

RESEARCH

Open Access



Investigating the parameter space of evolutionary algorithms

Moshe Sipper^{1,2*} , Weixuan Fu¹, Karuna Ahuja¹ and Jason H. Moore¹

*Correspondence:

sipper@upenn.edu

¹Institute for Biomedical Informatics,
University of Pennsylvania,
Philadelphia 19104-6021, PA, USA

²Department of Computer Science,
Ben-Gurion University, Beer Sheva
8410501, Israel

Abstract

Evolutionary computation (EC) has been widely applied to biological and biomedical data. The practice of EC involves the tuning of many parameters, such as population size, generation count, selection size, and crossover and mutation rates. Through an *extensive* series of experiments over multiple evolutionary algorithm implementations and 25 problems we show that parameter space tends to be rife with viable parameters, at least for the problems studied herein. We discuss the implications of this finding in practice for the researcher employing EC.

Keywords: Evolutionary algorithms, Genetic programming, Meta-genetic algorithm, Parameter tuning, Hyper-parameter

Introduction

Evolutionary computation (EC) has been widely applied to biological and biomedical data [1–4]. One of the crucial tasks of the EC practitioner is the tuning of parameters. The fitness-select-vary paradigm comes with a plethora of parameters relating to the population, the generations, and the operators of selection, crossover, and mutation. It seems natural to ask whether the myriad parameters can be obtained through some clever methodology (perhaps even an evolutionary one) rather than by trial and error; indeed, as we shall see below, such methods have been previously devised. Our own interest in the issue of parameters stems partly from a desire to better understand evolutionary algorithms (EAs) and partly from our recent investigation into the design and implementation of an accessible artificial intelligence system [5].

In this paper we examine key parameters, asking whether we might find new insight into the parameter-seeking process. EC practitioners often employ commonly used parameters “selected by conventions, ad hoc choices, and very limited experimental comparisons” [6] (see also [7]). We sought parameters that met a reasonable minimal performance level over an entire set of several problems. We experimented with a large and variegated assortment of problems in what is arguably one of the most extensive EC experiments, concluding that parameter space, in fact, tends to be rife with viable parameters, at least for the problems studied herein. This does not mean that the EC practitioner’s job is over, given the many desiderata that still remain, including the representation of solutions in the search space, defining the fitness function, designing crossover and mutation operators, and more. However, our experiments suggest that one can at least find good parameters with a bit more ease.

We begin in the next section by delineating previous work on finding parameters through intelligent means. In “[Software and datasets](#)” section we turn to our own experiments by first describing the software and datasets we used. “[Searching for parameters using a meta-genetic algorithm](#)” section presents our first approach for obtaining good parameter sets, which was based on a meta-genetic algorithm, followed by a random-search approach in “[Searching for parameters using random search](#)” section. We discuss our findings and conclude in “[Concluding remarks](#)” section.

Previous work: Seeking parameters

Before delving into a detailed discussion, let us make the distinction between *parameters*, which are part of the model being evolved, and *hyper-parameters* (also called *meta-parameters*), which are not part of the model and need to be set by the user before running the evolutionary process, either manually or algorithmically. Examples of parameters are synaptic weights in deep neural networks and ephemeral random constants (ERCs) in genetic programming (GP), while examples of hyper-parameters include the number of hidden layers in deep neural networks and several standard parameters in GP: population size, generation count, crossover rate, mutation probability, and selection type.

Brest et al. [8] noted that there are two major forms of setting (hyper-)parameter values: *parameter tuning* and *parameter control*. Tuning refers to the common practice of seeking good values for the parameters before running the algorithm, then running the algorithm using these values, which remain fixed during the run. Parameter control means that values for the parameters are changed during the run.

An early work by [9] looked into the problem of VLSI layout, aiming to minimize the overall connection length between the components (cells) of a circuit. They devised a genetic algorithm that used three operators: crossover (order, cycle, or partially mapped—PMX), mutation (pairwise interchange of genes), and inversion (taking a random segment in a solution string and flipping it). They used a meta-GA to optimize crossover rate, inversion rate, and mutation rate. The individuals in the meta-GA population consisted of three integers in the range [0,20] (the overall search space was thus quite small, comprising 8000 combinations). An individual’s fitness was defined as the quality of the best layout found when a GA was run with the parameters it embodied. The meta-GA was run on four test circuits, had a population size of 20, ran for 100 generations, and used uniform crossover (select parameter from either parent at random) and mutation (add random noise; no inversion used). They noted that crossover rate converged to values in the range 20–40%, with little difference in best fitness as long as the rate was within this range. The mutation rate evolved by the meta-GA was 0.5–1.5%, and the inversion rate was 0–30%. The authors then adopted a crossover rate of 33%, an inversion rate of 15%, and a mutation rate of 0.5%. These were used to run the optimizing GA on circuits of interest (different than those used in the meta-GA phase) and compare it with other techniques.

Another early work by [10] described meta-evolutionary programming. Their flavor of evolutionary programming used mutation only to evolve solutions to two functions in \mathbb{R}^2 . Mutation took a single parent and generated an offspring by adding Gaussian random noise with zero mean and variance equal to $F(x, y)$ —the function under investigation; this was applied to all members of the population. The meta-algorithm attached a perturbation term to each individual in the population, which was used as variance during mutation. This term was then evolved along with the solution. They compared both these

algorithms (meta- and non-meta) with each other, and with a standard, Holland-style genetic algorithm, concluding that both versions of evolutionary programming outperformed the GA. They also concluded that the standard evolutionary-programming method attained better results over the two functions studied, but that the meta-version was more robust to changes in the function.

Wu and Chow [11] applied a genetic algorithm to nonlinear constrained mixed discrete-integer optimization problems, using a meta-GA to optimize population size, crossover probability, mutation probability, and crossover operator. The total number of parameter combinations was 19,200. The fitness of an individual in the meta-GA population was taken as the optimum found by a GA run with the parameters defined by the individual. Their findings showed insensitivity to crossover rate but high sensitivity to mutation rate. Four-point crossover outperformed one-, two-, and three-point crossover.

Hinterding et al. [12] (see also [13]) noted that, “it is natural to expect adaptation to be used not only for finding solutions to a problem, but also for tuning the algorithm to the particular problem.” Their paper provided a short survey of adaptation techniques in evolutionary computation. They defined four categories of adaptation: static—constant throughout run and tuned externally; and dynamic, refined further into deterministic—parameter altered by a deterministic rule, adaptive—some feedback from the evolutionary algorithm determines the change in the parameter, and self-adaptive—the parameters to be adapted are encoded into the chromosomes and undergo crossover and mutation. They also defined four levels of adaptation: environment—such as changes to the fitness function (e.g., weights), population—parameters that apply to the entire population are adapted, individual—parameters held within an individual affecting only that individual, and component—parameters specific to a component or gene within an individual (such as self-adaptation of component-level mutation steps sizes and rotation angles in evolution strategies).

Ong and Keane [14] presented meta-Lamarckian learning in the context of memetic algorithms (MA), which incorporate local improvement procedures within traditional GAs. Their paper investigated the adaptive choice of local search (LS) methods to ensure robustness in MA search. In addition to Darwinian evolution they also studied Lamarckian learning, where the genotype reflected the result of improvement through placing the locally improved individual back into the population to compete for reproductive opportunities. They studied two adaptive meta-Lamarckian learning strategies, a heuristic approach based on subproblem decomposition, and a stochastic approach based on a biased roulette wheel. They tested their system on continuous parametric benchmark test problems and on a wing-design problem. They concluded that, “the strategies presented are effective in producing search performances that are close to the best traditional MA with a LS chosen to suit the problem in hand. Given that such knowledge is often not available a priori, this ability to tackle new problems in a robust way is of significant value.”

Ramos et al. [15] proposed the utilization of logistic regression for tuning the parameters of a “transgenetic” algorithm—an evolutionary algorithm that deals, basically, with a population of chromosomes and a population of transgenetic vectors. They cited symbiogenesis as their inspiration, a theory of evolution according to which new cell organelles, new bodies, new organs, and new species arise from symbiosis, wherein independent organisms merge to form composites. The chromosomes of a transgenetic algorithm do not share genetic material directly. There are no crossover and mutation operations but

rather transgenetic vectors that obtain and insert information into the chromosomes. They used logistic regression to set two main parameters of their algorithm (population size and maximum length of transgenetic vector that met certain constraints), their problem of interest being Traveling Salesman. They showed that their algorithm outperformed a standard memetic algorithm.

Brest et al. [8], mentioned above, described an efficient technique for adapting control parameter settings associated with differential evolution (DE). DE uses a floating-point encoding for global optimization over continuous spaces, creating new candidate solutions by combining the parent individual and several other individuals of the same population. A candidate replaces the parent only if it has better fitness. DE has three parameters: amplification factor of the difference vector, crossover control parameter, and population size. In [8], the parameter control technique was based on the self-adaptation of the first two parameters, which were encoded within an individual's genome. Their testbed consisted of twenty-one benchmark functions from [16]. They concluded that self-adaptive DE is better or comparable to the original DE and some other evolutionary algorithms they examined.

De Jong [17]—in his chapter in the book *Parameter Setting in Evolutionary Algorithms* ([18])—provided a thirty-year perspective of parameter setting in evolutionary computation. He wrote that, “It is not surprising, then, that from the beginning EA practitioners have wanted to know the answers to questions like:

- Are there optimal settings for the parameters of an EA in general?
- Are there optimal settings for the parameters of an EA for a particular class of fitness landscapes?
- Are there robust settings for the parameters of an EA that produce good performance over a broad range of fitness landscapes?
- Is it desirable to dynamically change parameter values during an EA run?
- How do changes in a parameter affect the performance of an EA?
- How do landscape properties affect parameter value choices?”

He went on to review static parameter-setting strategies, where he mentioned a two-level EA, the top level of which evolved the parameters of the lower-level EA. [17] stated that the, “key insight from such studies is the robustness of EAs with respect to their parameter settings. Getting ‘in the ball park’ is generally sufficient for good EA performance.” Our study herein not only confirms this observation through numerous experiments, but also presents the novel finding that the ballpark can be quite large. Of dynamic parameter-setting strategies he opined that, “it is difficult to say anything definitive and general about the performance improvements obtained through dynamic parameter setting strategies.” He added, interestingly, “My own view is that there is not much to be gained in dynamically adapting EA parameter settings when solving static optimization problems. The real payoff for dynamic parameter setting strategies is when the fitness landscapes are themselves dynamic ...” [17] also discussed the different aspects of setting various standard parameters: parent population size, offspring population size, selection, reproductive operators; adapting the representation; and parameterless EAs.

Kramer [19] provided a survey of self-adaptive parameter control in evolutionary computation, where—as noted above—control parameters are added to the (evolving)

genome. He complemented the taxonomy offered by [13], dividing parameter setting into two main categories: tuning and control. Tuning was further divided into tuning by hand, tuning by design of experiments, and tuning by meta-evolution; while control was divided as previously into deterministic, adaptive, and self-adaptive. This paper mainly focused on function optimization techniques, such as the covariance matrix self-adaptation evolution strategy (CMSA-ES). A mention of meta-evolution noted that they, “have to be chosen problem-dependent, which is the obvious drawback of the approach.” He concluded with the observation that most theoretical work on self-adaptation concentrated on mutation, stating that, “A necessary condition for the success of self-adaptation is a tight link between strategy parameters and fitness.”

Eiben and Smit [6] (see also [20, 21]) presented a conceptual framework for parameter tuning based on a three-tier hierarchy of: a problem, an evolutionary algorithm (EA), and a tuner. They argued that parameter tuning could be considered from two different perspectives, that of configuring an evolutionary algorithm by choosing parameter values that optimize its performance, and that of analyzing an evolutionary algorithm by studying how its performance depends on its parameter values. Furthermore, they distinguished between analyzing an evolutionary algorithm by studying how its performance depends on the problems it is solving, and analyzing an evolutionary algorithm by studying how its performance varies when executing independent repetitions of its run. They noted the existence of two types of parameters, *qualitative* (e.g., crossover type) and *quantitative* (e.g., crossover rate). They opined that, “using tuning algorithms is highly rewarding. The efforts are moderate and the gains in performance can be very significant. Second, by using tuning algorithms one does not only obtain superior parameter values, but also much information about parameter values and algorithm performance. This information can be used to obtain a deeper understanding of the algorithm in question.” The paper discussed a wide range of tuning algorithms, which they classified as sampling methods, model-based methods, screening methods, and meta-evolutionary algorithms. Of interest in their discussion of meta-evolutionary GAs was [22], an early (possibly first) though limited work; and the description of multi-objective meta-GAs, which tuned for more than a single objective, e.g., speed and accuracy. They opined that, “parameter tuning in EC has been a largely ignored issue for a long time ... *In the current EC practice parameter values are mostly selected by conventions, ad hoc choices, and very limited experimental comparisons.*” [italics added] This latter observation—with which we wholly concur—forms part of our motivation for the current study.

Arcuri and Fraser [23] carried out “the largest empirical analysis so far on parameter tuning in search-based software engineering.” They performed experiments in the domain of test generation for object-oriented software using genetic algorithms. The objective was to derive sets of test cases (suites) for a given class, such that the test suite maximized a chosen coverage criterion while minimizing the number of tests and their length. A test case in this domain was a sequence of method calls that constructed objects and called methods on them. Because their goal was to study the effects of tuning, they analyzed all the possible combinations of the selected parameter values. They concluded that, “tuning can improve performance, but default values coming from the literature can be already sufficient.”

Veček et al. [24] introduced a new tuning method—CRS-Tuning—that is based on meta-evolution and their novel method for comparing and ranking evolutionary

algorithms, Chess Rating System for Evolutionary Algorithms (CRS4EAs). They discussed the approach's advantages over other tuning methods.

Bergstra and Bengio [25] studied neural networks, showing that random experiments were more efficient than grid experiments for hyper-parameter optimization in the case of several learning algorithms on several datasets. They wrote that, “random experiments are more efficient because not all hyperparameters are equally important to tune ... Random experiments are also easier to carry out than grid experiments for practical reasons related to the statistical independence of every trial.” This paper partly motivated our choice of random search in “[Searching for parameters using random search](#)” section.

Smit and Eiben [26] is perhaps the most relevant paper to our current research, presenting a meta-EA called REVAC (Relevance Estimation and Value Calibration), which they used on a suite of 25 real-valued benchmark functions (real-parameter optimization functions defined for the CEC 2005 Special Session on Real-Parameter Optimization, including five unimodal functions and twenty multimodal functions [27]). They chose to improve G-CMA-ES, which they considered a hard-to-improve evolutionary algorithm, cycling through *parent selection-recombination-mutation-survivor selection-evaluation* over a population of G-CMA-ES parameter vectors. They were indeed successful in improving the algorithm's performance.

Our aim is to go further, casting our net much wider in terms of problem domains, seeking to better understand parameter space.

Software and datasets

We chose to work with two very different evolutionary-algorithm packages: Distributed Evolutionary Algorithms in Python (DEAP) [28]—which uses tree-based GP, and M4GP [29]—which is a stack-based evolutionary algorithm. We ran our experiments on a cluster of 224 cores (Intel® Xeon® E5-2650L), with 2 threads per core.

DEAP, available at github.com/DEAP, comes with five sample problems:

1. Symbolic Regression, with data points generated from the quartic polynomial $x^4 + x^3 + x^2 + x$.
2. Even-Parity: find the parity, even or odd, of n Boolean inputs (we set n to 8).
3. Multiplexer 3–8: reproduce the behavior of an electronic multiplexer with 3 address bits (inputs) and 8 data lines (outputs).
4. Artificial Ant: evolve simple controllers—“artificial ants”—that are able to eat all the food located in a given two-dimensional, grid environment.
5. Spambase: return true if an email is spam, false otherwise.

We performed a preliminary investigation of these five problems, essentially running DEAP with parameters tuned by hand (a common-enough undertaking in the EC community). We found that for the first four problems one can attain an accuracy level of close to 1, while for Spambase our preliminary investigation set the attainable accuracy level at 0.93. More detailed descriptions of each problem can be found on the DEAP website.

M4GP is entirely different, based on stack-based GP. This serves to reinforce our conclusions by running our experiments on two very different types of EC algorithms. M4GP uses a nearest centroid distance metric to make classifications, with each program producing multiple outputs. We ran M4GP over problems from PMLB, a new publicly available dataset suite (accessibly hosted on GitHub) initialized with 165 real-world,

simulated, and toy benchmark datasets for evaluating supervised classification methods [30]. Note that PMLB focuses on classification benchmarks, whereas of the DEAP sample problems above only Spambase involves classification. Thus, our study includes different types of problems.

The preliminary investigation in this case consisted of delving into the data provided by the PMLB authors. “Once the datasets were scaled,” wrote [30], “we performed a comprehensive grid search of each of the ML method’s parameters using 10-fold cross-validation to find the best parameters (according to mean cross-validation balanced accuracy) for each ML method on each dataset. This process resulted in a total of over 5.5 million evaluations of the 13 ML methods over the 165 data sets.”

This available data saved us the need to run initial investigative experiments over PMLB. Each row of the 5.5-million table of results represents a single run, of a single ML algorithm, using a specific set of parameters; a row contains six columns: dataset name, classifier (machine learning algorithm), parameters, accuracy, macro-averaged F1 score, balanced accuracy. Given that one usually wants first and foremost to *solve* a problem we focused on accuracy, specifically, balanced accuracy (the rightmost value of each row), which is a normalized version of accuracy that accounts for class imbalance by calculating accuracy on a per-class basis, then averaging the per-class accuracies [31, 32].

We composed two suites of ten datasets each: 1) 10 datasets for which a balanced accuracy of 1 was attained most frequently (Table 1(a)), and 2) 10 datasets whose average balanced accuracy was in the range [0.9, 0.95] (Table 1(b)).

Table 1 PMLB results by [30]

Problem	Features	Classes	Samples
(a) “Easier” problems			
mofn-3-7-10	10	2	1324
Clean2	168	2	6598
Clean1	168	2	476
Mushroom	22	2	8124
Irish	5	2	500
Agaricus-lepiota	22	2	8145
Corral	6	2	160
xd6	9	2	973
mux6	6	2	128
ThreeOf9	9	2	512
(b) “Harder” problems			
Breast-cancer-wisconsin	30	2	569
wdbc	30	2	569
Tokyo1	44	2	959
New-thyroid	5	3	215
Spambase	57	2	4601
Vote	16	2	435
Soybean	35	18	675
House-votes-84	16	2	435
Breast-w	9	2	699
Molecular-biology_promoters	58	2	106

(a) “Easier” problems: 10 datasets for which a balanced accuracy of 1 was attained most frequently.

(b) “Harder” problems: 10 datasets whose average balanced accuracy was in the range [0.9, 0.95]. Shown for each problem: number of features, number of classes, and number of samples

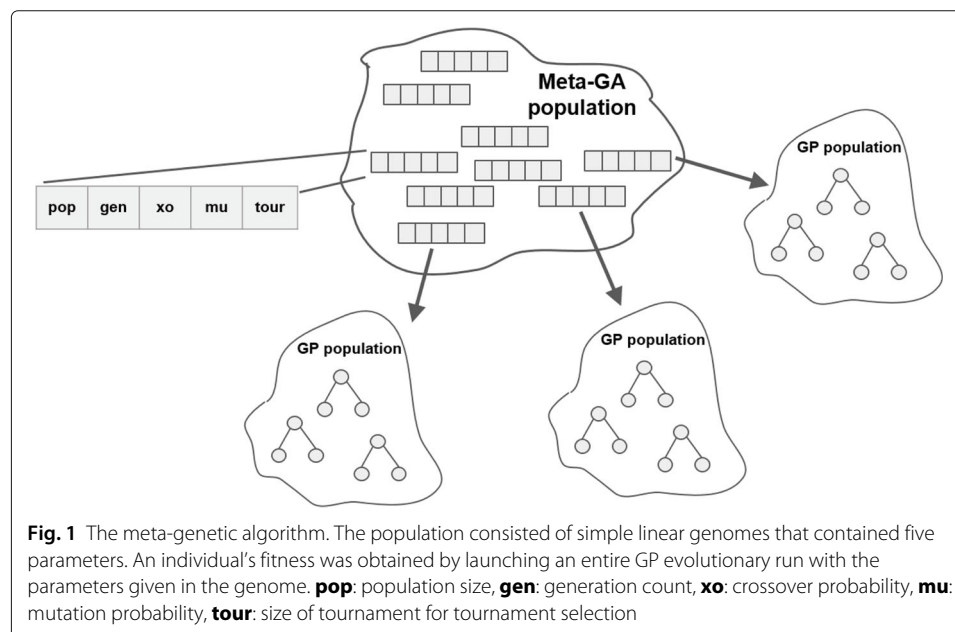
Searching for parameters using a meta-genetic algorithm

As done in a number of previous works discussed in “[Previous work: Seeking parameters](#)” section, we ran a meta-level genetic algorithm over the space of EC parameters, of which we identified five major ones:

1. Population size ($\in \mathbb{N}, [100, 3000]$),
2. Number of generations ($\in \mathbb{N}, [100, 2000]$),
3. Crossover rate ($\in \mathbb{R}, [0, 1]$),
4. Mutation rate ($\in \mathbb{R}, [0, 1]$),
5. Tournament size ($\in \mathbb{N}, [3, 100]$).

Figure 1 shows a schematic of the meta-GA’s workings. The meta-GA population comprised individuals with simple linear genomes that encoded the above five parameters as either integers or real values, respectively. An individual’s fitness was obtained by launching entire GP evolutionary runs with the parameters given in the genome. The GP in question was either DEAP or M4GP. DEAP’s goal was to solve the five problems listed in “[Software and datasets](#)” section (regression, parity, mux, ant, spam), while M4GP was set lose on the PMLB datasets. Each GP was run n times, where n was the number of problems it was set to solve (5 for DEAP, 10 for M4GP). Fitness was then a simple average of the n highest fitness values obtained during the n GP runs.

The meta-GA had a population size ranging from 100–300, ran for 50–160 generations, and used tournament selection with tournament size 3. Variation operators included: two-point crossover—with probability $p_{xo} = 0.5$ perform crossover on a pair of parents by selecting two (of four possible) parameter boundaries and switching the parameters between them; and value mutation—with probability $p_m = 0.2$ perform mutation on a single parent by selecting one of the five values at random, and mutating it by generating a random integer (population size, generation count, size of tournament) or real value (crossover probability, mutation probability) in the range specified above.



We experimented with the meta-GA for approximately two months, performing *tens of thousands of evolutionary runs*. We noted that numerous good parameter sets kept emerging, quite often appearing at random generation zero. Hence, we designed several putative improvements to the algorithmic process. First, we tweaked parameters such as crossover and mutation, and also added elitism (2%). It seemed that easier problems were causing the GP to move into local minima. To correct for this effect we introduced a weighted fitness function, where fitness was computed not as a simple average of the n GP runs but rather as a weighted one, with weights learned adaptively: examine every $m \in [1, 5]$ generations the average fitness per problem attained by all GP runs over that particular problem; then increase weights of below-average problems and decrease weights of above-average problems.

Other tweaks are not described here for brevity. Suffice it to say that after numerous runs of the various versions of the meta-GA, we eventually concluded that there seemed to be numerous successful parameter sets.

Searching for parameters using random search

How rife is parameter space with good parameters, i.e., ones whose use by a GP run results in success (which needs to be carefully defined)? Given the findings by [25] and the high cost grid search would incur in our case, we opted for random search.

We generated parameter sets at random, ran the following suites of experiments, and recorded the successful parameter sets:

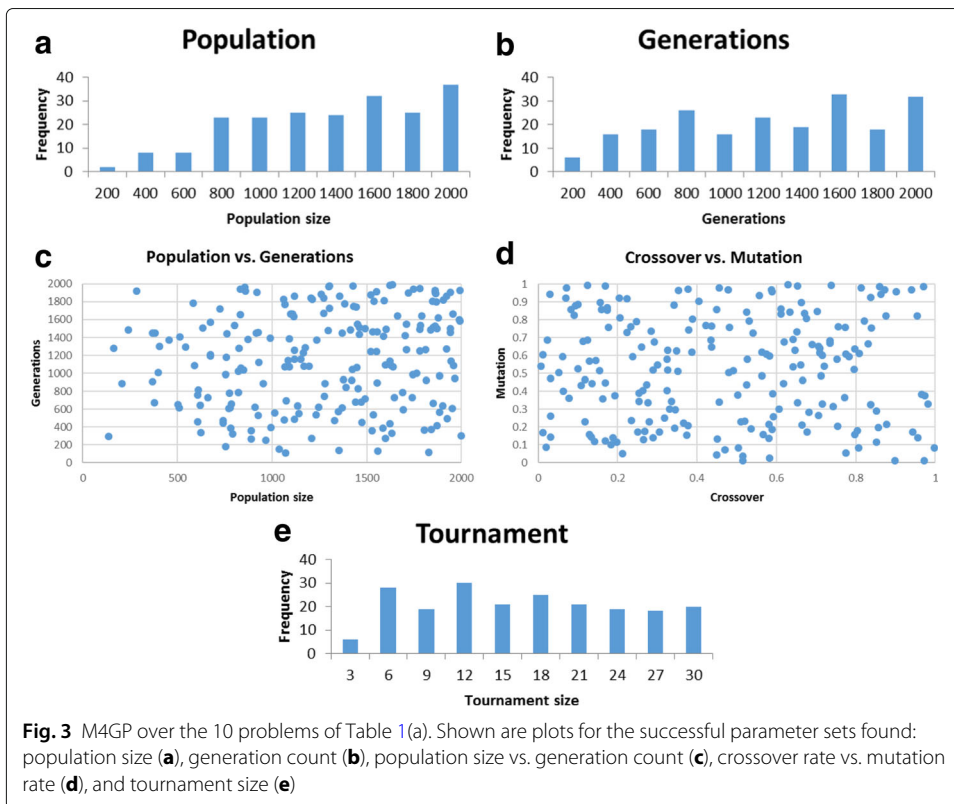
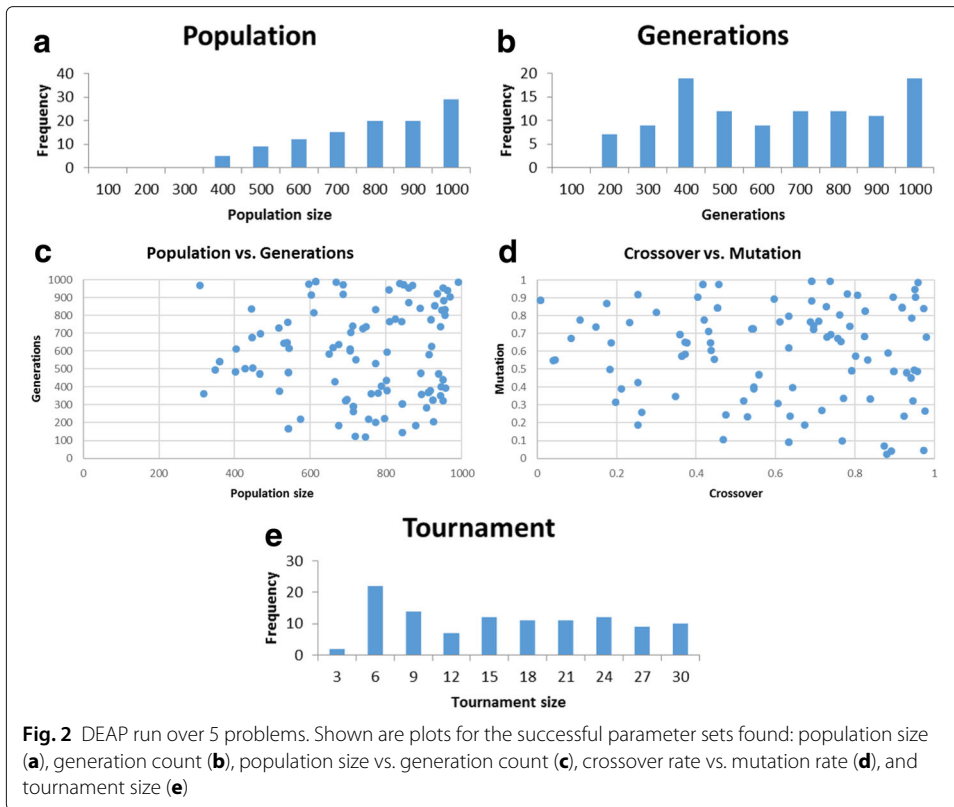
1. DEAP over the 5 sample problems (regression, parity, mux, ant, spam).
 - Generate random parameter sets with parameters in the following ranges: population size – [100, 1000]; generation count – [100, 1000]; crossover rate – [0, 1]; mutation rate – [0, 1]; tournament size – [3, 30].
 - Total runs (i.e., random parameter sets generated, with 5 GP runs launched per each, attempting to solve all five problems): 2693; number of successful parameter sets found: 110.
 - Success criterion of a parameter set: accuracy of 0.97 attained for all 5 problems but Spambase, where the accuracy threshold was set to 0.93.

Figure 2 shows our results.

2. M4GP over the 10 problems of Table 1(a).
 - Generate random parameter sets with parameters in the following ranges: population size – [100, 2000]; generation count – [100, 2000]; crossover rate – [0, 1]; mutation rate – [0, 1]; tournament size – [3, 30].
 - Total runs: 2610; number of successful parameter sets found: 207.
 - Success criterion of a parameter set: accuracy of 0.97 attained for all 10 problems.

Figure 3 shows our results.

3. M4GP over the 10 problems of Table 1(b).
 - Generate random parameter sets with parameters in the following ranges: population size – [100, 2000]; generation count – [100, 2000]; crossover rate – [0, 1]; mutation rate – [0, 1]; tournament size – [3, 30].



- Total runs: 5432; number of successful parameter sets found: 48.
- Success criterion of a parameter set: accuracy of 0.88 attained for all 10 problems.

Figure 4 shows our results.

The above experiments involved a total of 93,615 GP runs, each with a population size and generation count that could both be as high as 1000 or 2000.

Perhaps once given a parameter set, evolution is not providing any added value and simple random search would suffice? This is a standard question one should ask, and to answer it we used the good parameter sets found, in conjunction with random search (instead of evolution). Specifically, we performed the following:

1. DEAP: For each of the 110 good parameter sets found, generate $pop_size \times gen_count \times [1, 5]$ random solutions and check how many of them pass the same 5-problem criterion employed above.
Result: none passed.
The total number of random solutions examined was 47,028,011.
2. M4GP: For each of the 48 good Table 1(b) parameter sets found, generate $pop_size \times gen_count \times [1, 10]$ random solutions and check how many of them pass the same 10-problem criterion employed above¹.
Result: none passed.
The total number of random solutions examined was 205,377,967.

Random solutions were generated by using the respective software package to generate generation zero of size $pop_size \times gen_count$. This not only saved programming time but also prevented any bias vis-a-vis the GP experiments.

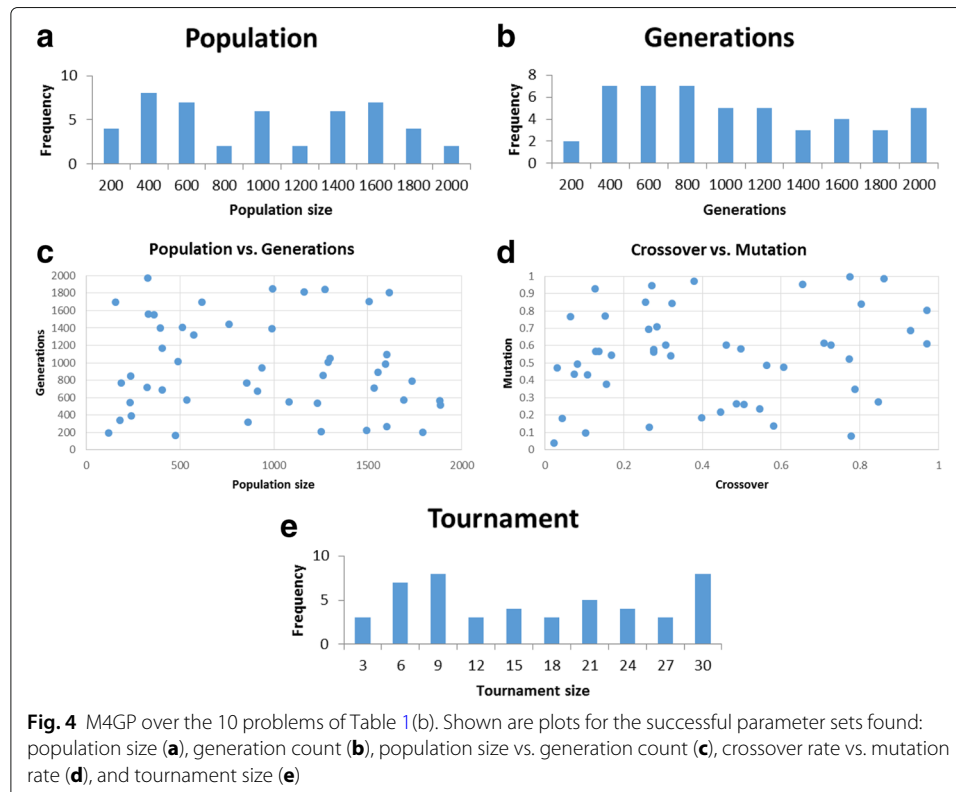


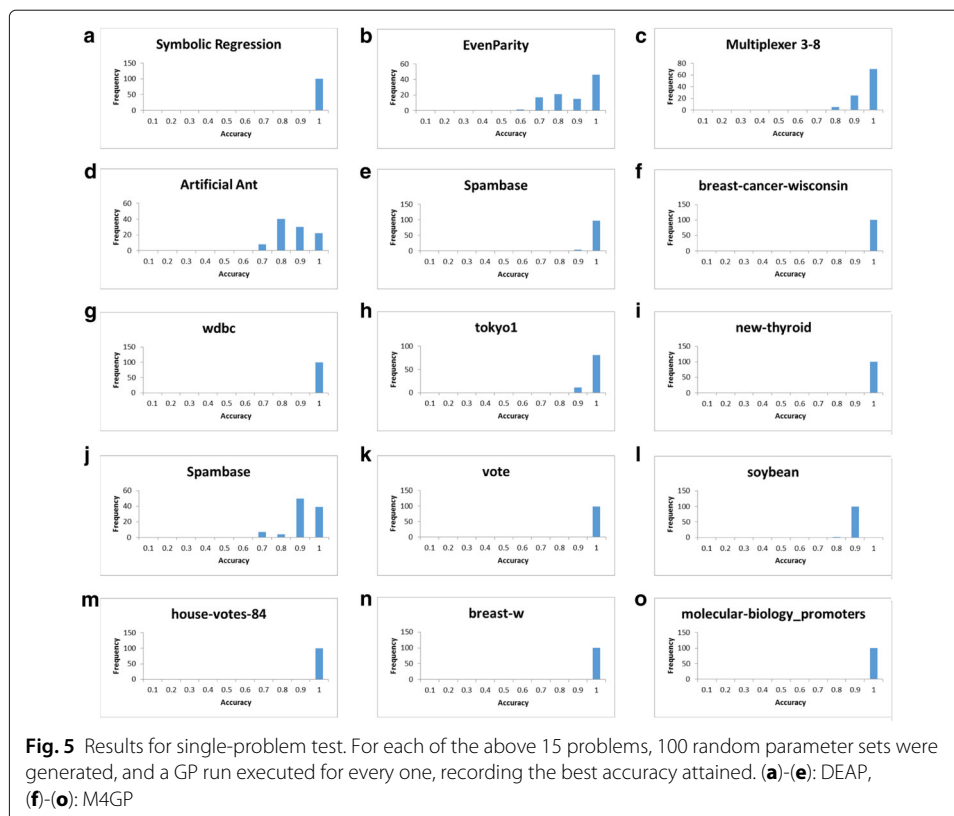
Fig. 4 M4GP over the 10 problems of Table 1(b). Shown are plots for the successful parameter sets found: population size (a), generation count (b), population size vs. generation count (c), crossover rate vs. mutation rate (d), and tournament size (e)

Concluding remarks

We performed what is arguably one of the most extensive EC experiments conducted. Studying our results, the only trend that seems to emerge is there being very little trend (at least for the problems studied herein). We can further make a number of interesting observations:

- Good parameters range over the entire spectrum, somewhat in contrast with common lore, which tends to focus on ad-hoc “good” values.
- While one is usually inclined to increase population size and generation count we note that this need not be so. At most, the two should not both be very low. We remark that recent evolutionary findings suggest the use of fewer generations [33].
- A commonly used range for tournament size is 3-7, but our experiments have shown that many more values work just as well.
- Crossover and mutation rates can take on widely diverse values, departing from the oft-used high or intermediate crossover rate and low mutation rate. Moreover, crossover-mutation pairs showed no tendency to aggregate anywhere. At most, such pairs should not both be low (which makes sense given that an evolutionary algorithm requires inter-generational variation).

The percentage of good random parameter sets found may seem relatively low, but that is not the point, rather, our major observation concerns the *diversity* in parameter space over a *wide range of problems*. Moreover, this percentage will likely be far higher in most situations encountered by evolutionary-algorithm practitioners, as we were aiming at a broad spectrum of problems—quite a high and somewhat unconventional bar—rather



than at the more common case of a single problem. To test this latter hypothesis we ran two additional sets of experiments:

1. DEAP: For each of the 5 problems, generate 100 random parameter sets and perform a GP run for each, recording the best accuracy attained.
2. M4GP: For each of the 10 problems of Table 1(b), generate 100 random parameter sets and perform a GP run for each, recording the best accuracy attained.

Figure 5 shows the results of these 1500 runs, which were very satisfactory.

We emphasize again what we stated in “Introduction” section—the EC practitioner’s job is far from over, what with solution representation, fitness function, crossover and mutation operators, and so forth. But perhaps one need not always spend too much time and resources on tuning hyper-parameters, with random search being a good choice for such tuning (after which EC will use these random hyper-parameters). Robustness to hyper-parameter tuning is a desired quality of an evolutionary algorithm and if one’s algorithm requires very specific parameters, the chance of finding them is slim; this would essentially be a needle-in-a-haystack situation in hyper-parameter space.

Endnote

¹Note: The number of actual solutions generated per test point could be less than 5 or 10, respectively, if the criterion failed early on.

Abbreviations

DE: Differential evolution; DEAP: Distributed evolutionary algorithms in python; EA: Evolutionary algorithm; EC: Evolutionary computation; GP: Genetic programming; MA: Memetic algorithm; PMLB: Penn machine learning benchmarks

Acknowledgments

We thank William La Cava, Randal Olson, and Ryan Urbanowicz for helpful discussions.

Funding

This work was supported by National Institutes of Health grants AI116794, DK112217, ES013508, HL134015, LM010098, LM011360, LM012601, and TR001263.

Availability of data and materials

The datasets supporting the conclusions of this article are available in the PMLB repository, <https://github.com/EpistasisLab/penn-ml-benchmarks> (see also [30]). We used two open-source software packages: DEAP [28] and M4GP [29].

Authors’ contributions

MS, WF, KA, and JHM conceived of and wrote the paper. All authors read and approved the final manuscript.

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

None of the authors have competing interests in this manuscript.

Publisher’s Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 12 October 2017 Accepted: 4 February 2018

Published online: 17 February 2018

References

1. Michal S, Ivry T, Cohen O, Sipper M, Barash D. Finding a common motif of RNA sequences using genetic programming: The GeRNAMo system. *IEEE/ACM Trans Comput Biol Bioinform (TCBB)*. 2007;4(4):596–610.
2. Moore JH, Hahn LW. Petri net modeling of high-order genetic systems using grammatical evolution. *BioSystems*. 2003;72(1):177–86.

3. Moore JH, Barney N, Tsai CT, Chiang FT, Gui J, White BC. Symbolic modeling of epistasis. *Hum Hered.* 2007;63(2):120–33.
4. Ritchie MD, White BC, Parker JS, Hahn LW, Moore JH. Optimization of neural network architecture using genetic programming improves detection and modeling of gene-gene interactions in studies of human diseases. *BMC Bioinformatics.* 2003;4(1):28.
5. Olson RS, Sipper M, La Cava W, Tartarone S, Vitale S, Fu W, Holmes JH, Moore JH. A system for accessible artificial intelligence. In: *Genetic Programming Theory & Practice XV*. Switzerland: Springer; 2017. <https://arxiv.org/abs/1705.00594>. (Submitted on 1 May 2017 (v1), last revised 10 Aug 2017 (this version, v2)).
6. Eiben AE, Smit SK. Parameter tuning for configuring and analyzing evolutionary algorithms. *Swarm Evol Comput.* 2011;1(1):19–31.
7. Eiben AE, Smit SK. Evolutionary algorithm parameters and methods to tune them. In: Hamadi Y, Monfroy E, Saubion F, editors. *Autonomous Search*. Berlin, Heidelberg: Springer; 2012. p. 15–36.
8. Brest J, Greiner S, Boskovic B, Mernik M, Zumer V. Self-adapting control parameters in differential evolution: a comparative study on numerical benchmark problems. *IEEE Trans Evol Comput.* 2006;10(6):646–57.
9. Shahookar K, Mazumder P. A genetic approach to standard cell placement using meta-genetic parameter optimization. *IEEE Trans Comput-Aided Design Integr Circ Sys.* 1990;9(5):500–11.
10. Fogel DB, Fogel LJ, Atmar JW. Meta-evolutionary programming. In: *Signals, Systems and Computers, 1991. 1991 Conference Record of the Twenty-fifth Asilomar Conference On*. IEEE Publications; 1991. p. 540–5.
11. Wu SJ, Chow PT. Genetic algorithms for nonlinear mixed discrete-integer optimization problems via meta-genetic parameter optimization. *Eng Optim+ A35.* 1995;24(2):137–59.
12. Hinterding R, Michalewicz Z, Eiben AE. Adaptation in evolutionary computation: A survey. In: *Evolutionary Computation, 1997, IEEE International Conference On*. IEEE Publications; 1997. p. 65–9.
13. Eiben AE, Hinterding R, Michalewicz Z. Parameter control in evolutionary algorithms. *IEEE Trans Evol Comput.* 1999;3(2):124–41.
14. Ong YS, Keane AJ. Meta-lamarckian learning in memetic algorithms. *IEEE Trans Evol Comput.* 2004;8(2):99–110.
15. Ramos IC, Goldbarg MC, Goldbarg EG, Neto ADD. Logistic regression for parameter tuning on an evolutionary algorithm. In: *2005 IEEE Congress on Evolutionary Computation*. IEEE Publications; 2005. p. 1061–8.
16. Yao X, Liu Y, Lin G. Evolutionary programming made faster. *IEEE Trans Evol Comput.* 1999;3(2):82–102.
17. De Jong K. Parameter setting in EAs: a 30 year perspective. In: *Parameter Setting in Evolutionary Algorithms*. Switzerland: Springer; 2007. p. 1–18.
18. Lobo F, Lima CF, Michalewicz Z, Vol. 54. *Parameter Setting in Evolutionary Algorithms*. Switzerland: Springer; 2007.
19. Kramer O. Evolutionary self-adaptation: a survey of operators and strategy parameters. *Evol Intel.* 2010;3(2):51–65.
20. Smit SK, Eiben AE. Comparing parameter tuning methods for evolutionary algorithms. In: *2009 IEEE Congress on Evolutionary Computation*. IEEE Publications; 2009. p. 399–406.
21. Smit SK, Eiben A. Parameter tuning of evolutionary algorithms: Generalist vs. specialist. In: *European Conference on the Applications of Evolutionary Computation*. Switzerland: Springer International Publishing; 2010. p. 542–51.
22. Mercer RE, Sampson J. Adaptive search using a reproductive meta-plan. *Kybernetes.* 1978;7(3):215–28.
23. Arcuri A, Fraser G. Parameter tuning or default values? an empirical investigation in search-based software engineering. *Empir Softw Eng.* 2013;18(3):594–623.
24. Veček N, Mernik M, Filipič B, Črepinšek M. Parameter tuning with chess rating system (CRS-Tuning) for meta-heuristic algorithms. *Inf Sci.* 2016;372:446–69.
25. Bergstra J, Bengio Y. Random search for hyper-parameter optimization. *J Mach Learn Res.* 2012;13(Feb):281–305.
26. Smit SK, Eiben AE. Beating the ‘world champion’ evolutionary algorithm via REVAC tuning. In: *IEEE Congress on Evolutionary Computation*. IEEE Publications; 2010. p. 1–8.
27. Suganthan PN, Hansen N, Liang JJ, Deb K, Chen YP, Auger A, Tiwari S. Problem definitions and evaluation criteria for the CEC 2005 special session on real-parameter optimization. *KanGAL report 2005005*. 2005;2005.
28. Fortin FA, De Rainville F-M, Gardner MA, Parizeau M, Gagné C. DEAP: Evolutionary algorithms made easy. *J Mach Learn Res.* 2012;13:2171–5.
29. La Cava W, Silva S, Vanneschi L, Specto L, Moore JH. Genetic programming representations for multi-dimensional feature learning in biomedical classification. In: *EvoStar 2017*. Switzerland: Springer; 2017.
30. Olson RS, La Cava W, Orzechowski P, Urbanowicz RJ, Moore JH. PMLB: a large benchmark suite for machine learning evaluation and comparison. *BioData Mining.* 2017;10(1):36.
31. Velez DR, et al. A balanced accuracy function for epistasis modeling in imbalanced datasets using multifactor dimensionality reduction. *Genetic Epidemiol.* 2007;31(4):306–15.
32. Urbanowicz RJ, Moore JH. ExSTraCS 2.0: description and evaluation of a scalable learning classifier system. *Evol Intell.* 2015;8(2–3):89–116.
33. Arnold C. Evolution Runs Faster on Short Timescales. 2017. *Quanta Magazine*. www.quantamagazine.org/20170314-time-dependent-rate-phenomenon-evolution-viruses. Accessed 14 Mar 2017.