

Jaws 30

W. B. Langdon

21 April 2023

Abstract It is 30 years since John R. Koza published “Jaws”, the first book on genetic programming [Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press (1992)]. I recount and expand the celebration at GECCO 2022, very briefly summarise some of what the rest of us have done and make suggestions for the next thirty years of GP research.

Keywords: genetic programming, genetic improvement, modularity, scaling, parallel computing

1 Introduction

An evening at the 2022 GECCO conference was devoted to celebrating the thirtieth anniversary of the publication of John Koza’s book “Genetic Programming: On the programming of computers by means of natural selection” [1]¹. Indeed that is the purpose of this special issue of Genetic Programming and Evolvable Machines. I hope to put my own spin on and fill out points raised in that panel discussion (which was recorded and is available on line²). I should stress this is not a survey of GP and that many valuable contributions are omitted. Similarly many digressions are placed in footnotes and there are hyper links to online articles in Wikipedia etc.

University College, London, Gower Street, London, WC1E 6BT, UK

¹ Named after Charles Darwin’s 1859 foundational book “On the Origin of Species by Means of Natural Selection” [2] which contains huge volumes of evidence (for example gathered on his five year voyage around the world [3]) in support of his scientific theory of evolution, which after a struggle was eventually accepted as the explanation for biology.

² A Conversation with John Koza, 30 years after the publication of Genetic Programming Sunday, July 10, 18:00-20:00 2022 <https://whova.com/portal/webapp/gecco.202207/Agenda/2516377>




Fig. 1 Prof. Dr. Wolfgang Banzhaf holding his copy of “Genetic Programming: On the programming of computers by means of natural selection” (Jaws) 834 pages [1] at the GECCO 2022 celebration of 30 years after its publication (Wolfgang says he was told that his copy was the first one sold by the bookshop in Boston.)




Fig. 2 At 834 pages, the first genetic programming book [1] weighs in at 4lb 2oz.

1.1 The Book

Dr. Amy Brand, Director of The MIT Press, was clearly delighted that John Koza had chosen MIT Press to publish the first book on genetic programming [1] (see Figures 1 and 2). She says it “was one of the seeds from which

sprang a whole ecosystems of books and journals at the intersection of computer and biological sciences for the MIT Press.” Adding it “is still available and selling in print-on-demand. That’s quite solid for a specialized and groundbreaking work in computer science from 1992.”³

John Koza said that the motivation for the book was his team in the preceding five years had published GP solutions to 81 diverse problems common to artificial intelligence, machine learning and knowledge based systems. They had shown that instead of, as had previously been done, using a solution technique devoted to each benchmark, a single evolutionary computing technique (now named Genetic Programming⁴) could solve them all^{5,6}. However the GP solutions were published in widely disperse conference venues. The goal of the book was to convince everyone that 1) a single technique could solve many diverse problems and 2) they could all be recast as the problem of searching for (and finding) a computer program. Whereas previous solutions had often used (non-evolutionary) search but used a representation, e.g. graph, grammar, network, often purpose built for each benchmark. The size of the book⁷ stems from the need to convince people that GP is a general solution. Whereas everyone who first comes to programming knows that programming languages are exceedingly picky about insisting they get everything, every comma, every semicolon, in the right place: so how could random stand a hope? Hence a substantial book, backed by a video, would be necessary to convince a skeptical public⁸.

³ John Koza’s publications have been at the top of the list of publications downloaded via the genetic programming bibliography since 2006, when download statistics were first gathered.

⁴ The name Genetic Programming was suggested by David E. Goldberg. John Koza said he was originally reluctant to use the name but came to realise it was a brilliant choice.

⁵ John Koza has previously likened GP’s success with early machine learning benchmarks with Sherman’s march through Georgia in 1864, which helped end the four year USA civil war.

⁶ In the late 1990’s Peter Nordin reported similar success with his linear genetic programming on the UCI machine learning benchmarks.

⁷ The first genetic programming book was colloquially known as “Jaws” after the 1975 Hollywood movie of the same name, where the shark appears to get progressively bigger throughout the film. In a similar way Koza remarked that as each new GP experiment was covered, the book got bigger, eventually exceeding 800 pages. The three succeeding GP books, are similarly known as Jaws 2 [4], Jaws 3 [5] and Jaws 4 [6], all four are each accompanied by an hour long video [7,8,9,10] (now available on YouTube and www.human-competitive.org). In 2009, John Koza gave a seminar at Stanford summarising his GP work which was recorded and is also available on YouTube [11].

⁸ There is a growing body of work, such as automatic bug fixing [12] and genetic improvement [13], that shows ordinary programs are not fragile [14,15,16,17,18,19,20,21]. The misplaced semicolon problem refers to the source code syntax as understood by the language compiler (another computer program). Since the syntax is formally defined, computer generated mutations can be automatically written to be syntactically correct. If mutated code compiles, it often runs and produces an answer which can be fed into a fitness function.

1.2 The Man

John R. Koza was born 1944 and did both his undergraduate degree and PhD at the University of Michigan in Ann Arbor, studying mathematics and, the then newfangled, computer science. He reports great interest in playing games including computer games, with students and faculty, for example, John H. Holland. As with John Holland's other students⁹, he was well versed in John Holland's genetic algorithms.

1.3 The Millionaire

John Koza graduated from the University of Michigan in December 1972, and using his mathematical skills in combinatorics, probability and game playing he joined a lottery company which printed games on paper which were sold at petrol stations and supermarkets. In 1974 he and a colleague formed their own company, Scientific Games Inc., to exploit John Koza's invention of a secure way of printing scratch off lottery tickets. They successfully lobbied various USA states to allow them to run the state's lottery¹⁰. By 1978 the technology of printing had moved on and they jettisoned their own technique in favour of more flexible computer based printing. In 1987, having made his fortune, he returned to research.

1.4 The Researcher

From about 1987 until 2005, John Koza devoted himself to research, applying genetic algorithms to the discovery of computer programs (GP). He published some 208 items, predominately papers but also book chapters, technical reports, proceedings, etc. and of course *Jaws* [1] and the three follow-up up door stoppers [4,5,6] and the four accompanying videos [7,8,9,10]. Initially the genetic programming systems were written in Lisp, although later implementations were in C, e.g. [22].

There were GP workshops associated with the International Conference on Genetic Algorithms, ICGA-93¹¹ and again in the summer of 1995 at ICGA-95 and ICML-95. In the fall, John R. Koza and Eric V. Siegel organised a GP event with the 1995 Fall Symposium of the AAAI in MIT. In 1994 Kim Kinnear had launched the "Advances in Genetic Programming" edited book series published by MIT Press [23,24,25]. But, since ICGA was a biannual conference, there was no ICGA conference in 1996, and instead it was the right time

⁹ John Holland's PhD students include: Stephanie Forrest, Tommaso F. Bersano-Begey, Melanie Mitchell, Tom Westerdale, Lashon Booker, Ted Codd, Clare Congdon, Dave Goldberg, Annie Wu, Ken De Jong, Leeann Fu, Rick Riolo, Chris Langton, Robert Reynolds, Bernie Zeigler and John Koza.

¹⁰ By 2009 the combined profits to the USA state governments which permitted lotteries had reached \$17.6 billion.

¹¹ ICGA had strong links with John Holland's students.

to launch the first GP conference [26]. One of the rules laid down at GP-96, was the absolute need for independent peer review.

July 1997 saw the return of ICGA-97, carefully scheduled a few days after the second GP conference, GP-97 [27], so attendance at both was encouraged. Again there was no ICGA in 1998, instead at GP-98 [28] there were serious discussions about combining the growing number of evolutionary computing conferences. John Koza in particular felt that the separate EC events were splitting EC into separate communities, and that the balkanisation of EC, did not make sense to people outside, particularly to funding bodies. And that this divergence was hurting the field. So at GP-98 there were negotiations about unifying, particularly: the Evolutionary Programming Society conference (EP), the IEEE's WCCI/ICEC, GP, ICGA, and the International Workshop on Learning Classifier Systems (IWLCS). These were only partially successful, leading in 1999 to the formation of the duopoly of CEC 1999 [29] and GECCO 1999 [30]. Of the european evolutionary computing conferences, only the IEE's Galesia elected to join CEC. PPSN¹², ICANNGA and the newly established EuroGP¹³ [31] continued as before¹⁴.

Again John Koza's organisational skills came to the for, with him helping to draft the byelaws for GECCO. These ensure it has a federal "big tent" structure, whereby none of its constituent groups would feel left out or put down by the others.

Having progressed genetic programming to the point where it could be described as a routine invention machine [6,32,33], John Koza turned to public service and electoral reform and in 2006 founded National Popular Vote.

1.5 The Public Benefactor

In 2004 John Koza started the annual "Humies" awards for human-competitive results produced by genetic and evolutionary computation. He continues to fund the cash prizes. The finals are held each year as part of the GECCO conference.

Since 2016 he has endowed Michigan State University with the first chair in genetic programming in the United States (held by Prof. Dr. Wolfgang Banzhaf).

¹² Parallel Problem Solving from Nature (PPSN) had started in Germany in 1990. It is a also a biannual conference on evolutionary computing and, although based in Europe, it was held on alternating years with ICGA. Like Genetic Programming, PPSN was also named by Dave Goldberg.

¹³ The First European Workshop on Genetic Programming had been held in 1998 in Paris, with the help of EvoNet, the EU Network of Excellence in Evolutionary Computing.

¹⁴ In 2003 John Koza listed 25 international conferences and workshops primarily devoted to the various forms of evolutionary computation. Many are still held annual or biannually, and some have combined. In most cases the proceedings are still available, often on line.

1.6 Pre-History

At GECCO-2022 the question of research before genetic programming was raised. John Koza pointed out that by 1987 the field of Genetic Algorithms was already well established¹⁵. There had been early experiments on machine learning in Columbia [34] and Manchester [35]¹⁶ universities. However John Koza traced Evolutionary Computing back to Alan Turing. He said Turing's 1948 paper on machine intelligence [36] suggested three routes to machine intelligence: 1) knowledge based, 2) based on logic (as would be expected of a mathematician), but John Koza highlighted the third: 3) in which machine intelligence was based on evolution. Although he pointed out it did not use crossover (which was added by John Holland).

1.7 Advice for the Future

Another question raised at GECCO-2022 was did John Koza have advice for new researchers. His answer was researchers must keep current, i.e., keep up to date with research, but not just in your area but with research in general. Take an interdisciplinary approach. He stressed be open to ideas from elsewhere, particularly from Biology.

John Koza's heuristic (perhaps common to all John Holland's students) was to ask himself "What would John Holland do?" to which the answer was often: John Holland would respond with his own question, "What does Nature do?" John Koza's particular example was: how did Nature evolve from microscopic organisms (like bacteria) which have genes for creating may be about 500 proteins to multicellular organisms (e.g. us) which have genes for creating about 20000 proteins. He reported asking this question around the Stanford School of Medicine.

The example John Koza quoted was the evolution of Myoglobin and Hemoglobin, which is thought to have occurred via gene duplication and subsequent specialisation. The idea being: "accidental" copying of parts of DNA sequences is common¹⁷. Once a species has two copies of a vital gene, it may be free to tinker with one. Since the other gene remains functional, the children with the duplicated gene remain viable and so some can survive long enough to carry both the working gene and the tinkered copy to the grand children. Over subsequent generations the two genes may diverge allowing the species to find new proteins which may help it survive. Susumu Ohno in his 1970 book [39] suggested that such gene duplication is a powerful mechanism in natural evolution. Indeed John Koza used it as inspiration [40] for his

¹⁵ In addition to genetic algorithms, there is early work on evolutionstrategie in Germany by Ingo Rechenberg and Hans-Paul Schwefel, and in the USA on evolutionary programming by Larry Fogel.

¹⁶ Kilburn, Grimsdale and Sumner ran their experiments in machine learning and thinking on the world's first digital stored program computer the Manchester Mark 1.

¹⁷ The evolution of repeating patterns in DNA due to crossover is common. Indeed crossover in GP can readily produce huge volumes of repeated code in trees [37,38].

architecture-altering operations. These GP operations allow, not just the code within automatically defined functions (ADFs) [4] to evolve, but also their structure (e.g. which ADF calls which ADF) evolves [41,5], [9, minute 10]. In terms of traditional AI, this can be thought of as dividing the whole problem into subcases and having an evolvable representation which facilitates not just the solution of the sub-problems but also their subsequent combination into a complete solution. Some form (or indeed many forms of) automatic problem decomposition is essential if any AI technique is to scale.

John Koza felt that in the 1960s the University of Michigan had had a wide ranging curriculum. He said computer scientists need to know about biology, language processing, psychology, information theory, electronic circuits, etc. However, this breadth has been lost from modern computer science curricula. Instead people should seek ideas from many places. He cited successful start ups in silicon valley, such as Adobe, which had come from co-working between two people with experience of newspaper publishing and another with a computer science background. Often in silicon valley success had come from partnerships of individuals with different experience. Alternatively, success may arise when different experience or many odd ideas are held by *one* person.

I would like to add, be ambitious in the problems you tackle. John Koza's impact, the impact of his book [1], stems from showing something widely viewed as *impossible* could be done. Before his work, the idea of automatically evolving a computer program was clearly ludicrous. Similarly, the idea of a computer fixing computer bugs was clearly impossible, until Stephanie Forrest et al. showed GP could do it [42]. Readers may remember Lewis Carroll's Alice and the White Queen [43] (Figure 3), Alice reproaches the White Queen for some nonsense, saying it is clearly impossible, to which the White Queen responds that Alice should practice believing the impossible. My suggestion would be to an ambitious researcher that she should *do* the impossible. Claire Le Goues was a PhD student in 2009 [44,45]. Fortunately her adviser did not tell her her idea was impossible. And so She and the team are famous, not because they completely solved the problem, but because they took something impossible and partially solved it. So that today the argument is not if it can be done, but what is the best way [12] to solve the previously impossible problem [46,47,48].

1.8 The Ones That Got Away: Missing Gaps

John Koza was asked to muse on his less successful experiments. Two came to mind: FPGAs and GPUs.

1.8.1 Genetic Programming and Field Programmable Gate Arrays, FPGAs

John Koza had hope to create a field programmable gate array (FPGA), which had all the likely to be useful program operations pre-loaded. An ultra fast



Fig. 3 When I was your age I could think of six impossible things before breakfast.

evolved GP program would then simply be an evolvable way of linking these together.

In some ways this seems similar to Juille's [49] way of running a GP interpreter on the hugely parallel MasPar MP-2 computer. Although it had thousands of processing units, they each did the same one thing at the same time. Juille's brainwave was to say: since computing is cheap, we will discard most of it. (Simplifying), Juille built a tiny interpreter which ran on all processing elements one of a handful of GP operations. The different members of the GP population were spread across the processing elements. Each with its own program counter. If the interpreter was currently executing a GP op code that was not the one the GP individual wanted, it did nothing but wait. However the interpreter cycled round all possible GP op codes. When it did reach the desired op code, that processor executed it and moved that GP individual's programme counter on by one. (The right hand side of Figure 4 shows the same idea in the context of GPUs.)

It sounds hideously inefficient, but bear in mind the GP is getting useful work done, whereas mostly human programmers could not handle the MasPar MP-2's SIMD architecture efficiently at all. Secondly often in many high performance computers (HPCs), most of the time the processing elements are waiting for data to arrive and so spend most of their time spinning in idle loops. This turns on its head our common conception of computers. In HPC (and indeed GPUs, see Section 1.8.2), computing is often cheap compared to moving data. Indeed sometimes it can be more efficient to compute a value a second time, rather than store it and retrieve it later when it is needed¹⁸.

In many cases FPGAs form the bed rock of evolvable hardware (EHW) [50, 51]. As well as offering a cheap and flexible alternative to dedicated integrated circuits (also known as application-specific integrated circuits, ASICs) they can be cost effective, particularly when only a limited number of chips will be

¹⁸ A second recommendation to the novice computer scientist, do not assume that a very old paper has no merit. Computer science is littered with examples of old ideas which returned, e.g. virtual memory, virtual machines and Maurice Wilkes' microcode.

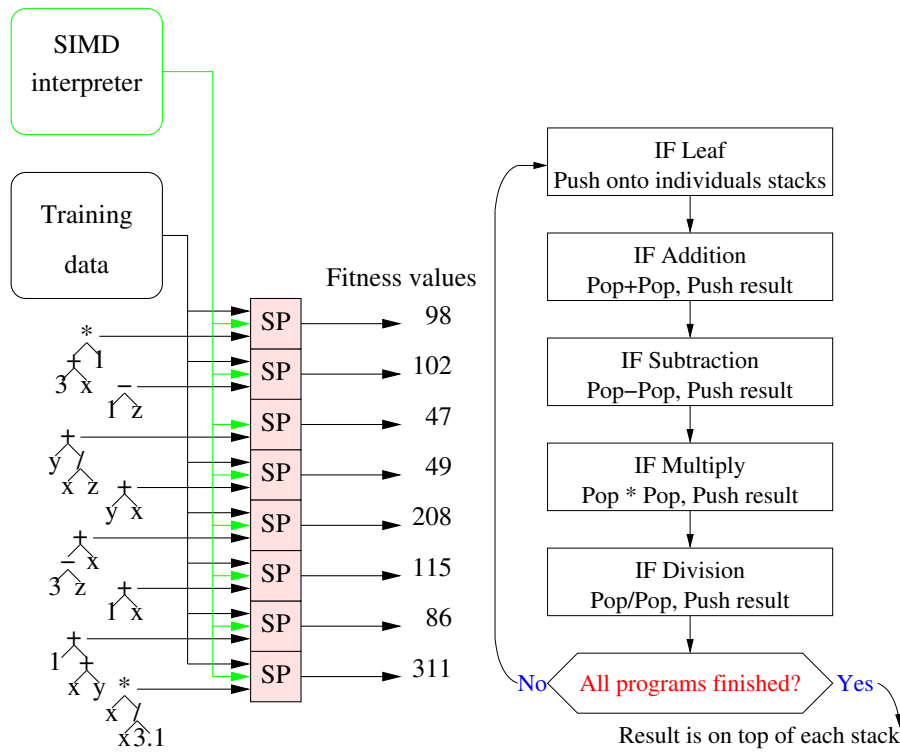


Fig. 4 Left: Avoid compilation overhead by interpreting GP trees. Run single SIMD interpreter on GPU’s stream processors (SP) on many trees. Right: Programs wait for the interpreter to offer an instruction they need evaluating. For example an addition. When the interpreter wants to do an addition, everyone in the whole population who is waiting for addition is evaluated. The operation is ignored by everyone else. The interpreter moves on to its next operation. The interpreter runs round its loop until the whole population has been interpreted. Fitness values can also be calculated in parallel.

needed. There are several examples where FPGAs have been used to run GP, e.g. [52,53,54].

1.8.2 Genetic Programming and Graphics Cards, GPUs

In the early 2000’s it was noticed that the graphics cards (GPUs) used to drive computer screens were becoming increasingly powerful parallel computing devices in their own right and so people started using them for other things.

Initially GPUs were designed just to rapidly render images on the computer’s screen. To do this quickly (in real time) they comprised many parallel components all doing the same thing but for different parts of the screen. As the computer video games market took off, GPUs rapidly ramped up their processing abilities and power. Each parallel component became a fully functional processor, often with special support for operations common in graphics applications (such as reciprocal square root [55]). This was so that more of the

parallel aspects of generating, rather than simply displaying, real time video could be devolved from the (serial) CPU to the (parallel) graphics card. As GPUs were often somewhat independent of the end users' computer mother board, keen video gamers could easily upgrade their GPU. This promoted rapid technological improvement, as rival GPU manufactures sought sales by offering better and/or cheaper hardware than their rivals. However even today, GPUs essentially (like the SIMD MasPar, page 8) require their parallel processing elements, to do the same thing at the same time.

Initially GPUs were very hard to program and their support software was only designed to be used by dedicated programmers employed by video game companies. However the abundant and cheap parallel processing the GPUs offered was taken up by scientific programming, leading to the field of General-Purpose Computing on GPUs (GPGPU) [56]. As GPGPU became more popular, the GPU manufactures, particularly nVidia provided much better software support.

At first in genetic programming GPUs were only used to speed up fitness evaluation, e.g. work by Simon Harding [57]. and Darren Chitty [58]. Indeed it was said that, due to the GPUs peculiar SIMD architecture, running the GP interpreter on the GPU was *impossible* (cf. Figure 3). Of course this was not true, and inspired by Juille's work with the MasPar SIMD supercomputer [59] (page 8), I built a SIMD interpreter for nVidia's GPUs (see Figure 4) [60, 61]¹⁹. See also [63, 64, 65, 66, 67, 68, 69]²⁰.

As the memory available on the GPU cards increased, it became possible to work with huge populations of small GP trees. In [70] I used a cascade of GP populations to winnow useful bioinformatic data from more than a million GeneChip features. The top level GP populations contained more than five million individuals trees. This GPU application could scale from a \$50 GPU to a top 500 super computer [71]. Figure 1 in [72] shows the dramatic improvement in nVidia GPU speed (2003 to 2012, which still continues), whilst Table 3 in [73] shows some high performance parallel GP implementations, almost all running on GPUs²¹.

1.8.3 Deep Learning and Accelerators: GPUs and TPUs

Due to the availability of internet scale data sets and GPGPU processing power, since 2010 the field of deep learning has taken off [76]. It is generally accepted that researchers need a GPU (possibly a whole cluster of GPUs) to

¹⁹ People also said that it was impossible to create random numbers on GPU. Again not true [62]. These days pseudo random number generators PRNG, (e.g. CuRAND) are supplied by nVidia with its CUDA software.

²⁰ The metric "Genetic Programming Operations per Second" (GPops) permits easy comparison of performance, e.g. across different implementations and hardware.

²¹ Recent extremely high performance on Intel multi-core CPU SIMD hardware [74] has been achieved, with unchanging fitness functions, in large trees, in converged populations [75], that do not have side effects (and so can be evaluated in any order). This is because a lot of work can be omitted, if it can be proved that a child has the same fitness as its parents.

do any form of competitive deep neural net learning. Even with the availability of cloud computing, this may soon have the effect of “pricing out” individual academic researchers from the future of deep learning [77].

Sometimes the whole notion of using a GPU to drive a computer’s screen (also called the computer’s monitor) may be disregarded. Often called “headless” GPUs, to save space and power, some GPUs dispensed with the screen interface altogether. An extreme examples of this is Google’s TPU, which is totally specialised to Artificial Neural Network (ANN) processing.

As gaming and now AI have become more important, the notion of a GPU as a cheap alternative to the computer’s CPU has also faded, and now a top end GPU can cost more than a CPU.

1.9 Other Gaps: Memory, Theory, Bloat

John Koza mentioned that even though Jaws [1] did not include much work on evolving memory, he regarded it as important because it provides another route to allow re-use. Since a value stored in memory can be re-used, potentially many times, without the code for it having to be evolved more than once. He mentioned my book [78], although using indexed memory in GP is due to Teller [79]. Surprisingly, there has been a steady stream of research on evolving memory within GP [80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115, 116, 117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 127, 128, 129, 130, 131, 132].

Genetic programming theory has a variety of forms [133]. Jaws [1] starts with adapting the then current explanations of how linear bit string genetic algorithms work, due to John Holland and Dave Goldberg. Such schema theories were also analysed by Una-May O’Reilly [134], Justin Rosca [135] and most notably by Riccardo Poli [136]. Another popular thread is to take ideas from biology about how evolution works and use them to understand GP [137], e.g. Price’s theorem [138] [139], population convergence [140, 141, 75] and neutral networks (plateaus) [142] in fitness landscapes [143] [144, 145, 146, 147]. Similarly biology has been an inspiration for other search operators, such as homologous crossover [148]. In recent years there has been a flowering of formal or rigorous run time analysis in evolutionary computing and some success applying mathematical techniques to GP problems [149, 150, 151, 152, 153, 154]. Of course it is difficult to make such theorems widely applicable and when using results we must remember the inevitable assumptions they require. For example, SAT has been proved to be NP-complete. Nevertheless in the last decade considerable progress has been made with practical SAT solvers and they are now routinely applied, e.g. in software engineering. Similarly, the No Free Lunch theorem [155] applies to GP (as with all optimisers) but fortunately (as in other branches of AI) that has not inhibited development of the

field. Although, as noted above, there are exceptions, but genetic programming as a whole remains a deeply empirical endeavour with many new ideas being reported. However it is difficult to persuade authors to carefully analyse their evolving populations of programs so as to be able to explain why their experiment succeeded (or even why it failed).

Although John Koza reports [1] bloat²² from the start of genetic programming, the tendency, indeed the name, for programs to be bigger than necessary is not unique to GP. Bloated human written programs are common. Indeed people writing computer programs with unnecessary instructions goes back to the very beginning of electronic digital computers, with bloat reported in programs run on the first stored program digital computer, the Manchester Mark I [35]. This human tendency is rampant, with some Internet code bases having grown to over a billion lines of code in less than 20 years. Bloat continues to be a well studied topic in GP with 426 entries in the GP bibliography mentioning it.

Although there are potential ways of mitigating bloat's impact on runtime [156] and reducing its memory requirements with DAGs [157] (indeed bloated trees produced by crossover [158] should be highly compressible), in practice bloated populations can quickly overwhelm the available computer resources and so the common approach is to shut bloat down. For example, by enforcing either depth or size limits on the evolving programs. However this is not risk free [159] and more sophisticated approaches may be wanted. For example, controls on selection, such as using multiple fitness objectives (e.g. a size versus performance Pareto trade-off [160,161]) or tighter controls on offspring generation [162,163,164,165]. In many cases bloat appears to be an unexpected aspect of early (even premature) convergence and so has some similarity with overfitting sometimes seen with artificial neural networks (ANNs), where prolonged search drives locally improved performance on the training data. This gives a more convoluted mapping between the ANN's inputs and outputs but at the possible expense of the ANN's ability to generalise to unseen data. Where the goal is to explain or predict, such complexity or overfitting is clearly unhelpful. In ANN anti-over fitting techniques are essential. These include stopping training early (i.e. in GP terms using fewer generations), regularization [166,167,168], changing the training data during training [169,170] and even expression simplification [171], either during evolution [172] or to increase comprehensibility and explainability, cf. XAI, after GP has finished [173]. Whilst Dale Hopper [172] and other authors, ensure their automatic rewrite of GP individuals gives a semantically equivalent but smaller replacement, in many cases this is not wanted. When a 100% correct program is not realistic, e.g. on many prediction tasks, it may be better to accept (or allow evolution to find) a similar but much simpler program, rather than spending a lot of effort creating an exactly equivalent program to what is essentially only an approximation.

²² Bloat is the tendency for programs to grow in size without commensurate increase in performance.

However, bear in mind that evolution is a hacker. It builds on what was there before. In biology evolution overfits. Classic examples include: 1) the Giraffe's left laryngeal nerve, which runs the whole length of its neck from its head, round the aorta in its chest and then returns to its throat at the top of its neck, because evolution did not find a shorter path, 2) the male peacock's heavy tail which helps secure a mate but impedes flight and 3) the human brain which consumes 20% of our food [174] but made our ancestors more appealing as mates to other members of their tribe [175].

2 A Brief Selection of Other Genetic Programming Work

In addition to continuing with evolving Lisp like trees, major branches of genetic programming include: linear genetic programming [176] cartesian genetic programming (CGP) [177] and grammatical evolution (GE) [178], all of which use a linear chromosome. Following John Koza's automatically defined functions, ADFs, see page 7, there were several attempts to encourage the evolution of modular programs using individuals with multiple trees or libraries of subtrees [179,180,181,182]. However, these seem not to have taken hold.

As with evolutionary computation in general, the major computational cost of GP is usually evaluating fitness [1, p783]. In tree GP this is usually the cost of interpreting the trees. When members of the population are going to be run many times²³ it may be worth the cost of compiling the population and then running the compiled programs²⁴ rather than interpreting them [187]. However, as Ronald Crepeau showed [188], for GP, it is not essential to run a full blown compiler, instead knowing the restricted set of primitives used by GP, he constructed a dedicated fast compiler which converted the evolved code into machine code and ran that directly.

Peter Nordin eliminated the compilation step entirely by using GP to evolve firstly Sun 32 bit SPARC RISC architecture machine code [84] and later Intel x86 binaries [189] (which in turn later became Discipulus [190]). He used tailored mutation operations which respected the layout of the machine code. Although perhaps first motivated by speed and simplicity, the idea of evolving variable length linear programs has taken off [191,192].

Grammatical Evolution (GE) [193,194] shows the virtues of trying ideas out. Michael O'Neill and Conor Ryan took the idea of a variable length linear chromosome, simplified it to become just an ordered list of byte sized integers (0..255) and married it to another favourite of computer scientists: the Backus-Naur form grammar (BNF). Pretty much anything which can be run on a computer can be expressed in a BNF grammar. They disregarded

²³ Considerable savings are sometimes possible by accepting fitness selection will be somewhat random and using a cheaper and approximate fitness function. After all, the goal of a fitness function, is not to measure performance (that can be done after the run) but to guide search. Why run thousands of fitness cases, when fitness will ultimately be reduced to a single bit: does this individual get a child or not? Even then, we typically add noise to this bit, e.g. via tournament selection, [183,184,185], see also [186, Sect. 10.1].

²⁴ Lisp provides compilation as an alternative to interpreting programs [1, p785].

that BNF is essentially tree shaped and trusted in evolution to find a way of putting them together. The linear stream of bytes is mapped using modulus to say which branch to take next in the grammar. If there are not enough bytes, we simply wrap round and start again from the first. If there are too many, we ignore the excess. The resulting grammar is treated as the individual's phenotype and in a problem dependent way converted into a trial solution with a fitness value. The sloppiness of the mapping from genotype to phenotype offended some and provoked wide discussion in a peer commentary issue of "Genetic Programming and Evolvable Machines" [195]. But as Conor Ryan says "GE works" [196]. Indeed the separation of genotype from BNF grammar makes grammatical evolution flexible and has been widely used. (The GP bibliography contains well over seven hundred entries relating to grammatical evolution.)

With Cartesian Genetic Programming (CGP) [177, 197, 198, 199, 200], Julian Miller turned to a fixed representation, more a kin to traditional bit string genetic algorithms (GAs). However the chromosome is a fixed sized two dimensional rectangle, rather than a single string, where each cell contains a digital computational unit, such as an XOR gate. Both the contents of the cells and crucially the connections between them are evolvable²⁵. Notice, like linear GP (but unlike GE), evolution directly sets the contents and connections of each cell (i.e. evolution acts directly on the phenotype). Also there is no explicit left-right flow of control. In CGP the chromosome is treated as a circuit and so its evaluation has to take note of where data enters and leaves. It is also not necessary to evaluate cells which are not connected. Cartesian GP has been widely used, including in the evolution of approximate computing [201, 202], where evolution can be well suited to finding good trade-offs between conflicting objectives, such as fidelity, size, number of components, power consumption and speed.

2.1 Inspired by Computer Science

In order for subtree crossover to freely mix subtrees from parents to create children, John Koza required the components of his GP trees to have closure [1, Sect. 6.1.1]. Meaning 1) any leaf or function in the tree can be an argument to any other function. Since components typically communicate via function return values, this often means GP trees use a single type, often float. 2) To ensure each function can deal with any combination of inputs, many functions have protected GP versions. Such as protected log RLOG [1, p83], which returns a defined value (rather than raising an exception) even if its input is zero or negative. Alternatives might be to allow evolution to deal with the exception, or simply assign poor fitness to individuals with illegal combinations.

²⁵ Nowadays in cartesian genetic programming people often set the width of the FPGA like rectangle of components to 1, making it effectively a string. Similarly it is common to allow only feed forward connections, so preventing recurrent loops.

However notice that ruling it out prevents GP exploring not only this tree but all the trees that might have evolved from it.

Perhaps the most famous extensions to closure are Dave Montana’s strongly typed GP [203] and Tina Yu’s polymorphic GP [204,205] which allow multiple types but ensure evolution explores only type safe expressions. Another approach is to use various types of grammar to try and keep evolution in the most productive parts of the search space [206]. For example, using context free grammars [207,208], using grammars to ensure the evolution of expressions which are dimensionally consistent [209], using tree-adjunct grammars to guide GP (TAG3) [210] and using GP with Lindenmayer Systems (often abbreviated to L-Systems) [211,212,213,214].

Whereas Lisp and most GP systems implicitly use the system stack, programs which explicitly use a stack [215,216], e.g. to pass vectors and matrices [217], are also possible. An explicit stack allows the evolution of Reverse Polish Notation (RPN) [61] and even infix expressions [218]. In PushGP [219] there are multiple stacks, one per type. These may include a code stack, so allowing GP to manipulate code, thus permitting GP to evolve its own genetic operators.

2.2 Non Genetic GP

John Koza’s GP [1] is clearly strongly influenced by his PhD supervisor, John Holland, and GP [1] is essentially the application of John Holland’s genetic algorithms to the evolution of Lisp s-expressions, i.e. tree shaped programs. But, as we have seen, the programs need not be trees, and similarly the search algorithm does not have to be a genetic algorithm. Other techniques include: local search, Simulated Annealing [220,221], Differential Evolution [222], Bayesian probability search [223], Estimation-of-Distribution Algorithms (EDAs) [224,225] Ken Stanley’s Neat [226,227,228] and even deterministic search, e.g. Trent McConaghy’s FFX [229]. Indeed search does not have to be guided only by fitness but can “look inside” the program [230] and its execution [231]. SR-bench [232] compares many GP and non-GP approaches to symbolic regression, including MRGP [233], M3GP [234], FEW [235] and Operon [236].

2.3 Less Explored

2.3.1 *Assembly Code, Byte Code*

In human terms assembly code is usually viewed as intermediate between high level languages and machine code. Offering the potential advantage of machine code (speed and compactness), and ease of use and readability of high level source code. There has been very little GP work on evolving assembly code. Exceptions include microcontroller assembly [237], nVidia GPU PTX [238,239] and the intermediate (IR) code used by LLVM [240], and again on GPUs [241].

Java, and some other interpreted languages, compile the source code into byte code which they then interpret. Eduard Lukschndl showed it is possible to run GP at the level of Java byte code [242].

2.3.2 Modularity, Recursion, Loops

Some of the work on encouraging the evolution of modular code was mentioned on page 13. In *Jaws*, John Koza described GP solving the Fibonacci problem [1, pp473–477] as an example requiring the evolution of recursion and several examples where GP evolved do-until loops and other forms of iteration, but again there has been relatively little work on either by others. Again a few exceptions. These include work by Peter Whigham [243,244] and Tom Castle [245].

2.3.3 Coevolution

As with many topics, there are examples of co-evolution [246,247] in *Jaws* [1] and many elsewhere in genetic programming [80], for example in agent learning [248]. However, it does feel like coevolution has not yet fulfilled its potential. In deep artificial neural networks there is interest in antagonistic adversarial learning and so perhaps this will stimulate renewed interest in coevolution in genetic programming.

3 The Future

At GECCO 2022 Erik Goodman asked if there were any applications of GP that had surprised John Koza. Amongst the many human competitive [6] results, perhaps one of the most encouraging is quantum computing. As with quantum physics, quantum computing has a deserved reputation for being difficult for people. However, the rules about quantum computing gates can be coded for GP to use without being an expert quantum physicist, and then GP can be left to evolve novel quantum circuit designs incorporating them [249,250,251]. Riccardo Poli, Leonardo Vanneschi and others have previously reported on the state of GP and in particular what remains to be done [252,253].

In genetic improvement [13] existing (human written) software is optimised (typically by using GP). Notice genetic improvement does not start from primordial ooze [1]. Instead search automates the potentially labour intensive, tedious and error prone task of finding modifications. For example, to repair bugs [42,46,12,254,48], including energy bugs [255], reducing memory consumption [256], reduce run time [173,257,258,259,260,261,262,263,264] improve existing functionality (e.g. to give better predictions [265]), porting to new hardware [266] including improving GPU applications [261,262,263,264,267,241] or even to incorporate existing functionality from outside the existing code base [268].

The idea of mixing evolutionary computing (including GP) with other optimisation tools to give hyperheuristics [269] has a long history. In particular, with the recent explosion of interest in deep artificial neural networks, combining evolutionary learning and artificial neural networks seems set to continue. One particularly encouraging trend is AutoML tools such as TPOT [270, 271] which automatically tune existing machine learning pipelines.

In GP, as in most optimisation problems, most of the computation effort is spent on evaluating how good the proposed solutions are. Various ideas for speeding up fitness evaluation have been proposed, for example surrogate fitness functions [272]. Colin Johnson’s Learned Guidance Functions [273] seem a particularly elegant approach to making best use of previously gained knowledge. It would be interesting to see Learned Guidance Functions applied to genetic programming or when using genetic improvement to adapt existing human written programs.

Since all digital computing progressively loses information, information about crossover and mutation gets progressively washed out the further it has to travel. In nested functions without side effects, deep genetic changes become invisible to the fitness function. Thus to evolve complex programs, they must remain shallow and so I propose that to evolve large complex programs, they be composed of many shallow trees, within a strong low entropy-loss data interconnect to and from the environment. This should ensure that the good and bad effects of most genetic code changes are externally measurable [274].

At GECCO John Koza pointed out that in both biology and in human design, modularity and reuse are ever present. Biology scales from a single cell to individuals containing billions of cells. It does this, like human engineers, not by solving many billions of individual problems but by reusing existing designs. We need to revisit the scaling problem.

4 Conclusions

We have seen that in the thirty years since John Koza published his first GP book, the field has blossomed. The genetic programming bibliography contains some 16 367 entries by 16 342 authors²⁶. Many of the genetic and evolutionary computation papers judged to be the best human competitive work of each year have used genetic programming. Clearly GP is doing well in its mission to help the world.

As mentioned at the end of the last section, although GP continues to flourish, perhaps we need to tackle the scaling problem. Are we evolving small things? Do we need to be more ambitious? Following Stephanie Forrest’s recent questions [275]: what could GP do with Google Deep AI scale resources?

As John Koza foresaw, 30 years of Moore’s law [276] (with component count doubling every 18 months) means 20 lots of doubling ($2^{20} = 1\,048\,576$). That is,

²⁶ The GP bibliography was started by John Koza. Although in recent years it has undoubtedly missed some work, in the five years before the pandemic (i.e. 2015–19) there were 3340 new entries and 5177 authors published at least one GP paper.

since the genetic programming field started, the computer power available to us has increased a million fold. What of the next 30 years? Perhaps Moore's Law will end? Certainly the death of Moore's Law has been confidently predicted many times. What seems certain is we will not see dramatic increases in silicon computing's clock speeds. Instead we anticipate the future of computing will be ever more parallel. But as John Koza says GP is *embarrassingly parallel*. Indeed the use of distributed parallel GP populations, not only makes good use of current and future compute resources but is in keeping with Sewall Wright's [277] model of natural evolution and as John Koza reports by keeping population diversity, the distributed population demes of the island model, improve GP results as well as speeding it up.

In 2052 will genetic programming researchers be using computers a million times faster than they use today? Certainly GP seems well placed to exploit them.

Acknowledgements

I would like to thank Sean Luke, and Andrew and Claire, and my anonymous reviewers. No competing interests. Funded by the Meta (formerly Facebook) Oops project (Award number 181551).

References

1. Koza, J.R.: Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press, Cambridge, MA, USA (1992), <http://mitpress.mit.edu/books/genetic-programming>
2. Darwin, C.: On the Origin of Species by Means of Natural Selection. John Murray, penguin classics, 1985 edn. (1859)
3. Darwin, C.: Voyage of the Beagle. Henry Colburn, London, penguin classics, 1989 edn. (1839)
4. Koza, J.R.: Genetic Programming II: Automatic Discovery of Reusable Programs. MIT Press, Cambridge Massachusetts (May 1994), <http://www.genetic-programming.org/gpbook2toc.html>
5. Koza, J.R., et al.: Genetic Programming III: Darwinian Invention and Problem Solving. Morgan Kaufmann (Apr 1999), <http://www.genetic-programming.org/gpbook3toc.html>
6. Koza, J.R., et al.: Genetic Programming IV: Routine Human-Competitive Machine Intelligence. Kluwer Academic Publishers (2003), http://dx.doi.org/10.1007/0-387-26417-5_1
7. Koza, J.R., Rice, J.P.: Genetic Programming: The Movie. MIT Press, Cambridge, MA, USA (1992), <https://youtu.be/tTMpKrKkYXo>
8. Koza, J.R.: Genetic Programming II Videotape: The next generation. MIT Press, 55 Hayward Street, Cambridge, MA, USA (1994), <http://www.genetic-programming.org/gpvideo2.html>
9. Koza, J.R., et al.: Genetic Programming III Videotape: Human Competitive Machine Intelligence. Morgan Kaufmann, 340 Pine Street - 6th Floor, San Francisco, CA 94104, USA (1999), <http://www.genetic-programming.org/gpvideo3.html>
10. Koza, J.R., et al.: Genetic Programming IV Video: Human-Competitive Machine Intelligence. Kluwer Academic Publishers (2003), <https://youtu.be/ordtzUccs2s>
11. Koza, J.: Automated design using Darwinian evolution and genetic programming. Stanford University, EE380: Computer Systems Colloquium (18 Feb 2009), <https://www.youtube.com/watch?v=xIoytwJWJP8>

12. Le Goues, C., et al.: Automated program repair. *Communications of the ACM* 62(12), 56–65 (Dec 2019), <http://dx.doi.org/10.1145/3318162>
13. Petke, J., et al.: Genetic improvement of software: a comprehensive survey. *IEEE Transactions on Evolutionary Computation* 22(3), 415–432 (Jun 2018), <http://dx.doi.org/doi:10.1109/TEVC.2017.2693219>
14. Langdon, W.B., Petke, J.: Software is not fragile. In: Parrend, P., et al. (eds.) *Complex Systems Digital Campus E-conference, CS-DC'15*. pp. 203–211. *Proceedings in Complexity*, Springer (Sep 30-Oct 1 2015), http://dx.doi.org/10.1007/978-3-319-45901-1_24, invited talk
15. Langdon, W.B., et al.: Efficient multi-objective higher order mutation testing with genetic programming. *Journal of Systems and Software* 83(12), 2416–2430 (Dec 2010), <http://dx.doi.org/10.1016/j.jss.2010.07.027>
16. Harrand, N., et al.: A journey among Java neutral program variants. *Genetic Programming and Evolvable Machines* 20(4), 531–580 (Dec 2019), <http://dx.doi.org/10.1007/s10710-019-09355-3>
17. Schulte, E., et al.: Software mutational robustness. *Genetic Programming and Evolvable Machines* 15(3), 281–312 (Sep 2014), <http://dx.doi.org/10.1007/s10710-013-9195-8>
18. Abou Assi, R., et al.: Coincidental correctness in the Defects4J benchmark. *Software Testing, Verification and Reliability* 29(3), e1696 (2019), <http://dx.doi.org/10.1002/stvr.1696>
19. Danglot, B., et al.: Correctness attraction: a study of stability of software behavior under runtime perturbation. *Empirical Software Engineering* 23(4), 2086–2119 (1 Aug 2018), <http://dx.doi.org/10.1007/s10664-017-9571-8>
20. Monperrus, M.: Principles of antifragile software. In: *Companion to the First International Conference on the Art, Science and Engineering of Programming*. pp. 32:1–32:4. *Programming '17*, ACM, New York, NY, USA (2017), <http://dx.doi.org/10.1145/3079368.3079412>
21. Petke, J., et al.: Software robustness: A survey, a theory, and some prospects. In: Avgeriou, P., Zhang, D. (eds.) *ESEC/FSE 2021, Ideas, Visions and Reflections*. pp. 1475–1478. ACM, Athens, Greece (23–28 Aug 2021), <http://dx.doi.org/10.1145/3468264.3473133>
22. Andre, D., Koza, J.R.: Parallel genetic programming on a network of transputers. In: Rosca, J.P. (ed.) *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*. pp. 111–120. Tahoe City, California, USA (9 Jul 1995), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/andre_1995_parallel.pdf
23. Kinnear, Jr., K.E. (ed.): *Advances in Genetic Programming*. MIT Press, Cambridge, MA (1994), <http://www.amazon.co.uk/Advances-Genetic-Programming-Complex-Adaptive/dp/0262111888>
24. Angeline, P.J., Kinnear, Jr., K.E. (eds.): *Advances in Genetic Programming 2*. MIT Press, Cambridge, MA, USA (1996), <http://dx.doi.org/10.7551/mitpress/1109.001.0001>
25. Spector, L., et al.: Quantum computing applications of genetic programming. In: Spector, L., et al. (eds.) *Advances in Genetic Programming 3*, chap. 7, pp. 135–160. MIT Press, Cambridge, MA, USA (Jun 1999), <http://dx.doi.org/10.7551/mitpress/1110.003.0010>
26. Koza, J.R., et al. (eds.): *Genetic Programming 1996: Proceedings of the First Annual Conference*. MIT Press, Stanford University, CA, USA (28–31 Jul 1996), <http://www.genetic-programming.org/gp96proceedings.html>
27. Koza, J.R., et al. (eds.): *Genetic Programming 1997: Proceedings of the Second Annual Conference*. Morgan Kaufmann, Stanford University, CA, USA (13–16 Jul 1997), <http://www.amazon.com/Genetic-Programming-2nd-Conference-Author/dp/1558604839>
28. Koza, J.R., et al. (eds.): *Genetic Programming 1998: Proceedings of the Third Annual Conference*. Morgan Kaufmann, University of Wisconsin, Madison, WI, USA (22–25 Jul 1998)
29. Angeline, P.J., et al. (eds.): *Proceedings of the 1999 Congress on Evolutionary Computation, CEC 1999*. IEEE Press, Washington, DC, USA (July 6–9 1999), <https://dblp.org/rec/conf/cec/1999.bib>

30. Banzhaf, W., et al. (eds.): GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference. Morgan Kaufmann, Orlando, Florida, USA (13-17 Jul 1999), <http://www.amazon.com/exec/obidos/ASIN/1558606114/qid%3D977054373/105-7666192-3217523>
31. Banzhaf, W., et al. (eds.): Genetic Programming, LNCS, vol. 1391. Springer-Verlag, Paris (14-15 Apr 1998), <http://dx.doi.org/10.1007/BFb0055923>
32. Koza, J.R., et al.: Evolving inventions. *Scientific American* 288(2), 52–59 (Feb 2003), <http://www.sciam.com/article.cfm?articleID=00073FCE-F36F-1E19-8B3B809EC588EEDF>
33. Koza, J.R.: Human-competitive results produced by genetic programming. *Genetic Programming and Evolvable Machines* 11(3/4), 251–284 (Sep 2010), <http://dx.doi.org/10.1007/s10710-010-9112-3>, tenth Anniversary Issue: Progress in Genetic Programming and Evolvable Machines
34. Friedberg, R.M.: A learning machine: I. *IBM Journal of Research and Development* 2(1), 2–13 (Jan 1958), <http://www.ams.org/mathscinet-getitem?mr=19%2c1085c>
35. Kilburn, T., et al.: Experiments in machine learning and thinking. In: *Information Processing, Proceedings of the 1st International Conference on Information Processing*. pp. 303–308. UNESCO, Paris (15-20 Jun 1959), <https://dblp.org/rec/conf/ifip/KilburnGS59.bib>
36. Turing, A.M.: *Intelligent machinery* (1948), <https://www.npl.co.uk/getattachment/about-us/History/Famous-faces/Alan-Turing/80916595-Intelligent-Machinery.pdf>, report for National Physical Laboratory. Reprinted in Ince, D. C. (editor). 1992. *Mechanical Intelligence: Collected Works of A. M. Turing*. Amsterdam: North Holland. Pages 107127. Also reprinted in Meltzer, B. and Michie, D. (editors). 1969. *Machine Intelligence 5*. Edinburgh: Edinburgh University Press [278].
37. Langdon, W.B., Banzhaf, W.: Repeated patterns in tree genetic programming. In: Keijzer, M., et al. (eds.) *Proceedings of the 8th European Conference on Genetic Programming*. *Lecture Notes in Computer Science*, vol. 3447, pp. 190–202. Springer, Lausanne, Switzerland (30 Mar - 1 Apr 2005), http://dx.doi.org/10.1007/978-3-540-31989-4_17
38. Langdon, W.B., Banzhaf, W.: Repeated patterns in genetic programming. *Natural Computing* 7(4), 589–613 (Dec 2008), <http://dx.doi.org/10.1007/s11047-007-9038-8>
39. Ohno, S.: *Evolution by Gene Duplication*. Springer (1970), <https://link.springer.com/book/10.1007/978-3-642-86659-3>
40. Koza, J.R., Andre, D.: A case study where biology inspired a solution to a computer science problem. In: Hunter, L., Klein, T.E. (eds.) *Pacific Symposium on Biocomputing '96*. pp. 500–511. World Scientific (1996), <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.368.6585>
41. Koza, J.R.: Architecture-altering operations for evolving the architecture of a multi-part program in genetic programming. Technical Report STAN-CS-94-1528, Dept. of Computer Science, Stanford University, Stanford, California 94305, USA (Oct 1994), <http://www.genetic-programming.com/jkpdf/tr1528.pdf>
42. Forrest, S., et al.: A genetic programming approach to automated software repair. In: Raidl, G., et al. (eds.) *GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*. pp. 947–954. ACM, Montreal (8-12 Jul 2009), <http://dx.doi.org/10.1145/1569901.1570031>, gECCO 2019 10-Year Most Influential Paper Award, Best paper
43. Carroll, L.: *Through the Looking-Glass, and What Alice Found There*. Macmillan (1871)
44. Weimer, W., et al.: Automatically finding patches using genetic programming. In: Fickas, S. (ed.) *International Conference on Software Engineering (ICSE) 2009*. pp. 364–374. Vancouver (May 16-24 2009), <http://dx.doi.org/10.1109/ICSE.2009.5070536>
45. Le Goues, C.: *Automatic Program Repair Using Genetic Programming*. Ph.D. thesis, Faculty of the School of Engineering and Applied Science, University of Virginia, USA (May 2013), <http://www.cs.virginia.edu/~weimer/students/claire-phd.pdf>

46. Haraldsson, S.O., et al.: Fixing bugs in your sleep: How genetic improvement became an overnight success. In: Petke, J., et al. (eds.) GI-2017. pp. 1513–1520. ACM, Berlin (15-19 Jul 2017), <http://dx.doi.org/10.1145/3067695.3082517>, best paper
47. Alshahwan, N.: Industrial experience of genetic improvement in Facebook. In: Petke, J., et al. (eds.) GI-2019, ICSE workshops proceedings. p. 1. IEEE, Montreal (28 May 2019), <http://dx.doi.org/10.1109/GI.2019.00010>, invited Keynote
48. Kirbas, S., et al.: On the introduction of automatic program repair in Bloomberg. *IEEE Software* 38(4), 43–51 (Jul-Aug 2021), <http://dx.doi.org/10.1109/MS.2021.3071086>
49. Juille, H., Pollack, J.B.: Massively parallel genetic programming. In: Angeline, P.J., Kinnear, Jr., K.E. (eds.) *Advances in Genetic Programming 2*, chap. 17, pp. 339–357. MIT Press, Cambridge, MA, USA (1996), <http://dx.doi.org/10.7551/mitpress/1109.003.0023>
50. Thompson, A.: *Hardware Evolution Automatic Design of Electronic Circuits in Reconfigurable Hardware by Artificial Evolution*. DISTDISS, Springer (1998), <http://dx.doi.org/10.1007/978-1-4471-3414-5>, winner 1997 BCS distinguished dissertation
51. Gordon, T.G.W.: *Exploiting Development to Enhance the Scalability of Hardware Evolution*. Ph.D. thesis, University College, London, UK (Jul 2005), <https://discovery.ucl.ac.uk/id/eprint/1444775/>
52. Martin, P.N.: *Genetic Programming in Hardware*. Ph.D. thesis, University of Essex, University of Essex, Wivenhoe Park, Colchester, UK (Mar 2003), <http://www.naiadhome.com/HardwareGeneticProgramming.pdf>
53. Sekanina, L., Vasiccek, Z.: Cgp acceleration using field-programmable gate arrays. In: Miller, J.F. (ed.) *Cartesian Genetic Programming*, chap. 7, pp. 217–230. Natural Computing Series, Springer (2011), http://dx.doi.org/10.1007/978-3-642-17310-3_7
54. Goribar-Jimenez, C., et al.: Towards the development of a complete gp system on an fpga using geometric semantic operators. In: Lozano, J.A. (ed.) *2017 IEEE Congress on Evolutionary Computation (CEC)*. pp. 1932–1939. IEEE, Donostia, San Sebastian, Spain (5-8 Jun 2017), <http://dx.doi.org/10.1109/CEC.2017.7969537>
55. Langdon, W.B., Krauss, O.: Genetic improvement of data for maths functions. *ACM Transactions on Evolutionary Learning and Optimization* 1(2), Article No.: 7 (Jul 2021), <http://dx.doi.org/10.1145/3461016>
56. Owens, J.D., et al.: A survey of general-purpose computation on graphics hardware. *Computer Graphics Forum* 26(1), 80–113 (March 2007), <http://dx.doi.org/10.1111/j.1467-8659.2007.01012.x>
57. Harding, S., Banzhaf, W.: Fast genetic programming on GPUs. In: Ebner, M., et al. (eds.) *Proceedings of the 10th European Conference on Genetic Programming*. Lecture Notes in Computer Science, vol. 4445, pp. 90–101. Springer, Valencia, Spain (11-13 Apr 2007), http://dx.doi.org/10.1007/978-3-540-71605-1_9
58. Chitty, D.M.: A data parallel approach to genetic programming using programmable graphics hardware. In: Thierens, D., et al. (eds.) *GECCO '07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*. vol. 2, pp. 1566–1573. ACM Press, London (7-11 Jul 2007), <http://dx.doi.org/10.1145/1276958.1277274>
59. Juille, H., Pollack, J.B.: Parallel genetic programming and fine-grained simd architecture. In: Siegel, E.V., Koza, J.R. (eds.) *Working Notes for the AAAI Symposium on Genetic Programming*. pp. 31–37. AAAI, MIT, Cambridge, MA, USA (10–12 Nov 1995), <http://www.aaai.org/Papers/Symposia/Fall/1995/FS-95-01/FS95-01-005.pdf>
60. Langdon, W.B.: A SIMD interpreter for genetic programming on GPU graphics cards. Tech. Rep. CSM-470, Department of Computer Science, University of Essex, Colchester, UK (3 Jul 2007), http://cswww.essex.ac.uk/technical-reports/2007/csm_470.pdf
61. Langdon, W.B., Banzhaf, W.: A SIMD interpreter for genetic programming on GPU graphics cards. In: O'Neill, M., et al. (eds.) *Proceedings of the 11th European Conference on Genetic Programming, EuroGP 2008*. Lecture Notes in Computer Science, vol. 4971, pp. 73–85. Springer, Naples (26-28 Mar 2008), http://dx.doi.org/10.1007/978-3-540-78671-9_7
62. Langdon, W.B.: A fast high quality pseudo random number generator for graphics processing units. In: Wang, J. (ed.) *2008 IEEE World Congress on Computational*

- Intelligence. pp. 459–465. IEEE, Hong Kong (1-6 Jun 2008), <http://dx.doi.org/10.1109/CEC.2008.4630838>
63. Robilliard, D., et al.: Genetic programming on graphics processing units. *Genetic Programming and Evolvable Machines* 10(4), 447–471 (Dec 2009), <http://dx.doi.org/10.1007/s10710-009-9092-3>, special issue on parallel and distributed evolutionary algorithms, part I
 64. Baumes, L.A., et al.: EASEA: a generic optimization tool for GPU machines in asynchronous island model. *Computer Methods in Materials Science* 11(3), 489–499 (2011), <http://icube-publis.unistra.fr/docs/7407/baumes.pdf>
 65. Vitola, J., et al.: Parallel algorithm for evolvable-based boolean synthesis on gpus. In: *Third IEEE Latin American Symposium on Circuits and Systems (LASCAS 2012)* (29 Feb-2 Mar 2012), <http://dx.doi.org/10.1109/LASCAS.2012.6180339>
 66. Maghazeh, A., et al.: General purpose computing on low-power embedded GPUs: Has it come of age? In: *Jeschke, H. (ed.) 2013 International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIII)*. IEEE, Samos, Greece (15-18 Jul 2013), <http://dx.doi.org/10.1109/SAMOS.2013.6621099>
 67. Chitty, D.M.: Faster GPU-based genetic programming using a two-dimensional stack. *Soft Computing* 21(14), 3859–3878 (Jul 2017), <http://dx.doi.org/10.1007/s00500-016-2034-0>
 68. Ono, K., Hanada, Y.: Self-organized subpopulation based on multiple features in genetic programming on GPU. *Journal of Advanced Computational Intelligence and Intelligent Informatics* 25(2), 177–186 (Mar 2021), <http://dx.doi.org/10.20965/jaciii.2021.p0177>
 69. Trujillo, L., et al.: GSGP-CUDA - a CUDA framework for geometric semantic genetic programming. *SoftwareX* 18, 101085 (Jun 2022), <http://dx.doi.org/10.1016/j.softx.2022.101085>
 70. Langdon, W.B., Harrison, A.P.: GP on SPMD parallel graphics hardware for mega bioinformatics data mining. *Soft Computing* 12(12), 1169–1183 (Oct 2008), <http://dx.doi.org/10.1007/s00500-008-0296-x>, special Issue on Distributed Bioinspired Algorithms
 71. Langdon, W.B.: Distilling GeneChips with genetic programming on the Emerald GPU supercomputer. *SIGEVolution newsletter of the ACM Special Interest Group on Genetic and Evolutionary Computation* 6(1), 15–21 (25 Jul 2012), <http://dx.doi.org/10.1145/2384697.2384699>
 72. Langdon, W.B.: Large scale bioinformatics data mining with parallel genetic programming on graphics processing units. In: *Tsutsui, S., Collet, P. (eds.) Massively Parallel Evolutionary Computation on GPGPUs*, chap. 15, pp. 311–347. *Natural Computing Series*, Springer (2013), http://dx.doi.org/10.1007/978-3-642-37959-8_15
 73. Langdon, W.B.: Large scale bioinformatics data mining with parallel genetic programming on graphics processing units. In: *Fernandez de Vega, F., Cantu-Paz, E. (eds.) Parallel and Distributed Computational Intelligence, Studies in Computational Intelligence*, vol. 269, chap. 5, pp. 113–141. Springer (Jan 2010), http://dx.doi.org/10.1007/978-3-642-10675-0_6
 74. Langdon, W.B., Banzhaf, W.: Long-term evolution experiment with genetic programming. *Artificial Life* 28(2), 173–204 (Summer 2022), http://dx.doi.org/10.1162/artl_a_00360, invited submission to *Artificial Life Journal* special issue of the ALIFE'19 conference
 75. Langdon, W.B.: Genetic programming convergence. *Genetic Programming and Evolvable Machines* 23(1), 71–104 (Mar 2022), <http://dx.doi.org/10.1007/s10710-021-09405-9>
 76. Goodfellow, I., et al.: *Deep Learning*. MIT Press (2016), <https://mitpress.mit.edu/9780262035613/>
 77. Weimer, W.: From deep learning to human judgments: Lessons for genetic improvement. *GI @ GECCO 2022* (9 Jul 2022), <http://geneticimprovementofsoftware.com/slides/gi2022gecco/weimer-keynote-gi-gecco-22.pdf>, invited keynote
 78. Langdon, W.B.: *Genetic Programming and Data Structures: Genetic Programming + Data Structures = Automatic Programming!*, *Genetic Programming*, vol. 1. Kluwer, Boston (1998), <http://dx.doi.org/10.1007/978-1-4615-5731-9>

79. Teller, A.: The evolution of mental models. In: Kinnear, Jr., K.E. (ed.) *Advances in Genetic Programming*, chap. 9, pp. 199–219. MIT Press (1994), <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/MentalModels.ps>
80. Jannink, J.: Cracking and co-evolving randomizers. In: Kinnear, Jr., K.E. (ed.) *Advances in Genetic Programming*, chap. 20, pp. 425–443. MIT Press (1994), <http://infolab.stanford.edu/pub/jannink/gp.ps>
81. Andre, D.: Evolution of mapmaking ability: Strategies for the evolution of learning, planning, and memory using genetic programming. In: *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*. vol. 1, pp. 250–255. IEEE Press, Orlando, Florida, USA (27–29 Jun 1994), <http://dx.doi.org/10.1109/ICEC.1994.350007>
82. Iba, H., et al.: Temporal data processing using genetic programming. In: Eshelman, L.J. (ed.) *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*. pp. 279–286. Morgan Kaufmann, Pittsburgh, PA, USA (15–19 Jul 1995), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/iba_1995_tdpdp.pdf
83. Haynes, T.D., Wainwright, R.L.: A simulation of adaptive agents in hostile environment. In: George, K.M., et al. (eds.) *Proceedings of the 1995 ACM Symposium on Applied Computing*. pp. 318–323. ACM Press, Nashville, USA (1995), <http://dx.doi.org/10.1145/315891.316007>
84. Nordin, P., Banzhaf, W.: Evolving Turing-complete programs for a register machine with self-modifying code. In: Eshelman, L.J. (ed.) *Genetic Algorithms: Proceedings of the Sixth International Conference (ICGA95)*. pp. 318–325. Morgan Kaufmann, Pittsburgh, PA, USA (15–19 Jul 1995), <http://www.cs.mun.ca/~banzhaf/papers/icga95-2.pdf>
85. Brave, S.: Evolving recursive programs for tree search. In: Angeline, P.J., Kinnear, Jr., K.E. (eds.) *Advances in Genetic Programming 2*, chap. 10, pp. 203–220. MIT Press, Cambridge, MA, USA (1996), <http://dx.doi.org/10.7551/mitpress/1109.003.0015>
86. Esparcia Alcazar, A.I., Sharman, K.C.: Some applications of genetic programming in digital signal processing. In: Koza, J.R. (ed.) *Late Breaking Papers at the Genetic Programming 1996 Conference Stanford University July 28–31, 1996*. pp. 24–31. Stanford Bookstore, Stanford University, CA, USA (28–31 Jul 1996), <http://www.iti.upv.es/~anna/papers/someappsgp96.ps>
87. Bruce, W.S.: *The Application of Genetic Programming to the Automatic Generation of Object-Oriented Programs*. Ph.D. thesis, School of Computer and Information Sciences, Nova Southeastern University, 3100 SW 9th Avenue, Fort Lauderdale, Florida 33315, USA (Dec 1995), https://nsuworks.nova.edu/gscis_etd/430/
88. Ronge, A., Nordahl, M.G.: Genetic programs and co-evolution developing robust general purpose controllers using local mating in two dimensional populations. In: Voigt, H.M., et al. (eds.) *Parallel Problem Solving from Nature IV, Proceedings of the International Conference on Evolutionary Computation*. LNCS, vol. 1141, pp. 81–90. Springer Verlag, Berlin, Germany (22–26 Sep 1996), http://dx.doi.org/10.1007/3-540-61723-X_972
89. Spector, L., Luke, S.: Cultural transmission of information in genetic programming. In: Koza, J.R., et al. (eds.) *Genetic Programming 1996: Proceedings of the First Annual Conference*. pp. 209–214. MIT Press, Stanford University, CA, USA (28–31 Jul 1996), <http://www.cs.gmu.edu/~sean/papers/culture-gp96.pdf>
90. Raik, S.E., Browne, D.G.: Evolving state and memory in genetic programming. In: *Simulated Evolution and Learning*. Springer (1997), <http://dx.doi.org/10.1007/BFb0028523>
91. Edmonds, B., Moss, S.: Modelling of boundedly rational agents using evolutionary programming techniques. In: Corne, D., Shapiro, J.L. (eds.) *Evolutionary Computing*, LNCS, vol. 1305, pp. 31–42. Springer-Verlag, University of Manchester, UK (7–8 Apr 1997), <http://dx.doi.org/10.1007/BFb0027164>
92. Bennett III, F.H.: A multi-skilled robot that recognizes and responds to different problem environments. In: Koza, J.R., et al. (eds.) *Genetic Programming 1997: Proceedings of the Second Annual Conference*. pp. 44–51. Morgan Kaufmann, Stanford University, CA, USA (13–16 Jul 1997), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/gp1997/bennet_1997_msrrrdpe.pdf

93. Angeline, P.J.: An alternative to indexed memory for evolving programs with explicit state representations. In: Koza, J.R., et al. (eds.) Genetic Programming 1997: Proceedings of the Second Annual Conference. pp. 423–430. Morgan Kaufmann, Stanford University, CA, USA (13–16 Jul 1997)
94. Lim, I.S., Thalmann, D.: Indexed memory as a generic protocol for handling vectors of data in genetic programming. In: Eiben, A.E., et al. (eds.) Fifth International Conference on Parallel Problem Solving from Nature. LNCS, vol. 1498, pp. 325–334. Springer-Verlag, Amsterdam (27–30 Sep 1998), <http://dx.doi.org/10.1007/BFb0056875>
95. Trenaman, A.: The Evolution of Autonomous Agents Using Concurrent Genetic Programming. Ph.D. thesis, Department of Computer Science, National University of Ireland, Maynooth, Ireland (Oct 1999), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/trenaman/at_thesis1.ps.gz
96. Silva, A., et al.: Evolving controllers for autonomous agents using genetically programmed networks. In: Poli, R., et al. (eds.) Genetic Programming, Proceedings of EuroGP'99. LNCS, vol. 1598, pp. 255–269. Springer-Verlag, Goteborg, Sweden (26–27 May 1999), http://dx.doi.org/10.1007/3-540-48885-5_22
97. Andersson, B., et al.: Reactive and memory-based genetic programming for robot control. In: Poli, R., et al. (eds.) Genetic Programming, Proceedings of EuroGP'99. LNCS, vol. 1598, pp. 161–172. Springer-Verlag, Goteborg, Sweden (26–27 May 1999), http://dx.doi.org/10.1007/3-540-48885-5_13
98. Martin, P.: Genetic programming for service creation in intelligent networks. In: Poli, R., et al. (eds.) Genetic Programming, Proceedings of EuroGP'2000. LNCS, vol. 1802, pp. 106–120. Springer-Verlag, Edinburgh (15–16 Apr 2000), http://dx.doi.org/10.1007/978-3-540-46239-2_8
99. Bearpark, K.: Learning and memory in genetic programming. Ph.D. thesis, School of Engineering Sciences, University of Southampton, UK (2000), <http://eprints.soton.ac.uk/45930/>
100. Karlsson, R., et al.: Sound localization for a humanoid robot using genetic programming. In: Cagnoni, S., et al. (eds.) Real-World Applications of Evolutionary Computing. LNCS, vol. 1803, pp. 65–76. Springer-Verlag, Edinburgh (17 Apr 2000), http://dx.doi.org/10.1007/3-540-45561-2_7
101. Martin, M.C.: Visual obstacle avoidance using genetic programming: First results. In: Spector, L., et al. (eds.) Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001). pp. 1107–1113. Morgan Kaufmann, San Francisco, California, USA (7–11 Jul 2001), <http://www.martinmartin.com/Dissertation/VisualObstacleAvoidanceGP.pdf>
102. Brumby, S.P., et al.: Evolving forest fire burn severity classification algorithms for multi-spectral imagery. In: Shen, S.S., Descour, M.R. (eds.) In Algorithms for Multispectral, Hyperspectral, and Ultraspectral Imagery VII, Proceedings of SPIE. vol. 4381, pp. 236–245 (2001), <http://dx.doi.org/10.1117/12.437013>
103. Howard, D., et al.: The boru data crawler for object detection tasks in machine vision. In: Cagnoni, S., et al. (eds.) Applications of Evolutionary Computing, Proceedings of EvoWorkshops2002: EvoCOP, EvoIASP, EvoSTim/EvoPLAN. LNCS, vol. 2279, pp. 222–232. Springer-Verlag, Kinsale, Ireland (3–4 Apr 2002), http://dx.doi.org/10.1007/3-540-46004-7_23
104. Imamura, K., et al.: *N*-version genetic programming via fault masking. In: Foster, J.A., et al. (eds.) Genetic Programming, Proceedings of the 5th European Conference, EuroGP 2002. LNCS, vol. 2278, pp. 172–181. Springer-Verlag, Kinsale, Ireland (3–5 Apr 2002), http://dx.doi.org/10.1007/3-540-45984-7_17
105. Johnson, M.: Sequence generation using machine language evolved by genetic programming. In: Wang, L., et al. (eds.) Proceedings of the 4th Asia-Pacific Conference on Simulated Evolution And Learning (SEAL'02). p. #1251. Orchid Country Club, Singapore (18–22 Nov 2002), <http://www.worldcat.org/title/seal02-proceedings-of-the-4th-asia-pacific-conference-on-simulated-evolution-and-learning-november-18-22-2002-orchid-oclc/51951214>
106. O'Neill, M., Ryan, C.: Investigations into memory in grammatical evolution. In: Barry, A.M. (ed.) GECCO 2002: Proceedings of the Bird of a Feather Workshops, Genetic and Evolutionary Computation Conference. pp. 141–144. AAAI, New York (8 Jul 2002), <http://www.grammatical-evolution.org/gews2002/oneill.ps>

107. Pillay, N.: Using genetic programming for the induction of novice procedural programming solution algorithms. In: SAC '02: Proceedings of the 2002 ACM symposium on Applied computing. pp. 578–583. ACM Press, Madrid, Spain (Mar 2002), <http://dx.doi.org/10.1145/508791.508903>
108. Quintana, M.L., et al.: Morphological algorithm design for binary images using genetic programming. *Genetic Programming and Evolvable Machines* 7(1), 81–102 (Mar 2006), <http://dx.doi.org/10.1007/s10710-006-7012-3>
109. Segond, M., et al.: Iterative filter generation using genetic programming. In: Collet, P., et al. (eds.) Proceedings of the 9th European Conference on Genetic Programming. Lecture Notes in Computer Science, vol. 3905, pp. 145–153. Springer, Budapest, Hungary (10 - 12 Apr 2006), http://dx.doi.org/10.1007/11729976_13
110. Kim, D.: A quantitative analysis of memory usage for agent tasks. In: Iba, H. (ed.) *Frontiers in Evolutionary Robotics*, chap. 14, pp. 247–274. IntechOpen, Rijeka (2008), <http://dx.doi.org/10.5772/5458>
111. Frias-Martinez, E., Gobet, F.: Automatic generation of cognitive theories using genetic programming. *Minds and Machines* 17(3), 287–309 (Oct 2007), <http://dx.doi.org/10.1007/s11023-007-9070-6>
112. McPhee, N.F., Poli, R.: Memory with memory: Soft assignment in genetic programming. In: Keijzer, M., et al. (eds.) GECCO '08: Proceedings of the 10th annual conference on Genetic and evolutionary computation. pp. 1235–1242. ACM, Atlanta, GA, USA (12-16 Jul 2008), <http://dx.doi.org/10.1145/1389095.1389336>
113. Katz, G., Peled, D.: Genetic programming and model checking: Synthesizing new mutual exclusion algorithms. In: *Automated Technology for Verification and Analysis*. Lecture Notes in Computer Science, vol. 5311, pp. 33–47. Springer (2008), http://dx.doi.org/10.1007/978-3-540-88387-6_5
114. Withall, M.S., et al.: An improved representation for evolving programs. *Genetic Programming and Evolvable Machines* 10(1), 37–70 (Mar 2009), <http://dx.doi.org/10.1007/s10710-008-9069-7>
115. Wilson, G.C., Banzhaf, W.: Soft memory for stock market analysis using linear and developmental genetic programming. In: Raidl, G., et al. (eds.) GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation. pp. 1633–1640. ACM, Montreal (8-12 Jul 2009), <http://dx.doi.org/10.1145/1569901.1570119>
116. Wolfson, K., Sipper, M.: Efficient list search algorithms. In: Collet, P., et al. (eds.) 9th International Conference, Evolution Artificielle, EA 2009. Lecture Notes in Computer Science, vol. 5975, pp. 158–169. Springer, Strasbourg, France (Oct 26-28 2009), http://dx.doi.org/10.1007/978-3-642-14156-0_14, revised Selected Papers
117. Hyde, M.: A genetic programming hyper-heuristic approach to automated packing. Ph.D. thesis, School of Computer Science, University of Nottingham, UK (Mar 2010), http://etheses.nottingham.ac.uk/1625/1/mvh_corrected_thesis.pdf
118. Suchorzewski, M.: Extending genetic programming to evolve perceptron-like learning programs. In: Rutkowski, L., et al. (eds.) 10th International Conference Artificial Intelligence and Soft Computing, ICAISC 2010, Part II. Lecture Notes in Computer Science, vol. 6114, pp. 221–228. Springer, Zakopane, Poland (Jun 13-17 2010), <http://dx.doi.org/10.1007/978-3-642-13232-2>
119. Agapitos, A., et al.: Learning environment models in car racing using stateful genetic programming. In: Proceedings of the 2011 IEEE Conference on Computational Intelligence and Games. pp. 219–226. IEEE, Seoul, South Korea (31 Aug - 3 Sep 2011), <http://dx.doi.org/10.1109/CIG.2011.6032010>
120. Weise, T., Tang, K.: Evolving distributed algorithms with genetic programming. *IEEE Transactions on Evolutionary Computation* 16(2), 242–265 (Apr 2012), <http://dx.doi.org/10.1109/TEVC.2011.2112666>
121. Kala, R.: Multi-robot path planning using co-evolutionary genetic programming. *Expert Systems with Applications* 39(3), 3817–3831 (2012), <http://dx.doi.org/10.1016/j.eswa.2011.09.090>
122. Yim, H., Kim, D.: Evolving internal memory strategies for the woods problems. In: 12th International Conference on Control, Automation and Systems (ICCAS 2012). pp. 366–369 (2012), <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=6393463>

123. Igwe, K., Pillay, N.: Automatic programming using genetic programming. In: Ngo, L.T., et al. (eds.) Proceedings of the 2013 Third World Congress on Information and Communication Technologies (WICT 2013). pp. 337–342. IEEE, Hanoi, Vietnam (15–18 Dec 2013), <http://dx.doi.org/10.1109/WICT.2013.7113158>
124. Qadir, O., et al.: Hardware architecture of the protein processing associative memory and the effects of dimensionality and quantisation on performance. *Genetic Programming and Evolvable Machines* 15(3), 245–275 (Sep 2014), <http://dx.doi.org/10.1007/s10710-014-9217-1>
125. Szczuko, P.: Genetic programming extension to APF-based monocular human body pose estimation. *Multimedia Tools and Applications* 68(1), 177–192 (2014), <http://dx.doi.org/10.1007/s11042-012-1147-4>
126. Yuan, X., et al.: Making lock-free data structures verifiable with artificial transactions. In: Proceedings of the 8th Workshop on Programming Languages and Operating Systems, PLOS 2015. pp. 39–45. ACM, Monterey, California, USA (4–7 Oct 2015), <http://dx.doi.org/10.1145/2818302.2818309>
127. Chaumont, N., Adami, C.: Evolution of sustained foraging in three-dimensional environments with physics. *Genetic Programming and Evolvable Machines* 17(4), 359–390 (Dec 2016), <http://dx.doi.org/10.1007/s10710-016-9270-z>
128. Smith, R., Heywood, M.: A model of external memory for navigation in partially observable visual reinforcement learning tasks. In: Sekanina, L., et al. (eds.) EuroGP 2019: Proceedings of the 22nd European Conference on Genetic Programming. LNCS, vol. 11451, pp. 162–177. Springer Verlag, Leipzig, Germany (24–26 Apr 2019), http://dx.doi.org/10.1007/978-3-030-16670-0_11
129. Kelly, S., et al.: Emergent tangled program graphs in partially observable recursive forecasting and ViZDoom navigation tasks. *ACM Transactions on Evolutionary Learning and Optimization* 1(3) (Sep 2021), <http://dx.doi.org/10.1145/3468857>
130. Real, E., et al.: AutoML-zero: Evolving machine learning algorithms from scratch. In: Daume, III, H., Singh, A. (eds.) Proceedings of the 37th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 119, pp. 8007–8019. PMLR (13–18 Jul 2020), http://www.human-competitive.org/sites/default/files/automl_zero_humies_competition_entry.txt, winner 2021 HUMIES
131. Sulyok, C., et al.: Evolving the process of a virtual composer. *Natural Computing* 18(1), 47–60 (Mar 2019), <http://dx.doi.org/10.1007/s11047-016-9561-6>
132. Al Masalma, M., Heywood, M.: Genetic programming with external memory in sequence recall tasks. In: Trautmann, H., et al. (eds.) Proceedings of the 2022 Genetic and Evolutionary Computation Conference Companion. pp. 518–521. GECCO ’22, Association for Computing Machinery, Boston, USA (9–13 Jul 2022), <http://dx.doi.org/10.1145/3520304.3528883>
133. Langdon, W.B., Poli, R.: Foundations of Genetic Programming. Springer-Verlag (2002), <http://dx.doi.org/10.1007/978-3-662-04726-2>
134. O’Reilly, U.M., Oppacher, F.: The troubling aspects of a building block hypothesis for genetic programming. In: Whitley, L.D., Vose, M.D. (eds.) Foundations of Genetic Algorithms 3. pp. 73–88. Morgan Kaufmann, Estes Park, Colorado, USA (31 Jul–2 Aug 1994), <http://dx.doi.org/10.1016/B978-1-55860-356-1.50008-X>, published 1995
135. Rosca, J.P., Ballard, D.H.: Rooted-tree schemata in genetic programming. In: Spector, L., et al. (eds.) Advances in Genetic Programming 3, chap. 11, pp. 243–271. MIT Press, Cambridge, MA, USA (Jun 1999), <http://dx.doi.org/10.7551/mitpress/1110.003.0015>
136. Poli, R.: Exact schema theory for genetic programming and variable-length genetic algorithms with one-point crossover. *Genetic Programming and Evolvable Machines* 2(2), 123–163 (Jun 2001), <http://dx.doi.org/10.1023/A:1011552313821>
137. Stephens, C.R., Poli, R.: EC theory – “in theory”: Towards a unification of evolutionary computation theory. In: Menon, A. (ed.) Frontiers of Evolutionary Computation, vol. 11, chap. 7, pp. 129–155. Kluwer, Boston, MA, USA (2004), http://dx.doi.org/10.1007/1-4020-7782-3_7
138. Price, G.R.: Selection and covariance. *Nature* 227, August 1, 520–521 (1970), <http://dx.doi.org/10.1038/227520a0>

139. Altenberg, L.: The evolution of evolvability in genetic programming. In: Kinnear, Jr., K.E. (ed.) *Advances in Genetic Programming*, chap. 3, pp. 47–74. MIT Press (1994), <http://dynamics.org/~altenber/PAPERS/EEGP/>
140. Ryan, C., et al.: A competitive building block hypothesis. In: Deb, K., et al. (eds.) *Genetic and Evolutionary Computation – GECCO-2004, Part II. Lecture Notes in Computer Science*, vol. 3103, pp. 654–665. Springer-Verlag, Seattle, WA, USA (26-30 Jun 2004), http://dx.doi.org/10.1007/978-3-540-24855-2_73
141. White, D.R., et al.: Modelling genetic programming as a simple sampling algorithm. In: Banzhaf, W., et al. (eds.) *Genetic Programming Theory and Practice XVII*, pp. 367–381. Springer, East Lansing, MI, USA (16-19 May 2019), http://dx.doi.org/10.1007/978-3-030-39958-0_18
142. Miller, J.: What bloat? cartesian genetic programming on Boolean problems. In: Goodman, E.D. (ed.) *2001 Genetic and Evolutionary Computation Conference Late Breaking Papers*, pp. 295–302. San Francisco, California, USA (9-11 Jul 2001), <http://www.elec.york.ac.uk/intsys/users/jfm7/gecco2001Late.pdf>
143. Jones, T.: One operator, one landscape. Tech. Rep. SFI TR 95-02-025, Santa Fe Institute (January 1995), <http://www.santafe.edu/sfi/publications/Working-Papers/95-02-025.ps>
144. O’Reilly, U.M.: Using a distance metric on genetic programs to understand genetic operators. In: *IEEE International Conference on Systems, Man, and Cybernetics, Computational Cybernetics and Simulation*, vol. 5, pp. 4092–4097. Orlando, Florida, USA (12-15 Oct 1997), <http://dx.doi.org/10.1109/ICSMC.1997.637337>
145. Vassilev, V.K., et al.: Smoothness, ruggedness and neutrality of fitness landscapes: from theory to application. In: Ghosh, A., Tsutsui, S. (eds.) *Advances in evolutionary computing: theory and applications*, pp. 3–44. Springer-Verlag New York, Inc. (2003), http://dx.doi.org/10.1007/978-3-642-18965-4_1
146. Langdon, W.B., Harman, M.: Fitness landscape of the Triangle program. In: Veerapen, N., Ochoa, G. (eds.) *PPSN-2016 Workshop on Landscape-Aware Heuristic Search*. Edinburgh (17 Sep 2016), http://www.cs.ucl.ac.uk/fileadmin/UCL-CS/research/Research_Notes/rn1605.pdf, also available as UCL RN/16/05
147. Langdon, W.B., et al.: Dissipative polynomials. In: Veerapen, N., et al. (eds.) *5th Workshop on Landscape-Aware Heuristic Search*, pp. 1683–1691. GECCO 2021 Companion, ACM, Internet (10-14 Jul 2021), <http://dx.doi.org/10.1145/3449726.3463147>
148. Francone, F.D., et al.: Homologous crossover in genetic programming. In: Banzhaf, W., et al. (eds.) *Proceedings of the Genetic and Evolutionary Computation Conference*, vol. 2, pp. 1021–1026. Morgan Kaufmann, Orlando, Florida, USA (13-17 Jul 1999), <http://gpbib.cs.ucl.ac.uk/gecco1999/GP-463.pdf>
149. Durrett, G., et al.: Computational complexity analysis of simple genetic programming on two problems modeling isolated program semantics. In: Beyer, H.G., Langdon, W.B. (eds.) *Foundations of Genetic Algorithms*, pp. 69–80. ACM, Schwarzenberg, Austria (5-9 Jan 2011), <http://dx.doi.org/10.1145/1967654.1967661>
150. Koetzing, T., et al.: The Max problem revisited: The importance of mutation in genetic programming. *Theoretical Computer Science* 545, 94–107 (2014), <http://dx.doi.org/10.1016/j.tcs.2013.06.014>
151. Nguyen, A., et al.: Single- and multi-objective genetic programming: New bounds for weighted order and majority. In: Neumann, F., De Jong, K. (eds.) *Foundations of Genetic Algorithms*, pp. 161–172. ACM, Adelaide, Australia (16-20 Jan 2013), <http://dx.doi.org/10.1145/2460239.2460254>
152. Lissovoi, A., Oliveto, P.S.: On the time and space complexity of genetic programming for evolving boolean conjunctions. *Journal of Artificial Intelligence Research* 66, 655–689 (2019), <http://dx.doi.org/10.1613/jair.1.11821>
153. Doerr, B., et al.: The impact of lexicographic parsimony pressure for ORDER/MAJORITY on the run time. *Theoretical Computer Science* 816, 144–168 (6 May 2020), <http://dx.doi.org/10.1016/j.tcs.2020.01.011>
154. Koetzing, T., et al.: Destructiveness of lexicographic parsimony pressure and alleviation by a concatenation crossover in genetic programming. *Theoretical Computer Science* 816, 96–113 (6 May 2020), <http://dx.doi.org/10.1016/j.tcs.2019.11.036>
155. Wolpert, D.H., Macready, W.G.: No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1(1), 67–82 (April 1997)

156. Langdon, W.B.: Incremental evaluation in genetic programming. In: Hu, T., et al. (eds.) EuroGP 2021: Proceedings of the 24th European Conference on Genetic Programming. LNCS, vol. 12691, pp. 229–246. Springer Verlag, Virtual Event (7-9 Apr 2021), http://dx.doi.org/10.1007/978-3-030-72812-0_15
157. Handley, S.: On the use of a directed acyclic graph to represent a population of computer programs. In: Proceedings of the 1994 IEEE World Congress on Computational Intelligence. vol. 1, pp. 154–159. IEEE Press, Orlando, Florida, USA (27-29 Jun 1994), <http://dx.doi.org/10.1109/ICEC.1994.350024>
158. Langdon, W.B., Banzhaf, W.: Repeated sequences in linear genetic programming genomes. *Complex Systems* 15(4), 285–306 (2005), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/wbl_repeat_linear.pdf
159. Soule, T., Foster, J.A.: Effects of code growth and parsimony pressure on populations in genetic programming. *Evolutionary Computation* 6(4), 293–309 (Winter 1998), <http://dx.doi.org/10.1162/evco.1998.6.4.293>
160. de Jong, E.D., Pollack, J.B.: Multi-objective methods for tree size control. *Genetic Programming and Evolvable Machines* 4(3), 211–233 (Sep 2003), <http://dx.doi.org/10.1023/A:1025122906870>
161. Bleuler, S., et al.: Multiobjective genetic programming: Reducing bloat using spea2. In: Proceedings of the 2001 Congress on Evolutionary Computation CEC2001. pp. 536–543. IEEE Press, COEX, World Trade Center, 159 Samseong-dong, Gangnam-gu, Seoul, Korea (27-30 May 2001), <http://dx.doi.org/10.1109/CEC.2001.934438>
162. Panait, L., Luke, S.: Alternative bloat control methods. In: Deb, K., et al. (eds.) Genetic and Evolutionary Computation – GECCO-2004, Part II. Lecture Notes in Computer Science, vol. 3103, pp. 630–641. Springer-Verlag, Seattle, WA, USA (26-30 Jun 2004), <http://dx.doi.org/10.1007/b98645>
163. Poli, R.: A simple but theoretically-motivated method to control bloat in genetic programming. In: Ryan, C., et al. (eds.) Genetic Programming, Proceedings of EuroGP'2003. LNCS, vol. 2610, pp. 204–217. Springer-Verlag, Essex (14-16 Apr 2003), http://dx.doi.org/10.1007/3-540-36599-0_19
164. Silva, S.: Controlling Bloat: Individual and Population Based Approaches in Genetic Programming. Ph.D. thesis, Coimbra University, Portugal (Apr 2008), <http://hdl.handle.net/10316/8542>, full author name is Sara Guilherme Oliveira da Silva
165. Dignum, S., Poli, R.: Operator equalisation and bloat free GP. In: O'Neill, M., et al. (eds.) Proceedings of the 11th European Conference on Genetic Programming, EuroGP 2008. Lecture Notes in Computer Science, vol. 4971, pp. 110–121. Springer, Naples (26-28 Mar 2008), http://dx.doi.org/10.1007/978-3-540-78671-9_10
166. Nikolaev, N.I., Iba, H.: Accelerated genetic programming of polynomials. *Genetic Programming and Evolvable Machines* 2(3), 231–257 (Sep 2001), <http://dx.doi.org/10.1023/A:1011949326249>
167. Kushchu, I.: Genetic programming and evolutionary generalization. *IEEE Transactions on Evolutionary Computation* 6(5), 431–442 (Oct 2002), <http://dx.doi.org/10.1109/TEVC.2002.805038>
168. Kowaliw, T., Doursat, R.: Bias-variance decomposition in genetic programming. *Open Mathematics* 14(1), 62–80 (Feb 2016), <http://dx.doi.org/10.1515/math-2016-0005>
169. Gathercole, C., Ross, P.: Dynamic training subset selection for supervised learning in genetic programming. In: Davidor, Y., et al. (eds.) Parallel Problem Solving from Nature III. LNCS, vol. 866, pp. 312–321. Springer-Verlag, Jerusalem (9-14 Oct 1994), http://dx.doi.org/10.1007/3-540-58484-6_275
170. Spector, L., et al.: Relaxations of lexibase parent selection. In: Banzhaf, W., et al. (eds.) Genetic Programming Theory and Practice XV. pp. 105–120. Genetic and Evolutionary Computation, Springer, University of Michigan in Ann Arbor, USA (May 18–20 2017), http://dx.doi.org/10.1007/978-3-319-90512-9_7
171. Javed, N., et al.: Simplification of genetic programs: a literature survey. *Data Mining and Knowledge Discovery* 36(4), 1279–1300 (Jul 2022), <http://dx.doi.org/10.1007/s10618-022-00830-7>, special Issue on Explainable and Interpretable Machine Learning and Data Mining
172. Hooper, D., Flann, N.S.: Improving the accuracy and robustness of genetic programming through expression simplification. In: Koza, J.R., et al. (eds.) Genetic Program-

- ming 1996: Proceedings of the First Annual Conference. p. 428. MIT Press, Stanford University, CA, USA (28–31 Jul 1996), <http://digital.cs.usu.edu/~flann/gp.pdf>
173. Langdon, W.B., Harman, M.: Optimising existing software with genetic programming. *IEEE Transactions on Evolutionary Computation* 19(1), 118–135 (Feb 2015), <http://dx.doi.org/10.1109/TEVC.2013.2281544>
 174. Raichle, M.E., Gusnard, D.A.: Appraising the brain’s energy budget. *Proceedings of the National Academy of Sciences* 99(16), 10237–10239 (6 aug 2002), <http://dx.doi.org/10.1073/pnas.172399499>
 175. Ridley, M.: *The Red Queen, Sex and the Evolution of Human Nature*. Penguin (1993), <http://www.penguin.co.uk/Penguin/Books/0140167722.html>
 176. Nordin, P.: A compiling genetic programming system that directly manipulates the machine code. In: Kinnear, Jr., K.E. (ed.) *Advances in Genetic Programming*, chap. 14, pp. 311–331. MIT Press (1994), <http://www.amazon.co.uk/Advances-Genetic-Programming-Complex-Adaptive/dp/0262111888>
 177. Miller, J.F.: An empirical study of the efficiency of learning Boolean functions using a cartesian genetic programming approach. In: Banzhaf, W., et al. (eds.) *Proceedings of the Genetic and Evolutionary Computation Conference*. vol. 2, pp. 1135–1142. Morgan Kaufmann, Orlando, Florida, USA (13–17 Jul 1999), <http://citeseer.ist.psu.edu/153431.html>
 178. Ryan, C., et al.: Grammatical evolution: Evolving programs for an arbitrary language. In: Banzhaf, W., et al. (eds.) *Proceedings of the First European Workshop on Genetic Programming*. LNCS, vol. 1391, pp. 83–96. Springer-Verlag, Paris (14–15 Apr 1998), <http://dx.doi.org/10.1007/BFb0055930>
 179. Angeline, P.J., Pollack, J.B.: The evolutionary induction of subroutines. In: *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. pp. 236–241. Lawrence Erlbaum, Bloomington, Indiana, USA (1992), <http://www.demon.cs.brandeis.edu/papers/glib92.pdf>
 180. Rosca, J.: Towards automatic discovery of building blocks in genetic programming. In: Siegel, E.V., Koza, J.R. (eds.) *Working Notes for the AAAI Symposium on Genetic Programming*. pp. 78–85. AAAI, MIT, Cambridge, MA, USA (10–12 Nov 1995), <http://www.aaai.org/Papers/Symposia/Fall/1995/FS-95-01/FS95-01-011.pdf>
 181. Spector, L.: Simultaneous evolution of programs and their control structures. In: Angeline, P.J., Kinnear, Jr., K.E. (eds.) *Advances in Genetic Programming 2*, chap. 7, pp. 137–154. MIT Press, Cambridge, MA, USA (1996), <http://dx.doi.org/10.7551/mitpress/1109.003.0012>
 182. Murphy, G., Ryan, C.: Seeding methods for run transferable libraries. In: Thierens, D., et al. (eds.) *GECCO ’07: Proceedings of the 9th annual conference on Genetic and evolutionary computation*. vol. 2, pp. 1755–1755. ACM Press, London (7–11 Jul 2007), <http://dx.doi.org/10.1145/1276958.1277305>
 183. Langdon, W.B.: *Data Structures and Genetic Programming*. Ph.D. thesis, University College, London, UK (27 Sep 1996), <http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/langdon.ps.gz>
 184. Teller, A., Andre, D.: Automatically choosing the number of fitness cases: The rational allocation of trials. In: Koza, J.R., et al. (eds.) *Genetic Programming 1997: Proceedings of the Second Annual Conference*. pp. 321–328. Morgan Kaufmann, Stanford University, CA, USA (13–16 Jul 1997), <http://www.cs.cmu.edu/afs/cs/usr/astro/public/papers/GR.ps>
 185. Spector, L.: Assessment of problem modality by differential performance of lexibase selection in genetic programming: A preliminary report. In: McClymont, K., Keedwell, E. (eds.) *1st workshop on Understanding Problems (GECCO-UP)*. pp. 401–408. ACM, Philadelphia, Pennsylvania, USA (7–11 Jul 2012), <http://dx.doi.org/10.1145/2330784.2330846>
 186. Poli, R., et al.: *A field guide to genetic programming*. Published via <http://lulu.com> and freely available at <http://www.gp-field-guide.org.uk> (2008), <http://www.gp-field-guide.org.uk>, (With contributions by J. R. Koza)
 187. Harding, S.L., Banzhaf, W.: Distributed genetic programming on GPUs using CUDA. In: Hidalgo, I., et al. (eds.) *Workshop on Parallel Architectures and Bioinspired Algorithms*. pp. 1–10. Universidad Complutense de Madrid, Raleigh, NC, USA (13 Sep 2009), <http://www.evolutioninmaterio.com/preprints/CudaParallelCompilePP.pdf>

188. Crepeau, R.L.: Genetic evolution of machine language software. In: Rosca, J.P. (ed.) Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications. pp. 121–134. Tahoe City, California, USA (9 Jul 1995), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/GEMS_Article.pdf
189. Nordin, P., et al.: Efficient evolution of machine code for CISC architectures using instruction blocks and homologous crossover. In: Spector, L., et al. (eds.) Advances in Genetic Programming 3, chap. 12, pp. 275–299. MIT Press, Cambridge, MA, USA (Jun 1999), <http://dx.doi.org/10.7551/mitpress/1110.003.0017>
190. Francone, F.D.: Discipulus Owner’s Manual. 11757 W. Ken Caryl Avenue F, PBM 512, Littleton, Colorado, 80127-3719, USA, version 3.0 draft edn. (2001), http://gpbib.cs.ucl.ac.uk/gp-html/francone_manual.html
191. Banzhaf, W., et al.: Genetic Programming – An Introduction; On the Automatic Evolution of Computer Programs and its Applications. Morgan Kaufmann, San Francisco, CA, USA (Jan 1998), <https://www.amazon.co.uk/Genetic-Programming-Introduction-Artificial-Intelligence/dp/155860510X>
192. Brameier, M., Banzhaf, W.: Linear Genetic Programming. No. XVI in Genetic and Evolutionary Computation, Springer (2007), <http://dx.doi.org/10.1007/978-0-387-31030-5>
193. O’Neill, M., Ryan, C.: Grammatical evolution. IEEE Transactions on Evolutionary Computation 5(4), 349–358 (Aug 2001), <http://dx.doi.org/10.1109/4235.942529>
194. O’Neill, M., Ryan, C.: Grammatical Evolution: Evolutionary Automatic Programming in a Arbitrary Language, Genetic programming, vol. 4. Kluwer Academic Publishers (2003), <http://dx.doi.org/10.1007/978-1-4615-0447-4>
195. Spector, L.: Introduction to the peer commentary special section on “on the mapping of genotype to phenotype in evolutionary algorithms” by Peter A. Whigham, Grant Dick, and James Maclaurin. Genetic Programming and Evolvable Machines 18(3), 351–352 (Sep 2017), <http://dx.doi.org/10.1007/s10710-017-9287-y>, special Peer Commentary on Mapping of Genotype to Phenotype in Evolutionary Algorithms
196. Ryan, C.: A rebuttal to whigham, dick, and maclaurin by one of the inventors of grammatical evolution: Commentary on “on the mapping of genotype to phenotype in evolutionary algorithms” by Peter A. Whigham, Grant Dick, and James Maclaurin. Genetic Programming and Evolvable Machines 18(3), 385–389 (Sep 2017), <http://dx.doi.org/10.1007/s10710-017-9294-z>, special Peer Commentary on Mapping of Genotype to Phenotype in Evolutionary Algorithms
197. Miller, J.F., et al.: Principles in the evolutionary design of digital circuits-part I. Genetic Programming and Evolvable Machines 1(1/2), 7–35 (Apr 2000), <http://dx.doi.org/10.1023/A:1010016313373>
198. Miller, J.F., et al.: Principles in the evolutionary design of digital circuits-part II. Genetic Programming and Evolvable Machines 1(3), 259–288 (Jul 2000), <http://dx.doi.org/10.1023/A:1010066330916>
199. Miller, J.F. (ed.): Cartesian Genetic Programming. Natural Computing Series, Springer (2011), <http://dx.doi.org/10.1007/978-3-642-17310-3>
200. Miller, J.F.: Cartesian Genetic Programming: its status and future. Genetic Programming and Evolvable Machines 21(1-2), 129–168 (Jun 2020), <http://dx.doi.org/10.1007/s10710-019-09360-6>, twentieth Anniversary Issue
201. Sekanina, L., Vasicek, Z.: Approximate circuit design by means of evolvable hardware. In: IEEE International Conference on Evolvable Systems (ICES 2013). pp. 21–28 (Apr 2013), <http://dx.doi.org/10.1109/ICES.2013.6613278>
202. Sekanina, L., et al.: Approximate circuits in low-power image and video processing: The approximate median filter. Radioengineering 26(3), 623–632 (Sep 2017), <http://dx.doi.org/10.13164/re.2017.0623>
203. Montana, D.J.: Strongly typed genetic programming. Evolutionary Computation 3(2), 199–230 (Summer 1995), <http://dx.doi.org/10.1162/evco.1995.3.2.199>
204. Yu, T.: Structure abstraction and genetic programming. In: Angeline, P.J., et al. (eds.) Proceedings of the Congress on Evolutionary Computation. vol. 1, pp. 652–659. IEEE Press, Mayflower Hotel, Washington D.C., USA (6-9 Jul 1999), <http://dx.doi.org/10.1109/CEC.1999.781995>

205. Yu, T.: Hierarchical processing for evolving recursive and modular programs using higher order functions and lambda abstractions. *Genetic Programming and Evolvable Machines* 2(4), 345–380 (Dec 2001), <http://dx.doi.org/10.1023/A:1012926821302>
206. McKay, R.I., et al.: Grammar-based genetic programming: a survey. *Genetic Programming and Evolvable Machines* 11(3/4), 365–396 (Sep 2010), <http://dx.doi.org/10.1007/s10710-010-9109-y>, tenth Anniversary Issue: Progress in Genetic Programming and Evolvable Machines
207. Whigham, P.A.: Grammatically-based genetic programming. In: Rosca, J.P. (ed.) *Proceedings of the Workshop on Genetic Programming: From Theory to Real-World Applications*. pp. 33–41. Tahoe City, California, USA (9 Jul 1995), <http://divcom.otago.ac.nz/sirc/Peterw/Publications/ml95.zip>
208. Whigham, P.A., et al.: On the mapping of genotype to phenotype in evolutionary algorithms. *Genetic Programming and Evolvable Machines* 18(3), 353–361 (Sep 2017), <http://dx.doi.org/10.1007/s10710-017-9288-x>, special Peer Commentary on Mapping of Genotype to Phenotype in Evolutionary Algorithms
209. Ratle, A., Sebag, M.: A novel approach to machine discovery: Genetic programming and stochastic grammars. In: Matwin, S., Sammut, C. (eds.) *Proceedings of Twelfth International Conference on Inductive Logic Programming*. LNCS, vol. 2583, pp. 207–222. Springer Verlag, Sydney, Australia (Jul 9–11 2002), http://dx.doi.org/10.1007/3-540-36468-4_14, revised Papers
210. Nguyen, X.H., et al.: Solving the symbolic regression problem with tree-adjunct grammar guided genetic programming: The comparative results. *The Australian Journal of Intelligent Information Processing Systems* 7(3/4), 114–121 (2001), <http://sc.snu.ac.kr/PAPERS/xuanetal.pdf>
211. Jacob, C.: Computer physics communications. *Evolution and coevolution of developmental programs* 121–122, 46–50 (Sep–Oct 1999), [http://dx.doi.org/10.1016/S0010-4655\(99\)00277-5](http://dx.doi.org/10.1016/S0010-4655(99)00277-5)
212. Jacob, C.: *Illustrating Evolutionary Computation with Mathematica*. Morgan Kaufmann (2001), <http://dx.doi.org/10.1016/B978-155860637-1/50020-5>
213. Hornby, G.S., Pollack, J.B.: Evolving L-systems to generate virtual creatures. *Computers and Graphics* 25(6), 1041–1048 (Dec 2001), [http://dx.doi.org/10.1016/S0097-8493\(01\)00157-1](http://dx.doi.org/10.1016/S0097-8493(01)00157-1), artificial Life
214. Hemberg, M., et al.: Genr8: Architects’ experience with an emergent design tool. In: Romero, J., Machado, P. (eds.) *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*, chap. 8, pp. 167–188. Springer (2008), http://dx.doi.org/10.1007/978-3-540-72877-1_8
215. Perkis, T.: Stack-based genetic programming. In: *Proceedings of the 1994 IEEE World Congress on Computational Intelligence*. vol. 1, pp. 148–153. IEEE Press, Orlando, Florida, USA (27–29 Jun 1994), <http://dx.doi.org/10.1109/ICEC.1994.350025>
216. Openshaw, S., Turton, I.: Building new spatial interaction models using genetic programming. In: Fogarty, T.C. (ed.) *Evolutionary Computing, AISB workshop*. Leeds, UK (11–13 Apr 1994), <http://dx.doi.org/10.1007/3-540-58483-8>, unpublished
217. Holladay, K., et al.: Fifth: A stack based gp language for vector processing. In: Ebner, M., et al. (eds.) *Proceedings of the 10th European Conference on Genetic Programming*. *Lecture Notes in Computer Science*, vol. 4445, pp. 102–113. Springer, Valencia, Spain (11–13 Apr 2007), http://dx.doi.org/10.1007/978-3-540-71605-1_10
218. Oltean, M., Grosan, C.: Solving classification problems using infix form genetic programming. In: Berthold, M.R., et al. (eds.) *Advances in Intelligent Data Analysis V*. *Lecture Notes in Computer Science*, vol. 2810, pp. 242–253. Springer, Berlin, Germany (Aug 28–30 2003), http://dx.doi.org/10.1007/978-3-540-45231-7_23
219. Spector, L., Robinson, A.: Genetic programming and autoconstructive evolution with the push programming language. *Genetic Programming and Evolvable Machines* 3(1), 7–40 (Mar 2002), <http://dx.doi.org/10.1023/A:1014538503543>
220. O’Reilly, U.M., Oppacher, F.: Program search with a hierarchical variable length representation: Genetic programming, simulated annealing and hill climbing. In: Davidor, Y., et al. (eds.) *Parallel Problem Solving from Nature – PPSN III*. pp. 397–406. No. 866 in *Lecture Notes in Computer Science*, Springer-Verlag, Jerusalem (9–14 Oct 1994), http://dx.doi.org/10.1007/3-540-58484-6_283

221. Esparcia-Alcazar, A.I., Sharman, K.C.: Genetic programming techniques that evolve recurrent neural networks architectures for signal processing. In: IEEE Workshop on Neural Networks for Signal Processing. pp. 139–148. IEEE, Seiko, Kyoto, Japan (4-6 Sep 1996), <http://dx.doi.org/10.1109/NNSP.1996.548344>
222. Moraglio, A., Silva, S.: Geometric differential evolution on the space of genetic programs. In: Esparcia-Alcazar, A.I., et al. (eds.) Proceedings of the 13th European Conference on Genetic Programming, EuroGP 2010. LNCS, vol. 6021, pp. 171–183. Springer, Istanbul (7-9 Apr 2010), http://dx.doi.org/10.1007/978-3-642-12148-7_15, best paper
223. Zhang, B.T.: Bayesian methods for efficient genetic programming. Genetic Programming and Evolvable Machines 1(3), 217–242 (Jul 2000), <http://dx.doi.org/10.1023/A:1010010230007>
224. Yanai, K., Iba, H.: Estimation of distribution programming based on Bayesian network. In: Sarker, R., et al. (eds.) Proceedings of the 2003 Congress on Evolutionary Computation CEC2003. pp. 1618–1625. IEEE Press, Canberra (8-12 Dec 2003), <http://dx.doi.org/10.1109/CEC.2003.1299866>
225. Bosman, P.A.N., de Jong, E.D.: Learning probabilistic tree grammars for genetic programming. In: Yao, X., et al. (eds.) Parallel Problem Solving from Nature - PPSN VIII. LNCS, vol. 3242, pp. 192–201. Springer-Verlag, Birmingham, UK (18-22 Sep 2004), <http://dx.doi.org/10.1007/b100601>
226. Rodriguez, A.: A Neat Approach To Genetic Programming. Master’s thesis, School of School of Electrical Engineering and Computer Science, University of Central Florida, Orlando, Florida, USA (2007), <https://stars.library.ucf.edu/etd/3323.pdf>
227. Buk, Z., et al.: NEAT in HyperNEAT substituted with genetic programming. In: Kolehmainen, M., et al. (eds.) 9th International Conference on Adaptive and Natural Computing Algorithms, ICANNGA 2009. Lecture Notes in Computer Science, vol. 5495, pp. 243–252. Springer, Kuopio, Finland (23-25 Apr 2009), http://dx.doi.org/10.1007/978-3-642-04921-7_25, revised selected papers
228. Trujillo, L., et al.: neat genetic programming: Controlling bloat naturally. Information Sciences 333, 21–43 (10 Mar 2016), <http://dx.doi.org/10.1016/j.ins.2015.11.010>
229. McConaghy, T.: FFX: Fast, scalable, deterministic symbolic regression technology. In: Riolo, R., et al. (eds.) Genetic Programming Theory and Practice IX, chap. 13, pp. 235–260. Genetic and Evolutionary Computation, Springer, Ann Arbor, USA (12-14 May 2011), http://dx.doi.org/10.1007/978-1-4614-1770-5_13
230. Moraglio, A., et al.: Geometric semantic genetic programming. In: Coello Coello, C.A., et al. (eds.) Parallel Problem Solving from Nature, PPSN XII (part 1). Lecture Notes in Computer Science, vol. 7491, pp. 21–31. Springer, Taormina, Italy (Sep 1-5 2012), http://dx.doi.org/10.1007/978-3-642-32937-1_3
231. Langdon, W.B.: Directed crossover within genetic programming. Research Note RN/95/71, University College London, Gower Street, London WC1E 6BT, UK (Sep 1995), http://www.cs.ucl.ac.uk/staff/W.Langdon/ftp/papers/directed_crossover.pdf
232. Orzechowski, P., et al.: Where are we now?: a large benchmark study of recent symbolic regression methods. In: Aguirre, H., et al. (eds.) GECCO ’18: Proceedings of the Genetic and Evolutionary Computation Conference. pp. 1183–1190. ACM, Kyoto, Japan (15-19 Jul 2018), <http://dx.doi.org/10.1145/3205455.3205539>
233. Arnaldo, I., et al.: Multiple regression genetic programming. In: Igel, C., et al. (eds.) GECCO ’14: Proceedings of the 2014 conference on Genetic and evolutionary computation. pp. 879–886. ACM, Vancouver, BC, Canada (12-16 Jul 2014), <http://dx.doi.org/10.1145/2576768.2598291>
234. Munoz, L., et al.: M3GP: multiclass classification with GP. In: Machado, P., et al. (eds.) 18th European Conference on Genetic Programming. LNCS, vol. 9025, pp. 78–91. Springer, Copenhagen (8-10 Apr 2015), http://dx.doi.org/10.1007/978-3-319-16501-1_7
235. La Cava, W., Moore, J.: A general feature engineering wrapper for machine learning using epsilon-lexicase survival. In: Castelli, M., et al. (eds.) EuroGP 2017: Proceedings of the 20th European Conference on Genetic Programming. LNCS, vol. 10196, pp. 80–95. Springer Verlag, Amsterdam (19-21 Apr 2017), http://dx.doi.org/10.1007/978-3-319-55696-3_6

236. Burlacu, B., et al.: Operon C++: An efficient genetic programming framework for symbolic regression. In: Allmendinger, R., et al. (eds.) Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion. pp. 1562–1570. GECCO '20, Association for Computing Machinery, internet (Jul 8-12 2020), <http://dx.doi.org/10.1145/3377929.3398099>
237. Mota Dias, D., et al.: Automatic synthesis of microcontroller assembly code through linear genetic programming. In: Nedjah, N., et al. (eds.) Genetic Systems Programming: Theory and Experiences, Studies in Computational Intelligence, vol. 13, pp. 193–227. Springer, Germany (2006), http://dx.doi.org/10.1007/3-540-32498-4_9
238. Lewis, T.E., Magoulas, G.D.: TMBL kernels for CUDA GPUs compile faster using PTX. In: Harding, S., et al. (eds.) GECCO 2011 Computational intelligence on consumer games and graphics hardware (CIGPU). pp. 455–462. ACM, Dublin, Ireland (12-16 Jul 2011), <http://dx.doi.org/10.1145/2001858.2002033>
239. Cupertino, L.F., et al.: Evolving CUDA PTX programs by quantum inspired linear genetic programming. In: Harding, S., et al. (eds.) GECCO 2011 Computational intelligence on consumer games and graphics hardware (CIGPU). pp. 399–406. ACM, Dublin, Ireland (12-16 Jul 2011), <http://dx.doi.org/10.1145/2001858.2002026>
240. Gregor, M., Spalek, J.: Using LLVM-based JIT compilation in genetic programming. In: 2016 ELEKTRO. pp. 406–411. IEEE, Strbske Pleso, Slovakia (16-18 May 2016), <http://dx.doi.org/10.1109/ELEKTRO.2016.7512108>
241. Liou, J.Y., et al.: GEVO: GPU code optimization using evolutionary computation. ACM Transactions on Architecture and Code Optimization 17(4), Article 33 (Dec 2020), <http://dx.doi.org/10.1145/3418055>
242. Lukschandl, E., et al.: Distributed java bytecode genetic programming. In: Poli, R., et al. (eds.) Genetic Programming, Proceedings of EuroGP'2000. LNCS, vol. 1802, pp. 316–325. Springer-Verlag, Edinburgh (15-16 Apr 2000), http://dx.doi.org/10.1007/978-3-540-46239-2_24
243. Whigham, P.A., McKay, R.I.: Genetic approaches to learning recursive relations. In: Yao, X. (ed.) Progress in Evolutionary Computation, Lecture Notes in Artificial Intelligence, vol. 956, pp. 17–27. Springer-Verlag (1995), http://dx.doi.org/10.1007/3-540-60154-6_44
244. Whigham, P.A.: A schema theorem for context-free grammars. In: 1995 IEEE Conference on Evolutionary Computation. vol. 1, pp. 178–181. IEEE Press, Perth, Australia (29 Nov - 1 Dec 1995), <http://dx.doi.org/10.1109/ICEC.1995.489140>
245. Castle, T., Johnson, C.G.: Evolving high-level imperative program trees with strongly formed genetic programming. In: Moraglio, A., et al. (eds.) Proceedings of the 15th European Conference on Genetic Programming, EuroGP 2012. LNCS, vol. 7244, pp. 1–12. Springer Verlag, Malaga, Spain (11-13 Apr 2012), http://dx.doi.org/10.1007/978-3-642-29139-5_1
246. Hillis, W.D.: Co-evolving parasites improve simulated evolution as an optimization procedure. In: Langton, C.G., et al. (eds.) Artificial Life II, Santa Fe Institute Studies in the Sciences of Complexity, vol. X, pp. 313–324. Addison-Wesley, Santa Fe Institute, New Mexico, USA (Feb 1990 1992)
247. Popovici, E., et al.: Coevolutionary principles. In: Rozenberg, G., et al. (eds.) Handbook of Natural Computing, chap. 31, pp. 987–1033. Springer (19 August 2012), http://dx.doi.org/10.1007/978-3-540-92910-9_31
248. Zhang, B.T., Cho, D.Y.: Coevolutionary fitness switching: Learning complex collective behaviors using genetic programming. In: Spector, L., et al. (eds.) Advances in Genetic Programming 3, chap. 18, pp. 425–445. MIT Press, Cambridge, MA, USA (Jun 1999), <http://dx.doi.org/10.7551/mitpress/1110.003.0023>
249. Leier, A., Banzhaf, W.: Exploring the search space of quantum programs. In: Sarker, R., et al. (eds.) Proceedings of the 2003 Congress on Evolutionary Computation CEC2003. vol. 1, pp. 170–177. IEEE Press, Canberra (8-12 Dec 2003), <http://dx.doi.org/10.1109/CEC.2003.1299571>
250. Spector, L.: Automatic Quantum Computer Programming: A Genetic Programming Approach, Genetic Programming, vol. 7. Kluwer Academic Publishers, Boston/Dordrecht/New York/London (Jun 2004), <http://dx.doi.org/10.1007/978-0-387-36791-0>

251. O'Brien, G., Clark, J.: Using genetic improvement to retarget quantum software on differing hardware. In: Petke, J., et al. (eds.) *GI @ ICSE 2021*. pp. 31–38. IEEE, internet (30 May 2021), <http://dx.doi.org/10.1109/GI52543.2021.00015>, winner Best Presentation
252. Poli, R., et al.: Theoretical results in genetic programming: The next ten years? *Genetic Programming and Evolvable Machines* 11(3/4), 285–320 (Sep 2010), <http://dx.doi.org/10.1007/s10710-010-9110-5>, tenth Anniversary Issue: Progress in Genetic Programming and Evolvable Machines
253. Vanneschi, L., Poli, R.: Genetic programming: Introduction, applications, theory and open issues. In: Rozenberg, G., et al. (eds.) *Handbook of Natural Computing*, vol. 2, chap. 24, pp. 709–739. Springer (19 Aug 2012), http://dx.doi.org/10.1007/978-3-540-92910-9_24
254. Marginean, A., et al.: SapFix: Automated end-to-end repair at scale. In: Atlee, J.M., Bultan, T. (eds.) *41st International Conference on Software Engineering*. pp. 269–278. ACM, Montreal (25–31 May 2019), <http://dx.doi.org/10.1109/ICSE-SEIP.2019.00039>
255. Bruce, B.R., et al.: Approximate oracles and synergy in software energy search spaces. *IEEE Transactions on Software Engineering* 45(11), 1150–1169 (Nov 2019), <http://dx.doi.org/10.1109/TSE.2018.2827066>
256. Wu, F., et al.: Deep parameter optimisation. In: Silva, S., et al. (eds.) *GECCO '15: Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. pp. 1375–1382. ACM, Madrid (11–15 Jul 2015), <http://dx.doi.org/10.1145/2739480.2754648>
257. Langdon, W.B., Harman, M.: Genetically improved CUDA C++ software. In: Nicolau, M., et al. (eds.) *17th European Conference on Genetic Programming*. LNCS, vol. 8599, pp. 87–99. Springer, Granada, Spain (23–25 Apr 2014), http://dx.doi.org/10.1007/978-3-662-44303-3_8
258. Langdon, W.B., et al.: Improving 3D medical image registration CUDA software with genetic programming. In: Igel, C., et al. (eds.) *GECCO '14: Proceeding of the sixteenth annual conference on genetic and evolutionary computation conference*. pp. 951–958. ACM, Vancouver, BC, Canada (12–15 Jul 2014), <http://dx.doi.org/10.1145/2576768.2598244>
259. Langdon, W.B., Harman, M.: Grow and graft a better CUDA pknotsRG for RNA pseudoknot free energy calculation. In: Langdon, W.B., et al. (eds.) *Genetic Improvement 2015 Workshop*. pp. 805–810. ACM, Madrid (11–15 Jul 2015), <http://dx.doi.org/10.1145/2739482.2768418>
260. Yeboah-Antwi, K., Baudry, B.: Embedding adaptivity in software systems using the ECSELR framework. In: Langdon, W.B., et al. (eds.) *Genetic Improvement 2015 Workshop*. pp. 839–844. ACM, Madrid (11–15 Jul 2015), <http://dx.doi.org/10.1145/2739482.2768425>
261. Langdon, W.B., et al.: Improving CUDA DNA analysis software with genetic programming. In: Silva, S., et al. (eds.) *GECCO '15: Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*. pp. 1063–1070. ACM, Madrid (11–15 Jul 2015), <http://dx.doi.org/10.1145/2739480.2754652>
262. Langdon, W.B., et al.: Genetic improvement of GPU software. *Genetic Programming and Evolvable Machines* 18(1), 5–44 (Mar 2017), <http://dx.doi.org/10.1007/s10710-016-9273-9>
263. Langdon, W.B.: Genetically improved software. In: Gandomi, A.H., et al. (eds.) *Handbook of Genetic Programming Applications*, chap. 8, pp. 181–220. Springer (2015), http://dx.doi.org/10.1007/978-3-319-20883-1_8
264. Langdon, W.B., Lam, B.Y.H.: Genetically improved BarraCUDA. *BioData Mining* 20(28) (2 Aug 2017), <http://dx.doi.org/10.1186/s13040-017-0149-1>
265. Langdon, W.B., et al.: Evolving better RNAfold structure prediction. In: Castelli, M., et al. (eds.) *EuroGP 2018: Proceedings of the 21st European Conference on Genetic Programming*. LNCS, vol. 10781, pp. 220–236. Springer Verlag, Parma, Italy (4–6 Apr 2018), http://dx.doi.org/10.1007/978-3-319-77553-1_14
266. Langdon, W.B., Harman, M.: Evolving a CUDA kernel from an nVidia template. In: Sobrevilla, P. (ed.) *2010 IEEE World Congress on Computational Intelligence*. pp.

- 2376–2383. IEEE, Barcelona (18–23 Jul 2010), <http://dx.doi.org/10.1109/CEC.2010.5585922>
267. Liou, J.Y., et al.: Genetic improvement of GPU code. In: Petke, J., et al. (eds.) GI-2019, ICSE workshops proceedings. pp. 20–27. IEEE, Montreal (28 May 2019), <http://dx.doi.org/10.1109/GI.2019.00014>, best Paper
268. Barr, E.T., et al.: Automated software transplanted. In: Xie, T., Young, M. (eds.) International Symposium on Software Testing and Analysis, ISSTA 2015. pp. 257–269. ACM, Baltimore, Maryland, USA (14–17 Jul 2015), <http://dx.doi.org/10.1145/2771783.2771796>, ACM SIGSOFT Distinguished Paper Award
269. Burke, E.K., et al.: Exploring hyper-heuristic methodologies with genetic programming. In: Mumford, C.L., Jain, L.C. (eds.) Computational Intelligence, Intelligent Systems Reference Library, vol. 1, chap. 6, pp. 177–201. Springer (2009), http://dx.doi.org/10.1007/978-3-642-01799-5_6
270. Olson, R.S., Moore, J.H.: TPOT: A tree-based pipeline optimization tool for automating data science. In: Hutter, F., et al. (eds.) AutoML 2016 workshop. New York City, USA (Jun 24 2016), <https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWwFpbXhdXRvbWwyMDE2fGd4OmFmYjMyNWU2NWl1YTBMZg>, collocated with ICML
271. Radecic, D.: Machine Learning Automation with TPOT. Packt Publishing (2021), <https://www.amazon.com/Machine-Learning-Automation-TPOT-automated/dp/180056788X?asin=180056788X&revisionId=&format=4&depth=1>
272. Krawiec, K., Liskowski, P.: Adaptive test selection for factorization-based surrogate fitness in genetic programming. *Foundations of Computing and Decision Sciences* 42(4), 339–358 (Dec 2017), <http://dx.doi.org/10.1515/fcds-2017-0017>
273. Johnson, C.G.: Solving the Rubik’s cube with stepwise deep learning. *Expert Systems: The Journal of Knowledge Engineering* 38(3) (May 2021), <http://dx.doi.org/10.1111/exsy.12665>
274. Langdon, W.B.: Evolving open complexity. *SIGEVolution newsletter of the ACM Special Interest Group on Genetic and Evolutionary Computation* 15(1) (Mar 2022), <http://dx.doi.org/10.1145/3532942.3532945>
275. Forrest, S.: Engineering and evolving software (30 May 2021), <http://dx.doi.org/10.1109/GI52543.2021.00008>, invited keynote
276. Moore, G.E.: Cramming more components onto integrated circuits. *Electronics* 38(8), 114–117 (April 19 1965)
277. Wright, S.: The roles of mutation, inbreeding, crossbreeding and selection in evolution. In: *Proceedings of the Sixth Annual Congress of Genetics*. pp. 356–366 (1932), <http://www.blackwellpublishing.com/ridley/classictexts/wright.pdf>
278. Turing, A.M.: Intelligent machinery. In: Meltzer, B., Michie, D. (eds.) *Machine Intelligence*, vol. 5, chap. 1, pp. 3–23. Edinburgh University Press, Edinburgh, UK (1969), <https://hashingit.com/elements/research-resources/1948-intelligent-machinery.pdf>