

A Hyper-heuristic Approach to Automated Generation of Mutation Operators for Evolutionary Programming

Libin Hong^a, John H. Drake^b, John R. Woodward^b, Ender Özcan^a

^a*School of Computer Science, University of Nottingham, Jubilee Campus, Wollaton Road, Nottingham, NG8 1BB, UK*

^b*Operational Research Group, School of Electronic Engineering and Computer Science, Queen Mary University of London, Mile End Road, London, E1 4NS, UK*

Abstract

Evolutionary programming can solve black-box function optimisation problems by evolving a population of numerical vectors. The variation component in the evolutionary process is supplied by a mutation operator, which is typically a Gaussian, Cauchy, or Lévy probability distribution. In this paper, we use genetic programming to automatically generate mutation operators for an evolutionary programming system, testing the proposed approach over a set of function classes, which represent a source of functions. The empirical results over a set of benchmark function classes illustrate that genetic programming can evolve mutation operators which generalise well from the training set to the test set on each function class. The proposed method is able to outperform existing human designed mutation operators with statistical significance in most cases, with competitive results observed for the rest.

Keywords: Evolutionary Programming; Genetic Programming; Automatic Design; Hyper-heuristics; Continuous Optimization

1. Introduction

Black-box function optimisation is the task of finding the optima of an objective function for which we do not have access to an analytical form. This paper

is concerned with evolutionary programming (EP) [1], which evolves a population of real-valued input vectors for a function, a technique widely applied to real-world problems [2, 3, 4]. As EP has an evolutionary basis, each vector undergoes selection, evaluation, and mutation, with the expectation that fitter and fitter vectors are obtained. Here we focus on the mutation component of EP, which in the past has been designed manually. A hyper-heuristic is a search method or learning mechanism for selecting or generating heuristics to solve computational search problems [5]. Genetic programming (GP) [6] is a population-based evolutionary computation method for evolving program trees, that has frequently been used as a generation hyper-heuristic in the literature [7].

In this paper, we use GP as an offline generation hyper-heuristic to automatically create mutation operators for EP operating on function classes. A mutation operator in EP is a probability distribution, represented as a random number generator. We present an algorithmic framework which can not only express a number of currently existing EP mutation operators, but also generate novel variants of EP mutation operators. Using a train-and-test approach, GP is used to evolve mutation operators for EP, using a training set drawn from a class of functions which is then validated on a larger set of unseen instances taken from the same class. We use the term *automatically designed mutation operators* (ADMs) to describe the mutation operators generated by GP. We demonstrate that the ADMs for EP are capable of comparable, and often superior, performance to existing human designed operators. An additional set of experiments that takes ADMs that are trained on one function class, but then tested on a different function class is also conducted to further examine the performance of the evolved ADMs.

This current paper builds on previous papers. Hong et al. [8] first demonstrated that GP could automatically construct random number generators which are typi-

cally used in EP. In a second paper, it was shown that ADMs could be trained on collections of functions classes, showing good performance across a broader range of functions [9], however a tradeoff between general training and specific performance was observed. This paper presents a study of the design of 23 ADMs, for 23 functions classes, and then tests each of the 23 ADMs on each of the function classes. Real-valued optimisation is an active research topic and several population-based metaheuristics have been applied to function optimisation, including differential evolution and its variants [10, 11], particle swarm optimisation [12], covariance matrix adaptation evolution strategy (CMA-ES) [13] and hybrid methods [14]. We acknowledge the existence of these algorithms. However, as this study is concerned specifically with the modification of EP, a full review of all of these algorithms is beyond the scope of this paper.

The outline of the remainder of this paper is as follows. In Section 2, we give the background to the proposed approach of automatically designing algorithms using GP-based hyper-heuristics. In Section 3 we consider the task of function optimisation and introduce the notion of a function class. Section 4 presents our experimental results, which are analysed in Section 5. In Section 6 we discuss the research presented and in Section 7 we summarise the article and outline potential further research directions.

2. Automated Design Using Hyper-Heuristics

2.1. Hyper-Heuristics and automated generation of heuristics

The key distinction between metaheuristics and hyper-heuristics is that the former operate directly on the solution search space, while the latter operate indirectly on the solution search space, working with a set of low-level heuristics or heuristic components. Hyper-heuristics come in two main types: *heuristics to choose*

heuristics and *heuristics to generate heuristics* [5]. In addition to the broad distinction between selection and generation hyper-heuristics, it is also possible to classify hyper-heuristics according to the source of feedback during learning. *Online* hyper-heuristics learn while solving a given instance of a problem, whilst hyper-heuristics which learn in an *offline* manner gather knowledge in the form of rules or programs, from a set of training instances, that would hopefully generalise to the process of solving unseen instances. In this paper we are concerned with heuristics to generate heuristics using offline learning [7].

The automated generation of heuristics has received much attention in the last few years. Woodward et al. [15] automatically search the space of genetic algorithm (GA) selection operators, which contain *fitness proportional* and *rank* selection, where bitstrings in the population are chosen in proportion to their fitness value or indexed position in the sorted population respectively. In a later paper by the same authors, novel mutation operators were automatically constructed using random search and multiple-restart hill-climbing to search the space of mutation operators [16]. The system was capable of expressing two well-known mutation operators: *one-point* and *uniform*. While random search and hill-climbing may not be considered to be particularly ‘sophisticated’, they were sufficient to discover new selection and mutation operators which outperformed their human designed counterparts. Diosan and Oltean [17] evolved crossover operators for genetic algorithms outperforming existing crossover operators on some function optimisation problems.

2.2. Genetic Programming as a Generation Hyper-heuristic

As one of the main applications areas of metaheuristics is combinatorial optimisation problems, it is not surprising that this type of problem has attracted the attention of automated design. GP has been widely adopted to generate heuristics

for a variety of problems. Nguyen et al. [18] used GP to automatically design algorithms for job-shop scheduling, Keller et al. [19] tackled the travelling salesman problem, while Bader El Den et al. [20, 21] evolved timetabling heuristics. Drake et al. [22] used GP to evolve a scoring mechanism to determine the order in which knapsack items should be considered by a constructive heuristic for the multidimensional knapsack problem. Hong et al. [8] used GP to automate the design of probability distributions as mutation operators for evolutionary programming. GP was used to generate new data mining algorithms which were tested on well-known machine learning benchmark datasets by Freitas and Pappa [23]. Their paper showed that GP could outperform random search in searching the space of rules for data mining. The rules evolved by GP were observed to be at least as good as human designed rules in terms of classification. GP has also been used to automatically design schedule policies for dynamic multi-objective job shop scheduling [24], to evolve ensembles of dispatching rules for the job shop scheduling problem [25], and to automate the design of production scheduling heuristics [26]. Both online and offline bin packing have attracted attention within the heuristic generation research community [27]. Heuristic functions have been evolved to determine in which bin to place a given item [28]. In this case, the evolved heuristic functions have been shown to perform well on problem instances drawn from the same problem class used in the training phase, while a degradation in performance is witnessed when heuristics are applied to problem instances drawn from different problem classes. Heuristics have been evolved on problem instances containing a small number of items, then applied to much larger problem instances containing many more items [29]. In these approaches, a heuristic function is evolved as a GP syntax tree. However, more recently, both a look-up table (referred to as a “matrix”) [30, 31] and function interpolation [32] have also been used to represent a heuristic function. Other metaheuristics have been

automatically designed using hyper-heuristics, such as particle swarm optimisation [33] and variable neighbourhood search [34].

3. Optimisation and Function Classes

In this section we discuss optimisation with Evolutionary Programming (EP), and then introduce function classes as probability distributions over functions. Function classes are central to this paper, differentiating our approach from the standard convention of benchmarking on arbitrary functions. Rather than demonstrating the utility of an optimisation algorithm for specific arbitrary functions, we demonstrate the utility of an ADM on a set of functions are drawn from a fixed probability distribution (i.e. a function class).

3.1. Evolutionary Programming and Optimisation

We follow the formulation of optimisation as stated by Yao et al. [1, 35], which we repeat here. Global minimisation can be formalised as a pair (X, f) , where $X \in \mathbb{R}^n$ is a bounded set in \mathbb{R}^n and $f : X \rightarrow \mathbb{R}$ is an n -dimensional real-valued function. The objective is to find a point $x_{min} \in X$ such that $f(x_{min})$ is a global minimum in X . More specifically, it is required to find an x_{min} such that

$$\forall x \in X : f(x_{min}) \leq f(x).$$

Here, f does not need to be continuous or differentiable. While the aim of optimisation is to identify global optima of the function, in practice we often settle for methods, such as EP, which identify near-optima. EP is a widely used evolutionary algorithm for continuous optimisation introduced by Bäck and Schwefel [36] as follows:

1. Generate the initial population of μ individuals. Each individual is taken as a pair of real-valued vectors, (x_i, η_i) , $\forall i \in \{1, \dots, \mu\}$, where x_i 's are objective variables and η_i 's are standard deviations for Gaussian mutations.

2. Evaluate the fitness value for each (x_i, η_i) , $\forall i \in \{1, \dots, \mu\}$, of the population based on the objective function, $f(x_i)$.

3. Each parent (x_i, η_i) , $\forall i \in \{1, \dots, \mu\}$, creates λ / μ offspring on average, so that a total of λ offspring are generated. Offspring are generated by: for $i \in \{1, \dots, \mu\}$, $j \in \{1, \dots, n\}$ and $p \in \{1, \dots, \lambda\}$,

$$\eta'_i(j) = \eta_i(j) \exp(\gamma' N(0, 1) + \gamma N_j(0, 1)) \quad (1)$$

$$x'_p(j) = x_i(j) + \eta'_p(j) D_j \quad (2)$$

where $x_i(j)$, $x'_p(j)$, $\eta_i(j)$ and $\eta'_p(j)$ denote the j -th component of the vectors \mathbf{x}_i , \mathbf{x}'_p , η_i and η'_p respectively. $N(0, 1)$ denotes a normally distributed one-dimensional random number with mean 0 and standard deviation 1. $N_j(0, 1)$ indicates that the random number is generated anew for each value of j . The factors γ and γ' are set to $(\sqrt{2\sqrt{n}})^{-1}$ and $(\sqrt{2n})^{-1}$ [36].

4. Calculate the fitness of each offspring (x'_p, η'_p) , $\forall p \in \{1, \dots, \lambda\}$, according to $f(x'_p)$.

5. Conduct pairwise comparison over the union of parents (x_i, η_i) and offspring (x'_p, η'_p) , $\forall i \in \{1, \dots, \mu\}$, $\forall p \in \{1, \dots, \lambda\}$. For each individual, q opponents are selected randomly from the parents and offspring. For each comparison, if the individuals' fitness is no smaller than the opponent's it receives a 'win'.

6. Select the μ individuals out of the parents and offspring $((x_i, \eta_i)$ and (x'_p, η'_p) , $\forall i \in \{1, \dots, \mu\}$, $\forall p \in \{1, \dots, \lambda\}$) that have the most wins to be the parents of the

next generation.

7. *Stop if the halting criterion is satisfied; otherwise return to Step 3.*

Different variants of EP can be obtained by using different probability distributions D_j in Step 3 above. If a Gaussian distribution is used, then the algorithm is classical evolutionary programming (CEP) [1]. The Cauchy distribution is used in Fast EP (FEP) [1] whereas the Lévy distribution is used in Lévy evolutionary programming (LEP) [37]. The Lévy distribution is parameterised with a single parameter α , and corresponds to the Cauchy distribution when $\alpha=1.0$ and the Gaussian distribution when $\alpha=2.0$. Where LEP is used in this paper, the Lévy $L_{\alpha,\gamma}(y)$ distribution is implemented according to Mantegna [38] as given by Lee and Yao [37], with γ fixed at 1 (note that γ here is independent from step 3 above):

$$L_{\alpha,\gamma}(y) = \frac{1}{\pi} \int_0^{\infty} e^{-\gamma q^\alpha} \cos(qy) dq, y \in \mathbb{R}. \quad (3)$$

In this paper we will use genetic programming to evolve distributions to replace D_j as a mutation operator in the EP system described above. In each case, we use the same parameters for EP as previous publications [1, 35], using tournament selection with tournament size 10 on a population of 100 individuals. The initial value of the strategy parameter η is set to 3.0. The number of generations is different for each function and is specified in Table 1 below.

3.2. *Functions and Function Classes*

A generate-and-test approach can be used to sample the domain of a function in order to identify potential optima. Different black-box optimisation techniques have been proposed; the majority of which are the result of much manual effort.

In addition, each optimisation algorithm is designed in isolation from a problem environment. A proposed algorithm is usually tested on a number of benchmark functions to demonstrate its ability to identify the optima of a function. Unlike an individual function, a *function class* represents a parameterised function in which the parameters have a certain range of values. Thus a function class represents an infinite source of functions, from which it is possible to draw sets of sample functions from the same distribution.

In previous papers [35, 37, 39, 40], methods have been compared on functions from an arbitrary set. Each method is executed multiple times on a *single* function to provide a statistical comparison. Our method deviates from this approach in that we evolve a mutation operator for a function class, representing a probability distribution over a set functions. An example of a function class is $y = ax^2$ where a is uniformly distributed in the range [1, 2]. In this case, $y = 1.3x^2$ is a function from this function class, while $y = 0.3x^2$ is not from this function class. For each of the 23 standard functions often used in EP research [1, 35], we have constructed a corresponding function class. The 23 benchmark function classes are given in Table 1. These can be classified as follows: f_1 – f_7 are unimodal functions, f_8 – f_{13} are multimodal functions with many local optima, and f_{14} – f_{23} are multimodal function with few local optima [35].

Table 1: Function classes with number of dimensions n , domain S , and number of generations Gen . The parameters a , b and c are uniformly distributed in [1, 2], [-1, 1] and [-1, 1], respectively. For the values of w in f_{19} to f_{23} , please see [35]

Function Class	n	S	Generations
$f_1(x) = \sum_{i=1}^n [(ax_i - b)^2 + c]$	30	$[-100, 100]^n$	1500
$f_2(x) = a \sum_{i=1}^n x_i + b \prod_{i=1}^n x_i $	30	$[-10, 10]^n$	2000
$f_3(x) = \sum_{i=1}^n [\sum_{j=1}^i (ax_j + b)]^2$	30	$[-100, 100]^n$	5000

Function Class	n	S	
$f_4(x) = \max_i \{a x_i , 1 \leq i \leq n\}$	30	$[-100, 100]^n$	5000
$f_5(x) = \sum_{i=1}^n [a(x_{i+1} - x_i^2)^2 + b(x_i - 1)^2 + c]$	30	$[-30, 30]^n$	1500
$f_6(x) = \sum_{i=1}^n ([ax_i + 0.5])^2$	30	$[-100, 100]^n$	1500
$f_7(x) = a \sum_{i=1}^n ix_i^4 + \text{random}[0, 1]$	30	$[-1.28, 1.28]^n$	3000
$f_8(x) = \sum_{i=1}^n -(x_i \sin(\sqrt{ x_i }) + a)$	30	$[-500, 500]^n$	1500
$f_9(x) = \sum_{i=1}^n [ax_i^2 + b(1 - \cos(2\pi x_i))]$	30	$[-5.12, 5.12]^n$	5000
$f_{10}(x) = -a \exp(-0.2 \sqrt{\frac{1}{n} \sum_{i=1}^n x_i^2}) - b \exp(\frac{1}{n} \sum_{i=1}^n \cos 2\pi x_i) + e + c$	30	$[-32, 32]^n$	1500
$f_{11}(x) = \frac{a}{4000} \sum_{i=1}^n x_i^2 - b \prod_{i=1}^n \cos(\frac{x_i}{\sqrt{i}}) + c$	30	$[-600, 600]^n$	1500
$f_{12}(x) = \frac{a\pi}{n} \{10 \sin^2(\pi y_i) + \sum_{i=1}^{n-1} (y_i - 1)^2 [1 + 10 \sin^2(\pi y_{i+1}) + (y_n - 1)^2]\} + \sum_{i=1}^n u(x_i, 10, 100, 4),$ $y_i = 1 + \frac{1}{4}(x_i + 1)$ $u(x_i, w, k, m) = \begin{cases} k(x_i - w)^m, & x_i > w, \\ 0, & -w \leq x_i \leq w, \\ k(-x_i - w)^m, & x_i < -w. \end{cases}$	30	$[-50, 50]^n$	1500
$f_{13}(x) = 0.1a \{\sin^2(3\pi x_1) + \sum_{i=1}^{n-1} (x_i - 1)^2 [1 + \sin^2(3\pi x_{i+1})] + (x_n - 1)[1 + \sin^2(2\pi x_n)]\} + \sum_{i=1}^n u(x_i, 5, 100, 4)$	30	$[-50, 50]^n$	1500
$f_{14}(x) = [\frac{1}{500} + a \sum_{i=1}^{25} \frac{1}{j + \sum_{i=1}^2 (x_i - w_{ij})^6}]^{-1}$	2	$[-65.536, 65.536]^n$	100
$f_{15}(x) = \sum_{i=1}^{11} [w_i - \frac{ax_1(y_i^2 + y_i x_2)}{b(y_i^2 + y_i x_3 + x_4)}]^2$	4	$[-5, 5]^n$	4000
$f_{16}(x) = a(4x_1^2 - 2.1x_1^4 + \frac{1}{3}x_1^6 + x_1x_2 - 4x_2^2 + 4x_2^4) + b$	2	$[-5, 5]^n$	100
$f_{17}(x) = a(x_2 - \frac{5.1}{4\pi^2}x_1^2 + \frac{5}{\pi}x_1 - 6)^2 + 10b(1 - \frac{1}{8\pi}) \cos x_1 + 10$	2	$[-5, 10] \times [0, 15]$	100
$f_{18}(x) = a[1 + (x_1 + x_2 + 1)^2(19 - 14x_1 + 3x_1^2 - 14x_2 + 6x_1x_2 + 3x_2^2)] \times [30 + (2x_1 - 3x_2)^2(18 - 32x_1 + 12x_1^2 + 48x_2 - 36x_1x_2 + 27x_2^2)] + b$	2	$[-2, 2]^n$	100
$f_{19}(x) = -\sum_{i=1}^4 y_i \exp[-\sum_{j=1}^4 aw_{ij}(x_j - p_{ij})^2 + b]$	3	$[0, 1]^n$	100
$f_{20}(x) = -\sum_{i=1}^4 y_i \exp[-\sum_{j=1}^6 aw_{ij}(x_j - p_{ij})^2 + b]$	6	$[0, 1]^n$	200

Function Class	n	S	
$f_{21}(x) = -\sum_{i=1}^5 [(x - w_i)^T (x - w_i) + y_i]^{-1} + a$	4	$[0, 10]^n$	100
$f_{22}(x) = -\sum_{i=1}^7 [a(x - w_i)^T (x - w_i) + y_i + b]^{-1}$	4	$[0, 10]^n$	100
$f_{23}(x) = -\sum_{i=1}^{10} [a(x - w_i)^T (x - w_i) + y_i + b]^{-1}$ where $y_i =$ 0.1	4	$[0, 10]^n$	100

4. Experimental Design

In this section we describe the experimental set-up of GP and EP. With hyper-heuristic approaches, it is important to identify the two levels at which the heuristics operate. In a typical hyper-heuristic, a metaheuristic operates on a space of (meta) heuristics, which operate directly on the space of solutions. Here we use GP as a mutation operator generator at the hyper-level to manipulate the mutation operators within a population of EP algorithms working at the base level. The overall framework used is given in Figure 1.

Previous approaches manually build optimisation methods and test them on benchmark functions [1, 35]. Here we employ a train-and-test approach, in which a first set of functions is used to train a method before it is tested on a second independent set of functions drawn from the same probability distribution to analyse performance.

4.1. The Training Phase

We call a program generated by GP an ‘automatically designed mutation operator’ (ADM), which is in effect a random number generator. Each ADM is used as an EP mutation operator on 5 training functions drawn from a given function class. The fitness of an ADM is the average of the best values obtained in each of




Figure 1: Overview of the hyper-heuristic framework used

the individual 5 EP runs on a given function class. We use the same 5 functions from each function class for the entire run of GP on a given function class. For each function class, 10 functions are taken for training, 5 of which are used to calculate the fitness, and 5 of which are used to monitor overfitting.

4.2. The Testing Phase

When the training phase of EP is complete for a given function class, the output is an ADM intended *solely* for the function class on which it was trained. We then draw 50 new functions from the given function class and test the ADM, comparing against 6 existing EP variants (CEP, FEP and EP with 4 settings for the α parameter of the Lévy distribution). Note that training is expensive, as many ADMs are evaluated, so a small set of functions is used when training. In contrast, testing involves only a single ADM so is less costly, and a larger number of testing samples allows a better

comparison through statistical tests.

4.3. Parameter Settings for Genetic Programming

The GP implementation used in this paper is the genetic programming toolbox for Matlab (GPLAB) [41]. The parameters for GP are given in Table 2 and in the main are the default settings of GPLAB.

The maximum number of generations is set dynamically to try to ensure that a sufficient number of evaluations are made to achieve convergence, whilst minimising the time spent on evaluating poor quality runs. If the best individual in the population is found within the final three generations of a run, the maximum number of generations is increased by 5 in order to allow the evolutionary process to continue. As the evolutionary process is incredibly expensive in computational terms (each GP individual has to perform 5 EP runs), the maximum number of generations is initially set at a low value to reduce the amount of computational effort spent.

The maximum number of generations is capped at 100, with the evolutionary process terminated regardless of when the best individual in the population was found. We use subtree crossover, in which a node in each of the parent programs is chosen uniformly at random, and the respective subtrees are swapped, creating two new offspring. One point mutation is also used, where a node is chosen in the parent tree and substituted for a new tree created with the Grow initialisation method [6], obeying the size and depth restrictions imposed by the GP parameters. The fitness of each GP individual is calculated as the average best fitness values of the EP runs on each of the 5 training functions, as described in Section 4.1 above. At each generation, the best individual from both parents and offspring, along with the best offspring created during that generation are retained in the population. The selection method used is ‘lexictour’, which uses lexicographic parsimony pressure when

two individuals are compared [42], if two individuals are of equal fitness, the tree with a smaller number of nodes is chosen.

Table 2: Parameter settings for GP

Parameter	Value
Population Size	10
Initial Number of Generations	5
Upper Bound of Number of Generations	100
Crossover Proportion	45%
Mutation Proportion	45%
Reproduction Proportion	10%
Selection Method	lexictour [42]
Tournament Size	2
Maximum Initial Size of Tree	28
Maximum Size of Tree	512

Table 3: Function set for GP

Symbol	Function	Arity
+	addition	2
-	subtraction	2
×	multiplication	2
÷	protected division	2
<i>pow</i>	power	2
<i>log</i>	logarithm (base e)	1
<i>sin</i>	sin	1
<i>cos</i>	cos	1
<i>sqr</i>	square root	1

Table 4: Terminal set for GP

Symbol	Terminal
U	$\sim[0, 1]$
$N(\mu, \sigma^2)$	Normal Distribution
<i>chy</i>	Cauchy Distribution

The function and terminal sets for GP are given in Tables 3 and 4, respectively. U is the uniform distribution on $[0, 1]$. $N(\mu, \sigma^2)$ is the normal distribution with mean μ and variance σ^2 . Here we clarify that this μ is independent to that used in the EP descriptions in Section 3. In our experiments, μ lies within the range $[-2, 2]$ and σ^2 is in $[0, 10]$. *chy* is the Cauchy distribution. ‘÷’ is protected divide: if the numerator is divided by a zero denominator, then the numerator is returned. The square root function is also the protected variant, taking the square root of the absolute value of a single argument to ensure that no negative input is used. As the Lévy distribution can be constructed from the normal distribution and arithmetic operators [38], it is also within the search space that GP is operating in. In the terminal set there are no input variables, here we are using GP to construct mutation operators which are

effectively random number generators so do not require any input variable.

5. Analysis of the Performance of the Automatically Designed Mutation Operators

Table 5 reports the average best values obtained over 50 EP runs of each of the 23 function classes using a number of different mutation operators. The corresponding standard deviations are shown underneath each mean value in parentheses. These values are displayed for Cauchy (FEP, Lévy with $\alpha = 1.0$), Lévy with $\alpha = 1.2$, $\alpha = 1.4$, $\alpha = 1.6$ and $\alpha = 1.8$, and Gaussian (CEP, Lévy with $\alpha = 2.0$), as well as the best ADM evolved by GP for that function class. The best values (lowest, as we are minimising) are in boldface.

Table 5: ADMs compared to human designed mutation operators, means and (standard deviations), averaged over 50 runs. The best fitness values are in bold.

	<i>Cauchy</i>	$\alpha = 1.2$	$\alpha = 1.4$	$\alpha = 1.6$	$\alpha = 1.8$	<i>Gaussian</i>	ADM
f_1	6.012303 (15.5123)	6.011758 (15.5123)	6.011538 (15.5123)	6.011426 (15.5123)	6.011369 (15.5123)	6.011267 (15.5123)	6.011234 (15.5123)
f_2	0.140 (2.8E-02)	0.102 (2.0E-02)	0.084 (1.7E-02)	0.073 (1.5E-02)	0.065 (1.3E-02)	0.043 (8.6E-03)	0.017 (3.5E-03)
f_3	0.028 (2.0E-02)	0.018 (2.2E-02)	0.014 (1.8E-02)	0.015 (3.4E-02)	0.018 (5.3E-02)	0.018 (2.6E-02)	0.008 (1.4E-02)
f_4	1.88 (1.87)	5.30 (3.49)	9.10 (4.36)	10.58 (4.43)	13.78 (6.56)	17.31 (6.92)	0.13 (0.21)
f_5	-19.94 (26.13)	-19.87 (26.51)	-20.42 (26.94)	-20.21 (25.71)	-19.61 (26.44)	-20.22 (26.35)	-20.63 (27.15)
f_6	0.0336 (3.0E-02)	0.0134 (1.3E-02)	0.0076 (8.2E-03)	0.0724 (2.4E-01)	0.9058 (2.87)	322.7146 (820.2)	0.0074 (9.4E-03)

	<i>Cauchy</i>	$\alpha = 1.2$	$\alpha = 1.4$	$\alpha = 1.6$	$\alpha = 1.8$	<i>Gaussian</i>	ADM
f_7	0.0586 (9.3E-03)	0.0530 (8.6E-03)	0.0506 (7.9E-03)	0.0524 (6.9E-03)	0.0554 (9.9E-03)	0.0609 (1.1E-02)	0.0486 (6.3E-03)
f_8	-11058.28 (397.2)	-10642.83 (507.5)	-10009.65 (483.6)	-9530.88 (586.8)	-8818.26 (645.2)	-8053.96 (603.5)	-12469.12 (109.0)
f_9	-10.953 (15.82)	-10.955 (15.82)	-10.956 (15.82)	-10.956 (15.82)	-10.357 (16.21)	-8.269 (16.29)	-10.958 (15.82)
f_{10}	-27.8428 (6.86)	-27.8496 (6.87)	-27.8529 (6.87)	-27.8549 (6.87)	-27.8562 (6.87)	-27.7300 (6.93)	-27.8634 (6.87)
f_{11}	-0.4963 (6.4E-01)	-0.4839 (6.4E-01)	-0.4704 (6.5E-01)	-0.4534 (6.7E-01)	-0.4554 (6.4E-01)	-0.4465 (6.3E-01)	-0.5030 (6.4E-01)
f_{12}	1.72E-05 (8.8E-06)	2.13E-02 (5.9E-02)	1.82E-01 (4.7E-01)	1.77E-01 (3.7E-01)	1.01 (2.0)	2.37 (2.9)	3.57E-06 (2.5E-06)
f_{13}	2.20E-04 (5.5E-05)	5.20E-04 (2.7E-03)	2.87E-01 (1.5)	5.67E-01 (1.6)	2.23 (6.9)	7.04 (13.0)	6.46E-05 (2.3E-05)
f_{14}	1.32 (1.1)	0.98 (6.9E-01)	1.24 (8.1E-01)	1.08 (7.1E-01)	1.11 (6.0E-01)	0.90 (5.1E-01)	0.78 (2.9E-01)
f_{15}	5.68E-04 (3.6E-04)	5.13E-04 (3.2E-04)	5.90E-04 (3.8E-04)	6.03E-04 (3.5E-04)	6.34E-04 (3.9E-04)	5.39E-04 (3.3E-04)	4.66E-04 (3.1E-04)
f_{16}	-1.522775 (0.6049959)	-1.522775 (0.6049966)	-1.522776 (0.6049960)	-1.522776 (0.6049964)	-1.522776 (0.6049962)	-1.522777 (0.6049962)	-1.522779 (0.6049964)
f_{17}	5.8792588 (2.9140509)	5.8792572 (2.9140486)	5.8792571 (2.9140487)	5.8792674 (2.9140637)	5.8792569 (2.9140487)	5.8792566 (2.9140488)	5.8792563 (2.9140489)
f_{18}	4.615353 (0.9234)	4.615353 (0.9235)	4.615322 (0.9234)	4.615324 (0.9234)	4.615301 (0.9234)	4.615228 (0.9234)	4.615192 (0.9234)
f_{19}	-3.353969 (1.7371)	-3.354014 (1.7370)	-3.354021 (1.7370)	-3.354022 (1.7370)	-3.354025 (1.7370)	-3.354040 (1.7370)	-3.354058 (1.7371)
f_{20}	-3.67 (2.15)	-3.73 (2.30)	-3.83 (2.24)	-3.56 (2.13)	-3.76 (2.25)	-3.74 (2.18)	-3.86 (2.25)

	<i>Cauchy</i>	$\alpha = 1.2$	$\alpha = 1.4$	$\alpha = 1.6$	$\alpha = 1.8$	<i>Gaussian</i>	ADM
f_{21}	-4.33 (2.0)	-5.99 (2.7)	-5.59 (2.7)	-5.84 (2.8)	-6.44 (2.5)	-6.96 (2.4)	-7.26 (2.2)
f_{22}	-11910.70 (25106.9)	-23968.23 (70957.4)	-20517.80 (36414.1)	-20734.62 (35143.6)	-12662.55 (20568.8)	-68694.69 (193492.5)	-85515.41 (195260.3)
f_{23}	-20808.83 (86762.9)	-19096.31 (39304.9)	-14155.95 (23208.8)	-18304.26 (46037.7)	-20858.36 (36873.8)	-27047.33 (40003.3)	-111864.57 (178141.4)

For all 23 of the function classes, the GP designed ADM outperform the 6 human designed mutation operators. In Table 5 the evolved ADMs show the best performance on both unimodal and multimodal functions generated by all listed function classes in Table 1. We retain 2, 3 or 4 digits after the decimal point for most of the results. To distinguish the difference of testing results on f_1 , f_{16} and f_{18} , we retain 6 digits after the decimal point. We retain 7 digits after the decimal point for testing results on f_{17} .

To determine which of these performances differ with statistical significance, we perform a Wilcoxon signed-rank test, the results of which are presented in Table 6. Shown are the results of the Wilcoxon signed-rank test within a 95% confidence interval of an ADM compared with other mutation operators. In this table, ‘ \geq ’ indicates that the ADM performs better than another mutation operator on average. In the cases where this difference is statistically significant, ‘ $>$ ’ is used. In the majority of the cases, the ADMs outperform human designed mutation operators including Gaussian, Cauchy and Lévy, and this performance difference is statistically significant.

Table 6: Wilcoxon Signed-Rank Test of ADMs versus Gaussian, Cauchy and Lévy (with $\alpha = 1.2, 1.4, 1.6, 1.8$) on f_1 - f_{23} .

Function Class	Cauchy	$\alpha = 1.2$	$\alpha = 1.4$	$\alpha = 1.6$	$\alpha = 1.8$	Gaussian
f_1	>	>	>	>	>	>
f_2	>	>	>	>	>	>
f_3	>	>	>	\geq	\geq	>
f_4	>	>	>	>	>	>
f_5	\geq	\geq	\geq	\geq	\geq	\geq
f_6	>	>	\geq	>	\geq	>
f_7	>	>	>	\geq	>	>
f_8	>	>	>	>	>	>
f_9	>	>	>	>	>	>
f_{10}	>	>	>	>	>	>
f_{11}	>	>	>	>	>	>
f_{12}	>	>	>	>	>	>
f_{13}	>	>	>	>	>	>
f_{14}	\geq	>	>	>	>	>
f_{15}	>	>	>	>	>	>
f_{16}	>	>	>	>	>	>
f_{17}	>	>	>	>	>	>
f_{18}	>	>	>	>	>	>
f_{19}	>	>	>	>	>	>
f_{20}	>	>	>	\geq	>	>
f_{21}	>	>	>	>	>	\geq
f_{22}	>	>	>	>	>	>
f_{23}	>	>	>	>	>	>

There are only 2 (f_5, f_{14}), 1 (f_5), 2 (f_5, f_6), 4 (f_3, f_5, f_7, f_{20}), 3 (f_3, f_5, f_6) and 1 (f_5) function classes for which evolved ADMs perform slightly better than

Cauchy, $\text{Levy}_{\alpha=1.2}$, $\text{Levy}_{\alpha=1.4}$, $\text{Levy}_{\alpha=1.6}$, $\text{Levy}_{\alpha=1.8}$ and Gaussian respectively, with no statistical significance. This could be for a number of reasons. It is possible that EP has been run for so many iterations that it does not matter which mutation operator is used, and any difference in performance is negligible. It also may be the case that GP would be able to find better mutation operators if it were allowed to run for a longer period of time. One reason that GP can consistently find ADMs which perform at least as well as human designed mutation operators is that they are easily expressed within the function and terminal set used.

The best ADMs obtained with GP for each function class are listed in Table 7. Figure 2 provides histograms for a subset of ADMs, showing 3000 samples taken from the ADM to give an indication of the underlying probability distribution they represent.

Table 7: Evolved ADMs for all 23 function classes. These have been algebraically simplified where possible.

ADMs	Automatically designed mutation operators
ADM1	$(\sin(+(\text{chy chy})))$
ADM2	$(\sin(-(U U)))$
ADM3	$(\log(\text{chy}))$
ADM4	$(\div(-(\text{chy chy}) +(-(0 N(-1.8493, 2.288)) \cos(\log(N(-1.6403, 4.4607))))))$
ADM5	$(\times(\cos(+(\div(1 \times(\div(-(U U) N(0.94108, 7.9111)) N(-0.5776, 0.10706)))) N(-1.0504, 5.9002))) N(-0.89638, 7.9277))$
ADM6	$(N(-0.11984, 5.631))$
ADM7	$(N(-0.058664, 3.2512))$
ADM8	$(+(\times(\times(\text{chy} \times(\times(\text{chy chy} \text{chy})) \times(\text{chy} +(\text{chy} U))) +(\text{chy} U)))$
ADM9	$(\div(\sin(N(0.0078838, 0.17049)) -(\cos(\div(\text{chy chy}) \text{chy})))$
ADM10	$(\times(\times(U U) \times(U \text{chy})))$
ADM11	$(\div(\times(U \div(\text{chy} \log(U))) U))$

ADMs	Automatically designed mutation operator
ADM12	$(\div(chy - (N(-0.62528, 8.6422) \sin(N(-1.3941, 5.1622))))))$
ADM13	$(\div(U chy))$
ADM14	$(\div(-(chy \quad N(-0.77005, 1.7459)) \quad +$ $(chy \quad N(0.7052, 4.8637))))$
ADM15	$(\sin(N(-0.039909, 2.854)))$
ADM16	$(\times(\cos(\log(U)) \quad \sin(\sin(\log(pow(\cos(sqr(\log(U))) \quad \times$ $(\cos(\log(U))$ $\sin(\sin(\log(pow(\cos(\sin(\log(pow(\cos(sqr(\log(U))) \quad U)))) \quad U))))))))))$
ADM17	$(\times(\log(N(0.29948, 0.99092)) \quad U))$
ADM18	$(\cos(N(1.6565, 0.8667)))$
ADM19	$(\log(\cos(\div(\log(\cos(\div(\times(\sin(U) \quad U) \quad pow((-(sqr(\cos(chy)) \quad \cos(\sin(N(1.904, 2.002)$ $)))) \quad \sin(chy)))))) \quad pow((-(sqr(\cos(chy)) \quad \cos(\sin(N(1.3206, 2.6021)))))) \quad \sin(chy))))))$
ADM20	$(\div(U \quad + (N(1.1209, 9.3713) \quad N(-1.1291, 6.3921))))$
ADM21	$(\log(sqr(N(0.41597, 1.3872))))$
ADM22	$(\times(\log(chy) \quad U))$
ADM23	$(\div(+ (N(-0.041901, 0.11743) \quad \cos(N(1.5605, 0.044548))) \quad pow(sqr(U) \quad chy)))$

One might expect that for unimodal functions (e.g. $f_1 - f_8$) unimodal distributions make good mutation operators. The intuition being, as one moves around the domain of the function, there are corresponding changes in the objective value which can clearly guide the search. Similarly, one might expect, multimodal probability distributions to be more suitable to search multimodal functions than unimodal probability distributions. However, our results do not support this. For example, in Figure 2, ADM1 clearly shows peaks at +1 and -1, which is due to the final application of the trigonometric function. In contrast, ADM5 is unimodal. Conversely, some multimodal functions have resulted in symmetric unimodal probability distributions.

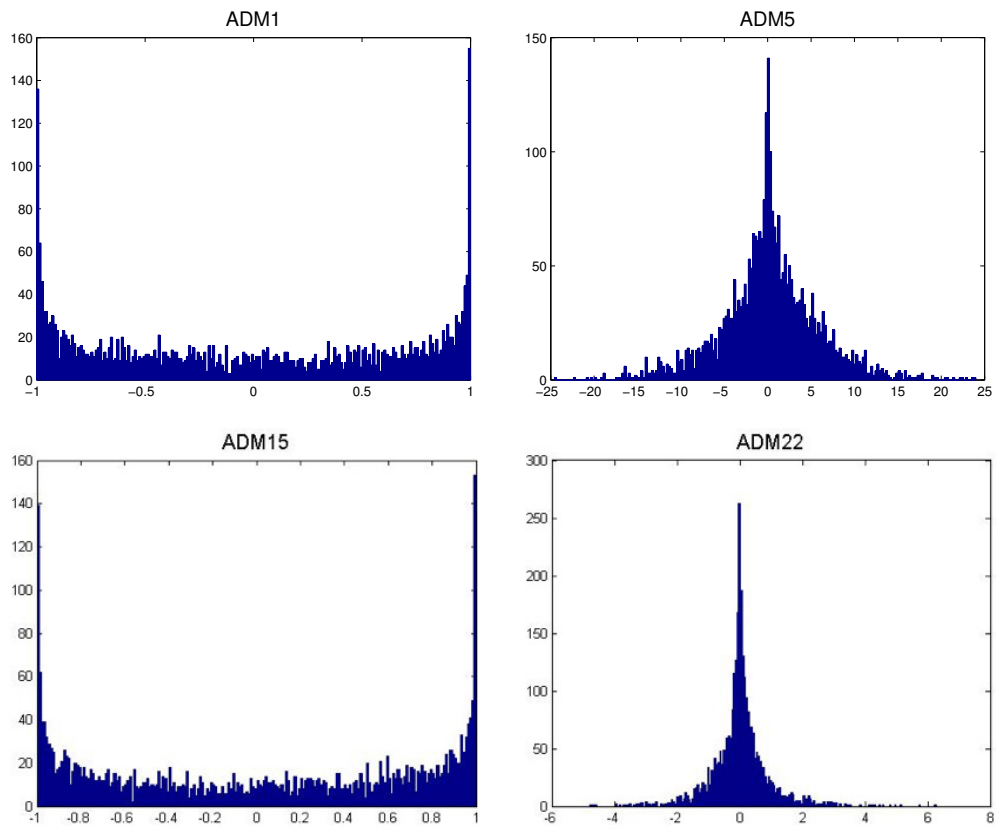


Figure 2: Histograms of the ADM1, ADM5, ADM15 and ADM22 for 3000 samples

For example, ADM15 is U-shaped and not the traditional bell-shape typically used with EP. Not all of the ADMs are symmetric, for example, ADM22, where the log function introduces asymmetry into the mutation operator.

In all but two cases (ADM6 and ADM7), the ADM is not a standard probability distribution (normal, Cauchy, Lévy) but something more complex. It is worth noting that these standard probability distributions are within the GP terminal set, however more ‘complex’ ADMs are the best found by GP for each function class. This supports the case for the automatic design of mutation operators. An alternative would be to automatically tune the numerical parameters of a Lévy distribution (the α parameter), or normal distribution (i.e. the mean and standard deviation), but this would only ever result in normal distributions which are a linearly scaled version of $N(0, 1)$. However, the automatic design process starts by defining a much broader search space than can be done with numerical parameters alone. This allows GP to find new probability distributions which perform better than the standard probability distributions as mutation operators for EP.

Figure 3 shows the performance of different mutation operators during an EP run, using a single function for eight of the function classes tested.

The ADMs show significant improvement on $f_2, f_3, f_4, f_{20}, f_{22}, f_{23}$. The figures of f_4, f_{19}, f_{20} and f_{23} shows that the ADMs not only perform well in early generations of the evolutionary process, but also in later generations. The figure for f_2 shows that ADM2 has good performance in the early generations, poor performance towards the middle, but outperforms all other methods by the end the run. This phenomenon was also found and discussed in previous references combining Cauchy and Gaussian mutation operators together. This suggests that, as the performance of a mutation operators varies at different stages of the search, using dynamic or multiple mutation operators may be preferable to using a single operator.

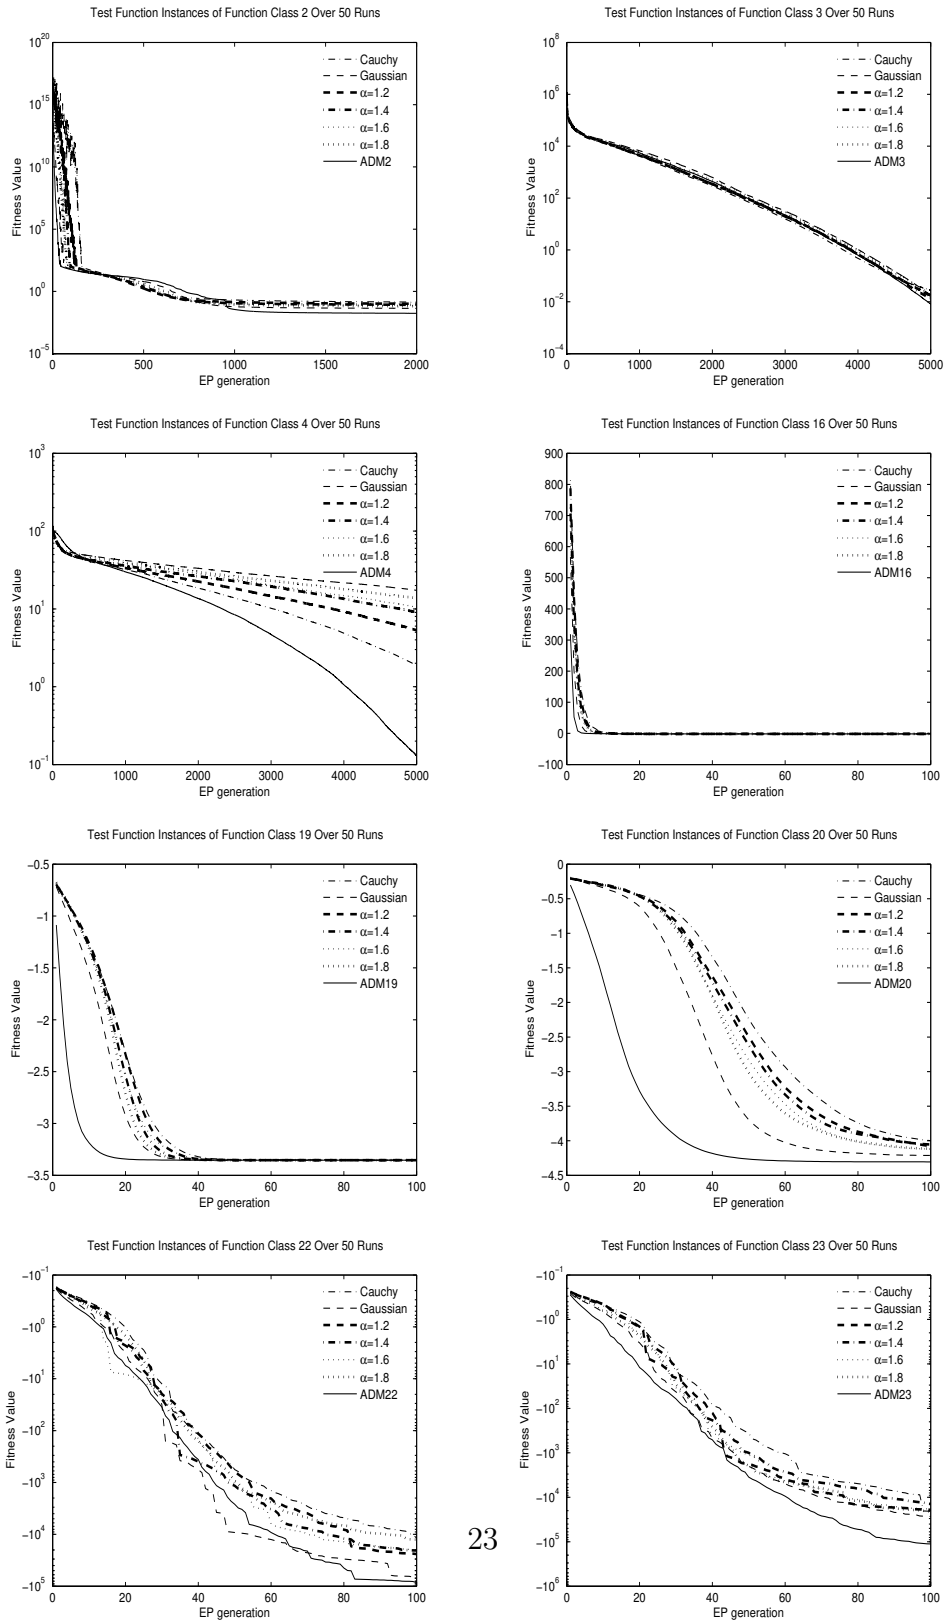


Figure 3: EP evolutionary process for ADM, Gaussian, Cauchy, and Lévy distribution with different value of α , the X axis represents EP generation, the Y axis represents fitness value of EP.

5.1. Performance of Evolved Mutation Operators on Other Function Classes

As a result of the train-and-test approach used, each ADM evolved by GP is designed for a specific function class. For example, ADM1 was designed for functions drawn from the function class f_1 . Specifically, it was trained on 10 functions (although only 5 were used for evaluation) drawn from f_1 and tested on a further 50 functions also taken from f_1 . This presents the following question: what happens when an ADM designed for one function class is used to optimise functions from a different function class? In other words, is the tailored ADM better than an arbitrary ADM?

We will now compare the performance of an ADM tailored to one function class with the other 22 ADMs which are intended for use on different function classes. The mean values and standard deviations achieved by each ADM are presented in Tables 8 and 9. The diagonal is in boldface and represents the performance of ADM i on the function class f_i . For a given row (f_i), the values in boldface indicate the ADMs that have beaten ADM i on f_i . For example, for f_1 , ADM1 is outperformed only by ADM10. From both tables we can see that ADM8, ADM11, ADM12 and ADM16 have the best performance on f_8 , f_{11} , f_{12} and f_{16} , respectively. ADM1, ADM4, ADM7, ADM9, ADM13, ADM19 and ADM20 have the second best performance on f_1 , f_4 , f_7 , f_9 , f_{13} , f_{19} and f_{20} , respectively. The worst performance of the tailored ADMs are ADM15, ADM18 and ADM22: their ranks are 11th, 11th, and 12th, on f_{15} , f_{18} and f_{22} , respectively. Overall, as expected an ADM tailored to a function class has better performance than the non-tailored ADMs. Table 10 shows the results of a Wilcoxon signed-rank test within a 95% confidence interval of a tailored ADM (TADM) compared with non-tailored ADMs (with statistically significant differences in boldface). The tailored ADM is an ADM trained for that specific function class. For example, ADM1 is the TADM for function class 1, but ADM1 is a non-tailored

ADM for function class 2. In this table, ' \geq ' and ' \leq ' indicate that an ADM performs better or worse than the other ADMs. In the case that this difference is statistically significant, ' $>$ ' and ' $<$ ' are used. Although ADM16 shows the best performance on f_{16} , in Table 10 ' $=$ ' indicates that the performances of ADM16 and ADM19 are equal on f_{16} . For f_8 , ' \geq ' indicates that the performance of ADM8 is better than that of ADM23, but that this is not significant. As can be seen from the last column in Table 10, in the majority of cases, the tailored ADM outperforms all of the other ADMs.

Table 9: Comparison (averaged over 20 runs) of each of the 23 ADMs on each of the 23 function classes (f_{20} - f_{23}).

	ADM1	ADM2	ADM3	ADM4	ADM5	ADM6	ADM7	ADM8	ADM9	ADM10	ADM11	ADM12	ADM13	ADM14	ADM15	ADM16	ADM17	ADM18	ADM19	ADM20	ADM21	ADM22	ADM23
f_{20}	-4.0774 (2.02)	-4.0748 (1.98)	-3.8878 (1.95)	-4.0932 (2.02)	-4.0110 (2.10)	-3.4181 (1.85)	-4.0560 (2.03)	-3.9522 (2.02)	-4.1095 (2.04)	-4.1141 (2.04)	-3.8083 (1.97)	-4.1051 (2.05)	-4.1100 (2.04)	-3.8682 (1.85)	-4.0393 (2.07)	-4.0958 (2.05)	-4.0869 (1.99)	-4.0615 (1.98)	-4.1102 (2.05)	-4.1107 (2.04)	-4.0990 (2.05)	-4.0948 (2.02)	-4.1099 (2.04)
f_{21}	-7.50 (2.4)	-8.26 (1.6)	-6.75 (2.9)	-3.69 (1.5)	-4.44 (1.9)	-4.58 (2.6)	-4.06 (1.7)	-3.64 (0.3)	-7.63 (2.3)	-7.12 (2.7)	-3.55 (0.7)	-7.50 (2.3)	-7.12 (2.7)	-5.96 (2.6)	-7.50 (2.1)	-3.99 (3.3)	-8.24 (1.5)	-6.87 (2.7)	-3.97 (2.7)	-8.14 (1.9)	-7.63 (2.4)	-7.24 (2.3)	-6.99 (2.5)
f_{22}	-22690.01 (30957.5)	-138764.16 (240207.8)	-16262.26 (22759.7)	-7564.46 (14484.9)	-4081.24 (7046.2)	-7710.71 (20971.9)	-8391.89 (15850.7)	-661.13 (1197.4)	-138490.10 (193997.9)	-156034.53 (256068.3)	-2163.99 (4065.4)	-121815.79 (272880.3)	-75620.18 (142754.5)	-21671.31 (34334.5)	-184449.89 (574931.9)	-20512221.26 (81568165.2)	71248.07 (109512.0)	-53103.24 (85281.3)	-3872071.64 (7903902.6)	-266869.50 (408822.5)	-88299.30 (146589.9)	-72258.79 (79909.2)	-156292.70 (227462.4)
f_{23}	-32152.70 (61337.3)	-52174.33 (93572.1)	-26607.57 (68340.3)	-21922.82 (54964.3)	-6153.03 (11938.5)	-5790.76 (11424.6)	-9617.22 (19983.7)	-380.46 (657.0)	-103530.92 (206684.0)	-132236.20 (244626.7)	-2388.34 (4202.0)	-50209.17 (100548.3)	-29264.20 (59321.4)	-14424.39 (31387.8)	-24997.68 (51026.8)	-3511793.67 (8105197.0)	-98707.40 (235422.7)	-211246.70 (642660.0)	-1024445.84 (2471291.7)	-290292.30 (839727.1)	-107802.76 (284455.3)	-36810.44 (74288.9)	-111941.27 (233075.7)

Table 10: Comparison of ADMs on different function classes.

FC	TADM -ADM1	TADM -ADM2	TADM -ADM3	TADM -ADM4	TADM -ADM5	TADM -ADM6	TADM -ADM7	TADM -ADM8	TADM -ADM9	TADM -ADM10	TADM -ADM11	TADM -ADM12	TADM -ADM13	TADM -ADM14	TADM -ADM15	TADM -ADM16	TADM -ADM17	TADM -ADM18	TADM -ADM19	TADM -ADM20	TADM -ADM21	TADM -ADM22	TADM -ADM23	TIMES
f_1	N/A	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	21
f_2	>	N/A	>	>	>	>	>	>	<	<	>	>	>	>	>	>	>	>	<	<	>	>	>	18
f_3	>	>	N/A	>	≤	≤	≤	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	19
f_4	>	>	>	N/A	>	>	>	>	>	≤	>	>	>	>	>	>	>	>	>	>	>	>	>	21
f_5	>	>	<	>	N/A	≤	≤	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	19
f_6	>	>	>	>	>	N/A	>	>	>	≤	>	<	<	<	>	>	>	>	>	≤	>	>	>	17
f_7	>	>	>	>	≤	>	N/A	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	21
f_8	>	>	>	>	>	>	>	N/A	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	22
f_9	>	>	>	>	>	>	>	>	N/A	>	>	>	>	>	>	>	>	>	>	>	>	>	>	21
f_{10}	>	>	>	>	>	>	>	>	<	N/A	>	>	>	>	>	>	>	>	>	<	>	>	>	20
f_{11}	>	>	>	>	>	>	>	>	>	>	N/A	>	>	>	>	>	>	>	>	>	>	>	>	22
f_{12}	>	>	>	>	>	>	>	>	>	>	>	N/A	>	>	>	>	>	>	>	>	>	>	>	22
f_{13}	>	>	>	>	>	>	>	>	>	>	>	>	<	N/A	>	>	>	>	>	>	>	>	>	21
f_{14}	>	>	>	>	>	>	>	>	>	≤	>	≤	≤	≤	N/A	>	>	>	>	≤	≤	≤	≤	19
f_{15}	≤	≤	≤	≤	≤	>	>	>	>	<	<	>	≤	≤	N/A	>	>	>	>	<	<	≤	<	12
f_{16}	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	N/A	>	>	=	>	>	>	>	22
f_{17}	>	>	>	>	>	>	>	>	>	<	<	>	>	>	>	<	N/A	>	>	<	>	≤	>	18
f_{18}	>	<	>	>	>	>	>	>	<	<	>	<	>	>	>	<	N/A	<	<	<	<	<	<	12
f_{19}	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	N/A	>	>	>	>	>	21
f_{20}	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	>	≤	N/A	>	>	>	21
f_{21}	>	<	>	>	>	>	>	>	>	>	>	>	>	>	>	>	≤	>	>	≤	N/A	≤	≤	19
f_{22}	>	≤	>	>	>	>	>	>	<	≤	≤	≤	≤	≤	≤	≤	≤	≤	<	≤	≤	N/A	≤	11
f_{23}	>	>	>	>	>	>	>	>	>	≤	≤	≤	≤	≤	≤	≤	≤	≤	≤	≤	≤	≤	N/A	17

6. Discussion

One of the advantages of the new method presented here is that it eliminates the need for human researchers to continually propose new distributions for use as mutation operators in EP. Instead, we have a search space which contains a rich set of mutation operators, and we can let a metaheuristic, such as GP, sample this space and select a suitable choice for the sample of functions at hand. In addition, it designs an ADM within the *context* of a function class. In other words, it tailors a mutation operator (random number generator) to a function class (probability distribution over functions). The suitability of mutation operator depends on the function class. Rather than tuning a *numerical parameter* to a function class, it tunes a *program* that generates random numbers to a function class.

One of the apparent disadvantages of the proposed system is the time needed to evolve the ADMs. This is because we have an EP algorithm at the base level, the mutation operator of which is being evolved by a GP algorithm at the hyper-level. While this may appear to be a superficial disadvantage, there are other advantages. Firstly, it is difficult to measure the amount of human effort required in designing a new mutation operator, and therefore it is difficult to directly compare the design phases of human and machine (GP in this case). We can only sensibly compare the performance of two mutation operators at the testing phase. Secondly, the system can be used to automatically generate new ADMs as and when needed to the demands of a new function class, whereas the human designer would have to start the whole process over again.

It is important to note that the fact that the training and testing are drawn from the same distribution is central to the train-and-test approach. In our case, this means that an ADM is developed to be used as a mutation operator *for a given*

function class, but also, importantly, *within a given EP algorithm*, which includes a fixed EP population size and number of generations.

One of the current limitations of the proposed method is that not only must the training set of functions be representative of the testing set of functions, but also the conditions under which they are sampled must be similar. For example, when testing a function from function class f_1 , we used the same population size and number of generations of EP as in the training phase. This is a limitation which needs to be addressed. One possibility would be to use GP to train EP using a *different* number of generations and population sizes. However, this would only partially address the issue, since if we attempted to use an ADM with a population size and number of generations outside the ranges seen in the training phase, there would be no guarantee of performance.

7. Summary and Future Work

In this paper we have used genetic programming (GP) as an offline hyper-heuristic to automatically evolve probability distributions, to use as mutation operators in evolutionary programming (EP). This is in contrast to existing operators in the literature which are human designed. The function and terminal set for GP was chosen to be able to express a number of currently existing human designed mutation operators, namely Cauchy, Gaussian and Lévy, and also express novel automatically designed mutation operators (ADMs). Each ADM is constructed from a function set including arithmetic and trigonometric functions, and a terminal set of probability distributions included as standard in many programming libraries and mathematical packages. Using a train-and-test approach, where two independent sets of functions are drawn from the same *function class* for training and testing, it is shown that GP is capable of generating ADMs which outperform existing EP variants over a

number of different function classes. As an additional validation exercise, we have also presented experiments to show that the ADM tailored to a given function class performs better than ADMs tailored to different function classes.

There are a number of possible directions for future work. As GP has been able to evolve good variation operators for EP, further work will explore the ability of GP to generate variation operators for other real-valued optimisation methods such as differential evolution (DE) and particle swarm optimisation (PSO), particularly for the case of function classes where a train-and-test approach can be used. Another direction within EP is to extend our hyper-heuristic approach beyond simple static mutation operators. It has been observed previously that both Cauchy and Gaussian mutation are effective in EP at different points of a search [35]. As both of these operators can be defined as a parameterised version of the Lévy distribution, through the use of different α values, we will evolve the value of α as a function over time, subsequently defining a family of adaptive mutation operators which can be trained to specialise in solving different classes of functions. Another possibility would be to use existing automated parameter tuning methods such as irace [43] to set this parameter, or indeed any of the other parameters within EP.

8. Acknowledgements

This work has been partially funded by the DAASE project, EPSRC grant EP/J017515/1, and the FAIME project, EPSRC grant EP/N002849/1.

References

- [1] X. Yao, Y. Liu, Fast evolutionary programming, in: Proceedings of the Fifth Annual Conference on Evolutionary Programming, MIT Press, 1996, pp. 451–460.

- [2] N. Sinha, R. Chakrabarti, P. Chattopadhyay, Evolutionary programming techniques for economic load dispatch, *IEEE Transactions on Evolutionary Computation* 7 (1) (2003) 83–94.
- [3] P. Shelokar, A. Quirin, O. Cordón, A multiobjective evolutionary programming framework for graph-based data mining, *Information Sciences* 237 (2013) 118–136.
- [4] M. Mutyalaraao, A. Sabarinath, M. Xavier James Raj, Taboo evolutionary programming approach to optimal transfer from earth to mars, in: *Proceedings of Swarm, Evolutionary, and Memetic Computing (SEMCCO 2011) - Part II*, Vol. 7077 of LNCS, Springer, Visakhapatnam, Andhra Pradesh, India, 2011, pp. 122–131.
- [5] E. Burke, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, J. Woodward, A classification of hyper-heuristic approaches, in: M. Gendreau, J.-Y. Potvin (Eds.), *Handbook of Metaheuristics*, Vol. 146 of International Series in Operations Research and Management Science, Springer US, 2010, pp. 449–468.
- [6] R. Poli, W. B. Langdon, N. F. McPhee, A field guide to genetic programming, Published via <http://lulu.com> and freely available at <http://www.gp-field-guide.org.uk>, 2008, (With contributions by J. R. Koza).
- [7] E. Burke, M. Hyde, G. Kendall, G. Ochoa, E. Özcan, J. Woodward, Exploring hyper-heuristic methodologies with genetic programming, in: C. Mumford, L. Jain (Eds.), *Computational Intelligence*, Vol. 1 of Intelligent Systems Reference Library, Springer Berlin Heidelberg, 2009, pp. 177–201.
- [8] L. Hong, J. Woodward, J. Li, E. Özcan, Automated design of probability dis-

tributions as mutation operators for evolutionary programming using genetic programming, in: Genetic Programming, Vol. 7831 of LNCS, Springer Berlin Heidelberg, 2013, pp. 85–96.

- [9] L. Hong, J. H. Drake, J. R. Woodward, E. Özcan, Automatically designing more general mutation operators of evolutionary programming for groups of function classes using a hyper-heuristic, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2016), ACM, New York, NY, USA, 2016, pp. 725–732.
- [10] A. K. Qin, V. L. Huang, P. N. Suganthan, Differential evolution algorithm with strategy adaptation for global numerical optimization, *IEEE transactions on Evolutionary Computation* 13 (2) (2009) 398–417.
- [11] R. Mallipeddi, P. N. Suganthan, Q.-K. Pan, M. F. Tasgetiren, Differential evolution algorithm with ensemble of parameters and mutation strategies, *Applied Soft Computing* 11 (2) (2011) 1679–1696.
- [12] J. Kennedy, R. Eberhart, Particle swarm optimization, in: Proceedings of the IEEE International Conference on Neural Networks, Vol. 4, Perth, Australia, 1995, pp. 1942–1948.
- [13] N. Hansen, The cma evolution strategy: A comparing review, in: Towards a New Evolutionary Computation, Vol. 192 of Studies in Fuzziness and Soft Computing, Springer Berlin Heidelberg, 2006, pp. 75–102.
- [14] S. M. Elsayed, R. A. Sarker, D. L. Essam, Adaptive configuration of evolutionary algorithms for constrained optimization, *Applied Mathematics and Computation* 222 (2013) 680–711.

- [15] J. R. Woodward, J. Swan, Automatically designing selection heuristics, in: Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO 2011), ACM, Dublin, Ireland, 2011, pp. 583–590.
- [16] J. R. Woodward, J. Swan, The automatic generation of mutation operators for genetic algorithms, in: Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO 2012), ACM, Philadelphia, Pennsylvania, USA, 2012, pp. 67–74.
- [17] L. Dioşan, M. Oltean, Evolving crossover operators for function optimization, in: P. Collet, M. Tomassini, M. Ebner, S. Gustafson, A. Ekárt (Eds.), Proceedings of the European Conference on Genetic Programming (EuroGP 2006), Vol. 3905 of LNCS, Springer, Budapest, Hungary, 2006, pp. 97–108.
- [18] S. Nguyen, M. Zhang, M. Johnston, K. Tan, Automatic design of scheduling policies for dynamic multi-objective job shop scheduling via cooperative coevolution genetic programming, *IEEE Transactions on Evolutionary Computation* 18 (2) (2013) 193–208.
- [19] R. E. Keller, R. Poli, Linear genetic programming of parsimonious metaheuristics, in: D. Srinivasan, L. Wang (Eds.), Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2007), IEEE Press, Singapore, 2007, pp. 4508–4515.
- [20] M. Bader El Den, R. Poli, Grammar-based genetic programming for timetabling, in: Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2009), IEEE Press, Piscataway, NJ, USA, 2009, pp. 2532–2539.
- [21] M. Bader El Den, R. Poli, S. Fatima, Evolving timetabling heuristics using a

- grammar-based genetic programming hyper-heuristic framework, *Memetic Computing* 1 (3) (2009) 205–219.
- [22] J. H. Drake, M. Hyde, K. Ibrahim, E. Özcan, A genetic programming hyper-heuristic for the multidimensional knapsack problem, *Kybernetes* 43 (9-10) (2014) 1500–1511.
- [23] G. L. Pappa, A. A. Freitas, Discovering new rule induction algorithms with grammar-based genetic programming, in: *Soft Computing for Knowledge Discovery and Data Mining*, Springer, 2008, pp. 133–152.
- [24] S. Nguyen, M. Zhang, M. Johnston, K. C. Tan, Automatic design of scheduling policies for dynamic multi-objective job shop scheduling via cooperative coevolution genetic programming, *IEEE Transactions on Evolutionary Computation* 18 (2) (2014) 193–208.
- [25] J. Park, S. Nguyen, M. Zhang, M. Johnston, Evolving ensembles of dispatching rules using genetic programming for job shop scheduling, in: *Genetic Programming*, Vol. 9025 of *Lecture Notes in Computer Science*, Springer International Publishing, 2015, pp. 92–104.
- [26] J. Branke, S. Nguyen, C. W. Pickardt, M. Zhang, Automated design of production scheduling heuristics: A review, *IEEE Transactions on Evolutionary Computation* 20 (1) (2016) 110–124.
- [27] R. Poli, J. Woodward, E. K. Burke, A histogram-matching approach to the evolution of bin-packing strategies., in: *Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2007)*, IEEE, 2007, pp. 3500–3507.

- [28] E. K. Burke, M. R. Hyde, G. Kendall, J. Woodward, Automatic heuristic generation with genetic programming: evolving a jack-of-all-trades or a master of one, in: Proceedings of the Genetic and Evolutionary Computation Conference Companion (GECCO 2007), ACM, London, England, 2007, pp. 1559–1565.
- [29] E. K. Burke, M. R. Hyde, G. Kendall, J. R. Woodward, The scalability of evolved on line bin packing heuristics, in: D. Srinivasan, L. Wang (Eds.), Proceedings of the IEEE Congress on Evolutionary Computation (CEC 2007), IEEE Press, Singapore, 2007, pp. 2530–2537.
- [30] A. Parkes, E. Özcan, M. Hyde, Matrix analysis of genetic programming mutation, in: A. Moraglio, S. Silva, K. Krawiec, P. Machado, C. Cotta (Eds.), Genetic Programming, Vol. 7244 of LNCS, Springer Berlin Heidelberg, 2012, pp. 158–169.
- [31] E. Özcan, A. J. Parkes, Policy matrix evolution for generation of heuristics, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2011), ACM, Dublin, Ireland, 2011, pp. 2011–2018.
- [32] J. H. Drake, J. Swan, G. Neumann, E. Özcan, Sparse, continuous policy representations for uniform online bin packing via regression of interpolants, in: European Conference on Evolutionary Computation in Combinatorial Optimization (EvoCOP 2017), Vol. 10197 of LNCS, Springer, 2017, pp. 189–200.
- [33] R. Poli, W. Langdon, O. Holland, Extending particle swarm optimisation via genetic programming, in: M. Keijzer, A. Tettamanzi, P. Collet, J. Hemert, M. Tomassini (Eds.), Genetic Programming, Vol. 3447 of LNCS, Springer Berlin Heidelberg, 2005, pp. 291–300.

- [34] J. Drake, N. Kililis, E. Özcan, Generation of vns components with grammatical evolution for vehicle routing, in: Genetic Programming, Vol. 7831 of LNCS, Springer Berlin Heidelberg, 2013, pp. 25–36.
- [35] X. Yao, Y. Liu, G. Lin, Evolutionary programming made faster, IEEE Transactions on Evolutionary Computation 3 (1999) 82–102.
- [36] T. Bäck, H.-P. Schwefel, An overview of evolutionary algorithms for parameter optimization, Evolutionary Computation 1 (1) (1993) 1–23.
- [37] C.-Y. Lee, X. Yao, Evolutionary programming using mutations based on the lévy probability distribution, IEEE Transactions on Evolutionary Computation 8 (1) (2004) 1–13.
- [38] R. N. Mantegna, Fast, accurate algorithm for numerical simulation of lévy stable stochastic processes, Physical Review E 49 (1994) 4677–4683.
- [39] H. Dong, J. He, H. Huang, W. Hou, Evolutionary programming using a mixed mutation strategy, Information Sciences 177 (1) (2007) 312 – 327.
- [40] R. Mallipeddi, S. Mallipeddi, P. Suganthan, Ensemble strategies with adaptive evolutionary programming, Information Sciences 180 (9) (2010) 1571–1581.
- [41] S. Silva, J. Almeida, Gplab-a genetic programming toolbox for matlab, in: Proceedings of the Nordic MATLAB conference, 2003, pp. 273–278.
- [42] S. Luke, L. Panait, Lexicographic parsimony pressure, in: Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002), Morgan Kaufmann Publishers, New York, NY, USA, 2002, pp. 829–836.

- [43] M. López-Ibáñez, J. Dubois-Lacoste, L. P. Cáceres, M. Birattari, T. Stützle, The irace package: Iterated racing for automatic algorithm configuration, *Operations Research Perspectives* 3 (2016) 43–58.