

A Classification of Hyper-heuristics Approaches – Revisited

Edmund K. Burke, Matthew R. Hyde, Graham Kendall, Gabriela Ochoa, Ender Özcan and John R. Woodward

Abstract Hyper-heuristics comprise a set of approaches that aim to automate the development of computational search methodologies, initially to address operational research problems but more recently venturing into new domains such as bioinformatics, strategies for games, and software engineering. This chapter overviews previous categorisations of hyper-heuristics and provides a unified classification and definition. We distinguish between two main hyper-heuristic categories: heuristic *selection* and heuristic *generation*. Some representative examples of each category are discussed in detail, and recent research trends are highlighted. Our goals are to clarify the main features of existing techniques and to suggest new directions for hyper-heuristic research.

Edmund K. Burke
Queen Mary University of London, UK.
e-mail: ekb@qmul.ac.uk
Graham Kendall
University of Nottingham Malaysia Campus, Malaysia.
Automated Scheduling, Optimisation and Planning (ASAP) Group, School of Computer Science,
University of Nottingham, UK.
Matthew Hyde at Automated Scheduling, Optimisation and Planning (ASAP) Group, School of
Computer Science, University of Nottingham, UK.
Gabriela Ochoa
Computing Science and Mathematics, University of Stirling, UK.
Ender Özcan
Automated Scheduling, Optimisation and Planning (ASAP) Group, School of Computer Science,
University of Nottingham, UK.
John R. Woodward
Queen Mary University of London, UK.

1 Introduction

The current state of the art in hyper-heuristic research comprises a set of approaches that share the common goal of automating the design and adaptation of heuristic methods in order to solve computational search problems. The motivation behind these approaches is to raise the level of generality at which search methodologies can operate [13]. The term hyper-heuristic was first used in 1997 [31] to describe a protocol that combines several artificial intelligence methods in the context of automated theorem proving. The term was independently used in 2000 [29] to describe ‘heuristics to choose heuristics’ in the context of combinatorial optimisation. In this context a hyper-heuristic is a high-level approach that, given a particular problem instance and a number of low-level heuristics, can select and apply an appropriate low-level heuristic at each decision point [13, 75]. The idea of automating the heuristic design process, however, is not new. Indeed, it can be traced back to the early 1960s [30, 37, 38], and was independently developed by a number of authors during the 1990s [35, 44, 46, 62, 91, 96]. Some historical notes, and a brief overview of early approaches can be found in [13] and [75], respectively. A more recent research trend in hyper-heuristics attempts to automatically *generate* new heuristics suited to a given problem or class of problems. This is typically done by combining, through the use of genetic programming for example, *components* or *building-blocks* of human designed heuristics [14].

We differentiate between the terms heuristic, metaheuristic and hyper-heuristic. A heuristic is a “rule of thumb” offering guidance for finding good solutions according to domain knowledge. Two main types of search heuristics can be distinguished, *perturbative* or local search heuristics, which operate on fully instantiated candidate solutions, and *constructive* heuristics which iteratively expand partial candidate solutions. Metaheuristics are search methodologies that coordinate local search and higher-level strategies to perform a robust search on a solution space and escape local optima. Hyper-heuristics are high-level strategies that operate on a search space of heuristics rather than directly on a search space of solutions.

A variety of hyper-heuristic approaches have been proposed in the literature. An introduction to the area appeared in the 2003 edition of the Handbook of Metaheuristics [13]. The present chapter updates our previous version [15], which suggested a unified classification and definition of hyper-heuristics. The proposed classification inspired a thorough survey article appearing in 2013 [12], and has been widely adopted by the research community.

The next section outlines previous classifications of hyper-heuristics. Section 3 proposes our unified classification and definition. Sections 4 and 5, describe the main categories of the proposed classification. They discuss some representative examples, and highlight current research trends. Finally, section 6 summarises our categorisation and suggests future research directions.

2 Previous classifications

In [90], hyper-heuristics are categorised into two types: (i) with learning, and (ii) without learning. Hyper-heuristics without learning include approaches that use several heuristics (neighbourhood structures), but select the heuristics to call according to a predetermined sequence. Therefore, this category contains approaches such as variable neighbourhood search. The hyper-heuristics with learning include methods that dynamically change the preference of each heuristic based on their historical performance, guided by some learning mechanism. As discussed in [90], hyper-heuristics can be further classified with respect to the learning mechanism employed, and a distinction is made between approaches which use a genetic algorithm, from those which use other mechanisms. This is because many hyper-heuristics to date have been based on genetic algorithms. In these genetic algorithm-based hyper-heuristics the idea is to evolve the solution methods, not the solutions themselves.

In [75], hyper-heuristics are classified into those which are *constructive* and those which are *local search* methods. Constructive hyper-heuristics build a solution incrementally by adaptively selecting heuristics, from a pool of constructive heuristics, at different stages of the construction process. Local search hyper-heuristics, on the other hand, start from a complete initial solution and iteratively select, from a set of neighbourhood structures, appropriate heuristics to lead the search in a promising direction.

When genetic programming started being used for hyper-heuristic research in the late 2000's (see [14] for an overview), a new class of hyper-heuristics emerged. This class was explicitly and independently mentioned in [5] and [20]. In the first class of heuristics, or 'heuristics to choose heuristics', the framework is provided with a set of pre-existing, generally widely known heuristics for solving the target problem. In contrast, in the second class, the aim is to generate new heuristics from a set of *building-blocks*, *components* or search trail of known heuristics, which are given to the framework. Therefore, the process requires, as in the first class of hyper-heuristics, the selection of a suitable set of heuristics known to be useful in solving the target problem. However, instead of supplying these directly to the framework, the heuristics are first decomposed into their basic components. Genetic programming hyper-heuristic researchers [5, 14, 20] have also made the distinction between 'disposable' and 'reusable' heuristics. A disposable heuristic is created just for one problem, and is not intended for use on unseen problems. Alternatively, the heuristic may be created for the purpose of re-using it on new unseen problems of a certain class.

In [27], hyper-heuristics are classified into four categories: (i) hyper-heuristics based on the random choice of low-level heuristics, (ii) greedy and peckish hyper-heuristics, which require preliminary evaluation of all or a subset of the heuristics in order to select the best performing one, (iii) metaheuristics based hyper-heuristics, and (iv) hyper-heuristics employing learning mechanisms to manage low level heuristics.

3 The proposed classification and definition

Building upon some of the previous classifications discussed above, and realising that hyper-heuristics lie at the interface of optimisation and machine learning research, we propose a general classification of hyper-heuristics according to two considerations: (i) the nature of the heuristic search space, and (ii) the source of feedback during learning. These considerations are orthogonal in that different heuristic search spaces can be combined with different sources of feedback, and thus different machine learning techniques.

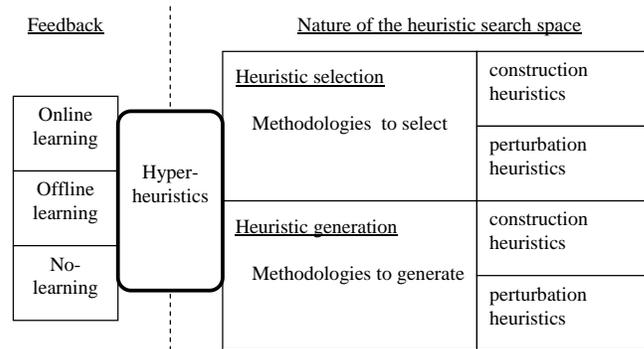


Fig. 1 A classification of hyper-heuristic approaches, according to two considerations: (i) the nature of the heuristic search space, and (ii) the source of feedback during learning.

The most fundamental hyper-heuristic categories from the previous classifications, are those represented by the processes of:

- *Heuristic selection*: Methodologies for choosing or selecting existing heuristics
- *Heuristic generation*: Methodologies for generating new heuristics (from primitive components or observed search trails of existing heuristics)

There is no reason why the higher level strategy (for selecting or generating heuristics) should be restricted to be a heuristic. Indeed, sophisticated knowledge-based techniques such as case-based reasoning have been employed in this way for university timetabling [25]. This leads us to propose the following more general definition of the term ‘hyper-heuristic’ which is intended to capture the idea of a method for automating the heuristic design and selection process:

A hyper-heuristic is an automated methodology for selecting or generating heuristics to solve computational search problems.

From this definition, there are two clear categories of hyper-heuristics: heuristic selection and heuristic generation. They form the first level in our first consideration (the nature of the search space). The second level in this consideration corresponds to the distinction between constructive and local search hyper-heuristics, also discussed in section 2. Notice that this categorisation is concerned with the nature of the low-level heuristics used in the hyper-heuristic framework. Our classification uses the terms *construction* and *perturbation* to refer to these classes of low-level heuristics. Sections 4 and 5 describe these categories in more detail, discussing some concrete examples of recent approaches that can be found in the literature.

There is an underlying theme to these two clear categories of hyper-heuristics. In both cases a high-level hyper-heuristic operates on a set of heuristics which in turn operate on the solution space. In the case of selection, we are choosing from a set of atomic predefined heuristics. In the case of generation, we are operating on a space of heuristic components.

We consider a hyper-heuristic to be a learning algorithm when it uses feedback information on the performance of the low-level heuristics from the search process. A non-learning hyper-heuristic selects a heuristic to apply uniformly at random from the existing pool, without keeping a record of their previous performance. According to the source of the feedback during learning about the low-level heuristics performance, we propose a distinction between *online* and *offline* learning. Notice that in the context of heuristic generation methodologies, an example of which is genetic programming-based hyper-heuristics (discussed in section 2), the notions of *disposable* and *reusable* have been used to refer to analogous ideas to those of online and offline learning, as described below:

Online learning hyper-heuristics: The learning takes place while the algorithm is solving an instance of a problem. Therefore, task-dependent local properties can be used by the high-level strategy to determine the appropriate low-level heuristic to apply.

Offline learning hyper-heuristics: The idea is to gather knowledge in the form of rules or programs, from a set of representative training instances, that we would expect to generalise to the process of solving unseen instances.

The proposed classification of hyper-heuristic approaches can be summarised as follows (also see figure 1):

- With respect to the nature of the heuristic search space
 - **Heuristic selection methodologies:** Produce combinations of pre-existing construction or perturbation heuristics.
 - **Heuristic generation methodologies:** Generate new heuristic methods using basic components/building-blocks or search trail of pre-existing construction or perturbation heuristics.
- With respect to the source of feedback during learning
 - **Online learning hyper-heuristics:** Learn while solving a given instance of a problem.

- **Offline learning hyper-heuristics:** Learn, from a set of training instances, a method that would generalise to unseen instances.
- **No-learning hyper-heuristics:** Do not use previous information from the search process on the low-level heuristic performance.

These categories reflect the most common research trends. However, there are methodologies that can cut across categories. For example, we can see hybrid methodologies that combine constructive with perturbation heuristics [42], or heuristic selection with heuristic generation [54, 61, 74, 84].

4 Heuristic selection methodologies

This section is not intended to be an exhaustive survey. The intention is to present a few examples to give the reader a flavour of the research that has been undertaken in this area. Some research trends are also highlighted.

4.1 Approaches based on construction low-level heuristics

These approaches build a solution incrementally. Starting with an empty solution, the goal is to intelligently select and use construction heuristics to gradually build a complete solution. The hyper-heuristic framework is provided with a set of pre-existing (generally problem specific) construction heuristics, and the challenge is to select the heuristic that is somehow the most suitable for the current problem state. This process continues until the final state (a complete solution) is obtained. Notice that there is a natural end to the construction process, that is, when a complete solution is reached. Therefore, the sequence of heuristic choices is finite and determined by the size of the underlying combinatorial problem. Furthermore, there is the interesting possibility of learning associations between partial solution stages and adequate heuristics for those stages.

Several approaches have been proposed to generate efficient hybridisations of existing construction heuristics in domains such as bin packing [60, 78], timetabling [23, 25, 71, 76, 77], production scheduling [97], and stock cutting [94, 95]. Both online and offline machine learning approaches have been investigated. Examples of online approaches are the use of metaheuristics in a search space of construction heuristics. For example, genetic algorithms [36, 46, 96, 97], tabu search [23] and other single-point based search strategies [73]. For this type of hyper-heuristic, the structure of the heuristic search space or hyper-heuristic landscape has been studied [66]. Examples of offline techniques are the use of learning classifier systems [60, 78], messy genetic algorithms [76, 77, 95] and case-based reasoning [25].

4.1.1 Representative examples

Two hyper-heuristics based on construction heuristics are described here in more detail. The first approach is online and is based on graph-colouring heuristics for timetabling problems, whilst the second is offline and is based on bin packing heuristics.

Graph-colouring hyper-heuristic for timetabling: In educational timetabling, a number of courses or exams need to be assigned to a number of timeslots, subject to a set of both hard and soft constraints. Timetabling problems can be modelled as graph colouring problems, where nodes in the graph represent events (e.g. exams), and edges represent conflicts between events. Graph heuristics in timetabling use the information in the graph to order the events by their characteristics (e.g. number of conflicts with other events or the degree of conflict), and assign them one by one into the timeslots. These characteristics suggest how difficult it is to schedule the events. Therefore, the most difficult event, according to the corresponding ordering strategy, will be assigned first. The graph-based hyper-heuristic developed in [23], implements the following five graph colouring-based heuristics, plus a random ordering heuristic:

- *Largest Degree (LD)*: Orders the events in decreasing order based on the number of conflicts the event has with the others events.
- *Largest Weighted Degree (LW)*: The same as *LD*, but the events are weighted by the number of students involved.
- *Colour Degree (CD)*: Orders the events in decreasing order in terms of the number of conflicts (events with common students involved) each event has with those already scheduled.
- *Largest Enrolment (RO)*: Orders the events in decreasing order based on the number of enrolments.
- *Saturation Degree (SD)*: Orders the events in increasing order based on the number of timeslots available for each event in the timetable at that time.

A candidate solution in the heuristic search space corresponds to a sequence (list) of these heuristics. The solution (timetable) construction is an iterative process where, at the i^{th} iteration, the i^{th} graph-colouring heuristic in the list is used to order the events not yet scheduled at that step, and the first e events in the ordered list are assigned to the first e least-cost timeslots in the timetable (see figure 2).

Tabu Search is employed as the high-level search strategy for producing good sequences of the low-level heuristics. Each heuristic list produced by tabu search is evaluated by sequentially using the individual heuristics to order the unscheduled events, and thus construct a complete timetable. Each heuristic in the list is used to schedule a number e of events. Therefore, the length of the heuristic list is n/e where n is the number of events to be scheduled. Values in the range of $e = 1, \dots, 5$ were tested (details can be found in [23]). This work also highlights the existence of two search spaces in constructive hyper-heuristics (the heuristic space and the problem solution). The approach was tested on both course and exam timetabling benchmark instances with competitive results. This graph-based hyper-heuristic was

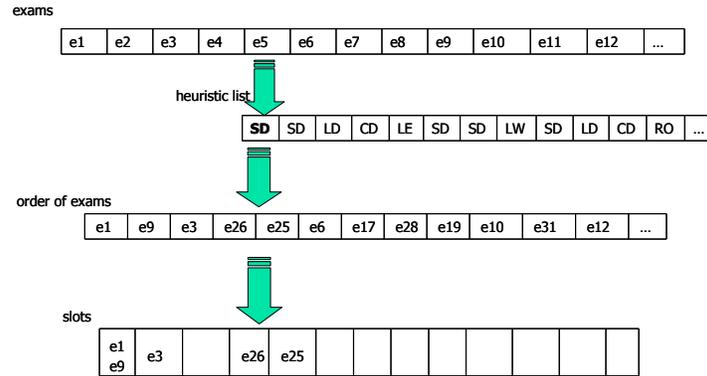


Fig. 2 A solution (timetable) is constructed by iteratively considering each heuristic in the list, and using it to order the events not yet scheduled. The first e events (in the figure $e = 5$) in the resulting ordering are assigned to the first e least-cost timeslots in the timetable.

later extended in [73] where a formal definition of the framework is presented. The authors also compare the performance of several high-level heuristics that operate on the search space of heuristics. Specifically, a steepest descent method, iterated local search and variable neighbourhood search are implemented and compared to the previously implemented tabu search. The results suggest that the choice of a particular neighbourhood structure on the heuristic space is not crucial to the performance. Moreover, iterative techniques such as iterated local search and variable neighbourhood search, were found to be more effective for traversing the heuristic search space than more elaborate metaheuristics such as tabu search. The authors suggest that the heuristic search space is likely to be smooth and to contain large plateaus (i.e. areas where different heuristic sequences can produce similar quality). The work also considers hybridisations of the hyper-heuristic framework with local search operating on the solution space. This strategy greatly improves the performance of the overall system, making it competitive with state-of-the-art approaches on the studied benchmark instances.

In a further study [66], the notion of fitness landscapes is used to analyse the search space of graph colouring heuristics. The study confirms some observations about the structure of the heuristic search space discussed in [73]. Specifically, these landscapes have a high level of neutrality (i.e. the presence of plateaus). Furthermore, although rugged, they have the encouraging feature of a globally convex or *big valley* structure, which indicates that an optimal solution would not be isolated but surrounded by many local minima. The study also revealed a *positional bias* in the search space comprising sequences of heuristics. Specifically, changes in the earlier positions of a heuristic sequence have a larger impact on the solution quality than changes in the later positions. This is because early decisions (heuristic

choices) in a construction process have a higher impact on the overall quality of the solution than later decisions.

Classifier system hyper-heuristic for bin packing: Classifier systems are rule-based learning systems that evolve fixed length stimulus-response rules. The rules are encoded as ternary strings, made of the symbols $\{0, 1, \#\}$, and have an associated strength. The system operates in two phases. First, the population of classification rules is applied to some task; and secondly, a genetic algorithm generates a new population of rules by selection based on strength, and by the application of standard genetic operators. Calculating the strength of each rule is a *credit assignment problem*, which refers to determining the contribution made by each sub-component or partial solution, in decomposable problems being solved collectively by a set of partial solutions.

In [78], a classifier system was used in the domain of one-dimensional bin packing, to learn a set of rules that associate characteristics of the current state of a problem with different low-level construction heuristics. In the one-dimensional bin packing problem, there is an unlimited supply of bins, each with capacity c . A set of n items is to be packed into the bins, the size of each item is given, and items must not overfill any bin. The task is to minimise the total number of bins required.

The set of rules evolved by the classifier system is used as follows: given the initial problem characteristics P , a heuristic H is chosen to pack an item, thus gradually altering the characteristics of the problem that remains to be solved. At each step a rule appropriate to the current problem state P' is selected, and the process continues until all items have been packed. For the training phase a total of 890 benchmark instances from the literature were used. The authors chose four bin packing heuristics from the literature, the selection being based on those that produced the best results on the studied benchmark set. These heuristics were as follows:

- *Largest-Fit-Decreasing*: Items are taken in order of size, largest first, and put in the first bin where they fit (a new bin is opened if necessary, and effectively all bins stay open).
- *Next-Fit-Decreasing*: An item is placed in the current bin if possible, or else a new bin is opened into which the piece is placed. This new bin becomes the current bin.
- *Djang and Finch's (DJD)*: A heuristic that considers combinations of up to three items to completely fill partially full bins.
- *A variation of DJD*: A variation of the previous heuristic that considers combinations of up to five items to completely fill partially full bins.

A simplified description of the current state of the problem is proposed. This description considers the number of items remaining to be packed, and calculates the percentage of items in each of four size ranges (huge, large, medium, and small); where the size of the items is judged in proportion to the bin size. The approach used single-step environments, meaning that rewards were available after each action had taken place. The classifier system was trained on a set of example problems and showed good generalisation to unseen problems. In [60], the classifier system approach is extended to multi-step environments. The authors test several reward

schemes in combination with alternate exploration/exploitation ratios, and several sizes and types of multi-step environments. Again, the approach was tested using a large set of one-dimensional benchmark bin packing problems. The classifier system was able to generalise well and create solution processes that performed well on a large set of NP-hard benchmark instances. The authors report that multi-step environments can obtain better results than single-step environments at the expense of a higher number of training cycles.

4.2 Approaches based on perturbation low-level heuristics

These approaches start with a complete solution, generated either randomly or using simple construction heuristics, and thereafter try to iteratively improve the current solution. The hyper-heuristic framework is provided with a set of neighbourhood structures and/or simple local searchers, and the goal is to iteratively select and apply them to the current complete solution. This process continues until a stopping condition has been met. Notice that these approaches differ from those based on construction heuristics, in that they do not have a natural termination condition. The sequence of heuristic choices can, in principle, be arbitrarily extended. This class of hyper-heuristics has the potential to be applied successfully to different combinatorial optimisation problems, since general neighbourhood structures or simple local searchers can be made available. Hyper-heuristics based on perturbation have been applied to personnel scheduling [22, 29], timetabling [6, 9, 21, 71, 82, 87, 88], shelf space allocation [7, 8], packing [6, 33] and vehicle routing problems [42, 72, 84].

So far, the approaches that combine perturbation low-level heuristics in a hyper-heuristic framework use online learning, in that they attempt to adaptively solve a single instance of the problem under consideration. Furthermore, the majority of the proposed approaches are single-point algorithms, in that they maintain a single incumbent solution in the solution space. Some approaches that maintain a population of points in the heuristic space have been attempted [28, 99].

As suggested in [68, 69] perturbation hyper-heuristics can be separated into two processes: (i) (low-level) heuristic selection, and (ii) move acceptance strategy. The authors classify hyper-heuristics with respect to the nature of the heuristic selection and move acceptance components. The heuristic selection can be done in a *non-adaptive* (simple) way: either randomly or along a cycle, based on a prefixed heuristic ordering [29]. No learning is involved in these approaches. Alternatively, the heuristic selection may incorporate an *adaptive* (or on-line learning) mechanism based on the probabilistic weighting of the low-level heuristics [22, 63, 72], or some type of performance statistics [29]. Both non-adaptive and adaptive heuristic selection schemes, are generally embedded within a single-point local search high-level heuristic.

The acceptance strategy is an important component of any local search heuristic. Many acceptance strategies have been explored within hyper-heuristics. Move acceptance strategies can be divided into two categories: *deterministic* and *non-*

deterministic. In general, a move is accepted or rejected, based on the quality of the move and the current solution during a single point search. At any point in the search, deterministic move acceptance methods generate the same result for the same candidate solution(s) used for the acceptance test. However, a different outcome is possible if a non-deterministic approach is used. If the move acceptance test involves other parameters, such as the current time, then these strategies are referred to as non-deterministic strategies. Well known meta-heuristic components are used as non-deterministic acceptance methods such and simulated annealing [33, 72].

4.2.1 Representative examples

Two hyper-heuristics based on perturbation heuristics are described here. The first is applied to a real-world packing problem, whilst the second uses large neighbourhood heuristics and is applied to five variants of the well known vehicle routing problem.

A simulated annealing hyper-heuristic for determining shipper sizes: In [33] the tabu search hyper-heuristic, originally presented in [22], is integrated within a simulated annealing framework. That is, a simulated annealing acceptance strategy is combined with the previously proposed heuristic selection mechanism. Figure 3 outlines the simulated annealing-based hyper-heuristic.

The tabu search hyper-heuristic [22], selects the low-level heuristics according to learned utilities or ranks. The framework also incorporates a dynamic tabu list of low-level heuristics that are temporarily excluded from the selection pool. The algorithm deterministically selects the low-level heuristic with the highest rank (not included in the tabu list), and applies it once regardless of whether the selected move causes an improvement or not (all moves acceptance). If there is an improvement, the rank is increased. If the new solution is worse, the rank of the low-level heuristic is decreased and it is made tabu. The rank update scheme is additive, and the tabu list is emptied each time a non-improvement move is accepted. This general tabu search approach was evaluated on various instances of two distinct timetabling and rostering (personal scheduling) problems, and the obtained results were competitive with those obtained using state-of-the-art problem-specific techniques. Apart from the simulated annealing acceptance criteria, some modifications are also introduced in [33]. In particular, a single application of a low-level heuristic h , is defined to be k iterations of h . Therefore, the decision points are set every k iterations, and the feedback for updating the quality of heuristic h is based on the best cost obtained during those k iterations. Additionally, a non monotonic cooling schedule is proposed to deal with the effects of having different neighbourhood sizes (given by the pool of low-level heuristics used). The methodology was applied to a packing problem in the cosmetics industry, where the shipper sizes for storage and transportation had to be determined. Real data was used for generating the instances, and the approach was compared with a simpler local search strategy (random descent), with favourable results.

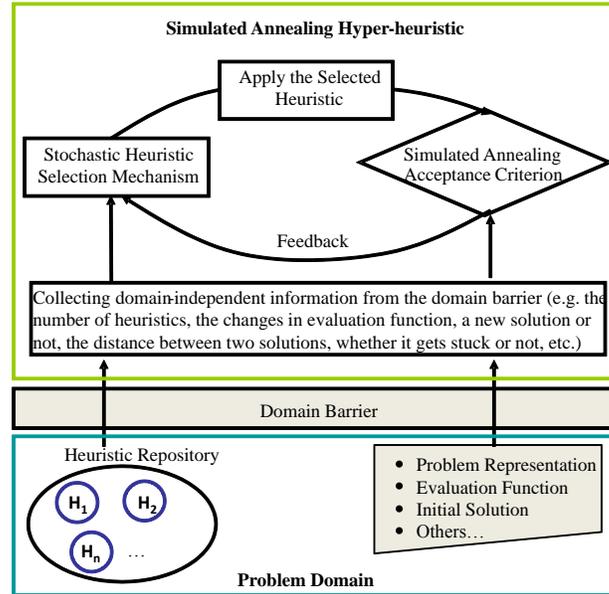


Fig. 3 A simulated annealing hyper-heuristic framework.

A general heuristic for vehicle routing problems: In [72], a unified methodology is presented, which is able to solve five variants of the vehicle routing problem: the vehicle routing problem with time windows, the capacitated vehicle routing problem, the multi-depot vehicle routing problem, the site-dependent vehicle routing problem and the open vehicle routing problem. All problem variants are transformed into a rich pickup and delivery model and solved using an adaptive large neighbourhood search methodology. The general framework is outlined in Fig.4, where the repeat loop corresponds to the local search framework at the master level. Implementing a simulated annealing algorithm is straightforward as one solution is sampled in each iteration of the algorithm. In each iteration of the main loop, the algorithm chooses one destroy ($N-$) and one repair neighbourhood ($N+$). An adaptive layer stochastically controls which neighbourhoods to choose according to their past performance (score, P_i). The more a neighbourhood N_i has contributed to the solution process, the larger score P_i it obtains, and hence it has a larger probability of being chosen. The adaptive layer uses roulette wheel selection for choosing a destroy and a repair neighbourhood.

The pickup and delivery model is concerned with serving a number of transportation requests using a limited number of vehicles. Each request involves moving a number of goods from a pickup location to a delivery location. The task is to construct routes that visit all locations such that the corresponding pickups and deliveries are placed on the same route and such that a pickup is performed before the corresponding delivery. Different constraints are added to model the different problem variants. The proposed framework adaptively chooses among a number of

insertion and removal heuristics to intensify and diversify the search. These competing sub-heuristics are selected with a frequency corresponding to their historic performance (stored as learned weights for each heuristic). The approach uses a simulated annealing acceptance strategy with a standard exponential cooling rate. A large number of tests were performed on standard benchmarks from the literature on the five variants of the vehicle routing problem. The results proved to be highly promising, as the methodology was able to improve on the best known solutions of over one third of the tested instances.

```

Construct a feasible solution  $x$ ; set  $x^* := x$ 
Repeat
  Choose a destroy and a repair neighbourhood:  $N^-$  and  $N^+$ 
    based on previously obtained scores ( $P_i$ )
  Generate a new solution  $x'$  from  $x$  using the heuristics
    corresponding to the chosen destroy and repair neighbourhoods
  If  $x'$  can be accepted then set  $x := x'$ 
  Update scores  $P_i$  of  $N^-$  and  $N^+$ 
  If  $f(x) < f(x^*)$  set  $x^* := x$ 
Until a stopping criteria is met
return  $x^*$ 

```

Fig. 4 Outline of the Adaptive Large Neighbourhood framework. N^- and N^+ correspond to destroy and repair neighbourhoods respectively, whilst P_i stands for the score associated to the heuristic i .

4.3 Recent research trends

The predominate approaches for heuristic selection have focused on a high-level search strategy resembling an existing metaheuristic such as simulated annealing, variable neighbourhood or genetic algorithms, which searches in a space of simple low-level heuristics. In recent years, new hyper-heuristic methodologies, such as the use of Monte Carlo Search [21, 83] have been explored. In addition, new domains such as software engineering [49, 99, 100], game playing [56], competitive travelling salesman problem [52], and DNA sequencing [10] have been studied. Attention has also been devoted to developing software frameworks, considering multi-objective problems and studying the theoretical foundations of these methods, as discussed below.

4.3.1 Software frameworks

The *HyFlex* (Hyper-heuristic Flexible) framework [65], features a common software interface for dealing with different combinatorial optimisation problems, and

provides the algorithm components that are problem specific. The algorithm designer does not require a detailed knowledge of the problem domains, and thus can concentrate his/her efforts on designing adaptive general-purpose optimisation algorithms. HyFlex provides six combinatorial problems implemented in Java, namely: boolean satisfiability, one dimensional bin packing, permutation flow shop, personnel scheduling, travelling salesman and vehicle routing. These domains supported an international research competition: the first Cross-Domain Heuristic Search Challenge (CHeSC) [64]. The framework has been widely used by the research community [12], and recently three new domains have been added: the 0-1 knapsack, quadratic assignment and max-cut problem [1].

In [92] a formulation of hyper-heuristics is introduced, which unifies the object-oriented programming type signature of hyper-heuristics. The formulation includes perturbative and constructive heuristics, as well as selection and generation heuristics. The goal is to extend the software design space of hyper-heuristic algorithms.

4.3.2 Multi-objective

There are emerging studies on multiobjective optimisation approaches mixing a set of existing multiobjective metaheuristics. For example, [98] describes a multialgorithm, genetically adaptive multiobjective method applied to benchmark function optimisation using NSGA-II, particle swarm optimisation, adaptive metropolis search, and differential evolution. This study demonstrates that combining standard metaheuristic algorithms can be better than using each one in isolation, as well as being competitive with other state-of-the-art methodologies on a set of benchmark functions. However, it has been once again observed that the choice of low level (meta)heuristics influences the overall performance of a hyper-heuristic in [57]. This latter study considers a variety of selection hyper-heuristics for multi-objective optimisation. The results showed that combining simple random choice and great deluge move acceptance from [59] is the best performing approach using NSGA-II, SPEA2 and IBEA as low level metaheuristics. This approach is additionally tested on a real world wind farm layout optimisation problem.

4.3.3 Theoretical and foundational studies

The structure of heuristic search spaces has been studied using the notion of fitness landscapes [66]. These landscapes are found to have a high level of neutrality (i.e. the presence of plateaus). Furthermore, although rugged, they have the encouraging feature of a globally convex or *big valley* structure, which indicates that there is a gradient guiding search towards good solutions. More recently, the search space of Boolean satisfiability (MAX-SAT) heuristics was analysed [26]. Using systematic enumeration of millions of heuristics, they gave evidence that the heuristic land-

scape had many clusters of good local optima, and was also amenable to other search methods, such as, (iterated) hill-climbing.

A series of formal run-time analysis of simplified hyper-heuristic algorithms has been conducted. In [48], the authors show the improved performance of a (1+1) evolutionary algorithms using a strategy combining multiple operators as compared to a strategy employing a single mutation operator with respect to a performance metric referred to *asymptotic hitting time*. In [55], a rigorous run-time analysis of hyper-heuristic mixing heuristics (operators) was performed, showing that there is some value to mixing heuristics leading to exponentially faster search than individual heuristics on some problems. In [2], theoretical analyses showed that alternatives to the additive updates in reinforcement learning, which is a commonly used scheme for heuristic selection, should be considered, since this strategy behaves in an asymptotically similar way to random choice based on the assumption that the probability of improving a solution at each step is less than 0.5. *Heuristic space diversity* is the focus of [45].

A fundamental open question in selection hyper-heuristics is which low-level heuristics (and how many) to use as part of the pool. In [3], the authors use tensor analysis as an advanced machine learning approach to decide on the subset of low level heuristics operating effectively with chosen deterministic move acceptance yielding improved results. A methodology for selecting a subset of effective heuristics from a given larger set is proposed in [86]. The approach considers non-parametric statistics and fitness landscape measurements from an available set of heuristics and benchmark instances, and it produces a compact subset of effective heuristics with improved performance for the underlying problem.

5 Heuristic generation methodologies

This section provides some examples of approaches that have the potential to automatically generate heuristics for a given problem. Many of the approaches in the literature to generate heuristics use genetic programming [14], a branch of evolutionary computation concerned with the automatic generation of computer programs. Genetic programming has been successfully applied to the automated generation of heuristics that solve hard combinatorial optimisation problems, such as boolean satisfiability, [5, 39, 40, 41, 53], bin packing [16, 18, 20, 58, 85], the traveling salesman problem [50, 51] and production scheduling [11, 32, 43, 47, 93]. In addition to the particular representation, using trees, graphs, grammars or linear program encodings, genetic programming differs from other evolutionary approaches in its application area. While most applications of evolutionary algorithms deal with optimisation problems, genetic programming could instead be positioned in the field of machine learning.

Some genetic programming-based hyper-heuristics have evolved local search heuristics [5, 17, 40, 41, 50, 51] or even evolutionary algorithms [67]. An alternative idea is to use genetic programming to evolve a program representing a function,

which is part of the processing of a given problem specific construction heuristic [16, 18, 32, 43, 93]. Most examples of using genetic programming as a hyper-heuristic are offline in that a training set is used for generating a program that acts as a heuristic, which is thereafter used on unseen instances of the same problem. That is, the idea is to generate *reusable* heuristics. However, research on *disposable* heuristics has also been conducted [5, 50, 51]. In other words, heuristics are evolved for solving a single instance of a problem. This approach is analogous to the online heuristic selection methodologies discussed in section 4, except that a new heuristic is generated for each instance, instead of choosing a sequence of heuristics from a predefined set.

The adaptation of heuristic orderings can also be considered as a methodology for heuristic generation. The adaptive approach proposed in [24], starts with one heuristic and adapts it to suit a particular problem instance ‘on the fly’. This method provides an alternative to existing forms of ‘backtracking’, which are often required to cope with the possible unsuitability of a heuristic. The adaptive method is more general, significantly easier to implement, and produces results that are at least comparable (if not better) than the current state-of-the-art examination timetabling algorithms.

5.1 Representative examples

Two representative examples of heuristic generation using genetic programming are discussed below. The first evolves packing heuristics that operate on a constructive framework. The second evolves complete local search algorithms, using components of successful, existing local search heuristics for boolean satisfiability.

Generation of construction heuristics for bin packing: As mentioned earlier, the one-dimensional bin packing problem involves a set of integer pieces L , which must be packed into bins of a certain capacity C , using the minimum number of bins possible. In the online version of the problem, the number of pieces and their sizes are not known in advance. This is in contrast to the offline version of the problem where the set of items to be packed is available at the start. An example of a construction heuristic used in online bin packing is first-fit, which packs a set of pieces one at a time, in the order that they are presented. The heuristic iterates through the open bins, and the current piece is placed in the first bin into which it fits.

In [16, 18], construction heuristics are evolved for the online bin packing problem. The evolved heuristics, represented as trees (see Fig. 5 for an example), operate within a fixed framework that resembles the operation of the first-fit heuristic discussed above. The key idea is to use the attributes of the pieces and bin capacities, that represent the state of the problem, in order to evolve functions (expressions) that would direct the process of packing. Each evolved function (GP tree) is applied in turn to the available bins, returning a value. If the value is zero or less then the system moves on to the next bin, but if the value is positive the piece is packed into

the bin. In this way, it is the expression which decides when to stop the search for a suitable bin and place the piece. The algorithm (depicted in Fig. 6) then repeats the process for each of the other pieces until all the pieces have been packed.

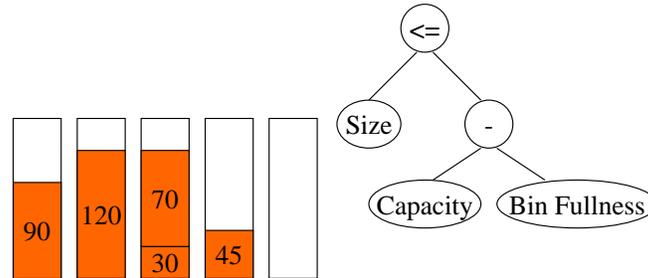


Fig. 5 Evolving one-dimensional packing heuristics with genetic programming.

In a genetic programming framework, the set of terminals and functions need to be specified. The hyper-heuristic framework for online bin packing uses some attributes that describe the state of the problem to define the terminals. In [16], the authors use the following terminals:

- S , the size of the current piece,
- C , the capacity of a bin (this is a constant for the problem) and,
- F , the fullness of a bin (i.e. the total size of all of the items occupying that bin).

Later [19], these three attributes were replaced by two: S , the size of the current item and, $E (= C - F)$, the emptiness of a bin (i.e. how much space is remaining in the bin, or how much more space can be allocated to it before it exceeds its capacity). The function set used in [16, 18] consisted of $\leq, +, -, \times, \%$, where $\%$ is the ‘protected divide function’. The results in [16] show that a simple genetic programming system can discover human designed heuristics such as first-fit, whilst in [19], heuristics that outperformed first-fit were evolved. In [18], it was also shown empirically that the choice of the training instances (categorised according to the piece size distribution), impacts on the trade-off between the performance and generality of the heuristics generated and their applicability to new problems.

Generation of local search heuristics for satisfiability testing: The boolean satisfiability problem consists of finding the true/false assignments of a set of boolean variables, to decide if a given propositional formula or expression (in conjunctive normal form) can be satisfied. The problem, denoted as SAT, is a classic NP-complete problem.

In [39, 40, 41] a genetic programming system, named CLASS (Composite Learned Algorithms for SAT Search), is proposed which automatically discovers new SAT local search heuristics. Figure 7 illustrates a generic SAT local search algorithm, where the ‘key detail’ is the choice of a *variable selection heuristic* in

```

For each piece p
  For each bin b
    output := evaluate Heuristic
    If (output > 0)
      place piece p in bin b
      break
    End If
  End For
End For

```

Fig. 6 Pseudo code showing the overall program structure within which an evolved packing heuristic operates.

the inner loop. Much research in the past decade has focused on designing a better variable selection heuristic, and as a result, local search heuristics have improved dramatically since the original method. The CLASS system was developed in order to automatically discover variable selection heuristics for SAT local search. It was noted in [39] that many of the best-known SAT heuristics (such as GSAT, HSAT, Walksat, and Novelty [41]) could be expressed as decision tree-like combinations of a set of primitives. Thus, it should be possible for a machine learning system to automatically discover new, efficient variable selection heuristics by exploring the space of combinations of these primitives. Examples of the primitives used in human designed SAT heuristics are the gain obtained by flipping a variable (i.e. the increase in the number of satisfied clauses in the formula) or the age of a variable (i.e. how long since it was last flipped).

The results using CLASS [41], show that a simple genetic programming system is able to generate local search heuristics that are competitive with efficient implementations of state-of-the-art heuristics (e.g. Walksat and Novelty variants), as well as previous evolutionary approaches. The evolved heuristics scale and generalise fairly well on random instances as well as more structured problem classes.

```

A:= randomly generated truth assignment
For j:= 1 to termination condition
  If A satisfies formula then return A
  v:= Choose variable using
    "variable selection heuristic"
  A:= A with value of v inverted
End If
End For
return FAILURE (no assignment found)

```

Fig. 7 A generic SAT local search algorithm. The “variable selection heuristic” is replaced by the evolved function.

5.2 *Some recent examples*

The most common approach so-far for generating heuristics has been tree-based genetic programming. Recently other genetic programming approaches have been used such as grammar-based [89], gene expression programming [80, 81], and grammatical evolution [17, 79].

Other machine learning techniques and representations have also been employed. In [4, 70], the authors applied a generic genetic algorithm which creates and searches heuristics in the form of policy matrices, communicating with an online bin packing simulator for evaluation. The empirical results show that the generated heuristics are specialised to the distributions and outperform the existing human designed ones.

A lifelong learning approach is proposed in [85], where an artificial immune system is combined with genetic programming in a system that continuously generates new heuristics and samples problems from its environment. The system is successfully tested on a large set of dynamically changing 1D bin-packing instances. A system that evolves an ensemble of heuristics for the job-shop scheduling problem is presented in [47]. The ensemble adopts a divide-and-conquer approach in which each heuristic solves a unique subset of the instance set considered. The system incorporates a heuristic generator that evolves heuristics composed of linear sequences of dispatching rules. Following a training period, the ensemble is shown to outperform both existing dispatching rules and a standard genetic programming algorithm on a large set of new test instances.

A new application domain includes generating strategies for games, such as the work in [34] where heuristics are evolved to guide staged deepening search for the hard game of FreeCell, obtaining solvers that outperform human players in this challenging puzzle.

6 Conclusions

The defining feature of hyper-heuristic research is that it investigates methodologies that operate on a search space of heuristics rather than directly on a search space of problem solutions. This feature provides the potential for increasing the level of generality of search methodologies. Several hyper-heuristic approaches have been proposed that incorporate different search and machine learning paradigms. We have suggested an updated definition of the term ‘hyper-heuristic’ to reflect recent work in the area.

With the incorporation of genetic programming and other machine learning methods, a new class of approaches can be identified; that is, heuristic generation methodologies. These approaches provide richer heuristic search spaces, and thus the freedom to create new methodologies for solving the underlying combinatorial problems. However, they are more difficult to implement than their coun-

terpart, heuristic selection methodologies, since they require the decomposition of existing heuristics, and the design of an appropriate framework.

We have further categorised the two main classes of hyper-heuristics (heuristic *selection* and heuristic *generation*), according to whether they use *construction* or *perturbation* low-level heuristics. These categories describe current research trends. However, the possibilities are open for the exploration of hybrid approaches. We also considered an additional orthogonal criterion for classifying hyper-heuristics with respect to the source of the feedback during the learning process, which can be either one instance (online approaches) or many instances of the underlying problem (offline approaches). Both online and offline approaches are potentially useful and therefore worth investigating. Although having a reusable method will increase the speed of solving new instances of problems, using online (or disposable) methods can have other advantages. In particular, searching over a space of heuristics may be more effective than directly searching the underlying problem space, as heuristics may provide an advantageous search space structure. Moreover, in newly encountered problems there may not be a set of related instances on which to train off-line hyper-heuristic methods.

Hyper-heuristic research lies at the the interface between search methodologies and machine learning methods. Machine learning is a well established artificial intelligence sub-field with a wealth of proven tools. The exploration of these techniques for automating the design of heuristics is only in its infancy. We foresee increasing interest in these methodologies in the coming years.

References

1. S. Adriaensen, G. Ochoa, and A. Nowé. A benchmark set extension and comparative study for the hyflex framework. In *IEEE Congress on Evolutionary Computation, CEC 2015*, pages 784–791, 2015.
2. F. Alanazi and P. K. Lehre. *Limits to Learning in Reinforcement Learning Hyper-heuristics*, pages 170–185. Springer International Publishing, Cham, 2016.
3. S. Asta and E. Özcan. A tensor-based selection hyper-heuristic for cross-domain heuristic search. *Information Sciences*, 299:412 – 432, 2015.
4. S. Asta., E. Özcan, and A. J. Parkes. Champ: Creating heuristics via many parameters for online bin packing. *Expert Systems with Applications*, 63:208 – 221, 2016.
5. M. Bader-El-Den and R. Poli. Generating SAT local-search heuristics using a GP hyper-heuristic framework. In *Evolution Artificielle, 8th International Conference*, volume 4926 of *Lecture Notes in Computer Science*, pages 37–49, Tours, France, 2007. Springer, Berlin.
6. R. Bai, J. Blazewicz, E. K. Burke, G. Kendall, and B. McCollum. A simulated annealing hyper-heuristic methodology for flexible decision support. *4OR - A Quarterly Journal of Operations Research*, 10(1):43–66, 2012.
7. R. Bai, E. K. Burke, and G. Kendall. Heuristic, meta-heuristic and hyper-heuristic approaches for fresh produce inventory control and shelf space allocation. *Journal of the Operational Research Society*, 59:1387 – 1397, 2008.
8. R. Bai, E.K. Burke, G. Kendall, and T. van Woensel. A new model and a hyper-heuristic approach for two-dimensional shelf space allocation. *4OR - A Quarterly Journal of Operations Research*, 11(1):31–55, 2013.

9. B. Bilgin, E. Özcan, and E. E. Korkmaz. An experimental study on hyper-heuristics and exam timetabling. In *Proceedings of the 6th Practice and Theory of Automated Timetabling (PATAT 2006)*, volume 3867 of *Lecture Notes in Computer Science*, pages 394–412, Brno, Czech Republic, 2007. Springer, Berlin.
10. J. Blazewicz, E. K. Burke, G. Kendall, W. Mruczkiewicz, C. Öguz, and A. Swiercz. A hyper-heuristic approach to sequencing by hybridization of dna sequences. *Annals of Operations Research*, 207(1):27–41, 2013.
11. J. Branke, S. Nguyen, C. Pickardt, and M. Zhang. Automated design of production scheduling heuristics: A review. *IEEE Transactions on Evolutionary Computation*, 20(1):110–124, 2016.
12. E. K. Burke, M. Gendreau, M. Hyde, G. Kendall, G. Ochoa, E. Ozcan, and R. Qu. Hyper-heuristics: a survey of the state of the art. *Journal of the Operational Research Society*, 64(12):1695–1724, 2013.
13. E. K. Burke, E. Hart, G. Kendall, J. Newall, P. Ross, and S. Schulenburg. Hyper-heuristics: An emerging direction in modern search technology. In F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, pages 457–474. Kluwer, 2003.
14. E. K. Burke, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, and J. Woodward. Exploring hyper-heuristic methodologies with genetic programming. In C. Mumford and L. Jain, editors, *Collaborative Computational Intelligence*, pages 177–201. Springer, Berlin, 2009.
15. E. K. Burke, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, and J. R. Woodward. *A Classification of Hyper-heuristic Approaches*, pages 449–468. Springer US, Boston, MA, 2010.
16. E. K. Burke, M. R. Hyde, and G. Kendall. Evolving bin packing heuristics with genetic programming. In *Proceedings of the 9th International Conference on Parallel Problem Solving from Nature (PPSN 2006)*, volume 4193 of *Lecture Notes in Computer Science*, pages 860–869, Reykjavik, Iceland, September 2006. Springer, Berlin.
17. E. K. Burke, M. R. Hyde, and G. Kendall. Grammatical evolution of local search heuristics. *IEEE Trans. Evolutionary Computation*, 16(3):406–417, 2012.
18. E. K. Burke, M. R. Hyde, G. Kendall, and J. Woodward. Automatic heuristic generation with genetic programming: Evolving a jack-of-all-trades or a master of one. In *Proceedings of the 9th ACM Genetic and Evolutionary Computation Conference (GECCO'07)*, pages 1559–1565, London, UK., July 2007. ACM, NY, USA.
19. E. K. Burke, M. R. Hyde, G. Kendall, and J. R. Woodward. The scalability of evolved on line bin packing heuristics. In *2007 IEEE Congress on Evolutionary Computation*, pages 2530–2537, Singapore, 2007. IEEE Computational Intelligence Society, IEEE Press.
20. E. K. Burke, M. R. Hyde, G. Kendall, and J. R. Woodward. A genetic programming hyper-heuristic approach for evolving 2-d strip packing heuristics. *IEEE Trans. Evolutionary Computation*, 14(6):942–958, 2010.
21. E. K. Burke, G. Kendall, M. Misir, and E. Özcan. Monte carlo hyper-heuristics for examination timetabling. *Annals of Operations Research*, 196(1):73–90, 2012.
22. E. K. Burke, G. Kendall, and E. Soubeiga. A tabu-search hyperheuristic for timetabling and rostering. *Journal of Heuristics*, 9(6):451–470, 2003.
23. E. K. Burke, B. McCollum, A. Meisels, S. Petrovic, and R. Qu. A graph-based hyper-heuristic for educational timetabling problems. *European Journal of Operational Research*, 176:177–192, 2007.
24. E. K. Burke and J. Newall. Solving examination timetabling problems through adaptation of heuristic orderings. *Annals of operations Research*, 129:107–134, 2004.
25. E. K. Burke, S. Petrovic, and R. Qu. Case based heuristic selection for timetabling problems. *Journal of Scheduling*, 9(2):115–132, 2006.
26. A. W. Burnett and A. J. Parkes. Exploring the landscape of the space of heuristics for local search in SAT. In *IEEE Congress on Evolutionary Computation (CEC 2017)*, pages 2518–2525. IEEE, 2017.
27. K. Chakhlevitch and P. I. Cowling. Hyperheuristics: Recent developments. In Carlos Cotta, Marc Sevaux, and Kenneth Sörensen, editors, *Adaptive and Multilevel Metaheuristics*, volume 136 of *Studies in Computational Intelligence*, pages 3–29. Springer, Berlin, 2008.

28. P. Cowling, G. Kendall, and L. Han. An investigation of a hyperheuristic genetic algorithm applied to a trainer scheduling problem. In *Proceedings of the Congress on Evolutionary Computation (CEC2002)*, pages 1185–1190, Hilton Hawaiian Village Hotel, Honolulu, Hawaii, USA, May 12-17 2002.
29. P. Cowling, G. Kendall, and E. Soubeiga. A hyperheuristic approach for scheduling a sales summit. In *Selected Papers of the Third International Conference on the Practice And Theory of Automated Timetabling, PATAT 2000*, Lecture Notes in Computer Science, pages 176–190, Konstanz, Germany, August 2001. Springer, Berlin.
30. W. B. Crowston, F. Glover, G. L. Thompson, and J. D. Trawick. Probabilistic and parametric learning combinations of local job shop scheduling rules. ONR research memorandum, Carnegie-Mellon University, Pittsburgh, 1963.
31. J. Denzinger, Matthias Fuchs, and Marc Fuchs. High performance ATP systems by combining several AI methods. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI 97)*, pages 102–107, CA, USA, 1997. Morgan Kaufmann.
32. C. Dimopoulos and A. M. S. Zalzala. Investigating the use of genetic programming for a classic one-machine scheduling problem. *Advances in Engineering Software*, 32(6):489–498, 2001.
33. K. A. Dowsland, E. Soubeiga, and E. K. Burke. A simulated annealing hyper-heuristic for determining shipper sizes. *European Journal of Operational Research*, 179(3):759–774, 2007.
34. A. Elyasaf, A. Hauptman, and M. Sipper. Evolutionary design of freecell solvers. *IEEE Trans. Comput. Intellig. and AI in Games*, 4(4):270–281, 2012.
35. H. L. Fang, P. Ross, and D. Corne. A promising genetic algorithm approach to job shop scheduling, rescheduling, and open-shop scheduling problems. In S. Forrest, editor, *Fifth International Conference on Genetic Algorithms*, pages 375–382, San Mateo, 1993. Morgan Kaufmann, CA, USA.
36. H. L. Fang, P. Ross, and D. Corne. A promising hybrid ga/ heuristic approach for open-shop scheduling problems. In A. Cohn, editor, *Eleventh European Conference on Artificial Intelligence*. John Wiley & Sons, NJ, USA, 1994.
37. H. Fisher and G. L. Thompson. Probabilistic learning combinations of local job-shop scheduling rules. In *Factory Scheduling Conference*, Carnegie Institute of Technology, May 10-12 1961.
38. H. Fisher and G. L. Thompson. Probabilistic learning combinations of local job-shop scheduling rules. In J. F. Muth and G. L. Thompson, editors, *Industrial Scheduling*, pages 225–251, New Jersey, 1963. Prentice-Hall, Inc.
39. A. S. Fukunaga. Automated discovery of composite SAT variable selection heuristics. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 641–648, Edmonton, Canada, 2002.
40. A. S. Fukunaga. Evolving local search heuristics for SAT using genetic programming. In *Genetic and Evolutionary Computation – GECCO-2004, Part II*, Lecture Notes in Computer Science, pages 483–494. Springer, Berlin, 2004.
41. A. S. Fukunaga. Automated discovery of local search heuristics for satisfiability testing. *Evol. Comput.*, 16(1):31–61, 2008.
42. P. Garrido and M.C. Riff. Dvrp: A hard dynamic combinatorial optimisation problem tackled by an evolutionary hyper-heuristic. *Journal of Heuristics*, 16:795–834, 2010.
43. C. D. Geiger, R. Uzsoy, and H. Aytüg. Rapid modeling and discovery of priority dispatching rules: An autonomous learning approach. *Journal of Scheduling*, 9:7–34, 2006.
44. J. Gratch and S. Chien. Adaptive problem-solving for large-scale scheduling problems: a case study. *Journal of Artificial Intelligence Research*, 4:365–396, 1996.
45. J. Grobler, A. P. Engelbrecht, G. Kendall, and V.S.S. Yadavalli. Heuristic space diversity control for improved meta-hyper-heuristic performance. *Information Sciences*, 300:49–62, 2015.
46. E. Hart, P. Ross, and J. A. D. Nelson. Solving a real-world problem using an evolving heuristically driven schedule builder. *Evolutionary Computing*, 6(1):61–80, 1998.

47. E. Hart and K. Sim. A hyper-heuristic ensemble method for static job-shop scheduling. *Evolutionary Computation*, 24(4):609–635, 2016.
48. J. He, F. He, and H. Dong. *Pure Strategy or Mixed Strategy?*, pages 218–229. Springer Berlin Heidelberg, Berlin, Heidelberg, 2012.
49. Y. Jia, M. B. Cohen, M. Harman, and J. Petke. Learning combinatorial interaction test generation strategies using hyperheuristic search. In *37th IEEE/ACM International Conference on Software Engineering, ICSE 2015*, pages 540–550, 2015.
50. R. E. Keller and R. Poli. Cost-benefit investigation of a genetic-programming hyperheuristic. In *Proceedings of Artificial Evolution (EA'07)*, pages 13–24, Tours, France, 2007.
51. R. E. Keller and R. Poli. Linear genetic programming of parsimonious metaheuristics. In *Proceedings of IEEE Congress on Evolutionary Computation (CEC 2007)*, Singapore, 2007.
52. G. Kendall and J. Li. Competitive travelling salesmen problem: A hyper-heuristic approach. *Journal of the Operational Research Society*, 64(2):208–216, 2013.
53. R. H. Kibria and Y. Li. Optimizing the initialization of dynamic decision heuristics in DPLL SAT solvers using genetic programming. In *Proceedings of the 9th European Conference on Genetic Programming*, volume 3905 of *Lecture Notes in Computer Science*, pages 331–340, Budapest, Hungary, 2006. Springer, Berlin.
54. N. Krasnogor and S. Maturana Gustafson. A study on the use of “self-generation” in memetic algorithms. *Natural Computing*, 3(1):53–76, 2004.
55. P. K. Lehre and E. Özcan. A runtime analysis of simple hyper-heuristics: To mix or not to mix operators. In *Proceedings of the Twelfth Workshop on Foundations of Genetic Algorithms XII, FOGA XII '13*, pages 97–104, New York, NY, USA, 2013. ACM.
56. J. Li and G. Kendall. A hyper-heuristic methodology to generate adaptive strategies for games. *IEEE Transactions on Computational Intelligence and AI in Games*, 9(1):1–10, 2017.
57. W. Li, E. Özcan, and R. John. Multi-objective evolutionary algorithms and hyper-heuristics for wind farm layout optimisation. *Renewable Energy*, 105:473 – 482, 2017.
58. Eunice Lopez-Camacho, Hugo Terashima-Marin, Peter Ross, and Gabriela Ochoa. A unified hyper-heuristic framework for solving bin packing problems. *Expert Systems with Applications*, 41(15):6876–6889, 2014.
59. M. Maashi, G. Kendall, and E. Özcan. Choice function based hyper-heuristics for multi-objective optimization. *Applied Soft Computing*, 28:312 – 326, 2015.
60. J. G. Marin-Blazquez and S. Schulenburg. A hyper-heuristic framework with XCS: Learning to create novel problem-solving algorithms constructed from simpler algorithmic ingredients. In *Learning Classifier Systems*, volume 4399 of *Lecture Notes in Computer Science*, pages 193–218, Berlin Heidelberg, 2007. Springer.
61. J. Maturana, F. Lardeux, and F. Saubion. Autonomous operator management for evolutionary algorithms. *Journal of Heuristics*, 16:881–909, 2010.
62. J. Mockus and L. Mockus. Bayesian approach to global optimization and applications to multi-objective constrained problems. *Journal of Optimization Theory and Applications*, 70(1):155–171, July 1991.
63. A. Nareyek. Choosing search heuristics by non-stationary reinforcement learning. In M. G. C. Resende and J. P. de Sousa, editors, *Metaheuristics: Computer Decision-Making*, chapter 9, pages 523–544. Kluwer, 2003.
64. G. Ochoa and M. Hyde. The Cross-domain Heuristic Search Challenge (CHeSC 2011). Website, 2011. <http://www.asap.cs.nott.ac.uk/chesc2011/>.
65. G. Ochoa, M. Hyde, T. Curtois, J. A. Vazquez-Rodriguez, J. Walker, M. Gendreau, G. Kendall, A. J. Parkes, S. Petrovic, and E. K. Burke. Hyflex: a benchmark framework for cross-domain heuristic search. In *Proceedings of the 12th European conference on Evolutionary Computation in Combinatorial Optimization, EvoCOP'12*, volume 7245 of *Lecture Notes in Computer Science*, pages 136–147. Springer-Verlag, 2012.
66. G. Ochoa, R. Qu, and E. K. Burke. Analyzing the landscape of a graph based hyper-heuristic for timetabling problems. In *Proceedings of the ACM Genetic and Evolutionary Computation Conference (GECCO 2009)*, pages 341–348, Montreal, Canada, 2009.
67. M. Oltean. Evolving evolutionary algorithms using linear genetic programming. *Evolutionary Computation*, 13(3):387–410, 2005.

68. E. Özcan, B. Bilgin, and E. E. Korkmaz. Hill climbers and mutational heuristics in hyper-heuristics. In *Proceedings of the 9th International Conference on Parallel Problem Solving from Nature (PPSN 2006)*, volume 4193 of *Lecture Notes in Computer Science*, pages 202–211, Reykjavik, Iceland, September 2006. Springer, Berlin.
69. E. Özcan, B. Bilgin, and E. E. Korkmaz. A comprehensive analysis of hyper-heuristics. *Intelligent Data Analysis*, 12(1):3–23, 2008.
70. E. Özcan and A. J. Parkes. Policy matrix evolution for generation of heuristics. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, GECCO '11, pages 2011–2018, New York, NY, USA, 2011. ACM.
71. N. Pillay. A review of hyper-heuristics for educational timetabling. *Annals of Operations Research*, 239(1):3–38, Apr 2016.
72. D. Pisinger and S. Ropke. A general heuristic for vehicle routing problems. *Computers and Operations Research*, 34:2403–2435, 2007.
73. R. Qu and E. K. Burke. Hybridisations within a graph based hyper-heuristic framework for university timetabling problems. *Journal of the Operational Research Society*, 60:1273–1285, 2009.
74. S. Remde, P. Cowling, K. Dahal, N. Colledge, and E. Selensky. An empirical study of hyper-heuristics for managing very large sets of low level heuristics. *Journal of the Operational Research Society*, 63(3):392–345, 2012.
75. P. Ross. Hyper-heuristics. In E. K. Burke and G. Kendall, editors, *Search Methodologies: Introductory Tutorials in Optimization and Decision Support Techniques*, chapter 17, pages 529–556. Springer, Berlin, 2005.
76. P. Ross and J. G. Marin-Blázquez. Constructive hyper-heuristics in class timetabling. In *IEEE Congress on Evolutionary Computation*, pages 1493–1500, Edinburgh, UK, 2005. IEEE, NJ, USA.
77. P. Ross, J. G. Marin-Blázquez, and E. Hart. Hyper-heuristics applied to class and exam timetabling problems. In *Proceedings of the 2004 IEEE Congress on Evolutionary Computation*, pages 1691–1698, Portland, Oregon, 2004. IEEE Press.
78. P. Ross, S. Schulenburg, J. G. Marin-Blázquez, and E. Hart. Hyper-heuristics: learning to combine simple heuristics in bin-packing problem. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO'02*, New York, USA, 2002. Morgan-Kaufman, CA, USA.
79. N. R. Sabar, M. Ayob, G. Kendall, and R. Qu. Grammatical evolution hyper-heuristic for combinatorial optimization problems. *IEEE Transactions on Evolutionary Computation*, 17(6):840–861, 2013.
80. N. R. Sabar, M. Ayob, G. Kendall, and R. Qu. Automatic design of hyper-heuristic framework with gene expression programming for combinatorial optimization problems. *IEEE Transactions on Evolutionary Computation*, 19(3):309–325, 2015. ISSN: 1089-778X.
81. N. R. Sabar, M. Ayob, G. Kendall, and R. Qu. A dynamic multiarmed bandit-gene expression programming hyper-heuristic for combinatorial optimization problems. *IEEE Transactions on Cybernetics*, 45(2):217–228, 2015. ISSN: 2168-2267.
82. N. R. Sabar, M. Ayob, R. Qu, and G. Kendall. A graph coloring constructive hyper-heuristic for examination timetabling problems. *Applied Intelligence*, 37(1):1–11, 2012.
83. N. R. Sabar and G. Kendall. Population based Monte Carlo tree search hyper-heuristic for combinatorial optimization problems. *Information Sciences*, 314:225–239, 2015.
84. K. Sim and E. Hart. A combined generative and selective hyper-heuristic for the vehicle routing problem. In *Genetic and Evolutionary Computation Conference*, pages 1093–1100, 2016.
85. K. Sim, E. Hart, and B. Paechter. A lifelong learning hyper-heuristic method for bin packing. *Evolutionary Computation*, 23(1):37–67, 2015.
86. J. A. Soria-Alcaraz, G. Ochoa, M. A. Sotelo-Figeroa, and E. K. Burke. A methodology for determining an effective subset of heuristics in selection hyper-heuristics. *European Journal of Operational Research*, 260(3):972 – 983, 2017.

87. J. A. Soria-Alcaraz, G. Ochoa, J. Swan, M. Carpio, H. Puga, and E. K. Burke. Effective learning hyper-heuristics for the course timetabling problem. *European Journal of Operational Research*, 238(1):77 – 86, 2014.
88. J. A. Soria-Alcaraza, E. Özcan, J. Swan, G. Kendall, and M. Carpio. Iterated local search using an add and delete hyper-heuristic for university course timetabling. *Applied Soft Computing*, 40:581–593, 2016.
89. Alejandro Sosa-Ascencio, Gabriela Ochoa, Hugo Terashima-Marin, and Santiago Enrique Conant-Pablos. Grammar-based generation of variable-selection heuristics for constraint satisfaction problems. *Genetic Programming and Evolvable Machines*, 17(2):119–144, 2016.
90. E. Soubeiga. *Development and Application of Hyperheuristics to Personnel Scheduling*. PhD thesis, School of Computer Science and Information Technology, University of Nottingham, June 2003.
91. R. H. Storer, S. D. Wu, and R. Vaccari. Problem and heuristic space search strategies for job shop scheduling. *ORSA Journal of Computing*, 7(4):453–467, 1995.
92. J. Swan, J. R. Woodward, E. Özcan, G. Kendall, and E. K. Burke. Searching the hyper-heuristic design space. *Cognitive Computation*, 6(1):66–73, 2014.
93. J. C. Tay and N. B. Ho. Evolving dispatching rules using genetic programming for solving multi-objective flexible job-shop problems. *Computers & Industrial Engineering*, 54:453–473, 2008.
94. H. Terashima-Marin, E. J. Flores-Alvarez, and P. Ross. Hyper-heuristics and classifier systems for solving 2D-regular cutting stock problems. In Hans-Georg Beyer and Una-May O’Reilly, editors, *Proceedings of the Genetic and Evolutionary Computation Conference GECCO 2005*, pages 637–643, Washington DC, USA, 2005. ACM, NY, USA.
95. H. Terashima-Marin, A. Moran-Saavedra, and P. Ross. Forming hyper-heuristics with GAs when solving 2D-regular cutting stock problems. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, volume 2, pages 1104–1110, Edinburgh, Scotland, UK, 2005. IEEE Press.
96. H. Terashima-Marin, P. Ross, and M. Valenzuela-Rendon. Evolution of constraint satisfaction strategies in examination timetabling. In *Genetic and Evolutionary Computation Conference, GECCO’99*, pages 635–642, 1999.
97. J. A. Vazquez-Rodriguez, S. Petrovic, and A. Salhi. A combined meta-heuristic with hyper-heuristic approach to the scheduling of the hybrid flow shop with sequence dependent setup times and uniform machines. In *Proceedings of the 3rd Multidisciplinary International Scheduling Conference: Theory and Applications (MISTA 2007)*, 2007.
98. J. A. Vrugt and B. A. Robinson. Improved evolutionary optimization from genetically adaptive multimethod search. *Proceedings of the National Academy of Sciences*, 104(3):708–711, 2007.
99. X. Wu, P. A. Consoli, L. L. Minku, G. Ochoa, and X. Yao. An evolutionary hyper-heuristic for the software project scheduling problem. In *Parallel Problem Solving from Nature - PPSN XIV - 14th International Conference, Edinburgh, UK, September 17-21, 2016, Proceedings*, pages 37–47, 2016.
100. K. Z. Zamil, F. Din, G. Kendall, and B. S. Ahmed. An experimental study of hyper-heuristic selection and acceptance mechanism for combinatorial t-way test suite generation. *Information Sciences*, 399:225–239, 2017.