

Genetic Programming

Riccardo Poli¹ and John R. Koza²

¹*School of Computer Science and Electronic Engineering, University of Essex, United Kingdom*

²*Stanford University, Stanford, California*

1 Introduction

The goal of getting computers to automatically solve problems is central to artificial intelligence, machine learning, and the broad area encompassed by what Turing called “machine intelligence” [161, 162].

In his 1983 talk entitled “AI: Where It Has Been and Where It Is Going”, machine learning pioneer Arthur Samuel stated the main goal of the fields of machine learning and artificial intelligence:

“[T]he aim [is] ... to get machines to exhibit behavior, which if done by humans, would be assumed to involve the use of intelligence.”

Genetic programming (GP) is a systematic method for getting computers to automatically solve a problem starting from a high-level statement of what needs to be done. GP is a domain-independent method that genetically breeds a population of computer programs to solve a problem. Specifically, GP iteratively transforms a population of computer programs into a new generation of programs by applying analogs of naturally occurring genetic operations. This process is illustrated in Figure 1.

The genetic operations include crossover (sexual recombination), mutation and reproduction. It may also include other analogues of natural operations such as gene duplication, gene deletion and developmental processes which transform an embryo into a fully developed structure. GP is an extension of the genetic algorithm [52] in which the *structures* in the population are not fixed-length character strings that encode candidate solutions to a problem, but *programs* that, when executed, *are* the candidate solutions to the problem.

Programs are expressed in GP as *syntax trees* rather than as lines of code. For example, the simple expression $\max(x*x, x+3*y)$ is represented as shown in Figure 2. The tree includes *nodes* (which we will also call *points*) and *links*. The nodes indicate the instructions to execute. The links indicate the arguments for each instruction. In the following the internal nodes in a tree will be called *functions*, while the tree’s leaves will be called *terminals*.

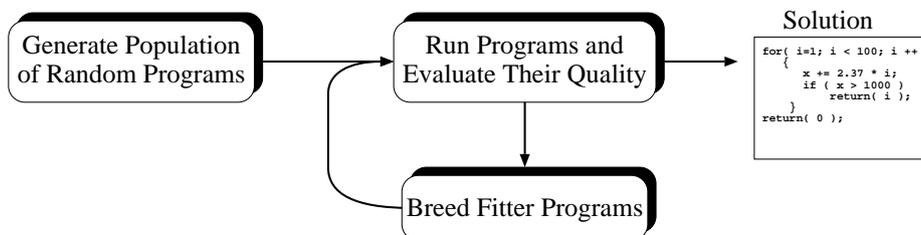


Figure 1: Main loop of genetic programming

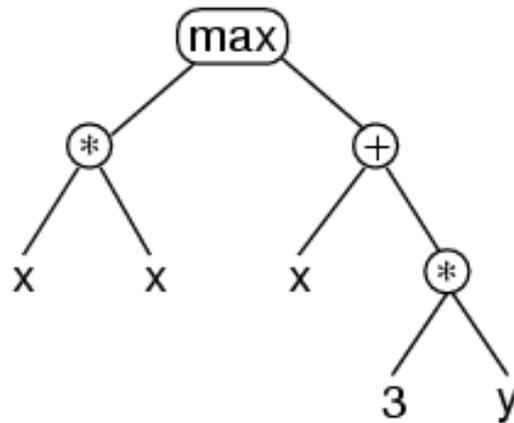


Figure 2: Basic tree-like program representation used in genetic programming

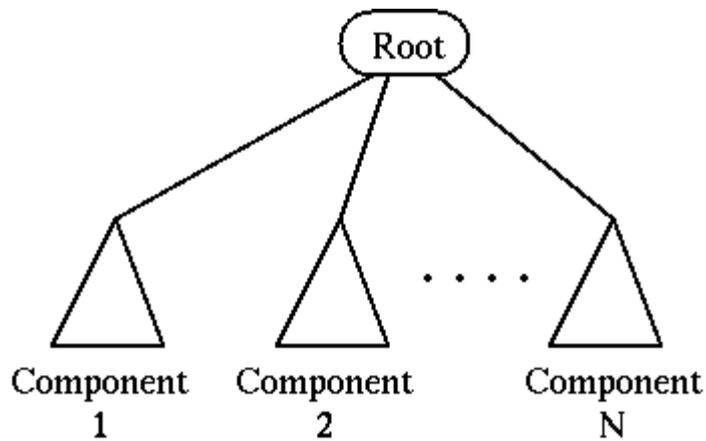


Figure 3: Multi-tree program representation

In more advanced forms of GP, programs can be composed of multiple components (e.g., subroutines). Often in this case the representation used in GP is a set of trees (one for each component) grouped together under a special node called *root*, as illustrated in Figure 3. We will call these (sub)trees *branches*. The number and type of the branches in a program, together with certain other features of the structure of the branches, form the *architecture* of the program.

GP trees and their corresponding expressions can equivalently be represented in *prefix notation* (e.g., as Lisp S-expressions). In prefix notation, functions always precede their arguments. For example, $\max(x*x, x+3*y)$ becomes $(\max (* x x) (+ x (* 3 y)))$. In this notation, it is easy to see the correspondence between expressions and their syntax trees. Simple recursive procedures can convert prefix-notation expressions into infix-notation expressions and vice versa. Therefore, in the following, we will use trees and their corresponding prefix-notation expressions interchangeably.

2 Preparatory Steps of Genetic Programming

Genetic programming starts from a high-level statement of the requirements of a problem and attempts to produce a computer program that solves the problem.

The human user communicates the high-level statement of the problem to the GP algorithm by performing certain well-defined preparatory steps.

The five major preparatory steps for the basic version of genetic programming require the human user to specify:

1. the set of terminals (e.g., the independent variables of the problem, zero-argument functions, and random constants) for each branch of the to-be-evolved program,
2. the set of primitive functions for each branch of the to-be-evolved program,
3. the fitness measure (for explicitly or implicitly measuring the quality of individuals in the population),
4. certain parameters for controlling the run, and
5. the termination criterion and method for designating the result of the run.

The first two preparatory steps specify the ingredients that are available to create the computer programs. A run of GP is a competitive search among a diverse population of programs composed of the available functions and terminals.

The identification of the function set and terminal set for a particular problem (or category of problems) is usually a straightforward process. For some problems, the function set may consist of merely the arithmetic functions of addition, subtraction, multiplication, and division as well as a conditional branching operator. The terminal set may consist of the program's external inputs (independent variables) and numerical constants.

For many other problems, the ingredients include specialized functions and terminals. For example, if the goal is to get GP to automatically program a robot to mop the entire floor of an obstacle-laden room, the human user must tell GP what the robot is capable of doing. For example, the robot may be capable of executing functions such as moving, turning, and swishing the mop.

If the goal is the automatic creation of a controller, the function set may consist of integrators, differentiators, leads, lags, gains, adders, subtractors, and the like and the terminal set may consist of signals such as the reference signal and plant output.

If the goal is the automatic synthesis of an analog electrical circuit, the function set may enable GP to construct circuits from components such as transistors, capacitors, and resistors. Once the human user has identified the primitive ingredients for a problem of circuit synthesis, the same function set can be used to automatically synthesize an amplifier, computational circuit, active filter, voltage reference circuit, or any other circuit composed of these ingredients.

The third preparatory step concerns the fitness measure for the problem. The fitness measure specifies what needs to be done. The fitness measure is the primary mechanism for communicating the high-level statement of the problem's requirements to the genetic programming system. For example, if the

goal is to get genetic programming to automatically synthesize an amplifier, the fitness function is the mechanism for telling GP to synthesize a circuit that amplifies an incoming signal (as opposed to, say, a circuit that suppresses the low frequencies of an incoming signal or that computes the square root of the incoming signal). The first two preparatory steps define the search space whereas the fitness measure implicitly specifies the search's desired goal.

The fourth and fifth preparatory steps are administrative. The fourth preparatory step entails specifying the control parameters for the run. The most important control parameter is the population size. Other control parameters include the probabilities of performing the genetic operations, the maximum size for programs, and other details of the run.

The fifth preparatory step consists of specifying the termination criterion and the method of designating the result of the run. The termination criterion may include a maximum number of generations to be run as well as a problem-specific success predicate. The single best-so-far individual is then harvested and designated as the result of the run.

3 Executional Steps of Genetic Programming

After the user has performed the preparatory steps for a problem, the run of genetic programming can be launched. Once the run is launched, a series of well-defined, problem-independent steps is executed.

GP typically starts with a population of randomly generated computer programs composed of the available programmatic ingredients (as provided by the human user in the first and second preparatory steps).

GP iteratively transforms a population of computer programs into a new generation of the population by applying analogs of naturally occurring genetic operations. These operations are applied to individual(s) selected from the population. The individuals are probabilistically selected to participate in the genetic operations based on their fitness (as measured by the fitness measure provided by the human user in the third preparatory step). The iterative transformation of the population is executed inside the main generational loop of the run of GP.

The executional steps of GP are as follows:

1. Randomly create an initial population (generation 0) of individual computer programs composed of the available functions and terminals.
2. Iteratively perform the following sub-steps (called a *generation*) on the population until the termination criterion is satisfied:
 - (a) Execute each program in the population and ascertain its fitness (explicitly or implicitly) using the problem's fitness measure.
 - (b) Select one or two individual program(s) from the population with a probability based on fitness (with reselection allowed) to participate in the genetic operations in (c).
 - (c) Create new individual program(s) for the population by applying the following genetic operations with specified probabilities:

- **Reproduction:** Copy the selected individual program into the new population.
 - **Crossover:** Create new offspring program(s) for the new population by recombining randomly chosen parts from two selected programs.
 - **Mutation:** Create one new offspring program for the new population by randomly mutating a randomly chosen part of one selected program.
 - **Architecture-altering operations:** If this feature is enabled, choose an architecture-altering operation from the available repertoire of such operations and create one new offspring program for the new population by applying the chosen architecture-altering operation to one selected program.
3. After the termination criterion is satisfied, the single best program in the population produced during the run (the best-so-far individual) is harvested and designated as the result of the run. If the run is successful, the result will be a solution (or approximate solution) to the problem.

Figure 4 is a flowchart of GP showing the genetic operations of crossover, reproduction, and mutation as well as the architecture-altering operations. This flowchart shows a two-offspring version of the crossover operation.

The preparatory steps specify what the user must provide in advance to the GP system. Once the run is launched, the executional steps as shown in the flowchart (Figure 4) are executed. GP is problem-independent in the sense that the flowchart specifying the basic sequence of executional steps is not modified for each new run or each new problem.

There is usually no discretionary human intervention or interaction during a run of GP (although a human user may exercise judgment as to whether to terminate a run).

GP starts with an initial population of computer programs composed of functions and terminals appropriate to the problem. The individual programs in the initial population are typically generated by recursively generating a rooted point-labeled program tree composed of random choices of the primitive functions and terminals (provided by the user as part of the first and second preparatory steps). The initial individuals are usually generated subject to a pre-established maximum size (specified by the user as a minor parameter as part of the fourth preparatory step). For example, in the “*Full*” *initialization method* nodes are taken from the function set until a maximum tree depth is reached. Beyond that depth only terminals can be chosen. Figure 5 shows several snapshots of this process. A variant of this, the “*Grow*” *initialization method*, allows the selection of nodes from the whole primitive set until the depth limit is reached. Thereafter, it behaves like the “*Full*” method.

In general, after the initialization phase, the programs in the population are of different size (number of functions and terminals) and of different shape (the particular graphical arrangement of functions and terminals in the program tree).

Each individual program in the population is either measured or compared in terms of how well it performs the task at hand (using the fitness measure provided in the third preparatory step). For many problems, this measurement

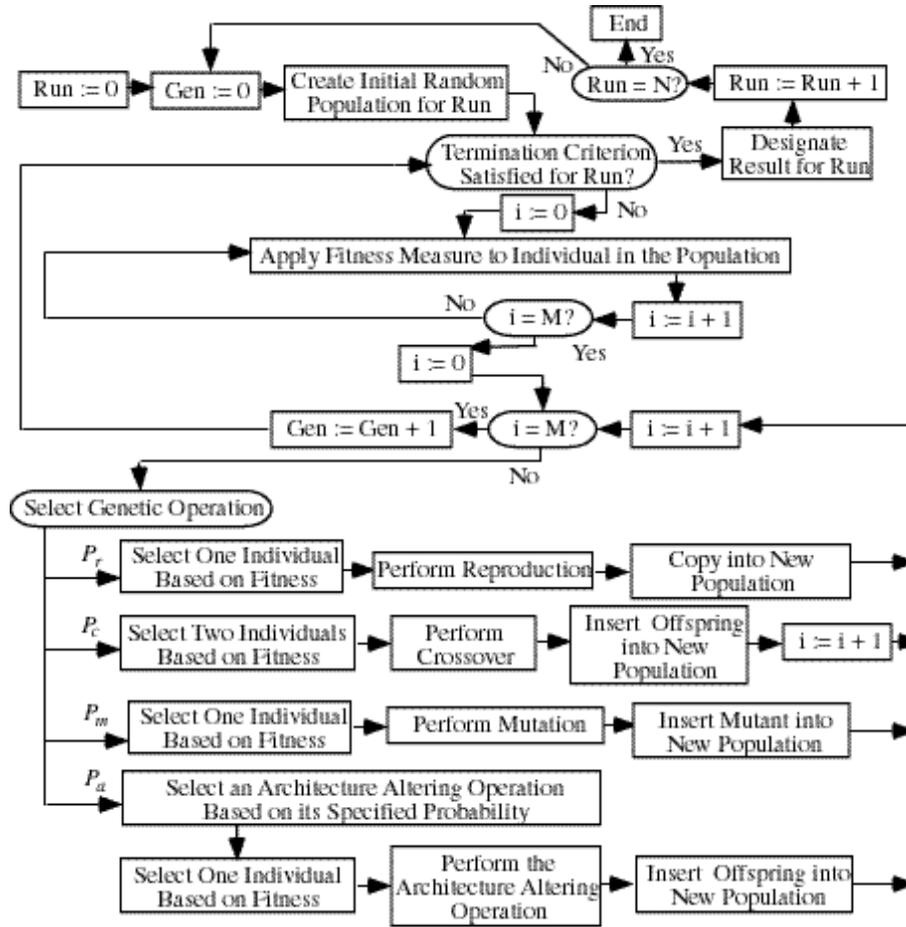


Figure 4: Flowchart of genetic programming

yields a single explicit numerical value, called *fitness*. Normally, fitness evaluation requires executing the programs in the population, often multiple times, *within* the GP system. A variety of execution strategies exist. The most common are virtual-machine-code compilation and interpretation. We will look at the latter.

Interpreting a program tree means executing the nodes in the tree in an order that guarantees that nodes are not executed before the value of their arguments (if any) is known. This is usually done by traversing the tree in a recursive way starting from the root node, and postponing the evaluation of each node until the value of its children (arguments) is known. This process is illustrated in Figure 6, where the numbers to the right of internal nodes represent the results of evaluating the subtrees rooted at such nodes. In this example, the independent variable x evaluates to -1 . Figure 7 gives a pseudo-code implementation of the interpretation procedure. The code assumes that programs are represented as prefix-notation expressions and that such expressions can be treated as lists of components (where a construct like $expr(i)$ can be used to read or set component i of expression $expr$).

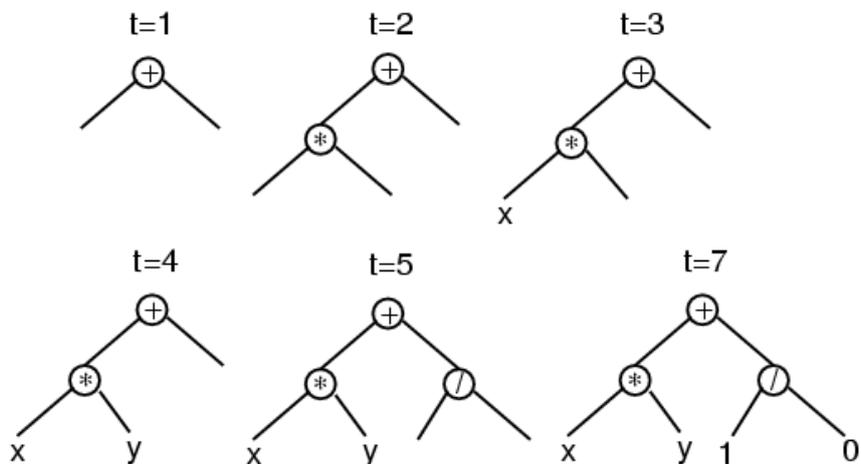


Figure 5: Creation of a seven-point tree using the “Full” initialization method (t=time)

Irrespective of the execution strategy adopted, the fitness of a program may be measured in many different ways, including, for example, in terms of the amount of error between its output and the desired output, the amount of time (fuel, money, etc.) required to bring a system to a desired target state, the accuracy of the program in recognizing patterns or classifying objects into classes, the payoff that a game-playing program produces, or the compliance of a complex structure (such as an antenna, circuit, or controller) with user-specified design criteria. The execution of the program sometimes returns one or more explicit values. Alternatively, the execution of a program may consist only of side effects on the state of a world (e.g., a robot’s actions). Alternatively, the execution of a program may yield both return values and side effects.

The fitness measure is, for many practical problems, multi-objective in the sense that it combines two or more different elements. In practice, the different elements of the fitness measure are in competition with one another to some degree.

For many problems, each program in the population is executed over a representative sample of different *fitness cases*. These fitness cases may represent different values of the program’s input(s), different initial conditions of a system, or different environments. Sometimes the fitness cases are constructed probabilistically.

The creation of the initial random population is, in effect, a blind random search of the search space of the problem. It provides a baseline for judging future search efforts. Typically, the individual programs in generation 0 all have exceedingly poor fitness. Nonetheless, some individuals in the population are (usually) fitter than others. The differences in fitness are then exploited by genetic programming. GP applies Darwinian selection and the genetic operations to create a new population of offspring programs from the current population.

The genetic operations include crossover (sexual recombination), mutation, reproduction, and the architecture-altering operations (when they are enabled). Given copies of two parent trees, typically, *crossover* involves randomly select-

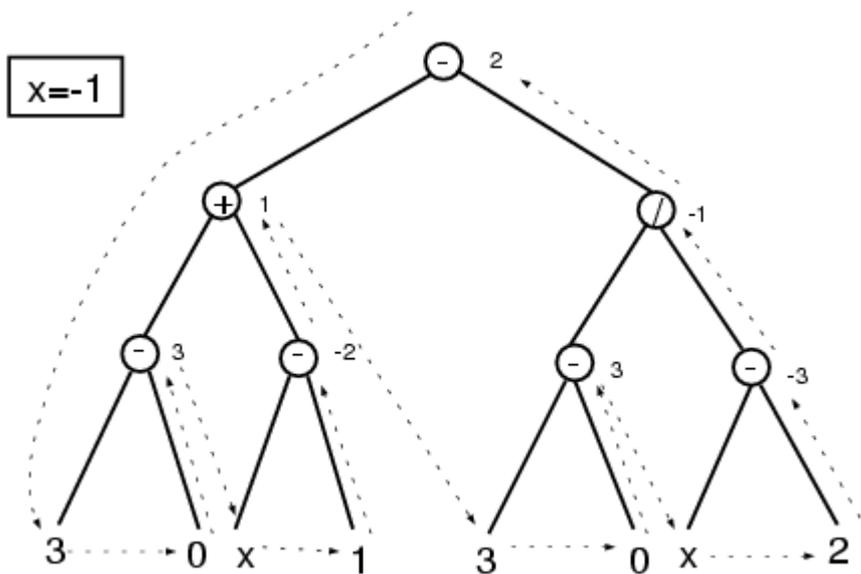


Figure 6: Example interpretation of a syntax tree (the terminal x is a variable holding the value -1)

ing a crossover point in each parent tree and swapping the sub-trees rooted at the crossover points, as exemplified in Figure 8. Often crossover points are not selected with uniform probability. A frequent strategy is, for example, to select internal nodes (functions) 90% of the times, and any node for the remaining 10% of the times. Traditional *mutation* consists of randomly selecting a mutation point in a tree and substituting the sub-tree rooted there with a randomly generated sub-tree, as illustrated in Figure 9. *Reproduction* involves simply copying certain individuals into the new population. Architecture altering operations will be discussed later in this chapter.

The genetic operations described above are applied to individual(s) that are probabilistically selected from the population based on fitness. In this probabilistic selection process, better individuals are favored over inferior individuals. However, the best individual in the population is not necessarily selected and the worst individual in the population is not necessarily passed over.

After the genetic operations are performed on the current population, the population of offspring (i.e., the new generation) replaces the current population (i.e., the now-old generation). This iterative process of measuring fitness and performing the genetic operations is repeated over many generations.

The run of GP terminates when the termination criterion (as provided by the fifth preparatory step) is satisfied. The outcome of the run is specified by the method of result designation. The best individual ever encountered during the run (i.e., the best-so-far individual) is typically designated as the result of the run.

All programs in the initial random population (generation 0) of a run of GP are syntactically valid, executable programs. The genetic operations that are performed during the run (i.e., crossover, mutation, reproduction, and the architecture-altering operations) are designed to produce offspring that are syn-

```

procedure: eval
arguments:
  expr    /* An expression in prefix notation */
results:
  value   /* A number */
begin
  if expr is a list then /* Non-terminal */
    proc = expr(1)
    value = proc(eval(expr(2)),eval(expr(3)),...)
  else /* Terminal */
    if expr is a variable or a constant then
      value = expr
    else /* 0-arity function */
      value = expr()
    endif
  endif
end

```

Figure 7: Typical interpreter for genetic programming

tactically valid, executable programs. Thus, every individual created during a run of genetic programming (including, in particular, the best-of-run individual) is a syntactically valid, executable program.

There are numerous alternative implementations of GP that vary from the preceding brief description. We will discuss some options in Section 5.

4 Example of a Run of Genetic Programming

To provide concreteness, this section contains an illustrative run of GP in which the goal is to automatically create a computer program whose output is equal to the values of the quadratic polynomial x^2+x+1 in the range from -1 to $+1$. That is, the goal is to automatically create a computer program that matches certain numerical data. This process is sometimes called *system identification* or *symbolic regression*.

We begin with the five preparatory steps.

The purpose of the first two preparatory steps is to specify the ingredients of the to-be-evolved program.

Because the problem is to find a mathematical function of one independent variable, the terminal set (inputs to the to-be-evolved program) includes the independent variable, x . The terminal set also includes numerical constants. That is, the terminal set, is $\mathbb{T} = \{x, \mathfrak{R}\}$, where \mathfrak{R} denotes constant numerical terminals in some reasonable range (say from -5.0 to $+5.0$).

The preceding statement of the problem is somewhat flexible in that it does not specify what functions may be employed in the to-be-evolved program. One possible choice for the function set consists of the four ordinary arithmetic functions of addition, subtraction, multiplication, and division. This choice is reasonable because mathematical expressions typically include these functions. Thus, the function set for this problem is $\mathbb{F} = \{+, -, *, /\}$, where

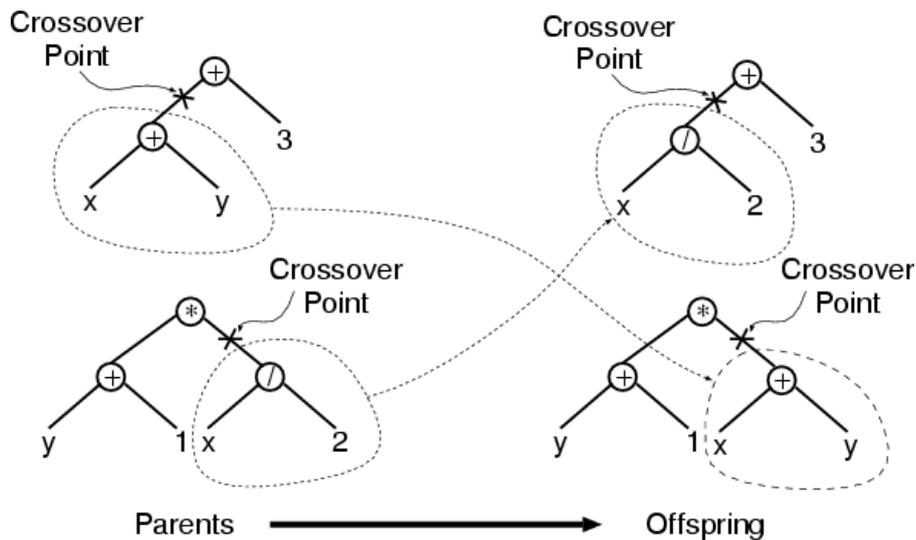


Figure 8: Example of two-child crossover between syntax trees

the two-argument $+$, $-$, $*$, and $\%$ functions add, subtract, multiply, and divide, respectively. To avoid run-time errors, the division function $\%$ is protected: it returns a value of 1 when division by 0 is attempted (including 0 divided by 0), but otherwise returns the quotient of its two arguments.

Each individual in the population is a composition of functions from the specified function set and terminals from the specified terminal set.

The third preparatory step involves constructing the fitness measure. The purpose of the fitness measure is to specify what the human wants. The high-level goal of this problem is to find a program whose output is equal to the values of the quadratic polynomial x^2+x+1 . Therefore, the fitness assigned to a particular individual in the population for this problem must reflect how closely the output of an individual program comes to the target polynomial x^2+x+1 . The fitness measure could be defined as the value of the integral (taken over values of the independent variable x between -1.0 and $+1.0$) of the absolute value of the differences (errors) between the value of the individual mathematical expression and the target quadratic polynomial x^2+x+1 . A smaller value of fitness (error) is better. A fitness (error) of zero would indicate a perfect fit.

For most problems of symbolic regression or system identification, it is not practical or possible to analytically compute the value of the integral of the absolute error. Thus, in practice, the integral is numerically approximated using dozens or hundreds of different values of the independent variable x in the range between -1.0 and $+1.0$.

The population size in this small illustrative example will be just four. In actual practice, the population size for a run of genetic programming consists of thousands or millions of individuals. In actual practice, the crossover operation is commonly performed on about 90% of the individuals in the population; the reproduction operation is performed on about 8% of the population; the mutation operation is performed on about 1% of the population; and the architecture-altering operations are performed on perhaps 1% of the popula-

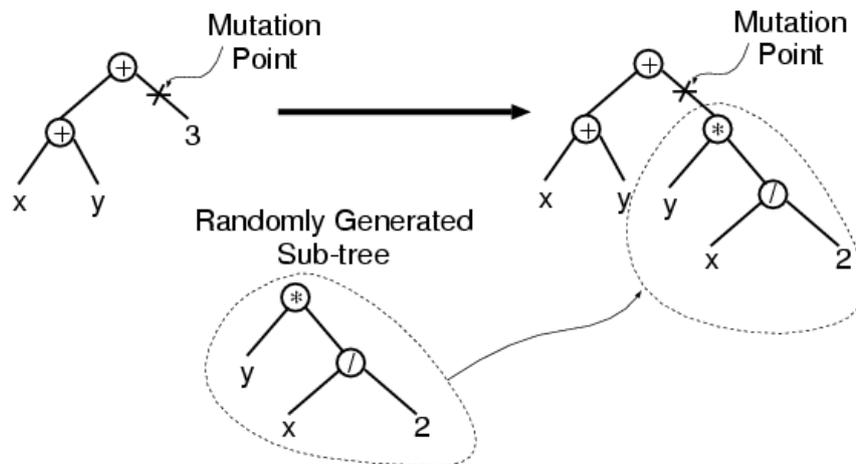


Figure 9: Example of sub-tree mutation

tion. Because this illustrative example involves an abnormally small population of only four individuals, the crossover operation will be performed on two individuals and the mutation and reproduction operations will each be performed on one individual. For simplicity, the architecture-altering operations are not used for this problem.

A reasonable termination criterion for this problem is that the run will continue from generation to generation until the fitness of some individual gets below 0.01. In this contrived example, the run will (atypically) yield an algebraically perfect solution (for which the fitness measure attains the ideal value of zero) after merely one generation.

Now that we have performed the five preparatory steps, the run of GP can be launched. That is, the executional steps shown in the flowchart of Figure 4 are now performed.

GP starts by randomly creating a population of four individual computer programs. The four programs are shown in Figure 10 in the form of trees.

The first randomly constructed program tree (Figure 10a) is equivalent to the mathematical expression $x+1$. A program tree is executed in a depth-first way, from left to right, in the style of the LISP programming language. Specifically, the addition function (+) is executed with the variable x and the constant value 1 as its two arguments. Then, the two-argument subtraction function (-) is executed. Its first argument is the value returned by the just-executed addition function. Its second argument is the constant value 0. The overall result of executing the entire program tree is thus $x+1$.

The first program (Figure 10a) was constructed, using the “Grow” method, by first choosing the subtraction function for the root (top point) of the program tree. The random construction process continued in a depth-first fashion (from left to right) and chose the addition function to be the first argument of the subtraction function. The random construction process then chose the terminal x to be the first argument of the addition function (thereby terminating the growth of this path in the program tree). The random construction process then chose the constant terminal 1 as the second argument of the addition

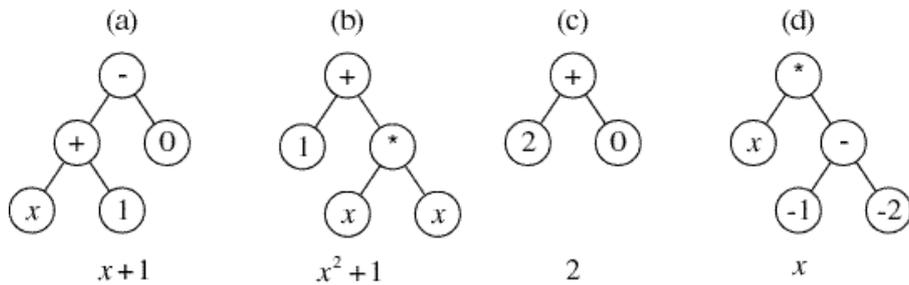


Figure 10: Initial population of four randomly created individuals of generation 0

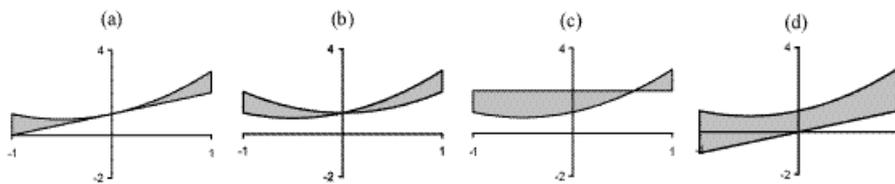


Figure 11: The fitness of each of the four randomly created individuals of generation 0 is equal to the area between two curves.

function (thereby terminating the growth along this path). Finally, the random construction process chose the constant terminal 0 as the second argument of the subtraction function (thereby terminating the entire construction process).

The second program (Figure 10b) adds the constant terminal 1 to the result of multiplying x by x and is equivalent to x^2+1 . The third program (Figure 10c) adds the constant terminal 2 to the constant terminal 0 and is equivalent to the constant value 2. The fourth program (Figure 10d) is equivalent to x .

Randomly created computer programs will, of course, typically be very poor at solving the problem at hand. However, even in a population of randomly created programs, some programs are better than others. The four random individuals from generation 0 in Figure 10 produce outputs that deviate from the

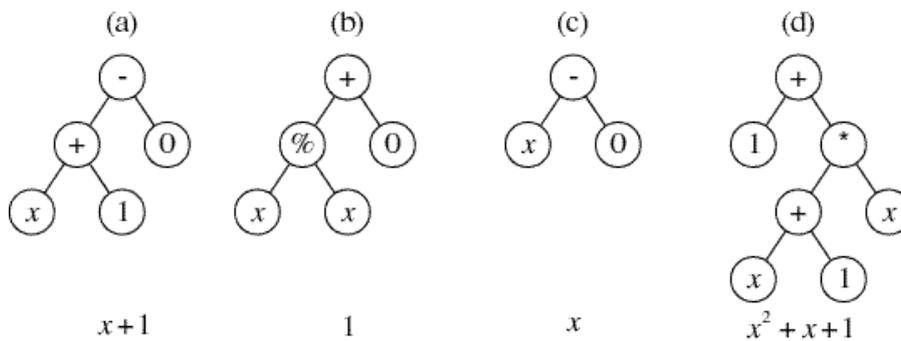


Figure 12: Population of generation 1 (after one reproduction, one mutation, and one two-offspring crossover operation)

output produced by the target quadratic function x^2+x+1 by different amounts. In this particular problem, fitness can be graphically illustrated as the area between two curves. That is, fitness is equal to the area between the parabola x^2+x+1 and the curve representing the candidate individual. Figure 11 shows (as shaded areas) the integral of the absolute value of the errors between each of the four individuals in Figure 10 and the target quadratic function x^2+x+1 . The integral of absolute error for the straight line $x+1$ (the first individual) is 0.67 (Figure 11a). The integral of absolute error for the parabola x^2+1 (the second individual) is 1.0 (Figure 11b). The integrals of the absolute errors for the remaining two individuals are 1.67 (Figure 11c) and 2.67 (Figure 11d), respectively.

As can be seen in Figure 11, the straight line $x+1$ (Figure 11a) is closer to the parabola x^2+x+1 in the range from -1 to $+1$ than any of its three cohorts in the population. This straight line is, of course, not equivalent to the parabola x^2+x+1 . This best-of-generation individual from generation 0 is not even a quadratic function. It is merely the best candidate that happened to emerge from the blind random search of generation 0. In the valley of the blind, the one-eyed man is king.

After the fitness of each individual in the population is ascertained, GP then probabilistically selects relatively more fit programs from the population. The genetic operations are applied to the selected individuals to create offspring programs. The most commonly employed methods for selecting individuals to participate in the genetic operations are tournament selection and fitness-proportionate selection. In both methods, the emphasis is on selecting relatively fit individuals. An important feature common to both methods is that the selection is not greedy. Individuals that are known to be inferior will be selected to a certain degree. The best individual in the population is not guaranteed to be selected. Moreover, the worst individual in the population will not necessarily be excluded. Anything can happen and nothing is guaranteed.

We first perform the reproduction operation. Because the first individual (Figure 10a) is the most fit individual in the population, it is very likely to be selected to participate in a genetic operation. Let's suppose that this particular individual is, in fact, selected for reproduction. If so, it is copied, without alteration, into the next generation (generation 1). It is shown in Figure 12a as part of the population of the new generation.

We next perform the mutation operation. Because selection is probabilistic, it is possible that the third best individual in the population (Figure 10c) is selected. One of the three nodes of this individual is then randomly picked as the site for the mutation. In this example, the constant terminal 2 is picked as the mutation site. This program is then randomly mutated by deleting the entire subtree rooted at the picked point (in this case, just the constant terminal 2) and inserting a subtree that is randomly grown in the same way that the individuals of the initial random population were originally created. In this particular instance, the randomly grown subtree computes the quotient of x and x using the protected division operation `%`. The resulting individual is shown in Figure 12b. This particular mutation changes the original individual from one having a constant value of 2 into one having a constant value of 1. This particular mutation improves fitness from 1.67 to 1.00.

Finally, we perform the crossover operation. Because the first and second individuals in generation 0 are both relatively fit, they are likely to be selected to

participate in crossover. The selection (and reselection) of relatively more fit individuals and the exclusion and extinction of unfit individuals is a characteristic feature of Darwinian selection. The first and second programs are mated sexually to produce two offspring (using the two-offspring version of the crossover operation). One point of the first parent (Figure 10a), namely the $+$ function, is randomly picked as the crossover point for the first parent. One point of the second parent (Figure 10b), namely its leftmost terminal x , is randomly picked as the crossover point for the second parent. The crossover operation is then performed on the two parents. The two offspring are shown in Figures 12c and 12d. One of the offspring (Figure 12c) is equivalent to x and is not noteworthy. However, the other offspring (Figure 12d) is equivalent to x^2+x+1 and has a fitness (integral of absolute errors) of zero. Because the fitness of this individual is below 0.01, the termination criterion for the run is satisfied and the run is automatically terminated. This best-so-far individual (Figure 12d) is designated as the result of the run. This individual is an algebraically correct solution to the problem.

Note that the best-of-run individual (Figure 12d) incorporates a good trait (the quadratic term x^2) from the second parent (Figure 10b) with two other good traits (the linear term x and constant term of 1) from the first parent (Figure 10a). The crossover operation produced a solution to this problem by recombining good traits from these two relatively fit parents into a superior (indeed, perfect) offspring.

In summary, GP has, in this example, automatically created a computer program whose output is equal to the values of the quadratic polynomial x^2+x+1 in the range from -1 to $+1$.

5 Further Features of Genetic Programming

Various advanced features of genetic programming are not covered by the foregoing illustrative problem and the foregoing discussion of the preparatory and executional steps of GP. In this section we will look at a few alternatives (a more complete and detailed survey is available in [115]).

5.1 Automatically Defined Functions and Libraries

Human programmers organise sequences of repeated steps into reusable components such as subroutines, functions and classes. They then repeatedly invoke these components, typically with different inputs. Reuse eliminates the need to “reinvent the wheel” every time a particular sequence of steps is needed. Reuse also makes it possible to exploit a problem’s modularities, symmetries and regularities (thereby potentially accelerating the problem-solving process). This can be taken further, as programmers typically organise these components into hierarchies in which top level components call lower level ones, which call still lower levels, etc. Automatically Defined Functions (ADFs) provide a mechanism by which the evolutionary process can evolve these kinds of potentially reusable components. We will review the basic concepts here, but ADFs are discussed in great detail in [66].

When ADFs are used, a program consists of multiple components. These typically consist of one or more function-defining branches (i.e., ADFs), as well

as one or more main result-producing branches (the *RPB*). The RPB is the “main” program that is executed when the individual is evaluated. It can, however, call the ADFs, which can in turn potentially call each other. A single ADF may be called multiple times by the same RPB, or by a combination of the RPB and other ADFs, allowing the logic that evolution has assembled in that ADF to be re-used in different contexts.

Typically, recursion is prevented by imposing an order on the ADFs within an individual and by restricting calls so that ADF_i can only call ADF_j if $i < j$. Also, in the presence of ADFs, recombination operators are typically constrained to respect the larger structure. That is, during crossover, a subtree from ADF_i can only be swapped with a subtree from another individual’s ADF_i .

The program’s result-producing branch and its ADFs typically have different function and terminal sets. For example, the terminal set for ADFs usually include arguments, such as `arg0`, `arg1`. Typically the user must decide in advance the primitive sets, the number of ADFs and any call restrictions to prevent recursion. However, these choices can be evolved using the architecture-altering operations described in Section 5.2.

There has also been proposals for the automated creation of libraries of functions within GP. For example, [7] and [131] studied the creation and use of dynamic libraries of subtrees taken from parts of the GP trees in the population.

Naturally, while including ADFs and automatically created libraries make it possible for modular re-use to emerge, there is no guarantee that they will be used that way. For example, it may be that the RPB never calls an ADF or only calls it once. It is also possible for an ADF to not actually encapsulate any significant logic.

5.2 Architecture-Altering Operations

The *architecture* of a program can be defined as the total number of trees, the type of each tree, the number of arguments (if any) possessed by each tree, and, finally, if there is more than one tree, the nature of the hierarchical references (if any) allowed among the trees (e.g., whether ADF_1 can call ADF_2) [66].

There are three ways to determine the architecture of the computer programs that will be evolved. Firstly, the user may specify in advance the architecture of the overall program, i.e., perform an *architecture-defining preparatory step* in addition to the five steps itemised in Section 2. Secondly, a run of GP may employ the *evolutionary design of the architecture* [66], thereby enabling the architecture of the overall program to emerge from a competitive process during the run. Finally, a run may employ a set of *architecture-altering operations* [66, 67, 64] which, for example, can create, remove or modify ADFs. Note that architecture changes are often designed not to initially change the semantics of the program and, so, the altered program often has exactly the same fitness as its parent. Nevertheless, the new architecture may make it easier to evolve better programs later.

5.3 Constraining Structures

Most GP systems require that all subtrees return data of the same type. This ensures that the output of any subtree can be used as one of the inputs to any node. The basic subtree crossover operator shuffles tree components entirely

randomly. Type compatibility ensures that crossover cannot lead to incompatible connections between nodes. This is also required to stop mutation from producing illegal programs.

There are cases, however, where this approach is not ideal. For example, there might be constraints on the structure of the acceptable solutions or a problem domain might be naturally represented with multiple types. To apply GP in these cases one needs to be able to use primitives with different type signatures. Below we will look at three approaches to constraining the syntax of the evolved expression trees in GP: simple structure enforcement, strongly typed GP and grammar-based constraints.

If a particular structure is believed or known to be important then one can modify the GP system to require that all individuals have that structure [65]. Enforcing a user specified structure on the evolved solutions can be implemented in a number of ways. For example, one can ensure that all the initial individuals have the structure of interest and then constrain crossover and mutation so that they do not alter any of the fixed regions of a tree. An alternative approach is to evolve the various (sub)components separately. A form of constraint-directed search in GP was also proposed in [160, 159].

Since constraints are often driven by or expressed using a type system, a natural approach is to incorporate types and their constraints into the GP system [99]. In *strongly typed GP*, every terminal has a type, and every function has types for each of its arguments and a type for its return value. The process that generates the initial, random expressions, and all the genetic operators are implemented so as to ensure that they do not violate the type system's constraints. For example, mutation replaces subtrees with new randomly generated trees ensuring that the root of the replacement tree has the same return type as the root of the excised tree. Similarly, crossover only allows the swap of subtrees having the same return type. This basic approach to types can be extended to more complex type systems [99, 49, 104, 169].

Another natural way to express constraints is via *grammars*, and these have been used in GP in a variety of ways [164, 45, 166, 105, 50]. In this sort of system, the grammar is typically used to ensure the initial population is made up of legal programs. The grammar is also used to guide the operations of the genetic operators. Thus we need to keep track not only of the program itself, but also the syntax rules used to derive it.

What actually is evolved in a grammar-based GP system depends on the particular system. [164], for example, evolved *derivation trees*, which effectively are a hierarchical representation of which rewrite rules must be applied, and in which order, to obtain a particular program. In this system, crossover is restricted to only swapping subtrees deriving from a common non-terminal symbol in the grammar. The actual program represented by a derivation tree can be obtained by reading out the leaves of the tree one by one from left to right.

Another approach is *grammatical evolution* (GE) [134, 105] which represents individuals as variable-length sequences of integers which are interpreted in the context of a user supplied grammar. For each rule in the grammar, the set of alternatives on the right hand side are numbered from 0 upwards. To create a program from a GE individual one uses the values in the individual to choose which alternative to take in the production rules. If a value exceeds the number of available options it is transformed via a modulus operation.

5.4 Developmental Genetic Programming

By using appropriate terminals, functions and/or interpreters, GP can go beyond the production of computer programs. In *cellular encoding* [43, 44, 42], programs are interpreted as sequences of instructions which modify (grow) a simple initial structure (embryo). Once the program has finished, the quality of the structure it has produced is measured and this is taken to be the fitness of the program.

Naturally, for cellular encoding to work the primitives of the language must be able to grow structures appropriate to the problem domain. Typical instructions involve the insertion and/or sizing of components, topological modifications of the structure, etc. Cellular encoding GP has successfully been used to evolve neural networks [43, 44, 42] and electronic circuits [72, 71, 70], as well as in numerous other domains. A related approach proposed by [51] combines tree adjoining grammars with L-systems [90] to create a system where each stage in the developmental process is a working program that respects the grammatical constraints.

One of the advantages of indirect representations such as cellular encoding is that the standard GP operators can be used to evolve structures (such as circuits) which may have nothing in common with standard GP trees. In many of these systems, the structures being “grown” are also still meaningful (and evaluable) at each point in their development. This allows fitness evaluation. Another important advantage is that structures resulting from developmental processes often have some regularity, which other methods obtain through the use of ADFs, constraints, types, etc.

5.5 Probabilistic Genetic Programming

Genetic programming typically uses an evolutionary algorithm as its main search engine. However, this is not the only option. This section considers recent work where the exploration is performed by *estimation of distribution algorithms* (EDAs).

EDAs [13, 89] are powerful population-based searchers where the variation operations traditionally implemented via crossover and mutation in EAs are replaced by the process of random sampling from a probability distribution. The distribution is modified generation after generation, using information obtained from the fitter individuals in the population. The objective of these changes in the distribution is to increase the probability of generating individuals with high fitness.

There have been several applications of probabilistic model-based evolution in the areas of tree-based and linear GP. The first EDA-style GP system was effectively an extension of the work in [13] to trees called *probabilistic incremental program evolution* (PIPE) [135]. In PIPE, the population is replaced by a hierarchy of probability tables organised into a tree. Each table represents the probability that a particular primitive will be chosen at that specific location in a newly generated program tree. At each generation a population of programs is created based on the current tree of probability tables. Then, the fitness of the new programs is computed and the probability hierarchy is updated on the basis of these fitnesses, so as to make the generation of above-average fitness programs more likely in the next generation. More recent work includes [168, 93, 94, 119].

A variety of other systems have been proposed which combine the use of grammars and probabilities [141]. For example, [127] used a stochastic context-free grammar to generate program trees where the probability of applying each rewrite rule was adapted using an EDA approach. A probabilistic L-system was used in [140] while a tree-adjunct grammar was used in [1, 139].

5.6 Bloat and Bloat Control

Starting in the early 1990s, researchers began to notice that in addition to progressively increasing their mean and best fitness, GP populations also showed another phenomenon: very often the average size (number of nodes) of the programs in a population after a certain number of generations would start growing at a rapid pace. Typically the increase in program size was not accompanied by any corresponding increase in fitness. This phenomenon is known as *bloat*.

Bloat has significant practical effects: large programs are computationally expensive to evolve and later use, can be hard to interpret, and may exhibit poor generalisation. Note that there are situations where one would expect to see program growth as part of the process of solving a problem. For example, GP runs typically start from populations of small random programs, and it may be necessary for the programs to grow in complexity for them to be able to comply with all the fitness cases. So, we should not equate growth with bloat and we should define *bloat as program growth without (significant) return in terms of fitness*.

Numerous empirical techniques have been proposed to control bloat [88, 144]. In the rest of this section we briefly review some of the most important. In section 7 we will review a subset of theoretical explanations for bloat. More information can be found in [115, 122].

Rather naturally, the first and simplest method to control code growth is the use of hard limits on the size or depth of the offspring programs generated by the genetic operators. Many implementations of this idea (e.g., [65]) apply a genetic operator and then check whether the offspring is beyond the size or depth limit. If it isn't, the offspring enters the population. If, instead, the offspring exceeds the limit, one of the parents is returned. A problem with this implementation is that parent programs that are more likely to violate the size limit will tend to be copied (unaltered) more often than programs that don't. That is, the population will tend to be filled up with programs that nearly infringe the size limit, which is typically not what is desired. However, the problem can be fixed by *not returning parents* if the offspring violates a constraint. Instead one should either return the oversize offspring, but give it a fitness of 0, so that selection will get rid of it at the next generation, or declare the genetic operation failed, and try again.

One can also control bloat by using genetic operators which directly or indirectly have an anti-bloat effect. *Size fair crossover* and *size fair mutation* [86, 26] achieve this by constraining the choices made during the execution of a genetic operation so as to actively prevent growth. In size-fair crossover, for example, the crossover point in the first parent is selected randomly, as in standard crossover. Then the size of the subtree to be excised is calculated. This is used to constrain the choice of the second crossover point so as to guarantee that the subtree chosen from the second parent will not be “unfairly” big. There are also several *mutation operators* that may help control the average tree size in the

population while still introducing new genetic material (e.g., see [61, 77, 63, 5]).

As will be clarified by the size evolution equation presented in Section 7.2, in systems with symmetric operators, bloat can only happen if there are some longer-than-average programs that are fitter than average or some shorter-than-average programs that are less fit than average, or both. So, it stands to reason that in order to control bloat one needs to somehow modulate the selection probabilities of programs based on their size.

A recent technique, the *Tarpeian method* [112], controls bloat by acting directly on selection probabilities. This is done by setting the fitness of randomly chosen longer-than-average programs to 0. This prevents them being parents. By changing how frequently this is done the anti-bloat intensity of Tarpeian control can be modulated. An advantage of the method is that the programs whose fitness is zeroed are never executed, thereby speeding up runs.

The well-known *parsimony pressure* method [65, 170, 171, 172] changes the selection probabilities by subtracting a value based on the size of each program from its fitness. Bigger programs have more subtracted and, so, have lower fitness and tend to have fewer children. That is, the new fitness function is $f(x) - c \times \ell(x)$, where $\ell(x)$ is the size of program x , $f(x)$ is its original fitness and c is a constant known as the *parsimony coefficient*. [171] showed some benefits of adaptively adjusting the coefficient c at each generation but most implementations actually keep the parsimony coefficient constant.

Recently, a theoretically sound method for setting the parsimony coefficient in a principled manner has been proposed [116]. This is called the *covariant parsimony pressure* method. The method is easy to implement. It recalculates the parsimony coefficient c at each generation using $c = \text{Cov}(\ell, f) / \text{Var}(\ell)$, where $\text{Cov}(\ell, f)$ is the covariance between program size ℓ and program fitness f in the population, and $\text{Var}(\ell)$ is the variance of program sizes. Using this equation ensures that the mean program size remains at the value set by the initialisation procedure. There is a variant of the method that allows the user to even decide what function the mean program size should follow over time. As shown in the figure this provides complete control over the population size dynamics.

6 Human-Competitive Results Produced by Genetic Programming

Samuel’s statement (quoted in Section 1) reflects the goal articulated by the pioneers of the 1950s in the fields of artificial intelligence and machine learning, namely to use computers to automatically produce human-like results. Indeed, getting machines to produce human-like results is *the* reason for the existence of the fields of artificial intelligence and machine learning.

To make the notion of human-competitiveness more concrete, we say that a result is “human-competitive” if it satisfies one or more of the eight criteria in Table 1.

As can be seen from Table 1, the eight criteria have the desirable attribute of being at arms-length from the fields of artificial intelligence, machine learning, and GP. That is, a result cannot acquire the rating of “human competitive” merely because it is endorsed by researchers *inside* the specialized fields that are attempting to create machine intelligence. Instead, a result produced by an

Table 1: Eight criteria for saying that an automatically created result is human-competitive

	Criterion
A	The result was patented as an invention in the past, is an improvement over a patented invention, or would qualify today as a patentable new invention.
B	The result is equal to or better than a result that was accepted as a new scientific result at the time when it was published in a peer-reviewed scientific journal.
C	The result is equal to or better than a result that was placed into a database or archive of results maintained by an internationally recognized panel of scientific experts.
D	The result is publishable in its own right as a new scientific result— independent of the fact that the result was mechanically created.
E	The result is equal to or better than the most recent human-created solution to a long-standing problem for which there has been a succession of increasingly better human-created solutions.
F	The result is equal to or better than a result that was considered an achievement in its field at the time it was first discovered.
G	The result solves a problem of indisputable difficulty in its field.
H	The result holds its own or wins a regulated competition involving human contestants (in the form of either live human players or human-written computer programs).

automated method must earn the rating of “human competitive” independent of the fact that it was generated by an automated method.

Since 2004, a competition has been held annually at ACM’s *Genetic and Evolutionary Computation Conference* (termed the Human-Competitive awards – the *Humies*). The \$10,000 prize is awarded to projects that have produced automatically-created human-competitive results according to the criteria in Table 1. Table 2 lists 71 human-competitive instances where GP has produced human-competitive results. Each entry in the table is accompanied by the criteria (from Table 1) that establish the basis for the claim of human-competitiveness or by the Humies competition where they won a prize or received a Honourable mention.

Table 2: Seventy one instances of human-competitive results produced by genetic programming

	Claimed instance	Basis for claim
1	Creation of a better-than-classical quantum algorithm for the Deutsch-Jozsa “early promise” problem [146]	B, F
2	Creation of a better-than-classical quantum algorithm for Grover’s database search problem [148]	B, F
3	Creation of a quantum algorithm for the depth-two AND/OR query problem that is better than any previously published result [147, 16]	D

	Claimed instance	Basis for claim
4	Creation of a quantum algorithm for the depth-one OR query problem that is better than any previously published result[16]	D
5	Creation of a protocol for communicating information through a quantum gate that was previously thought not to permit such communication [149]	D
6	Creation of a novel variant of quantum dense coding [149]	D
7	Creation of a soccer-playing program that won its first two games in the Robo Cup 1997 competition [95]	H
8	Creation of a soccer-playing program that ranked in the middle of the field of 34 human-written programs in the Robo Cup 1998 competition [3]	H
9	Creation of four different algorithms for the transmembrane segment identification problem for proteins [66, sections 18.8 and 18.10] and [70, sections 16.5 and 17.2]	B, E
10	Creation of a sorting network for seven items using only 16 steps [70, sections 21.4.4, 23.6, and 57.8.1]	A, D
11	Rediscovery of the Campbell ladder topology for lowpass and highpass filters [70, section 25.15.1] and [73, section 5.2]	A, F
12	Rediscovery of the Zobel “ M -derived half section” and “constant K ” filter sections [70, section 25.15.2]	A, F
13	Rediscovery of the Cauer (elliptic) topology for filters [70, section 27.3.7]	A, F
14	Automatic decomposition of the problem of synthesizing a crossover filter [70, section 32.3]	A, F
15	Rediscovery of a recognizable voltage gain stage and a Darlington emitter-follower section of an amplifier and other circuits [70, section 42.3]	A, F
16	Synthesis of 60 and 96 decibel amplifiers [70, section 45.3]	A, F
17	Automatic synthesis of asymmetric bandpass filter [72]	
18	Synthesis of analog computational circuits for squaring, cubing, square root, cube root, logarithm, and Gaussian functions [70, section 47.5.3]	A, D, G
19	Synthesis of a real-time analog circuit for time-optimal control of a robot [70, section 48.3]	G
20	Synthesis of an electronic thermometer [70, section 49.3]	A, G
21	Synthesis of a voltage reference circuit [70, section 50.3]	A, G
22	Automatic synthesis of digital-to-analog converter (DAC) circuit [17]	
23	Automatic synthesis of analog-to-digital (ADC) circuit [17]	
24	Creation of a cellular automata rule for the majority classification problem that is better than the Gacs-Kurdyumov-Levin (GKL) rule and all other known rules written by humans [4] and [70, section 58.4]	D, E
25	Creation of motifs that detect the D-E-A-D box family of proteins and the manganese superoxide dismutase family [70, section 59.8]	C

	Claimed instance	Basis for claim
26	Synthesis of topology for a PID-D2 (proportional, integrative, derivative, and second derivative) controller [73, section 3.7]	A, F
27	Synthesis of an analog circuit equivalent to Philbrick circuit [73, section 4.3]	A, F
28	Synthesis of NAND circuit [73, section 4.4]	A, F
29	Simultaneous synthesis of topology, sizing, placement, and routing of analog electrical circuits [73, chapter 5]	
30	Synthesis of topology for a PID (proportional, integrative, and derivative) controller [73, section 9.2]	A, F
31	Rediscovery of negative feedback [73, chapter 14]	A, E, F, G
32	Synthesis of a low-voltage balun circuit [73, section 15.4.1]	A
33	Synthesis of a mixed analog-digital variable capacitor circuit [73, section 15.4.2]	A
34	Synthesis of a high-current load circuit [73, section 15.4.3]	A
35	Synthesis of a voltage-current conversion circuit [73, section 15.4.4]	A
36	Synthesis of a cubic signal generator [73, section 15.4.5]	A
37	Synthesis of a tunable integrated active filter [73, section 15.4.6]	A
38	Creation of PID tuning rules that outperform the Ziegler-Nichols and Astrom-Hagglund tuning rules [73, chapter 12]	A, B, D, E, F, G
39	Creation of three non-PID controllers that outperform a PID controller that uses the Ziegler-Nichols or Astrom-Hagglund tuning rules [73, chapter 13]	A, B, D, E, F, G
40	An Evolved Antenna for Deployment on NASA's Space Technology 5 Mission [92]	Humies 2004
41	Automatic Quantum Computer Programming: A Genetic Programming Approach [145]	Humies 2004
42	Evolving Local Search Heuristics for SAT Using Genetic Programming [39]; Automated discovery of composite SAT variable-selection heuristics [38]	Humies 2004
43	How to Draw a Straight Line Using a GP: Benchmarking Evolutionary Design Against 19th Century Kinematic Synthesis [91]	Humies 2004
44	Organization Design Optimization Using Genetic Programming [60]	Humies 2004
45	Discovery of Human-Competitive Image Texture Feature Programs Using Genetic programming [76]	Humies 2004
46	Novel image filters implemented in hardware [138]	Humies 2004
47	Automated Re- Invention of Six Patented Optical Lens Systems using Genetic Programming: two telescope eyepieces, a telescope eyepiece system, an eyepiece for optical instruments, two wide angle eyepieces, and a telescope eyepiece [68, 69]	Humies 2005

	Claimed instance	Basis for claim
48	Evolution of a Human- Competitive Quantum Fourier Transform Algorithm Using Genetic Programming [97]	Humies 2005
49	Evolving Assembly Programs: How Games Help Microprocessor Validation [25]	Humies 2005
50	Attaining Human-Competitive Game Playing with Genetic Programming [143]; GP-Gammon: Using Genetic Programming to Evolve Backgammon Players [9]; GP-Gammon: Genetically Programming Backgammon Players [8]	Humies 2005
51	GP-EndChess: Using Genetic Programming to Evolve Chess Endgame [47]	Humies 2005
52	GP-Robocode: Using Genetic Programming to Evolve Robocode Players [142]	Humies 2005
53	Evolving Dispatching Rules for Solving the Flexible Job-Shop Problem [156]	Humies 2005
54	Solution of differential equations with Genetic Programming and the Stochastic Bernstein Interpolation [54]	Humies 2005
55	Determining Equations for Vegetation Induced Resistance using Genetic Programming [59]	Humies 2005
56	Sallen-Key filter [74]	
57	Using Evolution to Learn How to Perform Interest Point Detection [157]; Synthesis of Interest Point Detectors Through Genetic Programming [158]	Humies 2006
58	Evolution of an Efficient Search Algorithm for the Mate-In-N Problem in Chess [48]	Humies 2007
59	Evolving local and global weighting schemes in Information Retrieval [28]; An analysis of the Solution Space for Genetically Programmed Term-Weighting Schemes in Information Retrieval [27]; Term-Weighting in Information Retrieval using Genetic Programming: A Three-Stage Process [29]	Humies 2007
60	Real-Time, Non- Intrusive Evaluation of VoIP [126]	Humies 2007
61	Automated Reverse Engineering of nonlinear Dynamical Systems [20]	Humies 2007
62	Genetic programming approach for electron-alkalimetal atom collisions [125, 35]; Prediction of Non-Linear System in Optics Using Genetic Programming [124]; Genetic programming approach for flow of steady state fluid between two eccentric spheres [34]	Humies 2007
63	Genetic Programming for Finite Algebras [150]	Humies 2008
64	Automatic synthesis of quantum computing circuit for the two-oracle AND/OR problem [151]	
65	Automatic synthesis of quantum computing algorithms for the parity problem a special case of the hidden subgroup problem [153]	
66	Automatic synthesis of mechanical vibration absorbers [55]	
67	Automatically finding patches and automated software repair [100, 163]	Humies 2009

	Claimed instance	Basis for claim
68	GP-Rush: Using Genetic Programming to Evolve Solvers for the Rush Hour Puzzle [46]	Humies 2009
69	Learning Invariant Region Descriptor Operators with Genetic Programming and the F-measure [106]; Evolutionary Learning of Local Descriptor Operators for Object Recognition [107]	Humies 2009
70	Solution of matrix Riccati differential equation for nonlinear singular system using genetic programming [11]	
71	Distilling Free- Form Natural Laws from Experimental Data [136, 137]	

Clearly Table 2 shows GP’s potential as a powerful invention machine. There are 31 instances where the human-competitive result produced by GP duplicated the functionality of a previously patented invention, infringed a previously issued patent, or created a patentable new invention. These include one instance where genetic programming has created an entity that either infringes or duplicates the functionality of a previously patented 19th-century invention, 21 instances where GP has done the same with respect to previously patented 20th-century inventions, 7 instances where genetic programming has done the same with respect to previously patented 21st-century inventions, and two instances where GP has created a patentable new invention. The two new inventions are general-purpose controllers that outperform controllers employing tuning rules that have been in widespread use in industry for most of the 20th century.

7 Genetic Programming Theory

GP is a search technique that explores the space of computer programs. As discussed above, the search for solutions to a problem starts from a group of points (random programs) in this search space. Those points that are of above average quality are then used to generate a new generation of points through crossover, mutation, reproduction and possibly other genetic operations. This process is repeated over and over again until a termination criterion is satisfied.

If we could visualize this search, we would often find that initially the population looks a bit like a cloud of randomly scattered points, but that, generation after generation, this cloud changes shape and moves in the search space following a well defined trajectory. Because GP is a stochastic search technique, in different runs we would observe different trajectories. These, however, would very likely show clear regularities to our eye that could provide us with a deep understanding of how the algorithm is searching the program space for the solutions to a given problem. We could probably readily see, for example, why GP is successful in finding solutions in certain runs and with certain parameter settings, and unsuccessful in/with others.

Unfortunately, it is normally impossible to exactly visualize the program search space due to its high dimensionality and complexity, and so we cannot just use our senses to understand and predict the behavior of GP.

One approach to gain an understanding of the behavior of a genetic programming system and predict its behaviour in precise terms is to define and study

mathematical models of evolutionary search. There are a number of cases where this approach has been very successful in illuminating some of the fundamental processes and biases in GP systems. In this section we will review some theoretical approaches to understanding GP. The reader is referred to [122, 84, 115] for a more extensive review of GP theory.

7.1 Models of GP Search

Schema theories are among the oldest and the best known models of evolutionary algorithms [53, 165]. Schema theories are based on the idea of partitioning the search space into subsets, called *schemata*. They are concerned with modelling and explaining the dynamics of the distribution of the population over the schemata. Modern genetic algorithm schema theory [154, 155] provides exact information about the distribution of the population at the next generation in terms of quantities measured at the current generation, without having to actually run the algorithm. Exact schema theories are also available for GP systems with a variety of genetic operators (e.g., see [108, 109, 110, 84, 120, 117, 118]). Markov chain theory has also started being applied to GP [121, 120, 98], although so far this hasn't been developed as fully as the schema theory.

Exact mathematical models of GP, such as schema theories and Markov chains, are probabilistic descriptions of the operations of selection, reproduction, crossover and mutation. They explicitly represent how these operations determine which areas of the program space will be sampled by GP, and with what probability. These models treat the fitness function as a black box, however. That is, there is no representation of the fact that in GP, unlike in other evolutionary techniques, the fitness function involves the execution of computer programs on a variety of inputs. In other words, schema theories and Markov chains do not tell us how fitness is distributed in the search space. Yet, without this information, we have no way of closing the loop and fully characterising the behaviour of a GP systems which is always the result of the interaction between the fitness function and the search biases of the representation and genetic operations used in the system.

Fortunately, the characterisation of the space of computer programs explored by GP has been another main topic of theoretical research [84]. In this category are theoretical results showing that the distribution of functionality of non Turing-complete programs approaches a limit as program length increases. That is, although the number of programs of a particular length grows exponentially with length, beyond a certain threshold the fraction of programs implementing any particular functionality is effectively constant. There is a very substantial body of empirical evidence indicating that this happens in a variety of systems. In fact, there are also mathematical proofs of these convergence results for two important forms of programs: Lisp (tree-like) S-expressions (without side effects) and machine code programs without loops [84, 78, 79, 81, 80, 82]. Also, recently, [83, 113] started extending these results to Turing complete machine code programs.

7.2 Bloat

In Section 5.6 we introduced the notion of bloat and described some effective mechanisms for controlling it. Below we review a subset of theoretical models

and explanations for bloat. More information can be found in [122, 115, 84].

There have been some efforts to approximately mathematically model bloat. For example, [14] defined an *executable model of bloat* where only the fitness, the size of active code and the size of inactive code were represented (i.e., there was no representation of program structures). Fitnesses of individuals were drawn from a bell-shaped distribution, while active and inactive code lengths were modified by a size-unbiased mutation operator. Various interesting effects were reported which are very similar to corresponding effects found in GP runs. [130] proposed a similar, but slightly more sophisticated model which also included an analogue of crossover.

A strength of these executable models is their simplicity. A weakness is that they suppress or remove many details of the representation and operators typically used in GP. This makes it difficult to verify if all the phenomena observed in the model have analogues in GP runs, and if all important behaviours of GP in relation to bloat are captured by the model.

In [111, 118], a *size evolution equation* for GP was developed, which provided an exact formalisation of the dynamics of average program size. The equation has recently been simplified [116] giving:

$$E[\mu(t+1) - \mu(t)] = \sum_{\ell} \ell \times (p(\ell, t) - \Phi(\ell, t)), \quad (1)$$

where $\mu(t+1)$ is the mean size of the programs in the population at generation $t+1$, E is the expectation operator, ℓ is a program size, $p(\ell, t)$ is the probability of selecting programs of size ℓ from the population in generation t , and $\Phi(\ell, t)$ is the proportion of programs of size ℓ in the current generation. The equation applies to a GP system with selection and any form of symmetric subtree crossover.¹ Note that the equation constrains what can and cannot happen size-wise in GP populations. Any explanation for bloat has to agree with it.

In particular, Equation (1) shows that there can be bloat only if the selection probability $p(\ell, t)$ is different from the proportion $\Phi(\ell, t)$ for at least some ℓ . So, for bloat to happen there will have to be some small ℓ 's for which $p(\ell, t) < \Phi(\ell, t)$ and also some bigger ℓ 's for which $p(\ell, t) > \Phi(\ell, t)$ (at least on average).

We conclude this review on theories of bloat with a recent promising explanation for bloat called the *crossover bias theory* [114, 30], which is based on and is consistent with Equation (1). The theory goes as follows. On average, each application of subtree crossover removes as much genetic material as it inserts; consequently crossover on its own does not produce growth or shrinkage. While the *mean* program size is unaffected, however, *higher moments* of the distribution are. In particular, crossover pushes the population towards a particular distribution of program sizes, known as a *Lagrange distribution of the second kind*, where small programs have a much higher frequency than longer ones. For example, crossover generates a very high proportion of single-node individuals. In virtually all problems of practical interest, however, very small programs have no chance of solving the problem. As a result, programs of above average size have a selective advantage over programs of below average size, and the mean program size increases. Because crossover will continue to create small programs, which will then be ignored by selection (in favour of

¹In a symmetric operator the probability of selecting particular crossover points in the parents does not depend on the order in which the parents are drawn from the population.

the larger programs), the increase in average size will continue generation by generation.

8 Conclusions

In his seminal 1948 paper entitled “Intelligent Machinery,” Turing identified three ways by which human-competitive machine intelligence might be achieved. In connection with one of those ways, Turing [161] said:

“There is the genetical or evolutionary search by which a combination of genes is looked for, the criterion being the survival value.”

Turing did not specify how to conduct the “genetical or evolutionary search” for machine intelligence. In particular, he did not mention the idea of a population-based parallel search in conjunction with sexual recombination (crossover) as described in John Holland’s 1975 book *Adaptation in Natural and Artificial Systems*. However, in his 1950 paper “Computing Machinery and Intelligence,” Turing [162] did point out

“We cannot expect to find a good child-machine at the first attempt. One must experiment with teaching one such machine and see how well it learns. One can then try another and see if it is better or worse. There is an obvious connection between this process and evolution, by the identifications

“Structure of the child machine = Hereditary material

“Changes of the child machine = Mutations

“Natural selection” = Judgment of the experimenter”

That is, Turing perceived in 1948 and 1950 that one possibly productive approach to machine intelligence would involve an evolutionary process in which a description of a computer program (the hereditary material) undergoes progressive modification (mutation) under the guidance of natural selection (i.e., selective pressure in the form of what we now call “fitness”).

Today, many decades later, we can see that indeed Turing was right. Genetic programming has started fulfilling Turing’s dream by providing us with a systematic method, based on Darwinian evolution, for getting computers to automatically solve hard real-life problems. To do so, it simply requires a high-level statement of what needs to be done (and enough computing power).

Turing also understood the need to evaluate objectively the behaviour exhibited by machines, to avoid human biases when assessing their intelligence. This led him to propose an imitation game, now known as the *Turing test for machine intelligence*, whose goals are wonderfully summarised by Arthur Samuel’s position statement quoted in the introduction of this chapter.

At present GP is certainly not in a position to produce computer programs that would pass the full Turing test for machine intelligence, and it might not be ready for this immense task for centuries. Nonetheless, thanks to the constant technological improvements in GP technology, in its theoretical foundations and in computing power, GP has been able to solve tens of difficult problems with human-competitive results (see Table 2) in the recent past. These are a small step towards fulfilling Turing and Samuel’s dreams, but they are also early signs

of things to come. It is, indeed, arguable that in a few years' time GP will be able to *routinely* and *competently* solve important problems for us in a variety of specific domains of application, even when running on a single personal computer, thereby becoming an essential collaborator for many of human activities. This, we believe, will be a remarkable step forward towards achieving true, human-competitive machine intelligence.

Acknowledgements

We would like to thank the editors of this volume for giving us the opportunity to extend and revise this chapter. Their patience and flexibility has also been much appreciated.

We would also like to thank Bill Langdon and Nic McPhee for allowing us to draw some inspiration and to use some text from the *Field Guide to Genetic Programming* book [115].

A Tricks of the Trade and Resources

Newcomers to the field of GP often ask themselves (and/or other more experienced genetic programmers) questions such as the following:

1. Will GP be able to solve my problem?
2. What is the best way to get started with GP? Which books or papers should I read?
3. Should I implement my own GP system or should I use an existing package? If so, what package should I use?

In this appendix we will try to answer these questions (and many more) by considering the ingredients of successful GP applications (Section A.1) and reviewing some of the wide variety of available sources on GP which should assist readers who wish to explore further (Sections A.2–A.4).

A.1 Application Areas where GP is Likely to Excel

Based on the experience of numerous researchers over many years, it appears that GP and other evolutionary computation methods have been especially productive in areas having some or all of the following properties:

- The interrelationships among the relevant variables is unknown or poorly understood (or where it is suspected that the current understanding may possibly be wrong).
- Finding the size and shape of the ultimate solution is a major part of the problem.
- Significant amounts of test data are available in computer-readable form.
- There are good simulators to test the performance of tentative solutions to a problem, but poor methods to directly obtain good solutions.

- Conventional mathematical analysis does not, or cannot, provide analytic solutions.
- An approximate solution is acceptable (or is the only result that is ever likely to be obtained).
- Small improvements in performance are routinely measured (or easily measurable) and highly prized.

A.2 Key Books and Journals

There are more than 30 books written in English principally on GP or its applications with more being written. These start with Koza's 1992 book *Genetic Programming* (often referred to as Jaws). Koza has subsequently published three additional books on GP: *Genetic Programming II: Automatic Discovery of Reusable Programs* (1994) deals with ADFs; *Genetic Programming 3* (1999) covers, in particular, the evolution of analogue circuits; *Genetic Programming 4* (2003) uses GP for automatic invention. MIT Press published three volumes in the series *Advances in Genetic Programming* [62, 6, 152]. The joint GP / genetic algorithms Kluwer book series now contains over 10 books starting with *Genetic Programming and Data Structures* [85]. Apart from Jaws, these tend to be for the GP specialist. The late 1990s saw the introduction of the first textbook dedicated to GP [15].

The 2008 book *A Field Guide to Genetic Programming* [115] provides a gentle introduction to GP as well as a review of its different flavours and application domains. The book is freely available on the Internet in PDF and HTML formats.

Other titles include: *Genetic Programming* (in Japanese) [56], *Principia Evolvica – Simulerte Evolution mit Mathematica* (in German) [57] (English version [58]), *Data Mining Using Grammar Based Genetic Programming and Applications* [167], *Grammatical Evolution: Evolutionary Automatic Programming in a Arbitrary Language* [105], *Humanoider: Sjavlarande robotar och artificiell intelligens* (in Swedish) [103], and *Linear Genetic Programming* [22].

Readers interested in mathematical and empirical analyses of GP behaviour may find *Foundations of Genetic Programming* [84] useful.

Each of Koza's four books has an accompanying video. These videos are now available in DVD format. Also, a small set of videos on specific GP techniques and applications is available via on-line resources such as Google Video and YouTube.

In addition to GP's own *Genetic Programming and Evolvable Machines* journal, *Evolutionary Computation*, the *IEEE transaction on Evolutionary Computation*, *Complex Systems* (Complex Systems Publication, Inc.), and many others publish GP articles. The GP bibliography [87] lists several hundred different journals worldwide that have published articles related to GP.

A.3 GP Implementations

One of the reasons behind the success of GP is that it is easy to implement own versions, and implementing a simple GP system from scratch remains an excellent way to make sure one really understands the mechanics of GP. In addition to being an exceptionally useful exercise, it is often easier to customise

(e.g., adding new, application specific genetic operators or implementing unusual, knowledge-based initialisation strategies) a system one has built for new purposes than a large GP distribution. All of this, however, requires reasonable programming skills and the will to thoroughly test the resulting system until it behaves as expected.

This is actually an extremely tricky issue in highly stochastic systems such as GP. The problem is that almost any system will produce “interesting” behaviour, but it is typically very hard to test whether it is exhibiting the *correct* interesting behaviour. It is remarkably easy for small mistakes to go unnoticed for extended periods of time (even years). It is also easy to incorrectly assume that “minor” implementation decisions will not significantly affect the behaviour of the system.

An alternative is to use one of the many public domain GP implementations and adapt this for one’s purposes. This process is faster, and good implementations are often robust, efficient, well documented and comprehensive. The small price to pay is the need to study the available documentation and examples. These often explain how to modify the GP system to some extent. However, deeper modifications (such as the introduction of new or unusual operators) will often require studying the actual source code and a substantial amount of trial and error. Good publicly available GP implementations include: Lil-GP [123], ECJ [96], Open Beagle [40] and GPC++ [37]. The most prominent commercial implementation remains Discipulus [129]; see [36] for a review.

While the earliest GP systems were implemented in Lisp, people have since coded GP in a huge range of different languages, including C/C++, Java, Python, JavaScript, Perl, Prolog, Mathematica, Pop-11, MATLAB, Fortran, Occam and Haskell. Typically, these evolve expressions and programs which look like simplified Lisp. More complex target languages can be supported, however, especially with the use of more advanced tools such as grammars and type systems. Conversely, many successful programs in machine code or low-level languages have also climbed from the primordial ooze of initial randomness.

A.4 On-Line Resources

On-line resources appear, disappear, and move with great speed, so the addresses here, which were correct at the time of writing, are obviously subject to change without notice after publication. Hopefully, the most valuable resources should be readily findable using standard search tools.

One of the key on-line resources is the GP bibliography [87] available from <http://www.cs.bham.ac.uk/~wbl/biblio/> At the time of writing, this bibliography contains over 6,000 GP entries, roughly half of which can be downloaded immediately.

The GP bibliography has a variety of interfaces. It allows for quick jumps between papers linked by authors and allows one to sort the author list by the number of GP publications. Full references are provided in both BIB_TE_X and Refer formats for direct inclusion in papers written in L^AT_EX and Microsoft Word, respectively. The GP bibliography is also part of the Collection of Computer Sciences Bibliographies [2], which provides a comprehensive Lucerne syntax search engine.

From early on there has been an active, open email discussion list: the GP-list [41]. The [32] is a moderated list covering evolutionary computation more

broadly, and often contains GP related announcements.

Koza's <http://www.genetic-programming.org/> contains a ton of useful information for the novice, including a short tutorial on "What is Genetic Programming" and the Lisp implementation of GP from *Genetic Programming* [65].

References

- [1] H. Abbass, N. Hoai, and R. McKay. Anttag: A new method to compose computer programs using colonies of ants. In *IEEE Congress on Evolutionary Computation, 2002.*, 2002.
- [2] Alf-Christian Achilles and Paul Ortyl. The Collection of Computer Science Bibliographies, 1995-2008.
- [3] D. Andre and A. Teller. Evolving Team Darwin United. In M. Asada and H. Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *LNCS*, pages 346–351, Paris, France, July 1998 1999. Springer Verlag.
- [4] David Andre, Forrest H Bennett III, and John R. Koza. Discovery by genetic programming of a cellular automata rule that is better than any known rule for the majority classification problem. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 3–11, Stanford University, CA, USA, 28–31 July 1996. MIT Press.
- [5] Peter J. Angeline. An investigation into the sensitivity of genetic programming to the frequency of leaf selection during subtree crossover. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 21–29, Stanford University, CA, USA, 28–31 July 1996. MIT Press.
- [6] Peter J. Angeline and K. E. Kinnear, Jr., editors. *Advances in Genetic Programming 2*. MIT Press, Cambridge, MA, USA, 1996.
- [7] Peter J. Angeline and Jordan B. Pollack. The evolutionary induction of subroutines. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 236–241, Bloomington, Indiana, USA, 1992. Lawrence Erlbaum.
- [8] Yaniv Azaria and Moshe Sipper. GP-gammon: Genetically programming backgammon players. *Genetic Programming and Evolvable Machines*, 6(3):283–300, September 2005. Published online: 12 August 2005.
- [9] Yaniv Azaria and Moshe Sipper. GP-gammon: Using genetic programming to evolve backgammon players. In Maarten Keijzer, Andrea Tettamanzi, Pierre Collet, Jano I. van Hemert, and Marco Tomassini, editors, *Proceedings of the 8th European Conference on Genetic Programming*, volume 3447 of *Lecture Notes in Computer Science*, pages 132–142, Lausanne, Switzerland, 30 March - 1 April 2005. Springer.
- [10] Vladan Babovic. *Emergence, evolution, intelligence; Hydroinformatics - A study of distributed and decentralised computing using intelligent agents*. A. A. Balkema Publishers, Rotterdam, Holland, 1996.

- [11] P. Balasubramaniam and A. Vincent Antony Kumar. Solution of matrix Riccati differential equation for nonlinear singular system using genetic programming. *Genetic Programming and Evolvable Machines*, 10(1):71–89, March 2009.
- [12] Joze Balic. *Flexible Manufacturing Systems; Development - Structure - Operation - Handling - Tooling*. Manufacturing technology. DAAAM International, Vienna, 1999.
- [13] S. Baluja and R. Caruana. Removing the genetics from the standard genetic algorithm. In A. Prieditis and S. Russell, editors, *Machine Learning: Proceedings of the Twelfth International Conference*, pages 38–46. Morgan Kaufmann Publishers, San Francisco, CA, 1995.
- [14] W. Banzhaf and W. B. Langdon. Some considerations on the reason for bloat. *Genetic Programming and Evolvable Machines*, 3(1):81–91, March 2002.
- [15] Wolfgang Banzhaf, Peter Nordin, Robert E. Keller, and Frank D. Francone. *Genetic Programming – An Introduction; On the Automatic Evolution of Computer Programs and its Applications*. Morgan Kaufmann, San Francisco, CA, USA, January 1998.
- [16] Howard Barnum, Herbert J Bernstein, and Lee Spector. Quantum circuits for OR and AND of ORs. *Journal of Physics A: Mathematical and General*, 33(45):8047–8057, 17 November 2000.
- [17] Forrest H Bennett III, John R. Koza, Martin A. Keane, Jessen Yu, William Myrdlowec, and Oscar Stiffelman. Evolution by means of genetic programming of analog circuits that perform digital functions. In Wolfgang Banzhaf, Jason Daida, Agoston E. Eiben, Max H. Garzon, Vasant Honavar, Mark Jakiela, and Robert E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1477–1483, Orlando, Florida, USA, 13–17 July 1999. Morgan Kaufmann.
- [18] Bir Bhanu, Yingqiang Lin, and Krzysztof Krawiec. *Evolutionary Synthesis of Pattern Recognition Systems*. Monographs in Computer Science. Springer-Verlag, New York, 2005.
- [19] Tobias Blickle. *Theory of Evolutionary Algorithms and Application to System Synthesis*. PhD thesis, Swiss Federal Institute of Technology, Zurich, November 1996.
- [20] Josh Bongard and Hod Lipson. Automated reverse engineering of nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 104(24):9943–9948, June 12 2007.
- [21] Anthony Brabazon and Michael O’Neill. *Biologically Inspired Algorithms for Financial Modelling*. Natural Computing Series. Springer, 2006.
- [22] Markus Brameier and Wolfgang Banzhaf. *Linear Genetic Programming*. Number XVI in Genetic and Evolutionary Computation. Springer, 2007.

- [23] Miran Brezocnik. *Uporaba genetskega programiranja v inteligentnih proizvodnih sistemih*. University of Maribor, Faculty of mechanical engineering, Maribor, Slovenia, 2000. In Slovenian.
- [24] Shu-Heng Chen, editor. *Genetic Algorithms and Genetic Programming in Computational Finance*. Kluwer Academic Publishers, Dordrecht, July 2002.
- [25] F. Corno, E. Sanchez, and G. Squillero. Evolving assembly programs: how games help microprocessor validation. *Evolutionary Computation, IEEE Transactions on*, 9(6):695–706, 2005.
- [26] Raphael Crawford-Marks and Lee Spector. Size control via size fair genetic operators in the PushGP genetic programming system. In W. B. Langdon, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. A. Potter, A. C. Schultz, J. F. Miller, E. Burke, and N. Jonoska, editors, *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 733–739, New York, 9-13 July 2002. Morgan Kaufmann Publishers.
- [27] Ronan Cummins and Colm O’Riordan. An analysis of the solution space for genetically programmed term-weighting schemes in information retrieval. In D. A. Bell, editor, *17th Irish Artificial Intelligence and Cognitive Science Conference (AICS 2006)*, Queen’s University, Belfast, 11th-13th September 2006.
- [28] Ronan Cummins and Colm O’Riordan. Evolving local and global weighting schemes in information retrieval. *Information Retrieval*, 9(3):311–330, June 2006.
- [29] Ronan Cummins and Colm O’Riordan. Term-weighting in information retrieval using genetic programming: A three stage process. In Gerhard Brewka, Silvia Coradeschi, Anna Perini, and Paolo Traverso, editors, *The 17th European Conference on Artificial Intelligence, ECAI-2006*, pages 793–794, Riva del Garda, Italy, August 28th - September 1st 2006. IOS Press.
- [30] Stephen Dignum and Riccardo Poli. Generalisation of the limiting distribution of program sizes in tree-based genetic programming and analysis of its effects on bloat. In Dirk Thierens, Hans-Georg Beyer, Josh Bongard, Jurgen Branke, John Andrew Clark, Dave Cliff, Clare Bates Congdon, Kalyanmoy Deb, Benjamin Doerr, Tim Kovacs, Sanjeev Kumar, Julian F. Miller, Jason Moore, Frank Neumann, Martin Pelikan, Riccardo Poli, Kumara Sastry, Kenneth Owen Stanley, Thomas Stutzle, Richard A Watson, and Ingo Wegener, editors, *GECCO ’07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, volume 2, pages 1588–1595, London, 7-11 July 2007. ACM Press.
- [31] Dimitris C. Dracopoulos. *Evolutionary Learning Algorithms for Neural Adaptive Control*. Perspectives in Neural Computing. Springer Verlag, P.O. Box 31 13 40, D-10643 Berlin, Germany, August 1997.
- [32] EC-Digest, 1985-2008.

- [33] A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer, 2003.
- [34] Mostafa Y. El-Bakry and Amr Radi. Genetic programming approach for flow of steady state fluid between two eccentric spheres. *Applied Rheology*, 17, 2007.
- [35] Salah Yaseen El-Bakry and Amr. Radi. Genetic programming approach for electron-alkali-metal atom collisions. *International Journal of Modern Physics B*, 20(32):5463–5471, 2006.
- [36] James A. Foster. Review: Discipulus: A commercial genetic programming system. *Genetic Programming and Evolvable Machines*, 2(2):201–203, June 2001.
- [37] Adam Fraser and Thomas Weinbrenner. GPC++ Genetic Programming C++ Class Library, 1993-1997.
- [38] Alex Fukunaga. Automated discovery of composite SAT variable selection heuristics. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 641–648, 2002.
- [39] Alex S. Fukunaga. Evolving local search heuristics for SAT using genetic programming. In Kalyanmoy Deb, Riccardo Poli, Wolfgang Banzhaf, Hans-Georg Beyer, Edmund Burke, Paul Darwen, Dipankar Dasgupta, Dario Floreano, James Foster, Mark Harman, Owen Holland, Pier Luca Lanzi, Lee Spector, Andrea Tettamanzi, Dirk Thierens, and Andy Tyrrell, editors, *Genetic and Evolutionary Computation – GECCO-2004, Part II*, volume 3103 of *Lecture Notes in Computer Science*, pages 483–494, Seattle, WA, USA, 26-30 June 2004. Springer-Verlag.
- [40] Christian Gagné and Marc Parizeau. Open BEAGLE: A new C++ evolutionary computation framework. In W. B. Langdon, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. A. Potter, A. C. Schultz, J. F. Miller, E. Burke, and N. Jonoska, editors, *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, page 888, New York, 9-13 July 2002. Morgan Kaufmann Publishers.
- [41] Genetic Programming mailing list, 2001-2008.
- [42] F. Gruau. *Neural Network Synthesis using Cellular Encoding and the Genetic Algorithm*. PhD thesis, Laboratoire de l’Informatique du Parallélisme, Ecole Normale Supérieure de Lyon, France, 1994.
- [43] F. Gruau and D. Whitley. Adding learning to the cellular development process: a comparative study. *Evolutionary Computation*, 1(3):213–233, 1993.
- [44] Frederic Gruau. Genetic micro programming of neural networks. In Kenneth E. Kinneer, Jr., editor, *Advances in Genetic Programming*, chapter 24, pages 495–518. MIT Press, 1994.

- [45] Frederic Gruau. On using syntactic constraints with genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 19, pages 377–394. MIT Press, Cambridge, MA, USA, 1996.
- [46] Ami Hauptman, Achiya Elyasaf, Moshe Sipper, and Assaf Karmon. GP-rush: using genetic programming to evolve solvers for the rush hour puzzle. In Guenther Raidl, Franz Rothlauf, Giovanni Squillero, Rolf Drechsler, Thomas Stuetzle, Mauro Birattari, Clare Bates Congdon, Martin Middendorf, Christian Blum, Carlos Cotta, Peter Bosman, Joern Grahl, Joshua Knowles, David Corne, Hans-Georg Beyer, Ken Stanley, Julian F. Miller, Jano van Hemert, Tom Lenaerts, Marc Ebner, Jaume Bacardit, Michael O’Neill, Massimiliano Di Penta, Benjamin Doerr, Thomas Jansen, Riccardo Poli, and Enrique Alba, editors, *GECCO ’09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 955–962, Montreal, 8-12 July 2009. ACM.
- [47] Ami Hauptman and Moshe Sipper. GP-endchess: Using genetic programming to evolve chess endgame players. In Maarten Keijzer, Andrea Tettamanzi, Pierre Collet, Jano I. van Hemert, and Marco Tomassini, editors, *Proceedings of the 8th European Conference on Genetic Programming*, volume 3447 of *Lecture Notes in Computer Science*, pages 120–131, Lausanne, Switzerland, 30 March - 1 April 2005. Springer.
- [48] Ami Hauptman and Moshe Sipper. Evolution of an efficient search algorithm for the mate-in-N problem in chess. In Marc Ebner, Michael O’Neill, Anikó Ekárt, Leonardo Vanneschi, and Anna Isabel Esparcia-Alcázar, editors, *Proceedings of the 10th European Conference on Genetic Programming*, volume 4445 of *Lecture Notes in Computer Science*, pages 78–89, Valencia, Spain, 11-13 April 2007. Springer.
- [49] Thomas D. Haynes, Dale A. Schoenefeld, and Roger L. Wainwright. Type inheritance in strongly typed genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 18, pages 359–376. MIT Press, Cambridge, MA, USA, 1996.
- [50] Nguyen Xuan Hoai, R. I. McKay, and H. A. Abbass. Tree adjoining grammars, language bias, and genetic programming. In Conor Ryan, Terence Soule, Maarten Keijzer, Edward Tsang, Riccardo Poli, and Ernesto Costa, editors, *Genetic Programming, Proceedings of EuroGP’2003*, volume 2610 of *LNCS*, pages 335–344, Essex, 14-16 April 2003. Springer-Verlag.
- [51] Tuan-Hao Hoang, Daryl Essam, Robert Ian (Bob) McKay, and Xuan Hoai Nguyen. Building on success in genetic programming: adaptive variation & developmental evaluation. In *Proceedings of the 2007 International Symposium on Intelligent Computation and Applications (ISICA)*, Wuhan, China, September 21-23 2007. China University of Geosciences Press.
- [52] John H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI, 1975.
- [53] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial*

Intelligence. MIT Press, 1992. First Published by University of Michigan Press 1975.

- [54] D. Howard and K. Kolibal. Solution of differential equations with genetic programming and the stochastic bernstein interpolation. Technical Report BDS-TR-2005-001, University of Limerick, Biocomputing-Developmental Systems Group, June 2005.
- [55] Jianjun Hu, Erik D. Goodman, Shaobo Li, and Ronald Rosenberg. Automated synthesis of mechanical vibration absorbers using genetic programming. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 22(3):207–217, 2008.
- [56] Hitoshi Iba. *Genetic Programming*. Tokyo Denki University Press, 1996.
- [57] Christian Jacob. *Principia Evolvica – Simulierte Evolution mit Mathematica*. dpunkt.verlag, Heidelberg, Germany, August 1997. In German.
- [58] Christian Jacob. *Illustrating Evolutionary Computation with Mathematica*. Morgan Kaufmann, 2001.
- [59] Maarten Keijzer, Martin Baptist, Vladan Babovic, and Javier Rodriguez Uthurburu. Determining equations for vegetation induced resistance using genetic programming. In Hans-Georg Beyer, Una-May O’Reilly, Dirk V. Arnold, Wolfgang Banzhaf, Christian Blum, Eric W. Bonabeau, Erick Cantu-Paz, Dipankar Dasgupta, Kalyanmoy Deb, James A. Foster, Edwin D. de Jong, Hod Lipson, Xavier Llorca, Spiros Mancoridis, Martin Pelikan, Guenther R. Raidl, Terence Soule, Andy M. Tyrrell, Jean-Paul Watson, and Eckart Zitzler, editors, *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 2, pages 1999–2006, Washington DC, USA, 25–29 June 2005. ACM Press.
- [60] Bijan KHosraviani, Raymond E. Levitt, and John R. Koza. Organization design optimization using genetic programming. In Maarten Keijzer, editor, *Late Breaking Papers at the 2004 Genetic and Evolutionary Computation Conference*, Seattle, Washington, USA, 26 July 2004.
- [61] Kenneth E. Kinnear, Jr. Evolving a sort: Lessons in genetic programming. In *Proceedings of the 1993 International Conference on Neural Networks*, volume 2, pages 881–888, San Francisco, USA, 28 March–1 April 1993. IEEE Press.
- [62] Kenneth E. Kinnear, Jr., editor. *Advances in Genetic Programming*. MIT Press, Cambridge, MA, 1994.
- [63] Kenneth E. Kinnear, Jr. Fitness landscapes and difficulty in genetic programming. In *Proceedings of the 1994 IEEE World Conference on Computational Intelligence*, volume 1, pages 142–147, Orlando, Florida, USA, 27–29 June 1994. IEEE Press.
- [64] John Koza, Forrest H Bennett III, David Andre, and Martin A. Keane. The design of analog circuits by means of genetic programming. In Peter Bentley, editor, *Evolutionary Design by Computers*, chapter 16, pages 365–385. Morgan Kaufmann, San Francisco, USA, 1999.

- [65] John R. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA, 1992.
- [66] John R. Koza. *Genetic Programming II: Automatic Discovery of Reusable Programs*. MIT Press, Cambridge Massachusetts, May 1994.
- [67] John R. Koza. Two ways of discovering the size and shape of a computer program to solve a problem. In L. Eshelman, editor, *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*, pages 287–294, Pittsburgh, PA, USA, 15-19 July 1995. Morgan Kaufmann.
- [68] John R. Koza, Sameer H. Al-Sakran, and Lee W. Jones. Automated re-invention of six patented optical lens systems using genetic programming. In Hans-Georg Beyer, Una-May O’Reilly, Dirk V. Arnold, Wolfgang Banzhaf, Christian Blum, Eric W. Bonabeau, Erick Cantu-Paz, Dipankar Dasgupta, Kalyanmoy Deb, James A. Foster, Edwin D. de Jong, Hod Lipson, Xavier Llorca, Spiros Mancoridis, Martin Pelikan, Guenther R. Raidl, Terence Soule, Andy M. Tyrrell, Jean-Paul Watson, and Eckart Zitzler, editors, *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 2, pages 1953–1960, Washington DC, USA, 25-29 June 2005. ACM Press.
- [69] John R. Koza, Sameer H. Al-Sakran, and Lee W. Jones. Automated ab initio synthesis of complete designs of four patented optical lens systems by means of genetic programming. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 22(3):249–273, 2008.
- [70] John R. Koza, David Andre, Forrest H Bennett III, and Martin Keane. *Genetic Programming 3: Darwinian Invention and Problem Solving*. Morgan Kaufman, April 1999.
- [71] John R. Koza, David Andre, Forrest H Bennett III, and Martin A. Keane. Use of automatically defined functions and architecture-altering operations in automated circuit synthesis using genetic programming. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 132–149, Stanford University, CA, USA, 28–31 July 1996. MIT Press.
- [72] John R. Koza, Forrest H Bennett III, David Andre, and Martin A. Keane. Automated WYWIWYG design of both the topology and component values of electrical circuits using genetic programming. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 123–131, Stanford University, CA, USA, 28–31 July 1996. MIT Press.
- [73] John R. Koza, Martin A. Keane, Matthew J. Streeter, William Myrdlowec, Jessen Yu, and Guido Lanza. *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Kluwer Academic Publishers, 2003.
- [74] Keane Martin A. Koza, John R. and Matthew J. Streeter. Human-competitive automated engineering design and optimization by means of genetic programming. In Periaux J. Cerrolaza M. Annicchiarico, W. and

- G. Winter, editors, *Evolutionary Algorithms and Intelligent Tools in Engineering Optimization*. Wit Pr/Computational Mechanics, 2005.
- [75] Krzysztof Krawiec. *Evolutionary Feature Programming: Cooperative learning for knowledge discovery and computer vision*. Number 385 in . Wydawnictwo Politechniki Poznańskiej, Poznan University of Technology, Poznan, Poland, 2004.
- [76] Brian Lam and Vic Ciesielski. Discovery of human-competitive image texture feature extraction programs using genetic programming. In Kalyanmoy Deb, Riccardo Poli, Wolfgang Banzhaf, Hans-Georg Beyer, Edmund Burke, Paul Darwen, Dipankar Dasgupta, Dario Floreano, James Foster, Mark Harman, Owen Holland, Pier Luca Lanzi, Lee Spector, Andrea Tettamanzi, Dirk Thierens, and Andy Tyrrell, editors, *Genetic and Evolutionary Computation – GECCO-2004, Part II*, volume 3103 of *Lecture Notes in Computer Science*, pages 1114–1125, Seattle, WA, USA, 26-30 June 2004. Springer-Verlag.
- [77] W. B. Langdon. The evolution of size in variable length representations. In *1998 IEEE International Conference on Evolutionary Computation*, pages 633–638, Anchorage, Alaska, USA, 5-9 May 1998. IEEE Press.
- [78] W. B. Langdon. Convergence rates for the distribution of program outputs. In W. B. Langdon, E. Cantú-Paz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. A. Potter, A. C. Schultz, J. F. Miller, E. Burke, and N. Jonoska, editors, *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 812–819, New York, 9-13 July 2002. Morgan Kaufmann Publishers.
- [79] W. B. Langdon. How many good programs are there? How long are they? In Kenneth A. De Jong, Riccardo Poli, and Jonathan E. Rowe, editors, *Foundations of Genetic Algorithms VII*, pages 183–202, Torremolinos, Spain, 4-6 September 2002. Morgan Kaufmann. Published 2003.
- [80] W. B. Langdon. Convergence of program fitness landscapes. In E. Cantú-Paz, J. A. Foster, K. Deb, D. Davis, R. Roy, U.-M. O’Reilly, H.-G. Beyer, R. Standish, G. Kendall, S. Wilson, M. Harman, J. Wegener, D. Dasgupta, M. A. Potter, A. C. Schultz, K. Dowsland, N. Jonoska, and J. Miller, editors, *Genetic and Evolutionary Computation – GECCO-2003*, volume 2724 of *LNCS*, pages 1702–1714, Chicago, 12-16 July 2003. Springer-Verlag.
- [81] W. B. Langdon. The distribution of reversible functions is Normal. In Rick L. Riolo and Bill Worzel, editors, *Genetic Programming Theory and Practise*, chapter 11, pages 173–188. Kluwer, 2003.
- [82] W. B. Langdon. The distribution of amorphous computer outputs. In Susan Stepney and Stephen Emmott, editors, *The Grand Challenge in Non-Classical Computation: International Workshop*, York, UK, 18-19 April 2005.

- [83] W. B. Langdon and R. Poli. The halting probability in von Neumann architectures. In Pierre Collet, Marco Tomassini, Marc Ebner, Steven Gustafson, and Anikó Ekárt, editors, *Proceedings of the 9th European Conference on Genetic Programming*, volume 3905 of *Lecture Notes in Computer Science*, pages 225–237, Budapest, Hungary, 10 - 12 April 2006. Springer.
- [84] W. B. Langdon and Riccardo Poli. *Foundations of Genetic Programming*. Springer-Verlag, 2002.
- [85] William B. Langdon. *Genetic Programming and Data Structures: Genetic Programming + Data Structures = Automatic Programming!*, volume 1 of *Genetic Programming*. Kluwer, Boston, 1998.
- [86] William B. Langdon. Size fair and homologous tree genetic programming crossovers. *Genetic Programming and Evolvable Machines*, 1(1/2):95–119, April 2000.
- [87] William B. Langdon, Steven M. Gustafson, and John Koza. The Genetic Programming Bibliography, 1995-2008.
- [88] William B. Langdon, Terry Soule, Riccardo Poli, and James A. Foster. The evolution of size and shape. In Lee Spector, William B. Langdon, Una-May O’Reilly, and Peter J. Angeline, editors, *Advances in Genetic Programming 3*, chapter 8, pages 163–190. MIT Press, Cambridge, MA, USA, June 1999.
- [89] Pedro Larrañaga and Jose A. Lozano. *Estimation of Distribution Algorithms, A New Tool for Evolutionary Computation*. Kluwer Academic Publishers, 2002.
- [90] A. Lindenmayer. Mathematic models for cellular interaction in development, parts I and II. *Journal of Theoretical Biology*, 18:280–299 and 300–315, 1968.
- [91] Hod Lipson. How to draw a straight line using a GP: Benchmarking evolutionary design against 19th century kinematic synthesis. In Maarten Keijzer, editor, *Late Breaking Papers at the 2004 Genetic and Evolutionary Computation Conference*, Seattle, Washington, USA, 26 July 2004.
- [92] Jason Lohn, Gregory Hornby, and Derek Linden. Evolutionary antenna design for a NASA spacecraft. In Una-May O’Reilly, Tina Yu, Rick L. Riolo, and Bill Worzel, editors, *Genetic Programming Theory and Practice II*, chapter 18, pages 301–315. Springer, Ann Arbor, 13-15 May 2004.
- [93] Moshe Looks. Scalable estimation-of-distribution program evolution. In Hod Lipson, editor, *GECCO*, pages 539–546. ACM, 2007.
- [94] Moshe Looks, Ben Goertzel, and Cassio Pennachin. Learning computer programs with the bayesian optimization algorithm. In Hans-Georg Beyer, Una-May O’Reilly, Dirk V. Arnold, Wolfgang Banzhaf, Christian Blum, Eric W. Bonabeau, Erick Cantu-Paz, Dipankar Dasgupta, Kalyanmoy Deb, James A. Foster, Edwin D. de Jong, Hod Lipson, Xavier Llorca, Spiros Mancoridis, Martin Pelikan, Guenther R. Raidl, Terence Soule,

- Andy M. Tyrrell, Jean-Paul Watson, and Eckart Zitzler, editors, *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 1, pages 747–748, Washington DC, USA, 25-29 June 2005. ACM Press.
- [95] Sean Luke. Genetic programming produced competitive soccer softbot teams for robocup97. In John R. Koza, Wolfgang Banzhaf, Kumar Chellapilla, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max H. Garzon, David E. Goldberg, Hitoshi Iba, and Rick Riolo, editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 214–222, University of Wisconsin, Madison, Wisconsin, USA, 22-25 July 1998. Morgan Kaufmann.
- [96] Sean Luke, Liviu Panait, Gabriel Balan, Sean Paus, Zbigniew Skolicki, Elena Popovici, Joseph Harrison, Jeff Bassett, Robert Hubley, and Alexander Chircop. ECJ: A Java-based Evolutionary Computation Research System , 2000-2007.
- [97] Paul Massey, John A. Clark, and Susan Stepney. Evolution of a human-competitive quantum fourier transform algorithm using genetic programming. In Hans-Georg Beyer, Una-May O’Reilly, Dirk V. Arnold, Wolfgang Banzhaf, Christian Blum, Eric W. Bonabeau, Erick Cantu-Paz, Dipankar Dasgupta, Kalyanmoy Deb, James A. Foster, Edwin D. de Jong, Hod Lipson, Xavier Llorca, Spiros Mancoridis, Martin Pelikan, Guenther R. Raidl, Terence Soule, Andy M. Tyrrell, Jean-Paul Watson, and Eckart Zitzler, editors, *GECCO 2005: Proceedings of the 2005 conference on Genetic and evolutionary computation*, volume 2, pages 1657–1663, Washington DC, USA, 25-29 June 2005. ACM Press.
- [98] Boris Mitavskiy and Jon Rowe. Some results about the markov chains associated to GPs and to general EAs. *Theoretical Computer Science*, 361(1):72–110, 28 August 2006.
- [99] David J. Montana. Strongly typed genetic programming. *Evolutionary Computation*, 3(2):199–230, 1995.
- [100] ThanhVu Nguyen, Westley Weimer, Claire Le Goues, and Stephanie Forrest. Using execution paths to evolve software patches. In Phil McMinn and Robert Feldt, editors, *International Conference on Software Testing, Verification and Validation Workshops, ICSTW ’09*, pages 152–153, Denver, Colorado, USA, 1-4 April 2009.
- [101] Nikolay Nikolaev and Hitoshi Iba. *Adaptive Learning of Polynomial Networks Genetic Programming, Backpropagation and Bayesian Methods*. Number 4 in Genetic and Evolutionary Computation. Springer, 2006. June.
- [102] Peter Nordin. *Evolutionary Program Induction of Binary Machine Code and its Applications*. PhD thesis, der Universitat Dortmund am Fachereich Informatik, 1997.
- [103] Peter Nordin and Wilde Johanna. *Humanoider: Sjavlarande robotar och artificiell intelligens*. Liber, 2003.

- [104] J. R. Olsson. *Inductive functional programming using incremental program transformation and Execution of logic programs by iterative-deepening A* SLD-tree search*. Dr scient thesis, University of Oslo, Norway, 1994.
- [105] Michael O’Neill and Conor Ryan. *Grammatical Evolution: Evolutionary Automatic Programming in a Arbitrary Language*, volume 4 of *Genetic programming*. Kluwer Academic Publishers, 2003.
- [106] Cynthia B. Perez and Gustavo Olague. Learning invariant region descriptor operators with genetic programming and the F-measure. In *19th International Conference on Pattern Recognition (ICPR 2008)*, pages 1–4, Tampa, Florida, USA, December 8-11 2008.
- [107] Cynthia B. Perez and Gustavo Olague. Evolutionary learning of local descriptor operators for object recognition. In Guenther Raidl, Franz Rothlauf, Giovanni Squillero, Rolf Drechsler, Thomas Stuetzle, Mauro Bittantari, Clare Bates Congdon, Martin Middendorf, Christian Blum, Carlos Cotta, Peter Bosman, Joern Grahl, Joshua Knowles, David Corne, Hans-Georg Beyer, Ken Stanley, Julian F. Miller, Jano van Hemert, Tom Lenaerts, Marc Ebner, Jaume Bacardit, Michael O’Neill, Massimiliano Di Penta, Benjamin Doerr, Thomas Jansen, Riccardo Poli, and Enrique Alba, editors, *GECCO ’09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 1051–1058, Montreal, 8-12 July 2009. ACM.
- [108] R. Poli. Hyperschema theory for GP with one-point crossover, building blocks, and some new results in GA theory. In Riccardo Poli, Wolfgang Banzhaf, William B. Langdon, Julian F. Miller, Peter Nordin, and Terence C. Fogarty, editors, *Genetic Programming, Proceedings of EuroGP’2000*, volume 1802 of *LNCS*, pages 163–180, Edinburgh, 15-16 April 2000. Springer-Verlag.
- [109] Riccardo Poli. Exact schema theorem and effective fitness for GP with one-point crossover. In D. Whitley, D. Goldberg, E. Cantu-Paz, L. Spector, I. Parmee, and H.-G. Beyer, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 469–476, Las Vegas, July 2000. Morgan Kaufmann.
- [110] Riccardo Poli. Exact schema theory for genetic programming and variable-length genetic algorithms with one-point crossover. *Genetic Programming and Evolvable Machines*, 2(2):123–163, 2001.
- [111] Riccardo Poli. General schema theory for genetic programming with subtree-swapping crossover. In *Genetic Programming, Proceedings of EuroGP 2001*, LNCS, Milan, 18-20 April 2001. Springer-Verlag.
- [112] Riccardo Poli. A simple but theoretically-motivated method to control bloat in genetic programming. In Conor Ryan, Terrence Soule, Maarten Keijzer, Edward Tsang, Riccardo Poli, and Ernesto Costa, editors, *Genetic Programming, Proceedings of the 6th European Conference, EuroGP 2003*, LNCS, pages 211–223, Essex, UK, 14-16 April 2003. Springer-Verlag.

- [113] Riccardo Poli and William B. Langdon. Efficient markov chain model of machine code program execution and halting. In Rick L. Riolo, Terence Soule, and Bill Worzel, editors, *Genetic Programming Theory and Practice IV*, volume 5 of *Genetic and Evolutionary Computation*, chapter 13. Springer, Ann Arbor, 11-13 May 2006.
- [114] Riccardo Poli, William B. Langdon, and Stephen Dignum. On the limiting distribution of program sizes in tree-based genetic programming. In Marc Ebner, Michael O’Neill, Anikó Ekárt, Leonardo Vanneschi, and Anna Isabel Esparcia-Alcázar, editors, *Proceedings of the 10th European Conference on Genetic Programming*, volume 4445 of *Lecture Notes in Computer Science*, pages 193–204, Valencia, Spain, 11-13 April 2007. Springer.
- [115] Riccardo Poli, William B. Langdon, and Nicholas Freitag 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).
- [116] Riccardo Poli and Nicholas F. McPhee. Covariant parsimony pressure in genetic programming. Technical Report CES-480, Department of Computing and Electronic Systems, University of Essex, January 2008.
- [117] Riccardo Poli and Nicholas Freitag McPhee. General schema theory for genetic programming with subtree-swapping crossover: Part I. *Evolutionary Computation*, 11(1):53–66, March 2003.
- [118] Riccardo Poli and Nicholas Freitag McPhee. General schema theory for genetic programming with subtree-swapping crossover: Part II. *Evolutionary Computation*, 11(2):169–206, June 2003.
- [119] Riccardo Poli and Nicholas Freitag McPhee. A linear estimation-of-distribution GP system. In Michael O’Neill, Leonardo Vanneschi, Steven Gustafson, Anna Isabel Esparcia Alcazar, Ivanoe De Falco, Antonio Della Cioppa, and Ernesto Tarantino, editors, *Proceedings of the 11th European Conference on Genetic Programming, EuroGP 2008*, volume 4971 of *Lecture Notes in Computer Science*, pages 206–217, Naples, 26-28 March 2008. Springer.
- [120] Riccardo Poli, Nicholas Freitag McPhee, and Jonathan E. Rowe. Exact schema theory and markov chain models for genetic programming and variable-length genetic algorithms with homologous crossover. *Genetic Programming and Evolvable Machines*, 5(1):31–70, March 2004.
- [121] Riccardo Poli, Jonathan E. Rowe, and Nicholas Freitag McPhee. Markov chain models for GP and variable-length GAs with homologous crossover. In Lee Spector, Erik D. Goodman, Annie Wu, W. B. Langdon, Hans-Michael Voigt, Mitsuo Gen, Sandip Sen, Marco Dorigo, Shahram Pezeshk, Max H. Garzon, and Edmund Burke, editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 112–119, San Francisco, California, USA, 7-11 July 2001. Morgan Kaufmann.
- [122] Riccardo Poli, Leonardo Vanneschi, William B. Langdon, and Nicholas Freitag McPhee. Theoretical results in genetic programming:

The next ten years? *Genetic Programming and Evolvable Machines*, 11(3/4):285–320, September 2010. Tenth Anniversary Issue: Progress in Genetic Programming and Evolvable Machines.

- [123] Bill Punch and Douglas Zongker. *lil-gp Genetic Programming System*, 1998.
- [124] Amr Radi. Prediction of non-linear system in optics using genetic programming. *International Journal of Modern Physics C*, 18, March 2007.
- [125] Amr M. Radi and Salah Yaseen El-Bakry. Genetic programming approach for positron collisions with alkali-metal atom. In Dirk Thierens, Hans-Georg Beyer, Josh Bongard, Jurgen Branke, John Andrew Clark, Dave Cliff, Clare Bates Congdon, Kalyanmoy Deb, Benjamin Doerr, Tim Kovacs, Sanjeev Kumar, Julian F. Miller, Jason Moore, Frank Neumann, Martin Pelikan, Riccardo Poli, Kumara Sastry, Kenneth Owen Stanley, Thomas Stutzle, Richard A Watson, and Ingo Wegener, editors, *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*, volume 2, pages 1756–1756, London, 7-11 July 2007. ACM Press.
- [126] Adil Raja, R. Muhammad Atif Azad, Colin Flanagan, and Conor Ryan. Real-time, non-intrusive evaluation of voIP. In Marc Ebner, Michael O’Neill, Anikó Ekárt, Leonardo Vanneschi, and Anna Isabel Esparcia-Alcázar, editors, *Proceedings of the 10th European Conference on Genetic Programming*, volume 4445 of *Lecture Notes in Computer Science*, pages 217–228, Valencia, Spain, 11-13 April 2007. Springer.
- [127] Alain Ratle and Michele Sebag. Avoiding the bloat with probabilistic grammar-guided genetic programming. In P. Collet, C. Fonlupt, J.-K. Hao, E. Lutton, and M. Schoenauer, editors, *Artificial Evolution 5th International Conference, Evolution Artificielle, EA 2001*, volume 2310 of *LNCS*, pages 255–266, Creusot, France, October 29-31 2001. Springer Verlag.
- [128] Rick L. Riolo and Bill Worzel, editors. *Genetic Programming Theory and Practice*, volume 6 of *Genetic Programming*, Boston, MA, USA, 2003. Kluwer. Series Editor - John Koza.
- [129] RML Technologies. *Discipulus Genetic-Programming Software*, 1998-2007.
- [130] Justinian Rosca. A probabilistic model of size drift. In Rick L. Riolo and Bill Worzel, editors, *Genetic Programming Theory and Practice*, chapter 8, pages 119–136. Kluwer, 2003.
- [131] Justinian P. Rosca and Dana H. Ballard. Discovery of subroutines in genetic programming. In Peter J. Angeline and K. E. Kinnear, Jr., editors, *Advances in Genetic Programming 2*, chapter 9, pages 177–202. MIT Press, Cambridge, MA, USA, 1996.
- [132] Franz Rothlauf. *Representations for genetic and evolutionary algorithms*. Springer-Verlag, pub-SV:adr, second edition, 2006. First published 2002, 2nd edition available electronically.

- [133] Conor Ryan. *Automatic Re-engineering of Software Using Genetic Programming*, volume 2 of *Genetic Programming*. Kluwer Academic Publishers, 1 November 1999.
- [134] Conor Ryan, J. J. Collins, and Michael O’Neill. Grammatical evolution: Evolving programs for an arbitrary language. In Wolfgang Banzhaf, Riccardo Poli, Marc Schoenauer, and Terence C. Fogarty, editors, *Proceedings of the First European Workshop on Genetic Programming*, volume 1391 of *LNCS*, pages 83–95, Paris, 14-15 April 1998. Springer-Verlag.
- [135] R. P. Salustowicz and J. Schmidhuber. Probabilistic incremental program evolution. *Evolutionary Computation*, 5(2):123–141, 1997.
- [136] Michael Schmidt and Hod Lipson. Distilling free-form natural laws from experimental data. *Science*, 324(5923):81–85, 3 April 2009.
- [137] Michael D. Schmidt and Hod Lipson. Solving iterated functions using genetic programming. In Anna I. Esparcia, Ying ping Chen, Gabriela Ochoa, Ender Ozcan, Marc Schoenauer, Anne Auger, Hans-Georg Beyer, Nikolaus Hansen, Steffen Finck, Raymond Ros, Darrell Whitley, Garnett Wilson, Simon Harding, W. B. Langdon, Man Leung Wong, Laurence D. Merkle, Frank W. Moore, Sevan G. Ficici, William Rand, Rick Riolo, Nawwaf Kharma, William R. Buckley, Julian Miller, Kenneth Stanley, Jaume Bacardit, Will Browne, Jan Drugowitsch, Nicola Beume, Mike Preuss, Stephen L. Smith, Stefano Cagnoni, Jim DeLeo, Alexandru Floares, Aaron Baughman, Steven Gustafson, Maarten Keijzer, Arthur Kordon, Clare Bates Congdon, Laurence D. Merkle, and Frank W. Moore, editors, *GECCO-2009 Late-Breaking Papers*, pages 2149–2154, Montreal, 8-12 July 2009. ACM.
- [138] Lukas Sekanina. *Evolvable Components: From Theory to Hardware Implementations*. Natural Computing. Springer-Verlag, 2003.
- [139] Y. Shan, H. Abbass, R. I. McKay, and D. Essam. AntTAG: a further study. In Ruhul Sarker and Bob McKay, editors, *Proceedings of the Sixth Australia-Japan Joint Workshop on Intelligent and Evolutionary Systems*, Australian National University, Canberra, Australia, 30 November 2002.
- [140] Y. Shan, R. I. McKay, H. A. Abbass, and D. Essam. Program evolution with explicit learning: a new framework for program automatic synthesis. In Ruhul Sarker, Robert Reynolds, Hussein Abbass, Kay Chen Tan, Bob McKay, Daryl Essam, and Tom Gedeon, editors, *Proceedings of the 2003 Congress on Evolutionary Computation CEC2003*, pages 1639–1646, Canberra, 8-12 December 2003. IEEE Press.
- [141] Yin Shan, Robert I. McKay, Daryl Essam, and Hussein A. Abbass. A survey of probabilistic model building genetic programming. In M. Pelikan, K. Sastry, and E. Cantu-Paz, editors, *Scalable Optimization via Probabilistic Modeling: From Algorithms to Applications*, volume 33 of *Studies in Computational Intelligence*, chapter 6, pages 121–160. Springer, 2006.
- [142] Yehonatan Shichel, Eran Ziserman, and Moshe Sipper. GP-robocode: Using genetic programming to evolve robocode players. In Maarten Keijzer, Andrea Tettamanzi, Pierre Collet, Jano I. van Hemert, and Marco

- Tomassini, editors, *Proceedings of the 8th European Conference on Genetic Programming*, volume 3447 of *Lecture Notes in Computer Science*, pages 143–154, Lausanne, Switzerland, 30 March - 1 April 2005. Springer.
- [143] Moshe Sipper. Attaining human-competitive game playing with genetic programming. In Samira El Yacoubi, Bastien Chopard, and Stefania Bandini, editors, *Proceedings of the 7th International Conference on Cellular Automata, for Research and Industry, ACRI*, volume 4173 of *Lecture Notes in Computer Science*, page 13, Perpignan, France, September 20-23 2006. Springer. Invited Lectures.
- [144] Terence Soule and James A. Foster. Effects of code growth and parsimony pressure on populations in genetic programming. *Evolutionary Computation*, 6(4):293–309, Winter 1998.
- [145] Lee Spector. *Automatic Quantum Computer Programming: A Genetic Programming Approach*, volume 7 of *Genetic Programming*. Kluwer Academic Publishers, Boston/Dordrecht/New York/London, June 2004.
- [146] Lee Spector, Howard Barnum, and Herbert J. Bernstein. Genetic programming for quantum computers. In John R. Koza, Wolfgang Banzhaf, Kumar Chellapilla, Kalyanmoy Deb, Marco Dorigo, David B. Fogel, Max H. Garzon, David E. Goldberg, Hitoshi Iba, and Rick Riolo, editors, *Genetic Programming 1998: Proceedings of the Third Annual Conference*, pages 365–373, University of Wisconsin, Madison, Wisconsin, USA, 22-25 July 1998. Morgan Kaufmann.
- [147] Lee Spector, Howard Barnum, Herbert J. Bernstein, and Nikhil Swamy. Finding a better-than-classical quantum AND/OR algorithm using genetic programming. In Peter J. Angeline, Zbyszek Michalewicz, Marc Schoenauer, Xin Yao, and Ali Zalzala, editors, *Proceedings of the Congress on Evolutionary Computation*, volume 3, pages 2239–2246, Mayflower Hotel, Washington D.C., USA, 6-9 July 1999. IEEE Press.
- [148] Lee Spector, Howard Barnum, Herbert J. Bernstein, and Nikhil Swamy. Quantum computing applications of genetic programming. In Lee Spector, William B. Langdon, Una-May O’Reilly, and Peter J. Angeline, editors, *Advances in Genetic Programming 3*, chapter 7, pages 135–160. MIT Press, Cambridge, MA, USA, June 1999.
- [149] Lee Spector and H. J. Bernstein. Communication capacities of some quantum gates, discovered in part through genetic programming. In Jeffrey H. Shapiro and Osamu Hirota, editors, *Proceedings of the Sixth International Conference on Quantum Communication, Measurement, and Computing (QCMC)*, pages 500–503. Rinton Press, 2003.
- [150] Lee Spector, David M. Clark, Ian Lindsay, Bradford Barr, and Jon Klein. Genetic programming for finite algebras. In Maarten Keijzer, Giuliano Antoniol, Clare Bates Congdon, Kalyanmoy Deb, Benjamin Doerr, Nikolaus Hansen, John H. Holmes, Gregory S. Hornby, Daniel Howard, James Kennedy, Sanjeev Kumar, Fernando G. Lobo, Julian Francis Miller, Jason Moore, Frank Neumann, Martin Pelikan, Jordan Pollack, Kumara Sastry,

- Kenneth Stanley, Adrian Stoica, El-Ghazali Talbi, and Ingo Wegener, editors, *GECCO '08: Proceedings of the 10th annual conference on Genetic and evolutionary computation*, pages 1291–1298, Atlanta, GA, USA, 12-16 July 2008. ACM.
- [151] Lee Spector and Jon Klein. Machine invention of quantum computing circuits by means of genetic programming. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 22(3):275–283, 2008.
- [152] Lee Spector, W. B. Langdon, Una-May O’Reilly, and Peter J. Angeline, editors. *Advances in Genetic Programming 3*. MIT Press, Cambridge, MA, USA, June 1999.
- [153] Ralf Stadelhofer, Wolfgang Banzhaf, and Dieter Suter. Evolving blackbox quantum algorithms using genetic programming. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 22(3):285–297, 2008.
- [154] C. R. Stephens and H. Waelbroeck. Effective degrees of freedom in genetic algorithms and the block hypothesis. In Thomas Bäck, editor, *Proceedings of the Seventh International Conference on Genetic Algorithms (ICGA97)*, pages 34–40, East Lansing, 1997. Morgan Kaufmann.
- [155] C. R. Stephens and H. Waelbroeck. Schemata evolution and building blocks. *Evolutionary Computation*, 7(2):109–124, 1999.
- [156] Joc Cing Tay and Nhu Binh Ho. Evolving dispatching rules using genetic programming for solving multi-objective flexible job-shop problems. *Computers & Industrial Engineering*, 54(3):453–473, 2008.
- [157] L. Trujillo and G. Olague. Using evolution to learn how to perform interest point detection. In X. Y Tang et al., editor, *ICPR 2006 18th International Conference on Pattern Recognition*, volume 1, pages 211–214. IEEE, 20-24 August 2006.
- [158] Leonardo Trujillo and Gustavo Olague. Synthesis of interest point detectors through genetic programming. In Maarten Keijzer, Mike Catolico, Dirk Arnold, Vladan Babovic, Christian Blum, Peter Bosman, Martin V. Butz, Carlos Coello Coello, Dipankar Dasgupta, Sevan G. Ficici, James Foster, Arturo Hernandez-Aguirre, Greg Hornby, Hod Lipson, Phil McMinn, Jason Moore, Guenther Raidl, Franz Rothlauf, Conor Ryan, and Dirk Thierens, editors, *GECCO 2006: Proceedings of the 8th annual conference on Genetic and evolutionary computation*, volume 1, pages 887–894, Seattle, Washington, USA, 8-12 July 2006. ACM Press.
- [159] Edward Tsang and Nanlin Jin. Incentive method to handle constraints in evolutionary. In Pierre Collet, Marco Tomassini, Marc Ebner, Steven Gustafson, and Anikó Ekárt, editors, *Proceedings of the 9th European Conference on Genetic Programming*, volume 3905 of *Lecture Notes in Computer Science*, pages 133–144, Budapest, Hungary, 10 - 12 April 2006. Springer.
- [160] Edward P. K. Tsang and Jin Li. EDDIE for financial forecasting. In Shu-Heng Chen, editor, *Genetic Algorithms and Genetic Programming*

- in *Computational Finance*, chapter 7, pages 161–174. Kluwer Academic Press, 2002.
- [161] A. M. Turing. Intelligent machinery. Report for National Physical Laboratory. Reprinted in Ince, D. C. (editor). 1992. *Mechanical Intelligence: Collected Works of A. M. Turing*. Amsterdam: North Holland. Pages 107–127. Also reprinted in Meltzer, B. and Michie, D. (editors). 1969. *Machine Intelligence 5*. Edinburgh: Edinburgh University Press, 1948.
- [162] A. M. Turing. Computing machinery and intelligence. *Mind*, 49:433–460, January 01 1950.
- [163] Westley Weimer, ThanhVu Nguyen, Claire Le Goues, and Stephanie Forrest. Automatically finding patches using genetic programming. In Stephen Fickas, editor, *International Conference on Software Engineering (ICSE) 2009*, pages 364–374, Vancouver, May 16-24 2009.
- [164] P. A. Whigham. Search bias, language bias, and genetic programming. In John R. Koza, David E. Goldberg, David B. Fogel, and Rick L. Riolo, editors, *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 230–237, Stanford University, CA, USA, 28–31 July 1996. MIT Press.
- [165] L. Darrell Whitley. A Genetic Algorithm Tutorial. *Statistics and Computing*, 4:65–85, 1994.
- [166] Man Leung Wong and Kwong Sak Leung. Evolving recursive functions for the even-parity problem using genetic programming. In Peter J. Angeline and K. E. Kinneer, Jr., editors, *Advances in Genetic Programming 2*, chapter 11, pages 221–240. MIT Press, Cambridge, MA, USA, 1996.
- [167] Man Leung Wong and Kwong Sak Leung. *Data Mining Using Grammar Based Genetic Programming and Applications*, volume 3 of *Genetic Programming*. Kluwer Academic Publishers, January 2000.
- [168] Kohsuke Yanai and Hitoshi Iba. Estimation of distribution programming based on bayesian network. In Ruhul Sarker, Robert Reynolds, Hussein Abbass, Kay Chen Tan, Bob McKay, Daryl Essam, and Tom Gedeon, editors, *Proceedings of the 2003 Congress on Evolutionary Computation CEC2003*, pages 1618–1625, Canberra, 8-12 December 2003. IEEE Press.
- [169] Tina Yu. Hierarchical processing for evolving recursive and modular programs using higher order functions and lambda abstractions. *Genetic Programming and Evolvable Machines*, 2(4):345–380, December 2001.
- [170] Byoung-Tak Zhang and Heinz Mühlenbein. Evolving optimal neural networks using genetic algorithms with Occam’s razor. *Complex Systems*, 7:199–220, 1993.
- [171] Byoung-Tak Zhang and Heinz Mühlenbein. Balancing accuracy and parsimony in genetic programming. *Evolutionary Computation*, 3(1):17–38, 1995.

- [172] Byoung-Tak Zhang, Peter Ohm, and Heinz Mühlenbein. Evolutionary induction of sparse neural trees. *Evolutionary Computation*, 5(2):213–236, 1997.