Genetic Programming Practice and Theory

Riccardo Poli

Department of Computing and Electronic Systems University of Essex

Overview

- □Basics
- □ Examples and Applications

R. Pdi - University of Essex

- □Theory
- \square Conclusions

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.

R. Poli - University of Esse

□ GP is a *domain-independent* method that genetically breeds a population of computer programs to solve a problem.











Terminal Set (Step 1)

R. Poli - University of Essex

□ Steps 1 and 2 specify the ingredients that are available to create the computer programs (*primitive set*).

- □The terminal set may consist of
 - The program's external inputs (e.g. x, y),
 - ■0-arity functions (e.g. rand(), go_left()),
 - Numerical constants (e.g. 0.1, 3, *R*).

Function Set (Step 2)

□ For some problems, the function set may consist of merely the arithmetic functions (+, -, *, /) and a *conditional* branching operator.

□ But all sort of functions are allowed, e.g.

Kind	Example(s)
Arithmetic	+, *, /
Mathematical	sin, cos, exp
Boolean	AND, OR, NOT
Conditional	IF-THEN-ELSE
Looping	FOR, REPEAT

R. Poli-University of Essex

- □ For many other problems, the primitive set includes specialized functions and terminals. E.g.
 - If the goal is to program a robot to mop the floor Function set = {moving, turning, swishing the mop}
 - If the goal is the automatic creation of a controller

Function set = {*integrators*, *differentiators*, *leads*, *lags*, *gains*}

Terminal set = {reference signal, plant output}

If the goal is the synthesis of analog electrical circuits

Function set = {transistors, capacitors, resistors}

Syntax Errors Are Impossible

- □ **IF** all programs in the initial population of a run of GP are syntactically valid, executable programs,
- □ AND the genetic operations performed during the run are designed to produce offspring that are syntactically valid, executable programs,
- □ **THEN** every individual created during a run of GP is a syntactically valid, executable program.

Run-time Errors Can Be Avoided

- □ **IF** all functions in the primitive set can take as input the results produced by any other function or terminal (closure)
- □ **THEN** run-time errors are avoided.
- Sometime this requires modifying the primitive set appropriately, e.g. using protected versions of division, logarithm, square root, etc.

Fitness Measure (Step 3)

- The fitness measure is the mechanism for giving a high-level statement of the problem's requirements to the GP system.
- The first two preparatory steps define the search space whereas the fitness measure implicitly specifies the search's desired goal.

Fitness can be measured 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
- □ The compliance of a structure with userspecified design criteria, ...

R. Pol.: University of Estect

- The fitness measure is, for many practical problems, multi-objective, i.e. it combines two or more different elements that are often in competition with one another.
- For many problems, each program in the population is executed over a representative sample of different fitness cases.
- Fitness cases may represent different values of the program's input(s), different initial conditions of a system, or different environments.

Control Parameters (Step 4)

R. Poli - University of Essex

- □ An 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.

Termination Criterion (Step 5)

R. Poli - University of Essex

- □ We need to specify 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 best-so-far individual is then harvested and designated as the *result* of the run.

Executional Steps of GP

1. Randomly create an initial population of programs from the available primitives.

R. Pali - University of Essex

- 2. Iterate the following sub-steps until the termination criterion is satisfied:
 - i. Execute each program and ascertain its fitness.
 - ii. Select one or two program(s) from the population with a probability based on fitness to participate in genetic operations.
 - iii. Create new individual program(s) by applying genetic operations with specified probabilities.
- 3. Return the best-so-far individual

Genetic Operations

R. Poli - University of Essax

- □ **Reproduction:** copy the selected individual program to the new population.
- Crossover: create new offspring program(s) by recombining randomly chosen parts from two selected programs.
- Mutation: create one new offspring program by randomly mutating a randomly chosen part of one selected program.
- Architecture-altering operations: choose an architecture-altering operation from the available repertoire and create one new offspring using it.



Random Program Generation

The programs in the initial population are typically built by recursively generating a tree composed of random choices of functions and terminals.

The initial individuals are usually generated subject to a pre-established maximum size.





Fitness

5

Normally, fitness evaluation requires executing the programs in the population multiple times within the GP system.

R. Poli - University of Essex

- □ A variety of execution strategies exist, including:
 - off-line or on-line compilation and linking,
 - virtual-machine-code compilation,
 - interpretation.

Selection

- Genetic operators are applied to individual(s) that are probabilistically selected based on fitness.
- Better individuals are favoured over inferior individuals.
- The most commonly employed methods for selecting individuals are tournament selection and fitness-proportionate selection.
- □ Both methods are not greedy.

Program Interpretation Interpreting a program tree means recursively traversing the tree and executing its nodes only when their arguments are known. E.g.

Sub-tree Crossover

Given two parents crossover randomly selects a crossover point in each parent tree and swaps the sub-trees rooted at the crossover points.







Toy Example

□ Goal: 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.

R. Poli - University of Es

- □ Step 1 Definition of the Terminal Set:
 - The problem is to find a mathematical function of one independent variable, so the terminal set must include *x*.
 - In order to evolve any necessary coefficients, the terminal set also includes numerical constants.
 - That is: T = {X, ℜ}, where ℜ denotes constant numerical terminals in some range (e.g. [-5.0,+5.0]).



- One possible choice consists of the four ordinary arithmetic functions of addition, subtraction, multiplication, and division:
 - $\mathsf{F} = \{+, -, *, \%\}.$
- To avoid run-time errors, the division function % is protected: it returns a value of 1 when division by 0 is attempted, but otherwise returns the quotient of its two arguments.

R. Pol : University of Essex

- □ Step 3 Definition of the Fitness Function:
 - The fitness of an individual in the population must reflect how closely the output of a program comes to *x*²+*x*+1.
 - The fitness measure could be defined as the value of the integral of the errors between the value of the individual mathematical expression and *x*²+*x*+1.
 - Often this numerically approximated using dozens or hundreds of different values of the independent variable x.

□ Step 4 – Fixing GP Parameters:

R. Poli - University of Essex

- Population size: 4 (typically thousands or millions of individuals).
- Crossover probability: 50% (commonly about 90%).
- Reproduction probability: 25% (typically about 8%).
- Mutation probability: 25% (usually about 1%)
- Architecture-altering operation probability: 0% (frequently around 1%).



Example Run

Initial population of four randomly created individuals of generation 0















R. Poli H	Inversity of Essex 43
Real Symbolic	Regression Run
□ Problem : find the syr best fits the data: {x _i ,y _i } = {(-1.0,0.0) (-0.9, □ GP parameters:	nbolic expression that -0.1629) (-0.8,-0.2624)(1.0,4.0)}
Parameter	Value
Population size	1000
Function set	{+ - * plog pexp sin cos pdiv}
Terminal set	{x}
Initial max depth	4
Initialisation method	Full
Number of generations	50
Crossover probability	0.7
Mutation probability	0

Fitness

- Sum of absolute errors













Human-competitive Results

Getting machines to produce human-like results is *the* reason for the existence of the fields of artificial intelligence and machine learning.

R. Pdi - University of Essex

- 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.
- A result produced by an automated method must earn the rating of "human competitive" independent of the fact that it was generated by an automated method.

Criteria for Human-competitiveness

R. Poli - University of Essex

- 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

Criteria for Human-competitiveness

- 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
- н. 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)

R. Poli - Universi

Pre-2004 GP Human-competitive Results

- □ 36 human-competitive results
- 23 instances where GP has *duplicated* the functionality of a previously patented invention, *infringed* a previously *patented invention*, or created a patentable *new invention*
- □ 15 instances where GP has created an entity that either *infringes* or *duplicates* the functionality of a previously patented *20th-century invention*
- □ 6 instances where GP has done the same with respect to an *invention patented after January 1, 2000*
- □ 2 instances where GP has created a *patentable new invention* (general-purpose controllers).

A selection of results

- 1. Creation of a better-than-classical quantum algorithm for Grover's database search problem
- Creation of a quantum algorithm for the depth-two AND/OR query problem that is better than any previously published result

R. Poli - University of Essex

- 3. Creation of a soccer-playing program that won its first two games in the Robo Cup 1997 competition
- 4. Creation of four different algorithms for the transmembrane segment identification problem for proteins
- 5. Creation of a sorting network for seven items using only 16 steps
- 6. Synthesis of 60 and 96 decibel amplifiers
- Synthesis of analog computational circuits for squaring, cubing, square root, cube root, logarithm, and Gaussian functions
- Synthesis of a real-time analog circuit for time-optimal control of a robot

R. Poli - University of Essex

9. Synthesis of an electronic thermometer

10. 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

R. Poli - University of Essex

- 11. Synthesis of topology for a PID-D2 (proportional, integrative, derivative, and second derivative) controller
- 12. Synthesis of NAND circuit
- 13. Simultaneous synthesis of topology, sizing, placement, and routing of analog electrical circuits
- 14. Synthesis of topology for a PID (proportional, integrative, and derivative) controller
- 15. Synthesis of a voltage-current conversion circuit
- 16. Creation of PID tuning rules that outperform the Ziegler-Nichols and Astrom-Hagglund tuning rules
- 17. Creation of three non-PID controllers that outperform a PID controller that uses the Ziegler-Nichols or Astrom-Hagglund tuning rules

Human-competitive-result competition

- □ Held at GECCO 2004-2007
- □ Example winners:
 - Automated Quantum Programming, L. Spector
 - An Evolved Antenna for Deployment on NASA's Space Technology 5 Mission, J.

Lohn et al.





R. Poli - University of Essex

- We can perform many GP runs with a small set of problems and a small set of parameters
- □ We record the variations of certain numerical descriptors.
- Then, we suggest explanations about the behaviour of the system that are compatible with (and could explain) the empirical observations.

Motivation

Problem with Empirical Studies

R. Poli - University of Essex

- □ GP is a complex adaptive system with zillions of degrees of freedom.
- □ So, any small number of descriptors can capture only a fraction of the complexities of such a system.
- Choosing which problems, parameter settings and descriptors to use is an art form.
- Plotting the wrong data increases the confusion about GP's behaviour, rather than clarify it.























R. Pai : University of Essex 72
States, inputs and outputs
 Assume <i>n</i> bits of memory There are 2ⁿ states. At each time step the machine is in a state, <i>s</i>

















Distribution of behaviours					
Behaviour A	Behaviour B	Behaviour C	Identity		
0	0	0	1		
1⁄2	1⁄2	0	0		
1⁄4	1⁄4	1⁄2	0		
1/8	1/8	3⁄4	0		
1/16	1/16	7/8	0		
0	0	1	0		
	Ution 0 Behaviour A 0 1/2 1/4 1/8 1/16 0	Ution of behaviour Behaviour Behaviour 0 0 1/2 1/2 1/4 1/4 1/8 1/8 1/16 1/16 0 0	Ution of behaviours Behaviour Behaviour Behaviour 0 0 0 1/2 1/2 0 1/4 1/2 1/2 1/8 1/8 3/4 1/16 1/16 7/8 0 0 1		

Markov chain proofs of limiting distribution

R. Poli - University of Essex

- Using Markov chain theory we have proved that a limiting distributions of functionality exists for a large variety of CPUs
- □ There are extensions of the proofs from linear to tree-based GP.
- See Foundations of Genetic
 Programming book for an introduction to the proof techniques.

Yes....

 ...for this primitive set the distribution tends to a limit where only behaviour C has non-zero probability.

R. Poli - University of Essex

 Programs in this search space tend to copy the initial value of m1 into m0.

So what?

- Generally instructions lose information.
 Unless inputs are protected, almost all long programs are constants.
- Write protecting inputs makes linear GP more like tree GP.
- □ No point searching above threshold?
- Predict where threshold is? Ad-hoc or theoretical.

Implication of |solution space|/|search space|=constant

- $\hfill\square$ GP can succeed if
 - the constant is not too small or
 - there is structure in the search space to guide the search or
 - the search operators are biased towards searching solution-rich areas of the search space

or any combination of the above.

GP Search Characterisation



Schema Theories

- Divide the search space into subspaces (schemata)
- □ Characterise the schemata using *macroscopic* quantities
- Model how and why the individuals in the population move from one subspace to another (*schema theorems*).





Exact Schema Theory for GP with Subtree Crossover























♦ Intervention of the Interventint of the Interventint of the Interventint of the Interventi















R. Poli - University of Essex

So what?

□ A model is as good as the predictions and the understanding it can produce

R. Poli - University of Essex

□ So, what can we learn from schema theorems?

Lessons

- □ Operator biases
- □ Size evolution equation
- Bloat control
- Optimal parameter setting
- Optimal initialisation
- □ ...















Crossover attempts to push the population towards distributions of primitives where each primitive of a given arity is equally likely to be found in any position in any individual.
 The primitives in a particular individual tend not just to be swapped with those of other individuals in the population, but also to diffuse within the representation of each individual.
 Experiments with unary GP confirm the theory.











Conclusions

Theory

- □ In the last few years the theory of GP has seen a formidable development.
- Today we understand a lot more about the nature of the GP search space and the distribution of fitness in it.
- □ Also, schema theories explain and predict the syntactic behaviour of GAs and GP.
- We know much more as to where evolution is going, why and how.

A. Pdi - University of Essex

- Theory primarily provides explanations, but many recipes for practice have also been derived (initialisation, sizing, parameters, primitives, anti bloat, ...)
- So, theory can help design competent algorithms
- Theory is hard and slow: empirical studies are important to direct theory and to corroborate it.

Turing's Intuition

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 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, but in his 1950 paper "Computing Machinery and Intelligence," he wrote

"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"

□ So, *over 50 years ago* Turing perceived that one 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 (what we now call "fitness").

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 know as the *Turing test for machine intelligence*, whose goals are summarised by Arthur Samuel's position statement

"[T]he aim [is] ... to get machines to exhibit behavior, which if done by humans, would be assumed to involve the use of intelligence." [Arthur Samuel, 1983]

GP Has Started Fulfilling Turing and Samuel's Dreams

- GP has started fulfilling Turing's dream by providing us with a systematic method, based on Darwinian evolution, for getting computers to automatically solve problems.
- To do so, GP simply requires a high-level statement of what needs to be done (and enough computing power).

Today GP certainly cannot produce computer programs that would pass the full Turing test for machine intelligence.

- But GP has been able to solve tens of difficult problems with human-competitive results.
- □No other AI technique has done this

"John Koza's genetic programming approach to machine discovery can invent solutions to more complex specifications than any other I have seen." [John McCarthy]

- These are a small step towards fulfilling Turing and Samuel's dreams, but they are also early signs of things to come.
- 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, becoming an essential collaborator for many of human activities.
- This will be a remarkable step forward towards achieving true, human-competitive machine intelligence.

More information

- □ W.B. Langdon & R. Poli, *Foundations of Genetic Programming*, Springer, 2002.
- J. Koza & R. Poli, Genetic Programming, in INTROS tutorial book, Chapter 5, Kluwer, 2005
- □ R. Poli, Exact Schema Theory for Genetic Programming and Variable-Length Genetic Algorithms with One-Point Crossover, *Genetic Programming and Evolvable Machines* 2(2), 2001.
- □ R. Poli & N. F. McPhee, Schema Theory for GP with Subtree Crossover, Parts I & II, *Evolutionary Computation*, 11(1):53-66 & 11(2):169–206, 2003.
- □ R. Poli, N. F. McPhee & J.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, 2004.