## AI Programmer: Autonomously Creating Software Programs Using Genetic Algorithms

Kory Becker

kbecker@primaryobjects.com Bloomberg Justin Gottschlich

justin.gottschlich@intel.com Intel Labs

## Abstract

In this paper, we present AI Programmer, a machine learning (ML) system that can automatically generate full software programs, while requiring only minimal human guidance. At its core, AI Programmer uses a genetic algorithm (GA), coupled with a tightly constrained programming language that minimizes the overhead of its ML search space. Part of AI Programmer's novelty stems from (i) its unique system design, including an embedded, hand-crafted interpreter for efficiency and security and (ii) its augmentation of classic GA to include instruction-gene randomization bindings and programming language-specific genome construction and elimination techniques. We provide a detailed examination of AI Programmer's system design, several examples detailing how the system works, and experimental data demonstrating its software generation capabilities and performance using only mainstream CPUs.

*CCS Concepts* • Software and its engineering  $\rightarrow$  Genetic programming; General programming languages; • Theory of computation  $\rightarrow$  *Program analysis* 

*Keywords* Genetic algorithm, program synthesis, genetic programming, evolutionary computation, artificial intelligence, machine learning, programming languages, code generation and optimization

## 1. Introduction

Since the invention of the computer, having the ability to correctly and efficiently develop software programs has been a principle challenge [4]. To help address this, countless breakthroughs have been made in the field of software and hardware development. Some of these include safety and flexibility advances in static, dynamic, and gradual type sys-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '21 Companion, July 10-14, 2021 Lille, France

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM

ISBN 978-1-4503-8351-6/21/07...\$15.00

DOI: http://dx.doi.org/10.1145/3449726.3463125

tems [2, 22]; simplification, safety, and robustness advances using automatic memory management and garbage collection systems [5, 10]; generality and specificity progress in both general-purpose and domain specific languages [21, 26]; and, of course, a plethora of tools aimed at assisting programmers in nearly every way [8, 9, 19].

While such advances continue to broaden and deepen the space of what is computationally tractable, they have the fracturing side-effect of complicating and exacerbating the tension between the ease of developing software and the ability for humans to write maximally efficient code.

In this paper we explore an alternative approach to traditional human-driven software development; one that autonomously creates software programs using a genetic algorithm (GA) requiring only minimal human guidance.

## 1.1 The Evolution of Programming Languages

Over the last several decades, programming languages (PLs) have followed a steady path of providing higher-level programming abstractions, aimed at reducing the challenge of human-driven software development [1]. To this end, PLs, in general, have proliferated toward a design goal of simplifying human use. Although this trend is natural in an era where humans perform the majority of software development, as we will show, it is suboptimal in an environment where programming is performed predominantly by machines.

This paper makes the following technical contributions:

- 1. We present AI Programmer, a software generation framework, which constructs programs using an enhanced genetic algorithm coupled with a minimalistic programming language.
- 2. We present some observations, including an embedded interpreter and simulator solution, for security and optimization of ML-generated software.
- 3. We provide empirical results demonstrating the efficacy and efficiency of AI Programmer across several of its fully generated software programs on mainstream CPU hardware.

## 2. Background

A genetic algorithm (GA) is a type of artificial intelligence, modeled after biological evolution, that begins with no knowledge of the subject, aside from an encoding of genes that represent a set of instructions or actions [6]. In the concept of GA-driven computer programming, a series of programming instructions are selected at random to serve as an initial chain of DNA. The complete genome is executed as a program, with the resulting fitness score calculated according to how well the program can solve a given task. This is performed with a sufficiently large population size. Those that have the best fitness are mated together to produce offspring.

Each generation of programs receive extra diversity from evolutionary techniques including roulette selection, crossover, and mutation [17]. The process is repeated at each epoch with each child generation hopefully producing more favorable results than its parents' generation until a target solution is found. Through this process, applying GAs to computer programming automation enacts a survival of the fittest model for computer program generation [16]. A deeper examination of these GA principles are provided in Section 3.

## 3. The Design of AI Programmer

In this section we provide a high-level overview of the AI Programmer software architecture.

### 3.1 Programming Language Selection and Challenges

We chose a typeless programming language that contains only eight instructions to drive AI Programmer's software generation [18]. We briefly discuss the advantages of this approach and the modifications to the language that were required to integrate it into a GA solution.

Table 1. AI Programmer Instruction Set and Gene Map			
Instr	Gene Range	Operation	
>	(0, 0.125]	Increment the pointer	
<	(0.125, 0.25]	Decrement the pointer	
+	(0.25, 0.375]	Increment the byte at the pointer	
-	(0.375, 0.5]	Decrement the byte at the pointer	
	(0.5, 0.625]	Output the byte at the pointer	
,	(0.625, 0.75]	Input a byte and store it at the ptr	
[	(0.75, 0.875]	Jump to matching ] if current 0	
]	(0.875, 1.0]	Jump back to matching [ unless 0	

*Turing Completeness.* AI Programmer's programming language, listed in Table 1, is Turing complete. A *Turing complete* programming language is theoretically capable of completing any (single taped Turing machine) programming task given an unlimited amount of time and memory [23]. In essence, a programming language with this characteristic is capable of implementations of a vast number of pro-

gramming problems. Likewise, programs generated with AI Programmer are theoretically capable of expressing all tasks that one might want to accomplished with computers.

**GA Engine and Uniform Gene Distributions.** AI Programmer's genetic algorithm engine represents each generated program's instructions as an array of floating point values, which, when considered as a unit, is its genome. Each individual location within a given genome is called a *gene*. Each gene within a program's genome corresponds to a single instruction from Table 1.

AI Programmer binds a gene value range to each of its instructions across a *continuous uniform distribution* (or *rectangular distribution*) [3] (see Table 1), where each instruction's gene range is equal in size to each of the others. This was done so each instruction would have an equally random probability of being chosen at any location in a gene sequence when randomization was needed. <sup>1</sup>

*Simplified Instruction Set.* Each of AI Programmer's instructions manipulate a memory "tape" of byte values, ranging from 0-255. The language works by applying increment and decrement operations to the current memory cell, while shifting the memory cell up and down the tape, as instructed by the program. The values at the current memory pointer can be input from the user or output to the terminal. Primitive looping instructions also exist (e.g., '[' and ']'), offering a complete instruction set for creating software. An example program is shown in Figure 1.

+-+-+>-<[++++>+++++<+<>++]>[-[---.--[[-.++++ [+++..].]]]]

Figure 1. A generated program that outputs "hello".

The simplified instruction set reduces the search space in which a target program code can be found. As computational devices improve in speed, larger problem spaces can be searched. However, on less powerful devices, the search space needs to be constrained. As AI Programmer is intended for general purpose developers, limiting the programming instruction set to eight instructions enables the engine to execute in reasonable times on commodity hardware (see Section 5).

### 3.2 Genomes and Generations

To generate a software program using genetic algorithms, one must first create a genome. A *genome* is a set of genes that are grouped together as a single unit. For AI Programmer, the genome is encoded as an array of floating point values, with fixed value ranges per unique instruction ranging between 0 and 1, as shown in the Gene Range column of Table 1.

Once a genome is created, it is converted to a corresponding program, executed, and the resulting program is assigned

<sup>&</sup>lt;sup>1</sup> We did not examine the impact of weighted ranges for different programs, but note that it may be of interest as future work.

a fitness score based on the program's output. The closer a generated program comes to solving the provided task, the greater its fitness score and, the more likely it is to continue to the next evolutionary generation. At each generation epoch, AI Programmer utilizes roulette selection, along with crossover and mutation, to create child programs that contain slight random perturbations, and potentially better, genomes than their parents for solving the target task.

*Constructing a Genome* Figure 2 demonstrates an example of constructing a genome from an array of floating point values. Each value range maps to a specific instruction in the programming language. Initially, these values are random, resulting in generated programs that either won't function properly, throw errors, or simply fail <sup>2</sup>. However, one or two are bound to run and execute, at a minimum, some number of valid instructions. The more successful a program is at executing, the more likely it is to continue on and produce offspring with code that achieves more successful results.



Figure 2. Decoding a genome as a program.

**Crossover and Mutation** To create offspring, a parent genome contributes part of its genes to the child, a process called *crossover*, as shown in Figure 3. In addition to inheriting programming instructions from its parent, each child can also experience *mutation*, which is the process of adding controlled, but random perturbation, to specific genes. This results in modified behavior of the value of a particular gene, resulting in a change to the resulting programming instruction, and thus, the overall program.



**Figure 3.** An example of crossover and mutation. The child genome inherits the first 5 instructions from its parent. One instruction is mutated.

Crossover copies forward potentially beneficial parts of the parent, while mutation offers differing behaviors of instruction combinations, which may or may not, end up making the child programs more successful. *Survival of the Fittest* Executable programs are ranked according to how well they have performed. As shown in Figure 4, a particular program that has failed is often immediately removed from the pool of genomes. However, programs that succeed are carried forward to produce child programs.<sup>3</sup>



**Figure 4.** Programs are weighted by fitness, with the most successful used for child program generation.

## 3.3 The Fitness Test

To use GAs, a fitness test is needed to determine how well a generated solution performs. In the context of AI Programmer, this can involve scoring the byte-level output of the generated program, inspecting the generated program's internal state, or even analysis of intermediate state changes of the program throughout its execution. The score of the fitness test is calculated by analyzing these characteristics, and many others, and then comparing them against a userdefined target.

This concept is similar to test-driven development. When all unit tests pass, a program may be considered to be functionally correct. Likewise, a fitness test for a GA can be considered as a set of unit tests. In the case of AI Programmer, a fitness test typically contains a suite of tests for varying scenarios, which guide the genome selection, preserving only programs that evaluate well on the test suite. Further details about the construction of AI Programmer fitness tests are described in Section 4.

#### 3.4 AI Programmer's Sandboxed Interpreter

Once a program has been generated, it must be executed so it can be evaluated against human created fitness tests. However, the execution of ML generated programs may include potential security risks as well as performance degradations. Because of this and the need for complex fitness tests (Section 3.3), we developed our own interpreter. This interpreter is sandboxed within the AI Programmer system to provide a secure, efficient, and GA-appropriate execution environment. We explain the challenges and benefits of this system in the following subsections.

# 3.4.1 Execution of Generated Programs in a Controlled Environment

As generated programs are executed to evaluate their fitness, the results can often be undesirable and potentially danger-

<sup>&</sup>lt;sup>2</sup> Most initial programs in the gene pool fail immediately upon being executed. Others may result in endless loops. It is due to these reasons that exception handling and maximum iteration limits are imposed on the interpreter.

<sup>&</sup>lt;sup>3</sup> In Figure 4, the bottom program is a valid running program that takes one byte for input, increments it, and then displays it twice as output.

ous. Consider a program generated with I/O instructions, allowing for the modification of files on disk. A generated program could potentially overwrite critical system files, rendering the entire machine inoperable. Likewise, a program generated with instructions to support networking could inadvertently flood a computer network (e.g., denial of service attack [28]) or replicate itself across machines (e.g., worm [12]).

Normally, these types of behaviors are malicious, yet ML software generators happen upon these situations in an attempt to satisfy fitness goals. By executing programs within our own secure interpreter, which includes instruction-level protection checks, AI Programmer can provide the additional security measures that are needed to prevent ML generated software from causing harmful behaviors. Non-ML generating interpreters and compilers do not generally include these types of checks because management of such issues are not usually within their scope [27].

#### 3.4.2 Termination of Infinite Loops

Automatically generated software has the potential to create infinite loops. This can occur from unsatisfied loop termination constraints or unexpected looping instructions. In our experiments, this type of behavior often arises in early program generations due to the GA maximizing the goal fitness score at the cost of program execution time. As a result, unterminated programs have the potential to halt the generation process, resulting in a failure of further program evaluations.

In an attempt to mitigate this, one can add fitness constraints to prefer programs with fewer instruction counts over larger ones. However, the generation of infinite loops, especially in early generations of programs, remains a possibility. AI Programmer's interpreter includes a customizable maximum instruction count per execution. Programs that exceed the instruction count are terminated. A fitness penalty can then be applied, reducing the likelihood of future generations of programs carrying forward the infinite loop constructs. With this addition, AI Programmer is guaranteed to terminate all infinite loops. Although this solution has unintended consequences (such as terminating and discarding viable programs prematurely), it does ensure forward progress. We are exploring other possible solutions that do not have such obvious negative side-effects.

#### 3.4.3 Simulation of Complex Instructions

Optimizing program execution is a principle concern for ML program generators. This is because such systems may generate and execute dozens to millions of programs before one with a high enough fitness score is found. While simple operations, such as add, load, store, jmp, may take a single clock cycle to complete, more complex operations can require many. Examples include disk I/O, networking, and peripheral device access. These types of operations can significantly increase program execution time, as they often rely on accessing services or devices with increased latency.

AI Programmer can simulate the execution of these complex instructions. In doing so, the GA-based programs it generates can execute more efficiently, while still retaining the ability to check the program's fitness goals. Moreover, such simulation protects the devices themselves from overuse by the plethora of programs that may attempt to access them during exploratory evolution of GA program generation.

## 4. Using AI Programmer

AI Programmer consists of a modular framework, designed in C# .NET. It includes an engine for running genetic algorithms, an encoder and decoder for genomes, a sandboxed interpreter for simulated program execution, and a compiler to transform code into binary executables. While the initial design of AI Programmer uses C#, it is important to note that the principles employed by it are not bound to C#.

AI Programmer's software framework for fitness test construction is extensible and was developed so users can devise a myriad of customized fitness suites, which eventually guide the system's GA generation and evolution of software programs.

*Specifying Requirements of a Program* To generate a program, AI Programmer must be provided with the requirements for the desired input and output of the target program. For example, if a program should prompt a user for a numerical input and then subsequently output a line of text, this must be specified in the form of training data to AI Programmer. The following subsections detail the step-by-step process of how a program is specified and generated with AI Programmer.

**Creating Your First Program** To begin, a user creates a C# class within the AI Programmer project, inheriting from the FitnessBase base class. This class includes all the necessary requirements for specifying a solution to be built by AI Programmer, including fitness scoring functionality, program termination rules, and program generation capabilities.

*Specifying a Target Fitness Score* Next, the user indicates a target fitness score which is specified in the constructor of the class as shown in Figure 5. The score is typically based upon characteristics of the desired program. For example, if the target program is intended to output a string, such as "Hello World", the fitness score might be the number of characters in the string (i.e., 11). However, since AI Programmer generates programming code at the byte level, the fitness score should account for incremental differences in output characters. In this case, one should multiply each target character in the output by 256, resulting in 2816 (e.g., 11 \* 256) and use that as the resulting target fitness score.

AI Programmer is designed to continue execution, generating incrementally better programs that satisfy the fitness conditions, until the current fitness score reaches its target.

*Specifying Fitness Conditions* Next, the user must specify the rules that are used to score each of the generated

```
public StringFitness(GA ga, int maxIterCount)
  : base(ga, maxIterCount)
{
   _targetStr = "Hello World";
   _targetFit = _targetStr.Length * 256;
}
```

Figure 5. Example target fitness score for "Hello World."

programs (i.e., the fitness test). At each generation epoch, AI Programmer will favor programs that have fitness scores that are closer to the target fitness score. Therefore, careful crafting of the fitness conditions are required so fitness scores accurately represent the desired goal of the program.

In the "Hello World" example, the user must specify that the output of the program should match the target string. To achieve this, one can add to the fitness score according to how close each character in the generated output string is to the target string. In particular, the fitness test can simply loop over the characters in the string "Hello World", and compare each one against the characters produced in the output of the generated program, adding or subtracting accordingly, as shown in Figure 6.

```
for (int i = 0; i < _targetStr.Length; i++)
if (_console.Length > i)
Fit += 256-Math.Abs(_console[i]-_targetStr[i]);
```

Figure 6. Adding and subtracting the fitness score based upon program output.

After assigning a fitness score to each generated program within the current pool, a check is made to determine whether the target fitness score has been achieved by any of the generated programs. If so, AI Programmer halts and returns the solution program. Otherwise, it continues with the next generation of programs.

*Specifying Conditions for Variable Output* Previously, we presented an example on how to train AI Programmer to find an exact string. However, for more complex scenarios, such as variable outputs or calculated values, a series of training examples may be required. In such scenarios, training data can be created to serve as an initial set of examples to base the fitness score upon. Thereafter, AI Programmer can be guided with an evolutionary goal to generalize from the training data and provide correct results for new data.

As an example of variable output, consider the generation of a program to output the summation of two numbers. The target fitness score for this program would be the desired output, which, in this case, would be one byte, multiplied by the range of potential values (256). To construct the target fitness score, we can simply multiply the target fitness for a single result by the number of training examples. Therefore, in this case, the target fitness is trainingCount \* 256.

After specifying the target fitness score, we can implement the actual fitness check for adding two numbers by looping over each training set combination (consisting of two numbers), inputting those values to our program, and checking the output for the correct sum. An example of this is shown in Figure 7. By providing varying training input values we can help foster the generalization of the solution program, rather than the generation of a program that only solves the exact training examples provided.

```
int val;
if (Int32.TryParse(_console.ToString(), out val)) {
  Fit += 256 - Math.Abs(val - (input1 + input2));
}
```

Figure 7. Calculating the fitness of adding two numbers.

**Programmatic Sequences of Action** Because different programs require different sequences of actions (e.g., requesting input, outputting a result, etc.), AI Programmer provides users with a mechanism to specify the necessary programmatic sequence of actions within the fitness method.

Programmatic sequences can be provided in the form of a simple state machine within the fitness check method. When the generated program executes a command to request input from the user, a bonus score can be applied to the fitness if it is executed at the correct time in the sequence of actions. Likewise, when data is output, one can add or subtract from the fitness score according to the time the action is executed.

It is important for users to account for programmatic sequence bonuses when they are generating the initial target fitness. Doing so will ensure the generated solution will satisfy all required constraints, including sequences of events, before returning a viable solution program.

## 5. Results

Using AI Programmer, we were able to generate numerous complete software programs. A complete listing of these programs, their associated program generation time, and the total number of evolutionary generations used to build them are shown in Table 2. It is important to note that the number of evolutionary generations is not equivalent to the total computational time to generate a program. This is due, in large part, to varying genome size and fitness function computation, which is unique to each program.

For the remainder of this section, we highlight the details of some of the programs listed in Table 2 and discuss novel aspects that emerged when generating them.

### 5.1 Greetings

"Hello World" is usually one of the first programs human programmers create when they begin learning programming. As such, we found it fitting to guide AI Programmer to learn some basic greetings for its early programs. Rather than starting with "Hello World", we first had AI Programmer create a more simplistic program that simply output "hi." It

Table 2.         AI Programmer Results			
Name	Duration (s)	Generations	
hi	52	5,700	
Hi!	7,644	1,219,400	
hello	1,713	252,000	
hello world	7,702	580,900	
reddit	1,362	195,000	
Keep Calm Keep Coding	944	21,400	
I love all humans	36,000	6,057,200	
hello {user}	1,793	42,800	
Addition	2,698	92,400	
Subtraction	4,305	177,900	
Multiply x2	6,353	242,000	
Multiply x3	5,165	87,200	
XOR	2,095	146,400	
Fibonacci	21,862	151,900	
If/then conditionals	8,313	46,200	
cats are evil	10,209	814,400	
Bottles of Beer on the Wall	2,957	61,400	
Reverse string	49	2,600	
CSV parse	173	9,000	
Extract in quotes	6,478	212,100	
Extract in quotes 2	9,996	188,400	
Trim left of quote	9,030	341,700	
XML to JSON	6,866	820,900	
Warning countdown	48	900	

+[+++++-+>++>+++++++<<]>++.[+.]

Figure 8. Generated program: "hi"

was successfully after 5,700 generations and the generated code is shown in Figure 8.

The generated program fulfilled its requirement to output the target text, but interestingly included subsequent random characters, which contained parsing errors, including nonmatching brackets. However, AI Programmer's interpreter computes results until the program fails. In this manner, the syntax error (which is later on in the code, after a solution is reached) does not negatively impact its fitness score, and thus offers a working solution. In fact, the generated code can be executed in almost any third-party interpreter as a valid working program (provided, warnings are ignored).

Next, we guided AI Programmer to generate the famous "hello world" output which was successfully constructed after 580,900 generations and consists of the code shown in Figure 9.



**"I love all humans"** As a humorous aside, we asked AI Programmer to create the program to output "I love all humans," which was successfully generated after 6,057,200 generations. It consists of the code shown in Figure 10. The fitness method for this example includes a check on the output string length to ensure an exact matching output, without extraneous text.

To ensure an exact output string, the fitness score includes not just a check on the output characters, but also a check on the length of the string. In this case, the target fitness included an additional 10 points, of which a percentage of this amount is added to the resulting fitness, depending on how close the length of the output string matches the length of the target. This forces the generation of a program that outputs the exact target string, without extraneous output instructions, as the generation process will not halt until the target fitness is reached, of which, 10 points comprise having the correct output length.

+[>+<+++]+>+<+++++++++++++++++++++++++++++
++++++++.>+++++++++++++++++++++++++++++
+++.+++.+++++++
+++++++<.>++++++++++++
+++++++++++++++++++++++++++++++++

Figure 10. Generated program: "I love all humans"

```
// Assigning the target fitness.
_targetFitness = _targetString.Length * 256;
_targetFitness += 10;
```

••

```
// Calculating the fitness length bonus.
Fitness += 10 * ((_targetString.Length -
    Math.Abs(_console.Length -
    _targetString.Length)) /
    _targetString.Length);
```

**Figure 11.** A percentage of 10 points is added to the fitness, according to how exact the length of the output is to the target.

#### 5.2 Input-Output Computations

We next guided AI Programmer to generate programs that perform computations based on user input. In such programs, the user provides some input and the computer program then generates the appropriate output.

*Reversing a String* AI Programmer was able to generate the program to reverse any string after only 2,600 generations. The generated code is shown in Figure 12.

+->,>,[>+,],,,,-<[.+<]

Figure 12. Generated program for reversing a string.

When executed, the program prompts the user for input. The user then types one character at a time until a value of "0" is entered. A novelty of this program is that it is required to take variable size input first before performing the majority of its program logic. However, the program's internal memory state must manage the variable input, as the program must read all input first to locate the final character entered, which is the first character in the reversed string. The genetic algorithm was able to produce this logic automatically, based upon the fitness method.

*Addition and Subtraction* AI Programmer was able to generate programs for addition after 92,400 generations (Figure 13) and subtraction after 177,900 generations (Figure 14).

,>,-[-<+>]<+.

Figure 13. Generated program for performing addition.

,-->,-[-<->]<+.

Figure 14. Generated program for performing subtraction.

*If-Then Conditionals with User Input* Generating programs involving more complex programming logic, such as the ability to perform if-then decisions and actions, requires a more advanced type of fitness function. However, as described in Section 3.4, AI Programmer's embedded interpreter provides significantly more access to program state than just its output, which is essential for generating a large variety of more complex programs.

For example, AI Programmer was able to produce a program which prompts the user for input (e.g., 1, 2 or 3) and outputs text based on which value was entered, similar to selecting an option from a menu. By entering the value "1", the program would output "hi". Entering "2", resulted in the program output of "z". Entering "3", resulted in the output "bye". The program was generated in 446,200 generations.

The produced code was notably larger than previously generated programs, containing 650 instructions (although not all instructions are needed). The larger code was required, as the conditional branches are contained within individual blocks of the code.

#### 5.3 Complexity in Fitness Functions

As the complexity of the target program grows, so too does the fitness function. After all, the fitness function needs to guide the engine in determining how well a particular child program matches the targeted solution. For conditionals and branching, successful program generation required more advanced techniques within the fitness function.

In particular, a check was needed to examine the interpreter's memory register (i.e., current data pointer via shift operations), where the distinct number of memory registers being used by the program was counted, providing a bonus to fitness to favor more memory register usage over less. This aided in inspiring diversity amongst child programs. Additionally, the instruction pointer used for each print command was recorded and weighed against the fitness score. A penalty was applied for reuse of the same print command. This helped to foster diversity and achieve a successful ifthen result.

## 6. Optimizing Program Generation

We noticed that the program generation time increased significantly as the length of the target output increased. Furthermore, the need to extend AI Programmer beyond the basic instruction set was deemed a necessity if we were to have it produce programs with more interesting features, such as file I/O and networking capabilities.

As such, we extended AI Programmer to use an extended programming instruction set, which reduced code generation time and improved code compression due to an increased range of instruction specificity (i.e., fewer instructions to achieve the same result). However, a disadvantage of utilizing the extended instruction set is that the generated programs would be difficult to test in standard interpreters. As the extended instruction set for AI Programmer deviated from the traditional programming language, standard interpreters would no longer be able to run the produced code. In our case, AI Programmer's internally developed interpreter was modified to support the extended instruction set, so this was not a practical obstacle.

#### 6.1 Extended Instruction Set

Several extensions of the programming language used by AI Programmer exist, which are suitable to decrease program generation time. Specifically, the speed-enhancing extension set, Extended Type III [7], offers several programming instructions that aid generation. These instructions include the ability to immediately set the value of a particular cell to a multiple of 16, also called "fast cell initializers". This aids in allowing a generated program to quickly reach displayable ASCII range characters for output, thus, decreasing the number of individual increment programming instructions that would normally be required.

In addition to key instructions taken from Extended Type III, we added several new instructions to support calling functions from within a program, allowing for increasingly complex programs to be generated.

*Fibonacci Sequence* With these extensions in place, AI Programmer was able to generate a program to output the Fibonacci sequence up to 233<sup>4</sup>, which was was generated in approximately six hours. The program prompts the user for input of the two starting values in the sequence. It then outputs the next digits in the Fibonacci sequence. The generated

 $<sup>^4\,255</sup>$  is the max value for a byte, with the next Fibonacci sequence value being 377.

,>,\$[!>--\$<<a>>]4]+,,-[-<+>]<+.\$@

**Figure 15.** Generated program to output the Fibonacci sequence from two starting input values.

code for this, using the extended instruction set, is shown in Figure 15.

Advancing Complexity The ability of the GA to generate a program for solving the Fibonacci sequence was a profound advancement. The solution program contains several distinct programming tasks, including prompting the user for input of two numbers at the beginning of execution, calculating the addition of values, determining the correct mathematical sequence, outputting the result, and looping to repeat the process for each value in the sequence.

This combination of tasks, spreading across a range of programming abilities, might typically be given to human programmers in order to evaluate their programming proficiency. The capability of the GA to automatically generate this type of program demonstrates the potential for future expansion of the system.

## 7. Related Work

Genetic algorithms have previously been applied in some restricted cases. A key limitation in their broader application has been in the computational density of the search space involved in program generation, which exponentially increases as programs grow in complexity [13]. AI Programmer provides to novel mitigation of this inefficiency by using a minimalistic programming language, exploiting the natural parallelism of GAs, simulating complex instructions, and embedding an optimized interpreter for fast execution and fast-failure of defunct programs.

#### 7.1 Genetic Algorithms in Other Domains

Somewhat related to our work, is the use of program synthesis driven by genetic programming in hardware-based niche fields. Koza et al. used an automated process for creating analog circuits, involving genetically evolved designs with evolutionary computation to produce circuit components that typically require human-level intelligence to construct [14]. In addition, they used human constructed fitness methods to guide their circuit design. Although applied in different domains, the high-level machine learning approach of Koza et al.'s system is similar to AI Programmer.

One of the key components of our research is the usage of a minimalistic programming language to limit the computational complexity of generated programs. This approach been found useful in other areas of genetic programming, including the simulation of artificial life, as described in Ling's work [15]. In a simulation library based on genetic algorithms and biological hierarchy, the system, called Ragaraja, uses biological concepts to form an esoteric programming language, consisting of a set of 3-character instructions. In this manner, the system is able to simplify the genetic algorithm generation and mutation process by limiting the number of possible instruction combinations. Although applied in a completely different domain, the affects of this approach are similar to AI Programmer, specifically for optimizing the generation time and limiting the complexity of generated solutions to a constrained set of instructions.

#### 7.2 Different Approaches in Program Generation

AI Programmer has similarities to a program synthesis technique called *sketching* in that each approach attempts to automatically generate software by using some human guidance. However, the similarities between the two approaches ends there. On one hand, sketching is a program synthesis technique where a programmer provides only a minimalistic outline of an implementation and the compiler generates the remaining code [24, 25]. On the other hand, AI Programmerrequires no partial implementation, but instead requires human developers to design fitness tests which guide the evolutionary algorithm for the entire program construction.

Another slightly related work is verified lifting [11]. Verified lifting aims to lift algorithms written in one language and place a formally verified equivalent in another language. The benefits of verified lifting are highly practical, especially when considering the need for such systems as real software systems often migrate from one programming language to another. However, verified lifting and AI Programmer are only loosely similar in that both systems perform automatic code generation, but do not possess any other similarities in their approaches.

## 8. Conclusion and Future Work

The results presented in this paper, provide early notions about the power that machine learning techniques, specifically genetic algorithms, may offer a partial solution for automatic program generation. We showed that fully functional programs can indeed be automatically generated, provided they are supplied with some human guidance in the way of input parameters and training data.

In addition to correctness, efficient implementation of fitness methods are imperative to the practical application of AI Programmer. This is because each generated program is checked against the fitness method every time a new program is evaluated. An important open area of future work is the deep examination of how to implement fitness methods as efficiently as possible while still retaining a high degree of correctness. One possible solution is to build superoptimizers specifically for fitness test optimizations [20].

Another important open area in ML-based program generation is the need for specifically crafted programming languages that have strong alignment with the constraints of ML computation. By re-thinking the type of languages used for ML-based program generation, we can envision a new future of software development, driven by artificial intelligence systems, with human creativity guiding the way.

## References

- [1] The 2017 top programming languages. URL https://spectrum.ieee.org/computing/software/ the-2017-top-programming-languages.
- [2] L. Cardelli and P. Wegner. On understanding types, data abstraction, and polymorphism. ACM Comput. Surv., 17(4): 471-523, Dec. 1985. ISSN 0360-0300. URL http://doi. acm.org/10.1145/6041.6042.
- [3] K. Claessen, J. Duregård, and M. H. Pałka. Generating constrained random data with uniform distribution. *Journal of functional programming*, 25, 2015.
- [4] T. H. Cormen, C. Stein, R. L. Rivest, and C. E. Leiserson. *Introduction to Algorithms*. McGraw-Hill Higher Education, 2nd edition, 2001. ISBN 0070131511.
- [5] E. W. Dijkstra, L. Lamport, A. J. Martin, C. S. Scholten, and E. F. M. Steffens. On-the-fly garbage collection: An exercise in cooperation. *Commun. ACM*, 21(11):966–975, Nov. 1978. ISSN 0001-0782. URL http://doi.acm.org/10.1145/ 359642.359655.
- [6] P. Domingos. The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. 2015.
- [7] Esolangs.org. Extended type iii. URL https://goo.gl/ 9bS2gF.
- [8] J. E. Gottschlich, M. P. Herlihy, G. A. Pokam, and J. G. Siek. Visualizing transactional memory. In *Proceedings of the 21st International Conference on Parallel Architectures and Compilation Techniques*, PACT '12, pages 159–170, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1182-3. URL http://doi.acm.org/10.1145/2370816.2370842.
- [9] S. L. Graham, P. B. Kessler, and M. K. Mckusick. Gprof: A call graph execution profiler. *SIGPLAN Not.*, 17(6):120–126, June 1982. ISSN 0362-1340. URL http://doi.acm.org/ 10.1145/872726.806987.
- [10] R. Jones and R. Lins. Garbage Collection: Algorithms for Automatic Dynamic Memory Management. John Wiley & Sons, Inc., New York, NY, USA, 1996. ISBN 0-471-94148-4.
- [11] S. Kamil, A. Cheung, S. Itzhaky, and A. Solar-Lezama. Verified lifting of stencil computations. In *Proceedings of the* 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '16, pages 711–726, New York, NY, USA, 2016. ACM. ISBN 978-1-4503-4261-2. doi: 10.1145/2908080.2908117. URL http://doi.acm.org/ 10.1145/2908080.2908117.
- [12] P. K. Kerr, J. Rollins, and C. A. Theohary. *The stuxnet computer worm: Harbinger of an emerging warfare capability*. Congressional Research Service Washington, DC, 2010.
- [13] J. R. Koza. Human-competitive results produced by genetic programming. *Genetic Programming and Evolvable Machines*, 11(3-4):251–284, 2010.
- [14] J. R. Koza, F. H. Bennett, D. Andre, M. A. Keane, and F. Dunlap. Automated synthesis of analog electrical circuits by means of genetic programming. *IEEE Transactions on evolutionary computation*, 1(2):109–128, 1997.

- [15] M. H. Ling. An artificial life simulation library based on genetic algorithm, 3-character genetic code and biological hierarchy. *The Python Papers*, 7:5, 2012.
- [16] Z. Michalewicz. Genetic Algorithms Plus Data Structures Equals Evolution Programs. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2nd edition, 1994. ISBN 0387580905.
- [17] M. Mitchell. An Introduction to Genetic Algorithms. MIT Press, Cambridge, MA, USA, 1998. ISBN 0262631857.
- [18] U. Müller. Esoteric programming. URL https://goo.gl/ GnzaQe.
- [19] H. Patil, C. Pereira, M. Stallcup, G. Lueck, and J. Cownie. Pinplay: A framework for deterministic replay and reproducible analysis of parallel programs. In *Proceedings of* the 8th Annual IEEE/ACM International Symposium on Code Generation and Optimization, CGO '10, pages 2–11, New York, NY, USA, 2010. ACM. ISBN 978-1-60558-635-9. URL http://doi.acm.org/10.1145/1772954.1772958.
- [20] P. M. Phothilimthana, A. Thakur, R. Bodik, and D. Dhurjati. Scaling up superoptimization. *SIGPLAN Not.*, 51(4):297–310, Mar. 2016. ISSN 0362-1340. doi: 10.1145/2954679.2872387. URL http://doi.acm.org/10.1145/2954679.2872387.
- [21] M. L. Scott. Programming Language Pragmatics, Third Edition. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2009. ISBN 0123745144, 9780123745149.
- [22] J. G. Siek and W. Taha. Gradual typing for functional languages. In IN SCHEME AND FUNCTIONAL PROGRAM-MING WORKSHOP, pages 81–92, 2006.
- [23] R. I. Soare. Turing computability. Springer, 2016.
- [24] A. Solar-Lezama, R. Rabbah, R. Bodík, and K. Ebcioğlu. Programming by sketching for bit-streaming programs. *SIG-PLAN Not.*, 40(6):281–294, June 2005. ISSN 0362-1340. doi: 10.1145/1064978.1065045. URL http://doi.acm.org/ 10.1145/1064978.1065045.
- [25] A. Solar-Lezama, L. Tancau, R. Bodik, S. Seshia, and V. Saraswat. Combinatorial sketching for finite programs. *SIGOPS Oper. Syst. Rev.*, 40(5):404–415, Oct. 2006. ISSN 0163-5980. doi: 10.1145/1168917.1168907. URL http: //doi.acm.org/10.1145/1168917.1168907.
- [26] A. K. Sujeeth, K. J. Brown, H. Lee, T. Rompf, H. Chafi, M. Odersky, and K. Olukotun. Delite: A compiler architecture for performance-oriented embedded domain-specific languages. ACM Trans. Embed. Comput. Syst., 13(4s):134:1– 134:25, Apr. 2014. ISSN 1539-9087. URL http://doi. acm.org/10.1145/2584665.
- [27] L. Torczon and K. Cooper. *Engineering A Compiler*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2nd edition, 2011. ISBN 012088478X.
- [28] S. T. Zargar, J. Joshi, and D. Tipper. A survey of defense mechanisms against distributed denial of service (ddos) flooding attacks. *IEEE Communications Surveys and Tutorials*, 15 (4):2046–2069, 2013. URL http://dblp.uni-trier.de/ db/journals/comsur/comsur15.html#ZargarJT13.