

# Fully Autonomous Programming using Iterative Multi-Agent Debugging with Large Language Models

ANASTASIYA GRISHINA\*, Simula, Norway and University of Oslo, Norway

VADIM LIVENTSEV\*, TU Eindhoven, The Netherlands and Philips Research, The Netherlands

AKI HÄRMÄ, Philips Research, The Netherlands

LEON MOONEN†, Simula, Norway

Program synthesis with Large Language Models (LLMs) suffers from a “near-miss syndrome”: the generated code closely resembles a correct solution but fails unit tests due to minor errors. We address this with a multi-agent framework called Synthesize, Execute, Instruct, Debug, and Repair (SEIDR). Effectively applying SEIDR to instruction-tuned LLMs requires determining (a) optimal prompts for LLMs, (b) what ranking algorithm selects the best programs in debugging rounds, and (c) balancing the repair of unsuccessful programs with the generation of new ones. We empirically explore these trade-offs by comparing replace-focused, repair-focused, and hybrid debug strategies. We also evaluate lexibase and tournament selection to rank candidates in each generation. On Program Synthesis Benchmark 2 (PSB2), our framework outperforms both conventional use of OpenAI Codex without a repair phase and traditional genetic programming approaches. SEIDR outperforms the use of an LLM alone, solving 18 problems in C++ and 20 in Python on PSB2 at least once across experiments. To assess generalizability, we employ GPT-3.5 and Llama 3 on the PSB2 and HumanEval-X benchmarks. Although SEIDR with these models does not surpass current state-of-the-art methods on the Python benchmarks, the results on HumanEval-C++ are promising. SEIDR with Llama 3-8B achieves an average pass@100 of 84.2%. Across all SEIDR runs, 163 of 164 problems are solved at least once with GPT-3.5 in HumanEval-C++, and 162 of 164 with the smaller Llama 3-8B. We conclude that SEIDR effectively overcomes the near-miss syndrome in program synthesis with LLMs.

CCS Concepts: • **Software and its engineering** → *Software design engineering*; • **Computing methodologies** → *Neural networks; Model development and analysis; Search methodologies*.

Additional Key Words and Phrases: automatic programming, large language models, program repair

## ACM Reference Format:

Anastasiya Grishina, Vadim Liventsev, Aki Härmä, and Leon Moonen. 2025. Fully Autonomous Programming using Iterative Multi-Agent Debugging with Large Language Models. *ACM Trans. Evol. Learn. Optim.*, (February 2025), 37 pages. <https://doi.org/10.1145/3719351>

## 1 Introduction

Automatic programming has been an important goal in the field of artificial intelligence almost since its inception, promising to reduce the workload of software developers by automatically solving some of the tasks they face [Manna and Waldinger 1971]. More recently, program synthesis

\* Both authors contributed equally to this research.

† Corresponding author.

Authors' Contact Information: Anastasiya Grishina, [anastasiia@simula.no](mailto:anastasiia@simula.no), Simula, Oslo, Norway and University of Oslo, Oslo, Norway; Vadim Liventsev, [v.liventsev@tue.nl](mailto:v.liventsev@tue.nl), TU Eindhoven, Eindhoven, The Netherlands and Philips Research, Eindhoven, The Netherlands; Aki Härmä, [aki.harma@philips.com](mailto:aki.harma@philips.com), Philips Research, Eindhoven, The Netherlands; Leon Moonen, [leon.moonen@computer.org](mailto:leon.moonen@computer.org), Simula, Oslo, Norway.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2025 Copyright held by the owner/author(s).

ACM 2688-3007/2025/2-ART

<https://doi.org/10.1145/3719351>

has emerged as an interpretable alternative to black-box machine learning methods that lets human experts understand, validate, and edit the algorithms generated by artificial intelligence [Bastani et al. 2022]. In addition to the scientific benefits of such knowledge, it extends the benefits of machine learning to domains, such as embedded systems, where it is technically challenging [Dhar et al. 2021], or healthcare, where it is avoided for safety reasons [Connolly et al. 2023; Jia et al. 2022].

The predominant methodology in automatic programming has shifted from deductive programming [Alur et al. 2015; Manna and Waldinger 1992] to genetic and evolutionary methods [Ahvanooey et al. 2019] to, more recently, large autoregressive language models trained on corpora of source code [Xu et al. 2022; Zan et al. 2023] to benefit from their remarkable capability for zero-shot generalization [Chen et al. 2021b]. However, even state-of-the-art models fine-tuned on a specific class of programming tasks still require costly filtering to discard the outputs that fail to compile or pass tests [Li et al. 2022]. These outputs tend to be superficially similar to the correct solutions [Liu et al. 2023b; Ren et al. 2020; Shirafuji et al. 2023] despite failing to produce the expected output, an issue known as “near-miss syndrome” or “last mile problem” [Bavishi et al. 2022].

Given these challenges, research in machine learning on source code [Allamanis et al. 2018] tends to focus on restricted domain-specific languages [Chen et al. 2021a; Liventsev et al. 2021; Polozov and Gulwani 2015] or automating specific parts<sup>1</sup> of the software development process [Lu et al. 2021; Niu et al. 2023], such as code search [Di Grazia and Pradel 2023; Kim et al. 2023], code translation [Yan et al. 2023], detection of issues [Chakraborty et al. 2022; Fernandes et al. 2016], improvement [Petke et al. 2018] and repair [Le Goues et al. 2019] rather than fully autonomous programming in a programming language popular with human developers.<sup>2</sup> However, two recent innovations potentially make the latter task tractable.

One is *Synthesize, Execute, Debug* [Gupta et al. 2020], a framework that attempts to bridge the “last mile” gap by introducing program repair into the program synthesis algorithm. A programming task is specified using both a natural language description and a set of input/output (I/O) pairs that demonstrate what output is expected from the program, thereby combining text-to-code [Iyer et al. 2018] and programming by example [Gulwani 2016; Halbert 1984] paradigms typical for competitive programming [Zavershynskiy et al. 2018]. *Synthesize, Execute, Debug* creates a first draft program using a generative model, compiles and executes it with given input examples. This is followed by a program repair step to fix the identified errors.

Another relevant innovation is instruction fine-tuned large language models [Zhang et al. 2024a]. Instruction fine-tuned models use human feedback in their training process and are designed to explicitly or implicitly admit two inputs: a source text (or code) and a textual command instructing the model to edit the source in a particular way, e.g., “summarize” or “translate to Python.” These models have been shown to be highly successful in automatic program repair [Fan et al. 2023]. However, given the free-form nature of these instructions,<sup>3</sup> how one should engineer instructions that maximize repair performance is an open question.

These innovations have led to the proposal of a framework, *Synthesize, Execute, Instruct, Debug and Rank*, or SEIDR,<sup>4</sup> initially introduced in our GECCO-2023 paper [Liventsev et al. 2023]. SEIDR is a multi-agent iterative framework that uses feedback from code execution and failing test cases to update the initially generated buggy code. Our initial explorations cover SEIDR with GPT-3 as

<sup>1</sup> Similarly to autonomous driving [Grigorescu et al. 2020; Marcano et al. 2020].

<sup>2</sup> <https://www.tiobe.com/tiobe-index/>

<sup>3</sup> Throughout this paper we avoid other definitions of *instruction*, such as *an individual operation in code*, to prevent ambiguity.

<sup>4</sup> Seiðr also refers to a type of Norse magic [Blain 2002] pertaining to predicting and controlling the future, which we deem thematically appropriate.

the bug summarization model and Codex (or GPT-3 trained on code) as the program generation and debugging model. In addition, by modifying the tree arity parameter (see Section 2.2), we investigate the trade-off between generating and repairing only one program versus regenerating any program that does not pass all the tests, as well as intermediate configurations, where we build a tree of programs and update the best ones.

While Codex is an early code generation model, the emergence of new models that score better in programming and natural languages motivates further research into the use of SEIDR with newer models. To study the generalizability of the initial SEIDR results, we use two other LLMs, Llama 3<sup>5</sup> and GPT-3.5,<sup>6</sup> and an additional dataset, HumanEval-X, with different tree arity parameters [Brown et al. 2020; Chen et al. 2021b; Zheng et al. 2023]. Moreover, we build up on the initial experiments with Codex and zoom in on the area with the best-performing tree arities in a hyperparameter search for a better repair-replace trade-off resolution.

To reflect on the parent selection strategies used in the Rank agent of SEIDR, we also explore whether the programs should be chosen based on the average performance across all tests or whether SEIDR can benefit from keeping such programs in the loop that fully cover individual tests, but do not perform well on average. Therefore, as an alternative to the tournament selection, we test the best tree arity setups with lexicase selection-based ranking [Helmuth et al. 2015]. Moreover, since language models bring in stochasticity, we run the experiments several times to measure the variability of results obtained with fixed hyperparameters and reflect on the results repeatability.

Overall, the current paper contributes to the field of Large Language Models for Software Engineering (LLM4SE) with a framework for code generation and its iterative repair. Compared to the earlier GECCO-2023 paper, this article refines and extends the preliminary results by providing a more in-depth explanation of SEIDR as a multi-agent framework for autonomous programming, extending the original experiments and analysis to include HumanEval-X as an additional benchmark, evaluating GPT-3.5 and Llama 3 as additional LLMs, investigating lexicase selection as an alternative to tournament selection, and conducting a repeatability analysis of the framework’s performance over multiple independent runs. Overall, this extension adds three research questions and doubles the number of configurations considered in these questions compared to our preliminary work.

Section 2 presents a framework that adapts *Synthesize*, *Execute*, *Debug* to instruction-tuned Large Language Models in agents that can solve programming tasks in an autonomous fashion. We discuss related work in Section 3, introduce experiments to explore effective search strategies for this framework in Section 4. Finally, in Section 5, we demonstrate that our framework outperforms conventional automatic programming techniques, such as genetic programming and naive application of large language models that generate one solution per problem without updating it iteratively.

## 2 Methodology

The proposed five-agent SEIDR framework is summarized in Figure 1, which we discuss in detail in Section 2.1. In essence, to solve a programming task defined as a text description and a collection of I/O examples, we split I/O examples into prompt and validation sets and use the prompt set in a large language model to SYNTHESIZE a population of candidate solutions. We EXECUTE the solutions, test them against the validation set, generate a text description of the identified problems used to INSTRUCT a large language model to produce repaired candidate solutions similar to the way a human developer DEBUGs a program. We RANK the candidates by correctness, measured as matching I/O pairs, discard the worst candidates, and repeat until a fully correct solution is found.

<sup>5</sup> <https://ai.meta.com/blog/meta-llama-3/>

<sup>6</sup> <https://openai.com/form/researcher-access-program/>

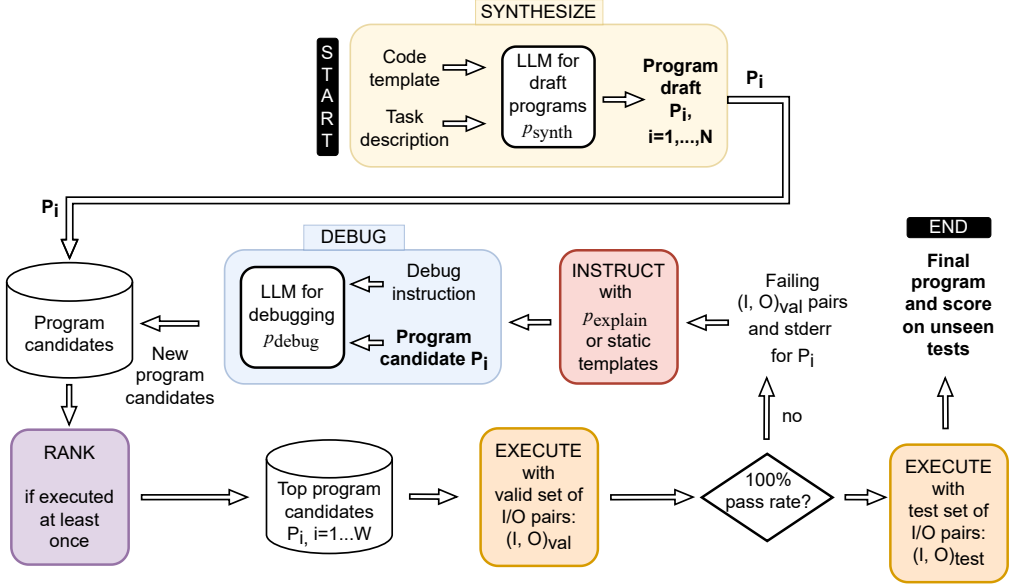


Fig. 1. Overview of SEIDR, a multi-agent iterative framework that uses LLMs to implement the Synthesize, Execute, Instruct, Debug, and Rank feedback loop.

## 2.1 Ingredients

SEIDR makes use of instruction fine-tuned large language models: a *synthesis* model  $p_{\text{synth}}$  (code, descr), a *debugging* model  $p_{\text{debug}}$  (code, descr), as well as a model  $p_{\text{explain}}$  (code, descr) that can be used for writing textual instructions, which are forwarded to the code generation model  $p_{\text{debug}}$  for code updates. Therefore, the design can be described as two agents communicating with each other, whereby one generates code and another provides critical or supervising comments on what should be changed in the generated code.

The models  $p_{\text{synth}}$ ,  $p_{\text{debug}}$ , and  $p_{\text{explain}}$  can be either separate or the same model. The prerequisites are that  $p_{\text{synth}}$  and  $p_{\text{debug}}$  models are able to “understand” natural language (descr) and partial or full programs (code) and generate code based on them. The model  $p_{\text{explain}}$  should be able to “understand” code and natural language and either autocomplete or generate the debugging instruction from scratch. Note that  $p_{\text{explain}}$  is optional, since alternatively the debugging instructions can be generated from failing tests using static pre-defined templates. In general, SEIDR requires sequence-to-sequence generative models for these agents. In our experiments, we select models for  $p_{\text{synth}}$ ,  $p_{\text{debug}}$ , and  $p_{\text{explain}}$  based on the state-of-the-art transformer architectures [Vaswani et al. 2017] (see Section 4.2).

Each LLM is a highly parameterised probability distribution over the space of (code, description)-tuples with parameters estimated on a large diverse (i.e., non-task-specific) corpus. This stochastic nature of language models is an important prerequisite for SEIDR, since it allows us to sample batches of diverse candidate solutions from  $p_{\text{synth}}$ ,  $p_{\text{debug}}$ , and  $p_{\text{explain}}$ . We denote the number of outputs generated with  $N_{\text{synth}}$ ,  $N_{\text{debug}}$ , and  $N_{\text{explain}}$ , correspondingly. Moreover, each model generates the most probable and less probable outputs in each batch, which helps diversify problem solving attempts. In the following implementation-related subsections, we explain how we vary

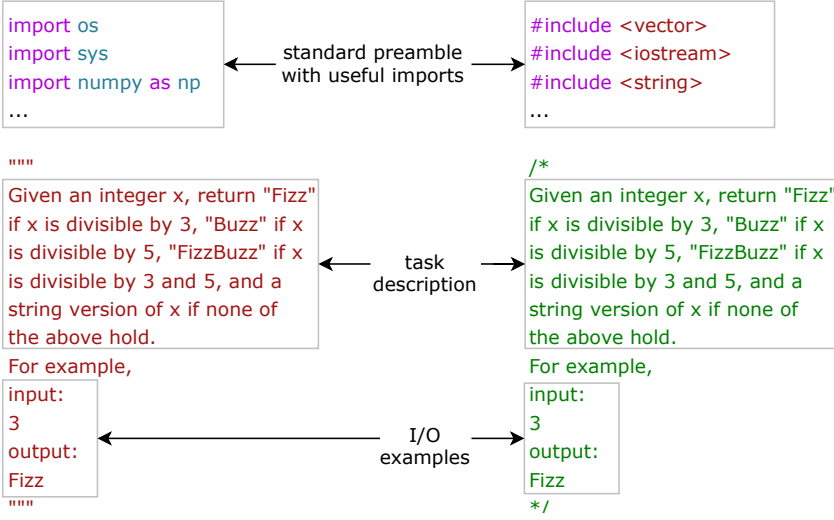


Fig. 2. Anatomy of SYNTHESIZE templates

the number of candidate solutions, debug instructions, and repairs generated in a batch by each LLM in SEIDR.

While SEIDR is described here as a multi-agent system, it can equally be seen as a form of evolutionary algorithm or genetic programming, where the initialization and mutation steps of the system are performed by LLMs,  $p_{\text{synth}}$  and  $p_{\text{debug}}$ , correspondingly. Throughout this work, we use agent-oriented programming terminology [Shoham 1993] and evolutionary optimization terminology interchangeably to try to bridge the gaps between these domains.

**2.1.1 Synthesize.** The framework starts with the SYNTHESIZE agent, which is responsible for generating initial draft solutions. We start with a basic template for a chosen programming language that contains a number of standard library imports, as shown in Figure 2.

We populate this template with a comment indicating a textual task description and several I/O examples from the training set. We design the templates with prompt engineering guidelines<sup>7</sup> and prior work [de Bruin et al. 2021] in mind. We then sample  $N_{\text{synth}}$  programs from  $p_{\text{synth}}$ , setting code to the populated template and description to the natural language description of what the model should generate. We use spring sampling:<sup>8</sup> a temperature-based sampling with a monotonically increasing temperature schedule where the  $i$ -th program is sampled with temperature  $t_i \approx \frac{i-1}{N_{\text{synth}}}$  (we use approximate equality to enable efficient implementation by means of batching). Thus, the sampling procedure for the first programs approximates a deterministic maximum-likelihood estimation. In combination with the naturalness principle of source code [Allamanis et al. 2018; Jiang et al. 2022], this approach ensures that the samples are diverse, but always contain the most likely programs for the given task.

**2.1.2 Execute.** The EXECUTE agent compiles the programs (if necessary) and launches them using the standard tools for the programming language. The program is run once for every I/O pair in the validation set. Its stdin stream receives all input lines in a given input pair, and its stdout and stderr streams are captured and saved. We then measure the *score* of the program defined as the

<sup>7</sup> <https://platform.openai.com/docs/guides/prompt-engineering/six-strategies-for-getting-better-results>

<sup>8</sup> <https://vadim.me/publications/spring/>

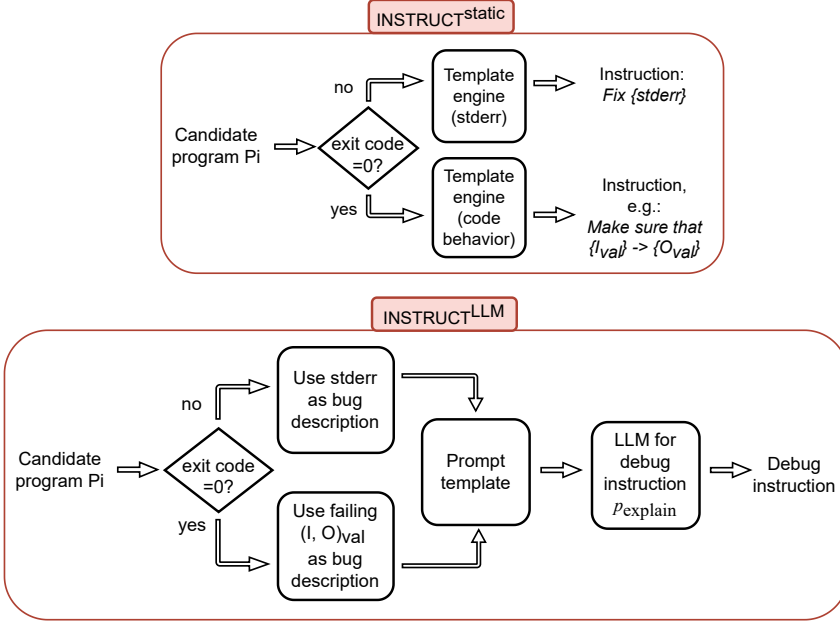


Fig. 3. Overview of the two designs for the INSTRUCT agent.

accuracy over the output lines, with  $O$  being the expected output, and  $n = \max\{|O|, |\text{stdout}|\}$ :

$$\text{score}(O, \text{stdout}) = \frac{\sum_i^n \mathbb{I}[\text{stdout}_i = O_i]}{n}$$

unless `stderr` is non-empty during compilation or execution, which is considered to indicate failure and is assigned a score of 0.

**2.1.3 Instruct.** The goal of the INSTRUCT agent is to provide instructions that summarize bugs in a candidate program and suggest a solution for  $p_{\text{debug}}$ . The resulting instructions with the bug summary should indicate what requirement is violated and instruct the LLM to edit the candidate program accordingly. The input to INSTRUCT consists of failing I/O pairs from the validation set and `stderr` output of the candidate execution. In order to represent this heterogeneous input as text that can be further processed by an LLM, we use template engines that replace placeholders in files or strings with input values and return a formatted string. In recent chat- and instruction-based LLMs, the terms *template engine* and *prompt template* are used interchangeably.

We consider two different designs of the INSTRUCT agent:  $\text{INSTRUCT}^{\text{static}}$  and  $\text{INSTRUCT}^{\text{LLM}}$  shown in Figure 3. In both cases, if `stderr` is not empty, i.e., execution exits with code 0 before getting any output to compare it with the expected output, the `stderr`-based template engine generates the instruction to fix the error. However, the designs differ in the way they transform failing I/O pairs to generate instructions in case `stderr` is empty.

$\text{INSTRUCT}^{\text{static}}$  uses a fixed template and substitutes placeholders for input and output with the corresponding strings of the first failing test case in its template engine. For example, we show the resulting instruction for an exemplar template in Figure 3. In contrast,  $\text{INSTRUCT}^{\text{LLM}}$  uses the failing I/O pair in the LLM for text completion, thereby prompting the text LLM to produce the bug explanation and a summary for debugging. In addition to providing a failing test case or `stderr`,

one may choose to give the model  $p_{\text{explain}}$  more context, such as the problem name, task description, and the code generated so far. Each call to  $p_{\text{explain}}$  can result in  $N_{\text{explain}} \geq 1$  instructions, a batch output of an LLM. The prompt templates used for the experiments are detailed in Section 4.3.

**2.1.4 Debug.** The main component of SEIDR that addresses the “near-miss syndrome” is the DEBUG agent. This agent iterates over all programs in the population to repair candidate programs and pass more tests. It uses the instructions written by INSTRUCT to sample from the  $p_{\text{debug}}$  model  $N_{\text{debug}}$  times to repair each candidate and create a new population of  $N$  candidates. For  $p_{\text{debug}}$ , the parameter code is set to the current version of the candidate solution and descr to the output of INSTRUCT and any additional context chosen for a specific implementation. The current generation of candidates is then replaced by  $N$  outputs of DEBUG.

**2.1.5 Rank.** The RANK agent implements what is known in genetic programming as *parent selection* [Koza 1994]: it selects the best  $W$  programs to be further improved by the DEBUG agent. We consider two different parent selection algorithms: tournament selection and lexicase selection. See Section 5.2.2 for their empirical comparison.

*Tournament selection* variant used in this work sorts the programs according to their average test score and selects top  $W$  candidates. The programs are selected based on the intuition that repair of the best so far (yet imperfect) programs begets good programs. The simple ranking is also referred to as *tournament selection*, where the best-performing candidates are chosen to participate in the next round. In the classical tournament selection, the first step is to draw  $n$  random candidates from the population. In our implementation, the first step is to generate exactly  $n$  candidates that are needed for the next round. This approach prioritizes candidates that perform the best on average over all tests.

*Lexicase selection* [Helmuth et al. 2015] is a ranking approach that maximizes the diversity of selected candidates in addition to their metric-based score. Lexicase selection ensures diversity by keeping the program candidates that perform the best on unique tests as opposed to the program candidates that perform best on average over all tests. The algorithm is as follows:

- (1) randomly shuffle the set of tests;
- (2) select a program with the best score on test 1;
- (3) if several programs are tied, resolve the tie by selecting the best program on test 2;
- (4) repeat for tests 3, 4,  $\dots$ , until only one program is left;
- (5) mark this program as “selected”;
- (6) if less than  $W$  programs are selected, go back to step 1.

This ensures that even if the average quality of the selected candidates is lower, the batch of  $W$  programs collectively contains a higher number of required “skills”, as measured by tests.

## 2.2 Meaning of Hyperparameters

After evaluating a given candidate solution in EXECUTE, SEIDR supports two approaches to address the candidate’s flaws:

- *Replace* the candidate with another sample from the current population.
- Use INSTRUCT and DEBUG to repair the candidate.

We refer to this problem as the *repair-replace trade-off*, by analogy with production economics [Jack and Van der Duyn Schouten 2000].

How does the choice of hyperparameters  $N$ , the total number of candidate programs in each generation, and  $W$ , the number of selected repairs to be preserved in a generation, influence the flow of SEIDR?  $N$  and  $W$  act as upper bounds on the *replace* option by limiting the size of the population. In the edge cases,  $N = W = 1$  corresponds to a repair-only process, while  $N = W = \infty$

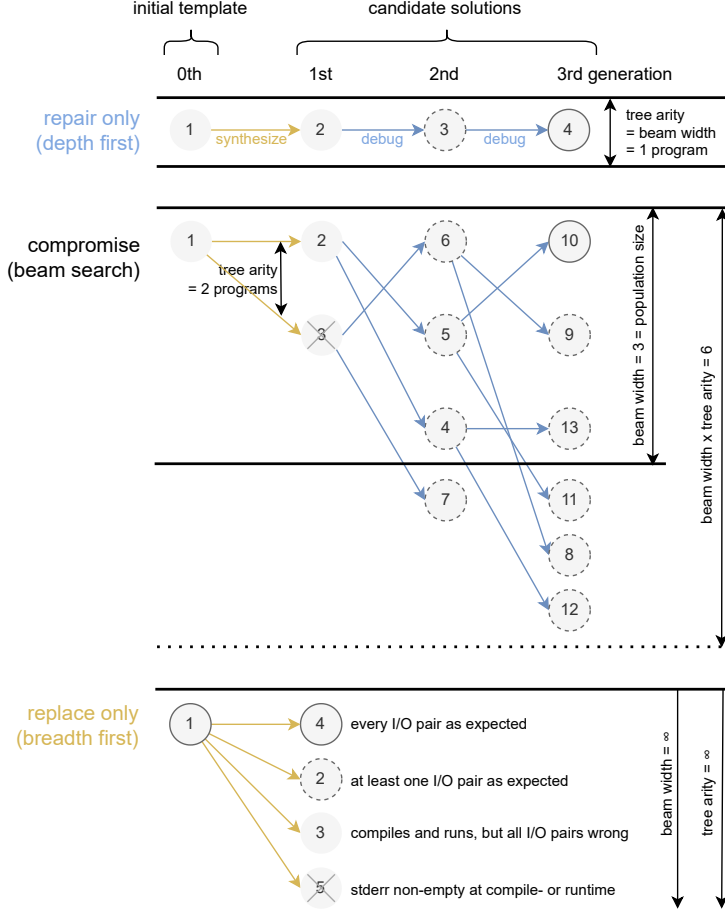


Fig. 4. Repair-replace trade-off as a tree search problem.

corresponds to replace-only, as illustrated in Figure 4. Here, the repair-only scenario can also be seen as an LLM-guided random walk [Xia et al. 2020] and replace-only as random sampling from the LLM. Strategies with tree arities between 1 and  $\infty$  are similar to population-based evolutionary algorithms. Note that  $N$  is defined by  $N_{\text{synth}}$  for the initial draft solutions in the first generation and  $N_{\text{explain}} \cdot N_{\text{debug}}$  for later generations.

Observe that a mutation-only genetic algorithm with tournament selection with fixed population size  $W$ , such as SEIDR, is equivalent to *local beam search* with beam width  $W$  on an  $N$ -ary tree [Russell 2010, Section 4.1.4]. This corresponds to a known property of local beam search: it degenerates into a depth-first search at  $W = 1$ , whereas setting  $W = \infty$  yields a breadth-first search.

### 3 Related Work

Until recently, the tasks of neural program synthesis [Gulwani et al. 2017] and program repair [Le Goues et al. 2012, 2019; Petke et al. 2018; Zhang et al. 2024b] have been considered separately. However, results from genetic programming [Sobania et al. 2021] suggest that evolution is a crucial step in synthesis. A number of important studies bridging this gap in the application of large language models have been carried out concurrently with this work, discussed below.



The use of large language models for a program repair step within a program synthesis pipeline has been studied by Joshi et al. [2022] and Gupta et al. [2020], while the specific case of instruction-driven LLMs has been explored by Fan et al. [2023], where the initial synthesis is done by the Codex model [Chen et al. 2021b] and for program repair both Codex and state-of-the-art specialized repair tools such as TBar and Recoder [Just et al. 2014] are considered and compared. Zhang et al. [2023b] do the same, but fine-tune PyCodeGPT-110M [Zan et al. 2022] to use it as a repair model. The resulting framework is a two-step process (1 draft step and 1 debug step), while iterative evolution and search are not explored.

Evolution through Large Models (ELM) [Lehman et al. 2022] proposes to use a language model in place of a mutation operator within a traditional genetic programming framework [Koza 1994]. They use a type of instruction fine-tuned model trained on git commit messages known as a diff model.<sup>9</sup> However, the model is directed neither to solve the programming problem at hand nor to fix bugs. Instead, it is provided with generic instructions such as “Change function f.” This approach is meant for the cases where creative freedom [Stanley and Lehman 2015] is encouraged rather than satisfying concrete requirements. Liu et al. [2023a] demonstrate the advantages of applying ELM to modular components of a solution, rather than the entire solution to the given problem.

Some related work is explicitly guided by the metaphor of conversation between agents. Zhang et al. [2023a] implements a repair-only program synthesis loop as a conversation between an analyst agent, a coder, and a tester, while [Dong et al. 2024] does the same for tester, developer, and reviewer. The coder and the developer are roughly equivalent to SYNTHESIZE, the tester to EXECUTE, and the reviewer to INSTRUCT, while the analyst is an agent that prepares a generation prompt for the coder. Both testers are language models that predict the output of a program without actual compilation, execution, or testing, which makes Dong et al. [2024] and Zhang et al. [2023a] specific cases of chain-of-thought prompting [Yu et al. 2023] for program synthesis.

Several studies explore an iterative approach to program synthesis in a manner similar to SEIDR [Chen et al. 2023; Shinn et al. 2023; Xia and Zhang 2023]. However, they do not explore the repair-replace trade-off and exclusively implement the repair-only approach that is prone to local minima. SelfEvolve [Jiang et al. 2023] is a repair-only version of a framework similar to SEIDR. SelfEvolve demonstrates the benefits of LLMs evolving not just the source code but an additional natural language text file that acts as the system’s knowledge base and is included in the model prompt when generating code. Finally, Self-Taught Optimizer (STOP) [Zelikman et al. 2023] takes the concept of self-improvement to the meta level and uses a large language model to edit the evolutionary algorithm itself (i.e., blocks or contents of Figure 1). Their reflections on the safety implications of such automated algorithm changes are of particular interest when considering this trajectory [Zelikman et al. 2023, Section 8]. These concerns do not hold in the context of SEIDR because the algorithm is fixed. In other words, SEIDR does not self-evolve but creates solutions that it improves, and the solutions that are synthesized are closely constrained by the validation set used by the EXECUTE agent.

These efforts to combine large language models and evolutionary techniques fall within the broader context of the quest for effective inductive bias in genetic programming [Whigham 1996]: a crucial property for any learning algorithm [Haussler 1988]. Reuter et al. [2023] suggest using graph neural networks for this purpose, while grammatical evolution methods, such as WHGE [Bartoli et al. 2020] and  $\pi$ -GE [O’Neill et al. 2004] use the known grammar of the programming language. Large language models are a novel and promising [Custode et al. 2024] alternative to grammatical evolution that incorporates semantics and idiom [Allamanis and Sutton 2014; Orlov 2020] of a programming language in addition to its grammar.

<sup>9</sup> <https://carper.ai/diff-models-a-new-way-to-edit-code/>

## 4 Experimental Design

To explore the capabilities of SEIDR and its generalizability, we test the framework on two benchmarks (PSB2 and HumanEval-X), a total of two different ranking strategies (tournament selection and lexicase selection), three models in the coding part of SEIDR (Codex, GPT-3.5, and Llama 3), two programming languages (Python and C++), and various branching factors. We use three models over the two parts of our experiments: one model in the initial exploration (also reported by [Liventsev et al. 2023]) and two models in the generalizability part. The problems in the benchmarks originate from coding competitions and human-written programming assignments.

During our empirical evaluation of SEIDR, we address the following research questions:

**RQ1. Repair-replace trade-off exploration:** What is the impact of using different tree search strategies in the autonomous programming setting? We experiment with six different tree arities but fix the tournament selection in the ranking part and one prompt. Here, we study the impact of tree arity on the number of resolved problems as well as the speed of obtaining solutions.

**RQ2. Generalizability of the approach to different LLMs and an additional dataset:** How does the choice of an LLM affect the performance of SEIDR? We vary the tree arity and experiment with two additional LLMs and one additional dataset. By default, we use tournament selection as the ranking strategy.

**RQ3. Repeatability of SEIDR in multiple runs with the same hyperparameters:** How does the non-deterministic nature of LLMs affect SEIDR performance when the method is restarted several times with the same hyperparameters? We study how SEIDR results vary in different restarts of the same experiments with unchanged hyperparameters as a result of the LLM non-determinism. The motivation for RQ3 is that LLMs exhibit stochastic behavior: the same prompt can yield different responses. Essentially, LLMs generate answers token-by-token and predict the next tokens based on the probability distribution over a vocabulary of tokens, which is sensitive to the precision of floating-point operations. If two tokens are predicted to be the next ones, with a very similar probability, it is likely that either of them will be chosen at each individual run. Further tokens are generated auto-regressively and depend on the previous tokens, so once one token diverges, the whole sequence is likely to diverge, too.

**RQ4. Effect of changing the parent selection strategy in the RANK agent:** How does the lexicase selection-based ranking strategy impact performance in comparison to tournament selection in the RANK agent? We use the best-performing tree arities from RQ3 to run the experiments with lexicase selection as the ranking strategy instead of tournament selection, to explore whether a different parent selection algorithm can further improve the results.

### 4.1 Data

Our experiments use the Program Synthesis Benchmark 2 (PSB2) [Helmuth and Kelly 2022] and HumanEval-X [Zheng et al. 2023] in C++ and Python. The key criteria for this choice are the availability of task descriptions in English and unit tests in Python and C++ or language-agnostic unit tests as well as the wide acceptance of these benchmarks in the areas of generative LLMs (HumanEval) and genetic programming (PSB2).

**4.1.1 PSB2.** The first dataset is a benchmark suite of 25 problems for program synthesis that resemble small real-world tasks. PSB2 was developed as a more realistic and challenging version of PSB1 [Helmuth and Spector 2015], the latter consisting of textbook problems and is widely used in genetic programming [Sobania et al. 2022]. The problems require different data structures and control flows to be used for effective solutions and are taken from sources, such as competitive programming platforms and educational courses. The problems have descriptions in English, as well as 1 million (M) tests for training and 1M testing-stage tests, including edge or corner cases

that test the resulting program on complicated inputs. The tests are provided as I/O pairs and are distributed together with the problem descriptions as a PyPI package.<sup>10</sup>

In PSB1, the training set consists of the edge test cases and is augmented by random test cases if the number of edge tests is not enough. The test set is formed by random test cases. This terminology is preserved in PSB2. We use the PSB2 training set for ranking and selection of programs (validation) within an experiment and the test set for reporting the result thereof (testing). Thus, we will refer to the PSB2 training set as the *validation set*, to be more consistent with how it is used in SEIDR.

**4.1.2 HumanEval-X.** The second dataset that we use is a development from the original set of human-written programming tasks in HumanEval [Chen et al. 2021b], which is a standard code generation benchmark for LLMs. HumanEval consists of 164 problems with a docstring representing problem description, a function signature, a correct solution, and unit tests in Python. HumanEval-X is the result of translating correct HumanEval programs and unit tests into five programming languages [Zheng et al. 2023]. We use HumanEval-Python for experiments in Python to ensure a comparison with other models in the setup without SEIDR. In addition, we test SEIDR on the HumanEval-C++ part of HumanEval-X.

The test functions of HumanEval-X contain all tests in one function. We split the aggregated test functions into separate tests so that the RANK agent can evaluate the score. On average, the number of tests in HumanEval-Python is 7.25 and 6.95 in HumanEval-C++, which is appointed to a repeated additional test present in some HumanEval-Python examples of the following type: `assert True, "This prints if this assert fails 1 (good for debugging!)"`. This test type is not present in HumanEval-C++.

To reiterate, we keep the original HumanEval-Python setup for direct comparison with the models tested on this benchmark without SEIDR. Because of the limited number of tests, we pass up to five tests to the draft prompt and make all tests visible to SEIDR for the debugging loop. In other words, we do not have a held-out test split for HumanEval-X in the same manner as we do for PSB2.

## 4.2 Models

SEIDR uses up to three LLMs —  $p_{\text{synth}}$ ,  $p_{\text{explain}}$ , and  $p_{\text{debug}}$  — in SYNTHESIZE, INSTRUCT<sup>LLM</sup>, and DEBUG, respectively. The main prerequisite is that  $p_{\text{synth}}$  and  $p_{\text{debug}}$  are a text-to-code models which take both a textual description and a draft code as input. Therefore,  $p_{\text{synth}}$  and  $p_{\text{debug}}$  can be a chat model, a code completion model, an instruction fine-tuned, or a foundation generative language model pre-trained on code in addition to text. By analogy, the text-to-text  $p_{\text{explain}}$  can be a chat, instruction fine-tuned, a text completion, or a text generation model pre-trained on text and code.

In our experiments, we use Open AI Generative Pre-trained Transformer (GPT) models by Open AI and Llama 3 by Meta [Rozière et al. 2023]. GPT models are auto-regressive transformer models that have the decoder-only architecture as opposed to the original full encoder-decoder transformer. They are pre-trained on both text and code and excel at sequence-to-sequence generative tasks, including code-to-code, text-to-code, and code-to-text.

In our initial experiments, we use Codex-edit<sup>11</sup> as the LLM for writing and debugging programs and GPT-3<sup>12</sup> for bug summarization via text completion [Brown et al. 2020] — both being 175B-parameter models. In our generalizability experiments, we use the GPT-3.5 model<sup>13</sup> for program

<sup>10</sup> <https://pypi.org/project/psb2/>

<sup>11</sup> code-davinci-edit-001

<sup>12</sup> text-davinci-003

<sup>13</sup> gpt-3.5-turbo

synthesis, bug summarization, and debugging. The GPT-3.5 model is an improvement over GPT-3 that is optimized for chat and available through an API. This switch is mainly motivated by rapid model updates, an OpenAI announcement that GPT-3 was due to become obsolete, i.e., not actively supported by the company, along with the company’s recommendation to switch to GPT-3.5.

To further evaluate the generalization capabilities of SEIDR, we have chosen Llama 3-8B [Rozière et al. 2023], an open-source alternative to GPT models with the same standard decoder-only transformer architecture. Compared to Llama 2 [Touvron et al. 2023], Llama 3 introduces improvements to the architecture and the training process, such as grouped query attention.

### 4.3 Prompts

**4.3.1 Prompts in the Initial Exploration of SEIDR.** The prompt for the LLM model is static and consists of the input for editing — candidate program generated so far — and a debug instruction to repair the candidate. The debug instructions are formulated as templates. The instructions describe the violated requirements in terms of the wrong output in a failing I/O test or summarize the bug to capture issues in code logic. We present debug instructions using the template engine format: the brackets  $\{ \}$  denote that the placeholder in the brackets will be replaced with the value generated during execution,  $\{I_{val}\}$  and  $\{O_{val}\}$  stand for values of the validation set I/O pair. As shown in Figure 3, the instruction to fix execution errors that abort the program before the resulting output is obtained with `stderr` lines: Fix  $\{stderr\}$ . Debug instruction that uses the output of candidate program execution is static, formulated as follows:

Make sure that  $I_{val} \rightarrow O_{val}$ . (S0)

**4.3.2 Prompts for Instruction Fine-tuned and Chat Models.** The prompts presented in this section are used in the experiments with GPT-3.5 and Llama 3. This implementation is optimized for instruction and chat models, which use prompts as inputs represented as text, partially with code fragments. The models use a system message that describes the “role” of the LLM and a regular message that works as an instruction or a chat message from a user. To provide more context, we always use a problem description, problem name, and programming language to the model as textual input (`descr`). As code, we also add an initial template depicted in Figure 2 to  $p_{synth}$  and the current program candidate to  $p_{explain}$  and  $p_{debug}$ . The resulting prompts are as follows:

The system message is:

---

```
You are an experienced software developer.
You write concise code in {language}.
The code must read input from user and return output corresponding to the task description.
```

---

The input to  $p_{synth}$  looks as follows:

---

```
Solve the following code contest problem: {problem_name}.
Problem description: {problem_description}.
{program_template}
Only complete the code, do not add triple quotes, do not give explanations.
```

---

Bug explanations are generated with  $p_{explain}$  using the following instructions:

Table 1. Initial exploration of SEIDR: hyperparameters in the tree arity experiments.

experiment #	1	2	3	4
population size (beam width), $W$	1	10	100	$\infty$ (1000)
tree arity, $N$	1	10	100	$\infty$ (1000)
max programs generated	1000			
prompt	S0			
models	Codex as $p_{\text{synth}}$ and $p_{\text{debug}}$			
# restarts (or runs)	1			
with the same hyperparameters				
datasets	PSB2			
languages	Python, C++			

---

I'm trying to solve the following code contest problem: {problem\_name}.

Problem description: {problem\_description}.

Currently, the code is

```
...
```

```
{program_candidate}
```

```
...
```

The issue is

{stderr} or "it must return {expected\_output} for input {input}, but it returns {output}".

Describe how I should fix the code in a very concise manner.

---

And the debugging model  $p_{\text{debug}}$  operates on the following instruction:

---

Solve the following code contest problem: {problem\_name}.

Problem description: {problem\_description}.

Currently, the code is

```
...
```

```
{program_candidate}
```

```
...
```

Modify the code as {bug\_summary}.

You must only return correct code.

Remove any triple quotes, language name or explanations.

---

#### 4.4 Repair-replace Trade-off Settings

The settings for tree arity will also be divided into two experiment sets: the ones for GPT-3 and Codex, and the ones for the GPT-3.5 and Llama 3 experiments.

**4.4.1 Tree Arity for the Initial Exploration of SEIDR.** As described in Section 2.2, the population size (number of parents to choose from in the tournament selection or beam width from the beam search perspective)  $W$  and tree arity  $N$  define the repair-replace trade-off, where higher  $W$  and  $N$  correspond to repair over replace. We evaluate four options for these hyperparameters as shown in Table 1. We only run the experiments once, due to the experimental timeline and the discontinuation of model support by the model provider.

Because we aim to compare tree search parameters, we fix one default debugging instruction S0 and use the INSTRUCT<sup>static</sup> agent. Moreover, we set the upper limit for the total number of generated program candidates to 1000 to limit the experimentation time. Although some solutions may not be found within the hard limit, we assume<sup>14</sup> that 1000 program candidates form a sufficiently large search space for our experiments.  $W = N = \infty$  is achieved in implementation by setting equal  $W$  and  $N$  equal to the upper limit of the program count of 1000. This ensures that a second generation of programs does not exist.

**4.4.2 Tree Arity for SEIDR Generalizability Experiments.** With the shift to chat and instruction models in the generalizability part of our study, we move from generating one bug explanation and one code draft or update to a batch of those. Specifically, each of the three LLMs in SYNTHESIZE, INSTRUCT, and DEBUG can generate sequences in batches. We generate  $N_{\text{synth}}$  programs in the first generation with  $p_{\text{synth}}$  model,  $N_{\text{explain}}$  bug explanations with  $p_{\text{explain}}$  for each program in a generation, and  $N_{\text{debug}}$  candidate repairs for each of the debugging instructions using  $p_{\text{debug}}$ . A new generation of  $N_{\text{explain}} \cdot N_{\text{debug}} \cdot W$  programs created from each of  $W$  parents in a previous generation is ranked and filtered to keep the best-performing  $W$  candidates for generating the next candidates.

To balance between a reasonable number of experiments and diverse sets of hyperparameters, we fix  $N_{\text{explain}} = 2$  to moderately vary the bug descriptions and set  $N_{\text{synth}} = N_{\text{debug}} = N$ . As a reference, in the experiments with GPT-3 and Codex, we generated only one bug explanation ( $N_{\text{explain}} = 1$ ) and used  $N_{\text{synth}} = N_{\text{debug}} = N$  setting, too. We evaluate six options of  $N$  as shown in Table 2 and use tournament selection as the ranking strategy.

The choice of these tree branching hyperparameters and the maximum number of generated programs is motivated by the experiments with GPT-3 and Codex, where the best results were obtained for  $N = 10$ . Therefore, we explore the area around this value more closely in the generalizability experiments. In the same experiments, the majority of problems in PSB2 were solved within the first 100 generated programs. Therefore, the upper limit for the total number of generated program candidates is set here to 100 to limit the experimentation time.

To account for the stochasticity of language models and the fact that OpenAI’s models do not support setting the effective sampling temperature to zero to force deterministic behavior,<sup>15</sup> we ran the experiments six times with each set of hyperparameters. This number of runs was selected to hit a sweet spot between the overall running time and costs of the experiments, while at the same time achieving confidence in the stability of the results in the presence of non-determinism. Note that reporting results over six runs is considerably better than the common practice of having only one run for every selection of hyperparameters, and it is in line with the best-of-class practice in the field of LLMs for code generation [Ouyang et al. 2023].

The total cost of running the experiments with GPT-3 and Codex were around 550 USD, and the experiments with GPT-3.5 amounted to 266 USD.<sup>16</sup> Moreover, the time to finish one run with several branching factors and all the tests amounts to ca. 42h for PSB2<sup>17</sup> and ca. 156h for HumanEval-X. Overall, we restart experiments with GPT-3.5 six times for each of the tree arities  $N_{\text{synth}} = N_{\text{debug}} = N^* \in \{1, 2, 4, 10, 16, 100\}$  and the same for Llama 3, with a total of  $6 \times 6 \times 2 = 72$  experiments. For the lexicase selection experiments, we have 6 runs per model but one best-performing tree arity, which adds 12 experiments to the total count.

<sup>14</sup> This assumption is later confirmed in Section 5.1.1.

<sup>15</sup> <https://community.openai.com/t/observing-discrepancy-in-completions-with-temperature-0/73380>

<sup>16</sup> For comparison, one run with GPT-4o cost us 315 USD (early July 2024), so further use of this model was discarded.

<sup>17</sup> Due to the local setup for testing, API call limits, and the number of tests.

Table 2. SEIDR generalizability experiments: hyperparameters in the tree arity grid search.

experiment #	16	17	18	19	20	21
population size (beam width), $W$	1	4	8	10	16	$\infty$ (100)
# programs in the 1st generation, $N_{\text{synth}}$	1	4	8	10	16	$\infty$ (100)
# bug explanations for candidate, $N_{\text{explain}}$	2	2	2	2	2	-
# repairs for each explanation, $N_{\text{debug}}$	1	4	8	10	16	-
max programs generated	100					
prompts	see Section 4.3.2					
models	(a) GPT-3.5 as $p_{\text{synth}}$ , $p_{\text{debug}}$ , $p_{\text{explain}}$ , (b) Llama 3 as $p_{\text{synth}}$ , $p_{\text{debug}}$ , $p_{\text{explain}}$					
# runs per experiment	6					
datasets	PSB2, HumanEval-X					
languages	Python, C++					

#### 4.5 Performance Indicators

In our experiments, we compare the number of fully solved programs obtained with SEIDR with different values of hyperparameters. For a more detailed analysis of results, we use *test pass rate* (TPR) and *Excess Programs Generated* (EPG). TPR reflects the percentage of fully passed test cases based on the exact match of program output and test output. The TPR metric is used for the final evaluation of generated programs and does not reflect partial passing of the I/O test as opposed to the *score* as calculated by the RANK agent (see Section 2.1.2).

We define *pass@k* as the number of problems that have  $TPR = 1$  if SEIDR is stopped after generating  $k$  programs, following Kulal et al. [2019] “*success rate at budget of  $k$  programs*.” Note that Jiang et al. [2023] and Chen et al. [2023] define  $k$  as the number of restarts of the iterative method, the budget in terms of trees of programs. We choose against this approach, since it threatens the validity of the comparison between iterative tree-based program synthesis and repair-only baseline by giving the iterative approach additional budget in terms of the number of programs it can generate.

Codex [Chen et al. 2021b] calculates *pass@n>k* and constructs an unbiased estimator of *pass@k* with lower variance, ensuring statistically robust results. We cannot apply this adjustment for SEIDR, since the adjustment assumes that the programs are independent and identically distributed, while SEIDR is a Markov chain with dependencies between iterations.

EPG reflects the number of programs generated before the first occurrence of the program that passes all validation test cases. EPG is indicative of the computational cost of solving a problem distributed in terms of LLM inferences and program compilations and executions. For a single execution of SEIDR, EPG is equivalent to the smallest  $k$  at which *pass@k*=1.

#### 4.6 Implementation Details

To summarize the setup, in this study, we have two groups of experiments. The first group of experiments is dedicated to the initial exploration of RQ1 with Codex-edit (code-davinci-edit-001) as the LLM for writing and debugging programs. We have referred to these experiments as *Initial*

*Exploration of SEIDR* with Codex and GPT-3, and test the hyperparameter choices only on PSB2 as detailed in Table 1. Here, we use wide steps between tree arity values (see Section 4.4.1).

The second set of experiments mainly focuses on the generalizability (RQ2) of SEIDR and its robustness to restarting experiments with the same hyperparameters (RQ3). The motivation here is to potentially improve on GPT-3 with a newer, generally more powerful version, GPT-3.5, and its smaller open-source competitor, Llama 3. GPT-3.5 and Llama 3 are used in more fine-grained repair-replace trade-off exploration and ranking experiments (RQ1). We have referred to these experiments as *SEIDR Generalizability Experiments* and test SEIDR both on PSB2 and HumanEval. Building on the findings of the initial exploration, we use more fine-grained tree arity values (see Section 4.4.2, Table 2) and use the prompts from Section 4.3.2. Thus, INSTRUCT is represented by the INSTRUCT<sup>LLM</sup> agent and creates  $N_{\text{debug}}$  bug summaries. Each program update creates  $N$  child programs from one parent with the SYNTHESIZE and DEBUG agents. We also compare the performance of SEIDR with the current state-of-the-art without SEIDR (see Section 5.2.1).

The second set of experiments is further updated with an alternative ranking strategy, lexicase selection (RQ4). For each model, dataset, and language, we choose the best-performing tree arity from RQ2 and exchange the tournament selection algorithm with the lexicase selection. This selection step chooses parents for debugging updates in each generation.

In all experiments, we set the limit to generate a maximum of  $M$  program candidates during the search for the candidate that passes all validation tests. If we reach  $M$  candidates and none of them pass all validation tests, we store the test pass rate for the last generated candidate and the best test pass rate achieved throughout the search. For the first set of experiments, we set  $M = 1000$ , and for the generalizability ones, we limit  $M$  to 100, after finding out that for the majority of problems, a solution is found among the first 100 programs or not found at all.

Following Helmuth and Kelly [2021], we use 2000 I/O pairs  $((I, O)_{\text{test}}$  in Figure 1) from the test split of PSB2 to evaluate the candidate program that has passed all validation test cases  $((I, O)_{\text{val}}$  in Figure 1) during debugging. Due to repetitive calls to EXECUTE, we have to resolve the speed of testing versus precision trade-off while choosing the number of validation test pairs. We resolve the trade-off by fixing the validation set size at 100 for the initial experiments and 50 for the generalizability exploration, which has more runs with the same hyperparameters. We have run a preliminary experiment to confirm that we do not lose the final test pass rate points on 2000 tests when we decreased the validation test set size from 100 (which was used in the GECCO-2023 paper) to 50 for the generalizability exploration. Due to a small number of tests in HumanEval-X, all tests are made visible to the debugging LLM and during the validation step. To operate with the chosen LLMs in SEIDR, we use ollama<sup>18</sup> and LangChain.<sup>19</sup> To ensure that the program candidates generated from the same parent program are different from each other, we change the temperature parameter of the LLMs.

## 5 Results and Discussion

In this section, we present the results of the initial exploration, where we investigate the repair-replace trade-off in SEIDR with Codex and GPT-3 (RQ1) using the PSB2 benchmark. We then continue with generalizability (RQ2) and repeatability (RQ3) experiments with GPT-3.5 and Llama 3 on PSB2 and HumanEval. Finally, we test lexicase selection as the ranking strategy (RQ4).

<sup>18</sup> <https://ollama.ai/>

<sup>19</sup> <https://www.langchain.com/>






Fig. 5. Number of solved PSB2 problems depending on tree arity in beam search for prompt type S0.

## 5.1 Initial Exploration

**5.1.1 Repair-replace Trade-off in the Initial Exploration of SEIDR.** We compare the number of solved problems in the experiments with tree arity of 1, 10, 100, and  $\infty$  and fixed debug instruction S0 in Python and C++ in Figure 5. The results of SEIDR are compared to the baseline performance of PushGP on the PSB2 benchmark, which solves 17 out of 25 problems. Note that experiments with  $N = 1$  and  $N = \infty$  can be considered as ablation studies, where the replace option and repair option are turned off correspondingly.

The results highlight the benefit of compromise strategies with tree arity of 10 and 100 over repair-only ( $N = 1$ ) and replace-only ( $N = \infty$ ) strategies. The results show that the repair-only scheme is outperformed by other strategies. We explain the poor performance of the repair-only strategy by the fact that the search space is under-explored. Specifically, the replace scenario ensures that the LLM for the debugging component represented by Codex-edit in our experiments generates different updates of program candidates using variable temperatures. The probability of finding a better fix is higher when more alternatives are generated to update the draft program at  $N > 1$  compared to  $N = 1$ . The search strategy with  $N = 10$  yields the best results: it performs on par with PushGP for C++ and outperforms the baseline during Python program synthesis by +2 problems, resulting in a total of 19 programs that pass all test cases. The results imply that generating a moderate number of programs in parallel during the DEBUG step works better than the policies in which more updates are generated for each program (100 or 1000), or those in which only one program is updated iteratively.

We present the analogy of the solution speed for all four arities and the fixed default debug instruction in Figure 6. In detail, we show the distribution of EPG values in all experiments to explore how many candidate updates are generated before the solution is found. We zoom in to the cases with solutions found with up to the first 10 program candidates in Figure 6a and show the EPG distribution with the step of 100 candidates in Figure 6b. In addition, we break down the results into each tree arity in Figure 7, showing the EPG on a heatmap scale and the TPR as a number between 0 and 1, or the signs “+” if a problem is solved and “-” if the final TPR is 0.

Out of 100 experiments for each language, in 21–24% of runs in Python and C++, the draft program is already the solution (EPG=0). For 31–33% of the experiments, the solution is found after discarding 5 candidates. Around half of the experiments do not generate more than 100 programs. However, 5 problems are solved with more than 500 generated programs in Python and 1 problem in C++ (with  $N = 10$ ). The results imply that the first steps in updating the draft program are crucial for solving the problem. The chances of solving the problem in later stages of the search, such as after 100 programs have been generated, are low. This confirms our initial assumption in Section 4.4 that 1000 programs are sufficient.

To briefly analyze Figure 7, we observe that some problems are solved in both languages, whereas some others — only in Python. In addition, only five problems are not solved in any SEIDR




Fig. 6. Distribution of the number of generated programs during each problem-solving attempt in the experiments with different tree arities where a problem solution is found.

configuration in Python (bouncing-balls, bowling, cut-vector, dice-game and leaders) and seven in C++ (bouncing-balls, cut-vector, dice-game, leaders, shopping-list, snow-day, and vector-distance). Upon closer inspection of generated programs, we have noticed that in bouncing-balls, the programs have logical errors and differ considerably between programming languages, as well as in the majority of unsolved problems. Test cases and debug instructions in bowling frequently skewed the resulting programs to return answers to individual bowling score strings instead of writing an algorithm to calculate the score based on each next character. The latter mistake happened in other unresolved problems, such as cut-vector. Qualitative analysis has also shown that some programs failed to read input from the user and instead defined input strings within the code, which limited the program to testing only one I/O pair, although the algorithm was correct.

**Repair-replace trade-off in the initial exploration (RQ1):** SEIDR with Codex as the coding LLM outperforms the PushGP baseline on PSB2 in Python and performs on par with it in C++ experiments with tree arity of 10. Search strategies with tree arity larger than one benefit from the replace possibility of the SEIDR framework as a consequence of using variable temperature for Codex-edit. The repair component is also crucial for the framework because the replace-only search policy (with tree arity of  $\infty$ ) performs worse than the policies that alternate between replace and repair during the program update (with tree arity of 10 or 100).

## 5.2 Generalizability Experiments

In this section, we present and discuss replace-repair trade-off results obtained with GPT-3.5 and Llama 3 and the effect of switching from tournament selection to lexicase selection with the best hyperparameter settings found for the trade-off. We count the number of fully solved problems




Fig. 7. Number of excess programs generated (in color) and test pass rate (as numbers) depending on tree arity. Higher EPG values are shown in darker shades. We denote solved problems with “+” (test pass rate = 1), unsolved problems with “-” (test pass rate = 0), and show the test pass rate for partially solved problems.

(i.e., reached  $TPR = 1$ ) in experiments with the hyperparameter settings described in Table 2 in six runs and present the language-specific results for PSB2 and HumanEval. We also explore the total number of programs that need to be generated before a solution is obtained, as well as in how many of the six runs each problem is fully solved. The figures are described in detail in the dedicated sections.

**5.2.1 Repair-replace Trade-off and Robustness to Restarts in the Generalizability Experiments.** We compare the number of solved problems in the experiments with  $N_{\text{synth}} = N_{\text{debug}}$  values of 1, 2, 4, 10, 16,  $\infty$  (100) and  $N_{\text{explain}} = 2$  in Python and C++ and tournament selection ranking strategy in Figure 8. We will refer to the equally set values of  $N_{\text{synth}}$  and  $N_{\text{debug}}$  as  $N^*$  hereafter. As before, experiments with  $N^* = 1$  and  $N^* = \infty$  correspond to ablation studies, where the replace option or the repair option is turned off. The results of SEIDR are compared to the baseline performance of PushGP on the PSB2 benchmark, which solves 17 out of 25 problems. The boxplots show the inter-quartile range between the first and third quartiles observed in six runs and the median values as horizontal lines within the boxes.

The trend present for all the datasets and models is that the results for Python are more condensed over six runs than for C++. Since access to statistics about the training data for Llama 3 and GPT-3.5 is not provided, reasons for more stable performance in Python across runs can possibly lie in the training data distribution but cannot be confirmed. However, the fact that HumanEval-Python is a

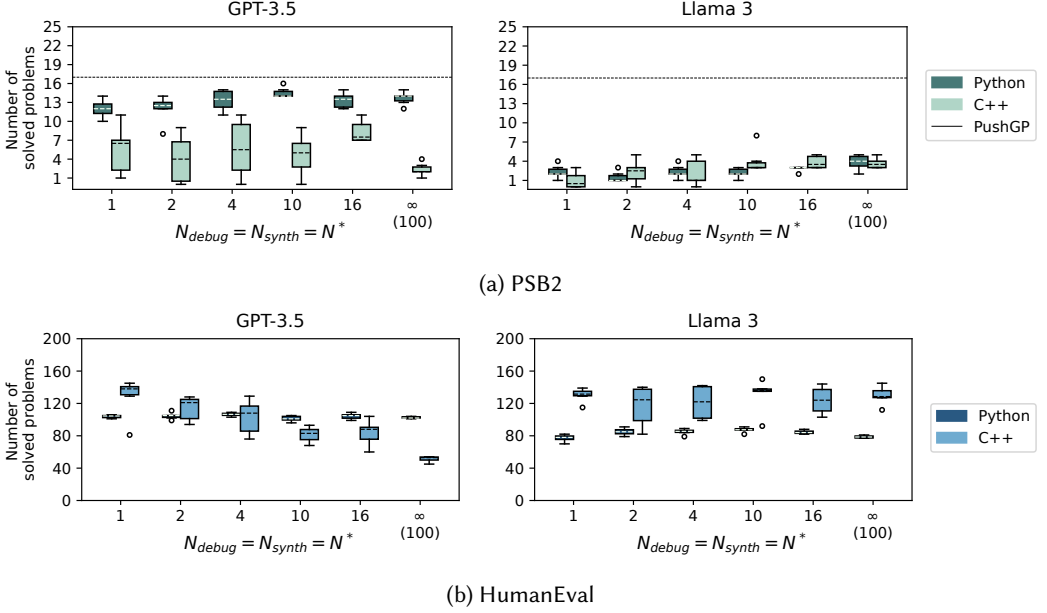


Fig. 8. Repair-replace trade-off as a tree search problem in SEIDR: the total number of solved problems as measured by  $TPR = 1$  using SEIDR with GPT-3.5 and Llama 3 depending on tree arity  $N^*$ .

popular code generation benchmark against which models are compared may affect the results on this and other Python benchmarks. In the same line of comparison of results between two programming languages, SEIDR with GPT-3.5 performs better in Python than in C++ on PSB2.

Llama 3 performs worse than GPT-3.5 on PSB2 in both languages. This dataset has more test cases in stock than HumanEval and can be considered a more thorough test of the coding and debugging capabilities of LLMs. Following a recent trend, where larger models outperform smaller ones, the results of SEIDR on PSB2 confirm that the smaller model (Llama 3) performs worse than the larger one (GPT-3.5).

Looking back at the initial experiments with Codex and PSB2, we notice the degradation of performance from Codex to GPT-3.5: SEIDR with Codex solved 19 problems in Python and 17 in C++ with tree arity 10, while the best-performing result of SEIDR GPT-3.5 is 16 (tree arity of 10, too) in Python and 11 in C++ (several tree arities, but not 10). This result can be explained by the focus of LLM builders on the generalization of knowledge and performance on a variety of tasks, while Codex specializes in code generation. Moreover, due to increased costs from Codex to GPT-3.5, we decrease the maximum number of generated program candidates from 1000 for Codex to 100 for GPT-3.5. However, only two problems are solved with Codex with  $EPG > 100$ : indices of substring (at program candidate #510) and substitution cipher (at candidate #210) in Python, bowling (candidate #912) and indices of substring (#410) in C++. Meanwhile, some problems are solved by Codex earlier than at the debugging attempt #100 and not solved by GPT-3.5 and vice versa. Therefore, the reduction of the maximum generated programs has only a partial effect on the difference between the results of the two models.

In addition to reporting the number of fully solved problems, Figure 9 also reports the mean Test Pass Rate measured in %. For HumanEval-C++, SEIDR with GPT-3.5 has better results with smaller tree arity values than with larger ones. In Python, experiments with  $1 < N^* < 100$  (i.e.,

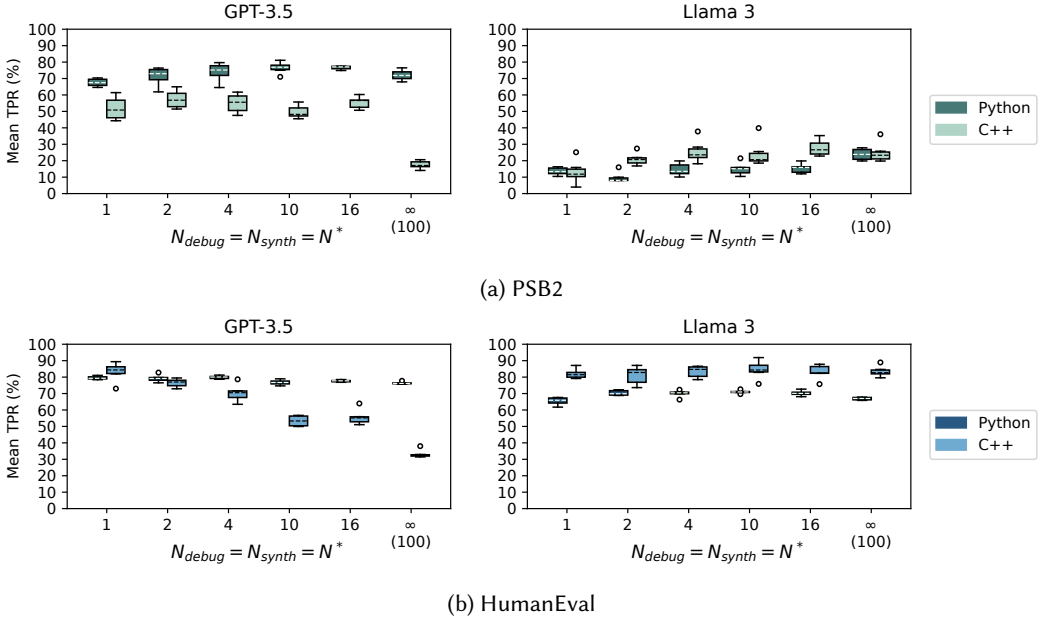


Fig. 9. Repair-replace trade-off as a tree search problem in SEIDR: mean  $TPR$  measured in % obtained using SEIDR with GPT-3.5 and Llama 3 depending on tree arity  $N^*$ .

non-corner case values of  $N^*$ ) yield better mean  $TPR$  for SEIDR with GPT-3.5 on PSB2 and with Llama 3 — for HumanEval-Python. SEIDR with Llama 3 performs with slightly better mean  $TPR$  towards larger  $N^*$  on PSB2-C++, and at the same time, it performs well on HumanEval-C++, so the difference between results with different  $N^*$  is small. The maximum of mean  $TPR$  for both datasets and models and max number of solved problems are obtained with  $N^* = 16$  or less, except for Llama 3 in PSB2-C++. From this part of the experiments, we notice that moderate or small tree arities ( $N^* \leq 16$ ) are preferred, but there is no one leading tree arity.

The number of runs in which a problem is solved and the average speed of finding those solutions are shown in Figure 10 for PSB2, Figure 11 for HumanEval-Python and Figure 12 for HumanEval-C++. These figures show which problems were not solved in any run of any experiment (a row with zeros colored white) and what problems are easier to solve than others (e.g., solved in all runs, or at least one with each tree arity  $N^*$ , or earlier in the search tree and shown in brighter rather than darker color but not white). We also show the results of lexicase selection runs marked with “lex.” in these three figures, but will discuss them in a separate section.

The majority of solved PSB2 problems are solved in more than one run per setting with GPT-3.5 and faster (with fewer attempts) in Python than in C++ as illustrated in Figure 10. For Llama 3, most of the solutions are obtained in 1–3 runs. The trend for both datasets is that, for problems where a solution is found, SEIDR with GPT-3.5 makes fewer attempts in Python (brighter colors prevail in the corresponding parts of Figures 10 and 11) than SEIDR with GPT-3.5 in C++ or SEIDR with Llama 3 in both languages (darker shades in the GPT-3.5 on C++ and Llama 3 parts of Figures 10 and 12).

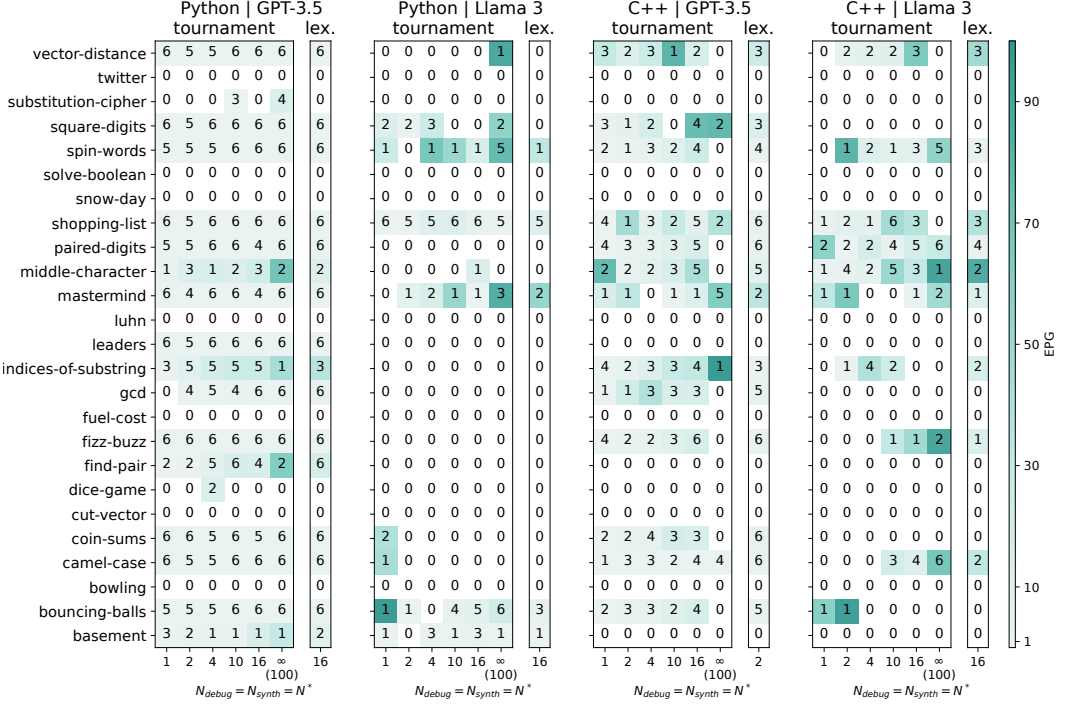


Fig. 10. PSB2: mean Excess Programs Generated (in color) and the number of runs in which a task is solved. The experiments in which a specific problem is not solved in any run are shown in white.

To explore the capabilities of SEIDR with the studied models, we can take a union over experiments and count the number of problems solved at least once across  $36^{20}$  restarts with a fixed dataset, language and model. In this way, in addition to reporting the number of problems solved in each run in the boxplot, we can derive the number of problems solved in any run with any set of hyperparameter using information in Figures 10–12. For example, SEIDR with GPT-3.5 does not solve only 7 problems out of 25 PSB2-Python tasks (see Figure 10). Namely, bouncing-balls, cut-vector, dice-game, indices-of-substring, middle-character, solve-boolean, and substitution-cipher are solved in 0 runs with all the variations of  $N^*$ . In other words, 18 unique problems are solved in the collection of all the experiments using SEIDR with GPT-3.5 on PSB2-Python. SEIDR with GPT-3.5 does not solve 12 out of 25 PSB2-C++ tasks with any hyperparameter settings. The number of solved problems in any run with SEIDR and Llama 3 are 10 for PSB2 in both languages. To support the finding about degradation of performance happening in a non-specialized model GPT-3.5 compared to the code-specialized Codex, we calculate the number of solved problems in any run with Codex and obtain 20 problems solved at least once in Python and 18 in C++.

Similarly, SEIDR with GPT-3.5 solves 141 out of 164 HumanEval-Python problems at least once in all runs, collectively, and 128 problems with Llama 3 (see Figure 11). In HumanEval-C++, SEIDR with GPT-3.5 solves 163 out of 164 problems (except CPP/137) in the union of all runs and 162 with Llama 3 (except CPP/137 and CPP/151) as follows from Figure 12.

In Figure 13, we present the analogy of the speed of obtaining solutions, EPG. We zoom in to the cases with solutions found with up to the first 10 program candidates in Figures 13a and 13c for

<sup>20</sup> We have 6 runs and 6 different values of  $N^*$ .

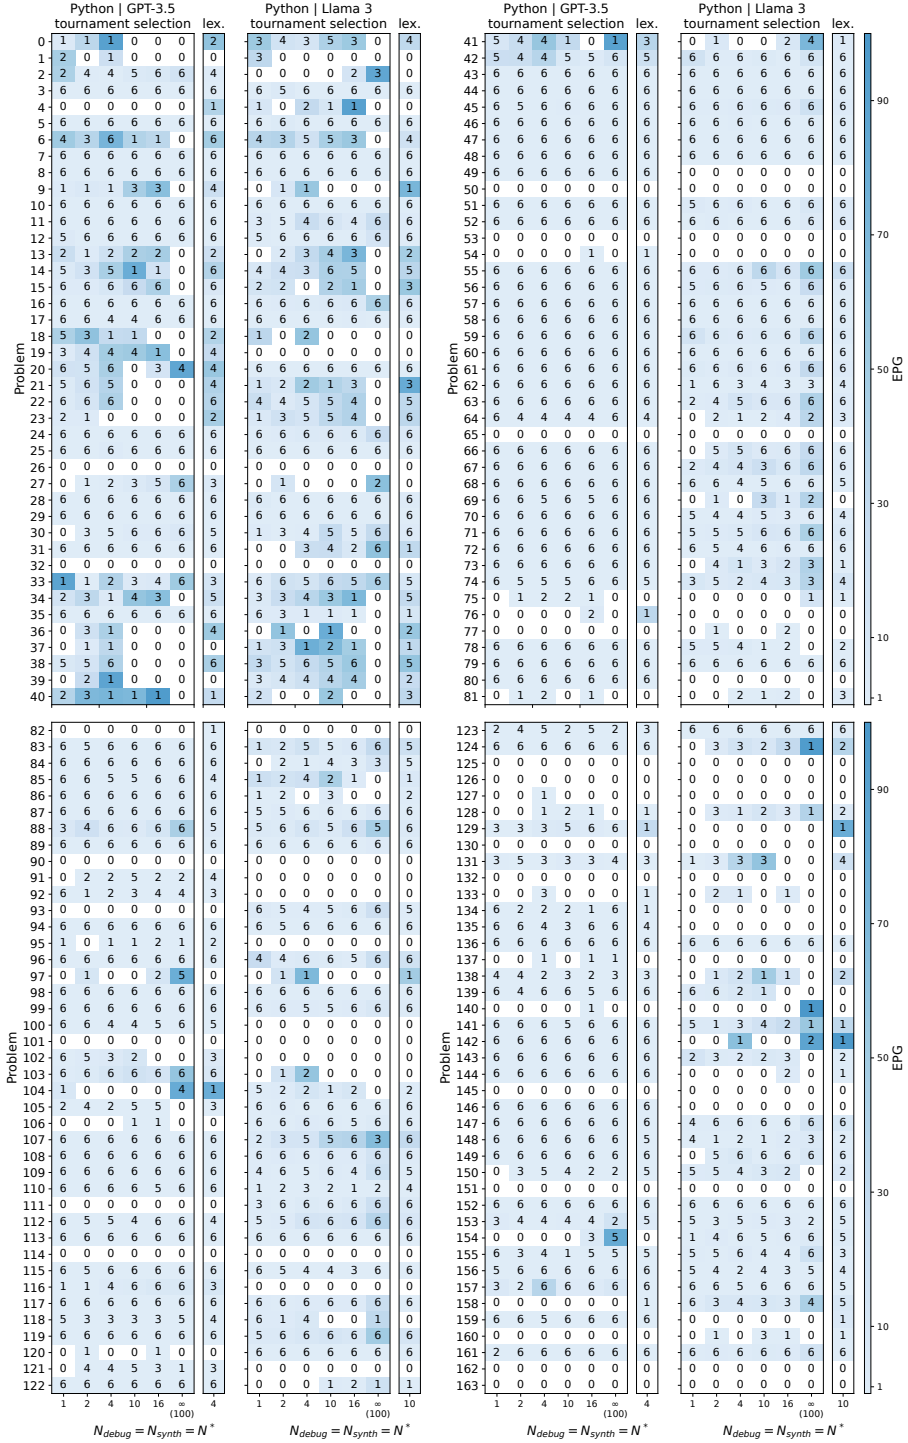


Fig. 11. HumanEval-Python: mean Excess Programs Generated (in color) and the number of runs in which a task is solved. The problems that are not solved in any run are colored white.

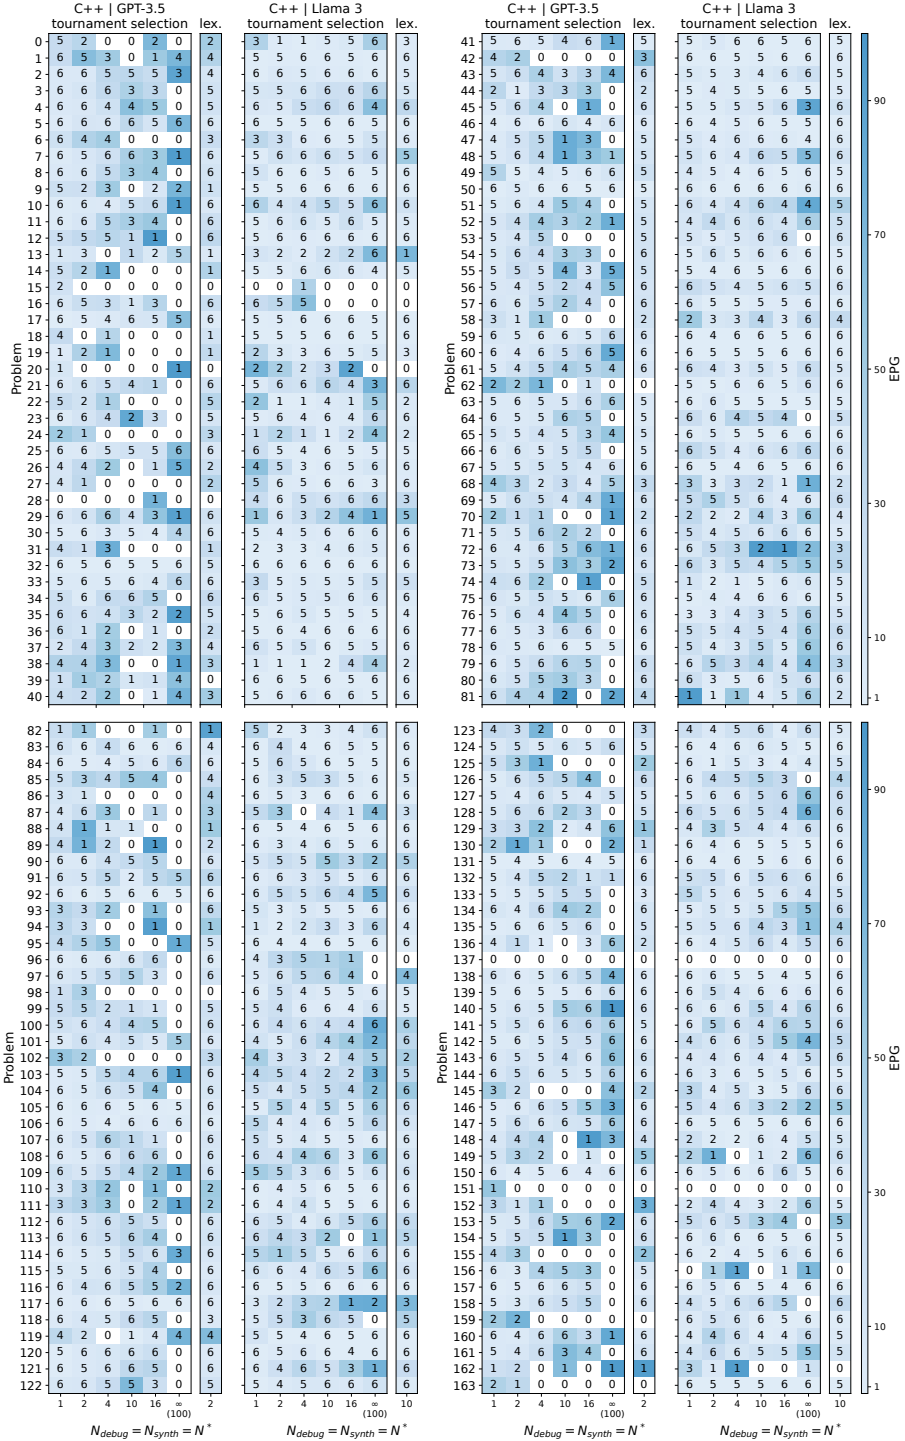


Fig. 12. HumanEval-C++: mean Excess Programs Generated (in color) and the number of runs in which a task is solved. The problems that are not solved in any run are colored white.



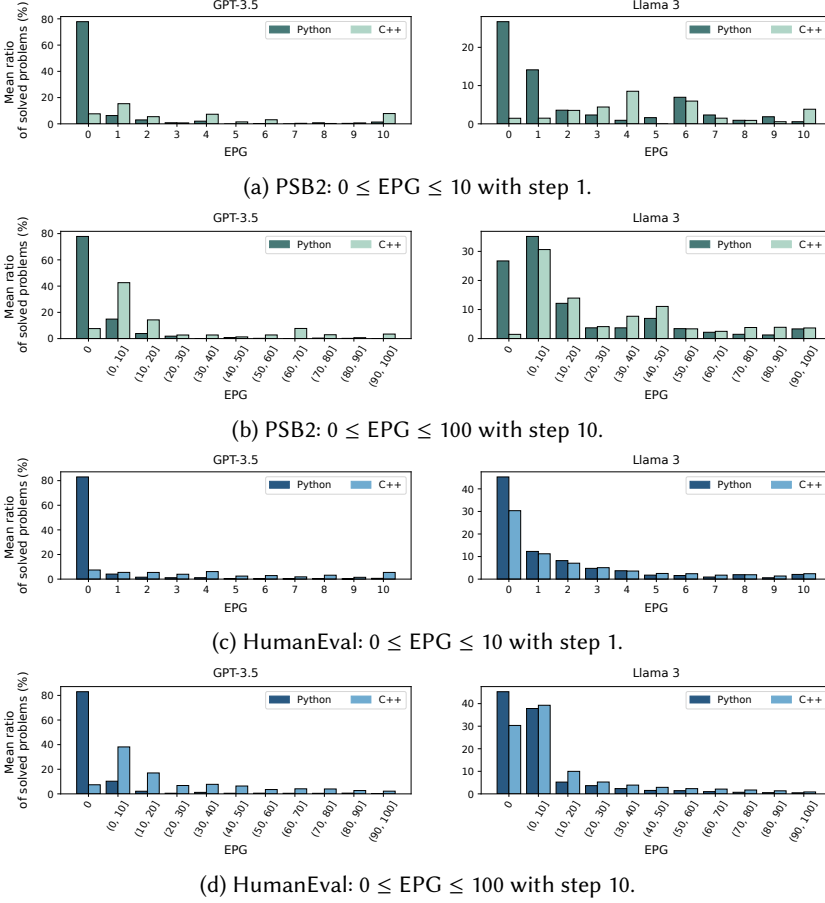


Fig. 13. Distribution (in %) of the number of generated programs with GPT-3.5 and Llama 3 during each problem-solving attempt on average over 6 runs with different tree arities  $N_{\text{synth}}$ ,  $N_{\text{debug}}$ .

PSB2 and HumanEval-X, respectively. The coarser-grained EPG distribution with the step of 10 candidates is shown in Figures 13b and 13d.

In detail, we show the distribution of EPG values in all experiments to explore what proportion of candidate updates is made before a solution is found. For example, on average, over 70% of Python solutions by SEIDR with GPT-3.5 are solved from the first attempt, i.e., have  $\text{EPG}=0$  (see Figure 13a). Python solutions are more frequently found from the first attempt by SEIDR with Llama 3 than later in the tree, although less frequently than with GPT-3.5. Most of solutions found by SEIDR benefit from the iterative repair and generate up to 10 extra programs with both models (see Figures 13b, 13d). However, some solutions are also found later in the tree search.

The EPG distribution results for correctly solved problems with  $\text{TPR} = 1$  imply that the first steps in the update of the draft program are crucial for solving the problem. The chances of solving the problem in the later stages of the search are low. This confirms our assumption in Section 4.4 that 100 programs are sufficient in the generalizability experiments.

Table 3. SEIDR generalizability experiments: hyperparameters in the lexicase selection experiments.

experiment #	22	23	24	25	26	27	28	29
datasets	PSB2				HumanEval			
models	GPT-3.5		Llama 3		GPT-3.5		Llama 3	
language	C++	Python	C++	Python	C++	Python	C++	Python
population size (beam width), $W$	2	16	16	16	2	4	10	10
# programs in the 1st generation, $N_{\text{synth}}$	2	16	16	16	2	4	10	10
# bug explanations for candidate, $N_{\text{explain}}$	2	2	2	2	2	2	2	2
# repairs for each explanation, $N_{\text{debug}}$	2	16	16	16	2	4	10	10
max programs generated	100							
prompts	see Section 4.3.2							
# runs per experiment	6							

**Repair-replace trade-off for SEIDR with GPT-3.5 and Llama 3 in the generalizability experiments (RQ1, RQ2, RQ3):** Unlike for SEIDR with Codex and a fixed debugging instruction (see Section 5.1.1), SEIDR with GPT-3.5 and Llama 3 does not show any distinct trend for all the languages and datasets in terms of the preferred tree arity value. Results over a number of runs with the same hyperparameter settings are more condensed for Python than for C++, which can be a result of optimizing LLMs for high performance on popular coding benchmarks in Python. If solutions to problems are found by SEIDR, it is done at an earlier tree search step for Python than for C++ (i.e., with a smaller EPG). The majority of solutions are found within the first 10 updates of program candidates. SEIDR solves 163 out of 164 HumanEval-C++ problems with GPT-3.5 at least once over all runs with all restarts and hyperparameter sets, and 162 — with Llama 3.

### 5.2.2 Parent Selection Strategies within the Ranking Agent in the Generalizability Experiments.

Based on the mean TPR over all runs reported in Figure 9, we fix the best-performing  $N^*$  for each dataset, language and model and run the experiments with lexicase selection instead of tournament selection as a parent selection strategy. The hyperparameters are shown in Table 3. We compare the number of problems solved with these settings and two types of parent selection algorithms in Figure 14 and mean TPR in Figure 15. Mean TPR with lexicase selection and the best tournament selection configuration are similar for two selection strategies for most of the experiments. Similar results are obtained for the number of fully solved problems, with the exception of C++ results of SEIDR with GPT-3.5 on both datasets, where lexicase selection improved the results.

To zoom in on the improvement details, we refer back to Figures 10–12, where separate columns are dedicated to lexicase selection experiments (marked with “lex.”). We do not observe different programs solved with lexicase selection than with tournament selection. Some problems are solved in all six runs with lexicase selection and in fewer runs with tournament selection, such as find-pair and fizz-buzz in the C++ experiments with GPT-3.5.

To confirm the effect of lexicase selection on the program candidate search, we have explored the test pass rate dynamics. An example run of SEIDR with GPT-3.5 on HumanEval is shown in Figure 16. We observe that in the vast majority of cases, the test score jumps from 0 to 1 directly, not as a result of reordering candidates, but as a result of SEIDR bug summarization and candidate

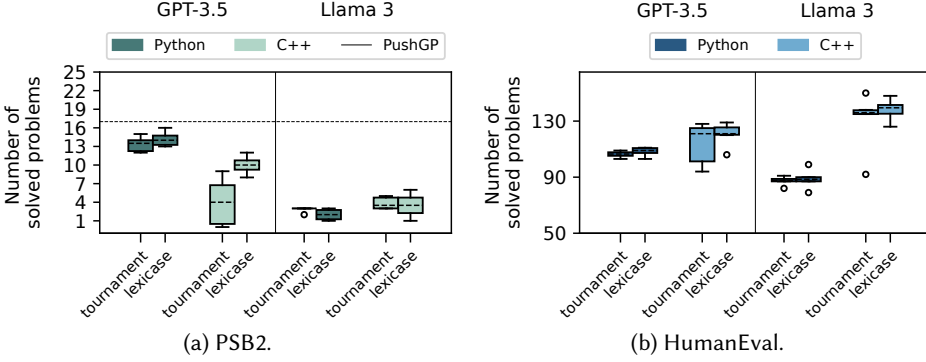


Fig. 14. Number of solved problems in lexicase and tournament selection experiments for SEIDR with GPT-3.5 and Llama 3 and fixed tree arity  $N^*$ .

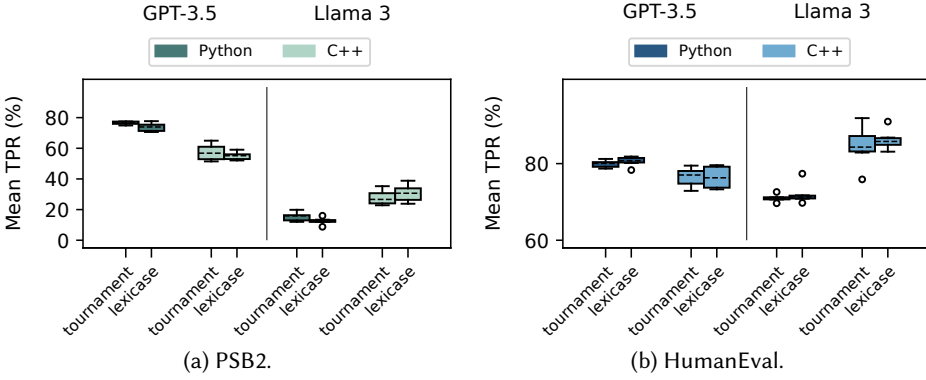


Fig. 15. Mean  $TPR$  (measured in %) in lexicase and tournament selection experiments for SEIDR with GPT-3.5 and Llama 3 and fixed tree arity  $N^*$ .

update. Therefore, we appoint differences in results between lexicase and tournament parent selection primarily to the LLMs themselves and the debugging loop rather than to the parent selection strategy. Specifically, when an LLM gets the prompt that fixes exactly the error in a program candidate, the score jumps to 1 because of the correct prompt provided to the model more frequently than because a better candidate is chosen on a previous step.

**Parent selection strategies for SEIDR with GPT-3.5 and Llama 3 in the generalizability experiments (RQ2, RQ3, RQ4):** No leading ranking strategy is found for SEIDR experiments with GPT-3.5 and Llama 3 as measured by the average test pass rate and the average number of solved problems over six runs. In the vast majority of experiments where a solution is found and at least one debugging attempt is made the score jumps from 0 to 1 as opposed to climbing up incrementally and makes the ranking strategy less impactful than effective prompting.

**5.2.3 Overall Generalizability of SEIDR with parent selection strategies and repair-replace trade-off.** We combine all results in Table 4 for PSB2 and Table 5 for HumanEval-X, where we show




Fig. 16. The best score on validation tests (y-axis) obtained up to an indicated logging step (x-axis) for problems solved from the second or later attempts with lexicase selection using SEIDR with GPT-3.5. Results are obtained for HumanEval. “Step” on the x-axis corresponds to logging settings and roughly indicates the speed of finding solutions to different problems. Each line stands for a unique solution.

the number of solved problems at different cutoffs ( $k$ ). An average is taken over all runs for the experiments with multiple restarts, and the results for one run are shown for SEIDR with Codex.

Note that at cutoff  $k = 1$ , iterations of SEIDR are not taken into account, because the first program candidate is the one LLMs generate from a template, and the trees at cutoff  $k = 1$  have only one program in each of them, regardless of the parent selection strategy or tree arity. Here, the difference between  $\text{pass}@1$  results with different parent selection and tree arity is explained by the stochastic behavior of the underlying LLMs. We keep  $\text{pass}@1$  results to inform the reader about the stochasticity of the models and for the ease of visual comparison of values in each row.

We compare the performance of PSB2 solutions synthesized with SEIDR to the PushGP genetic programming system with down-sampled lexicase selection [Helmuth and Spector 2022]. To compare the performance of SEIDR with other LLMs without SEIDR, we report the average percentage of solved problems for HumanEval-X. For HumanEval-X, we cite the performance of the state-of-the-art models as reported by their authors without SEIDR, such as CodeGeex [Zheng et al. 2023], GPT-3.5 and GPT-4 [OpenAI 2023], for reference, and compare SEIDR with GPT-3.5 and Code Llama [Rozière et al. 2023] results at different cutoffs ( $k$  for  $\text{pass}@k$ ). Furthermore, we report the number of solved problems in the union of all experiments for a given dataset, model, and language, because different problems are solved in different restarts of the method. We refer to these results as “solved at least once” hereafter.

In PSB2-Python experiments, SEIDR with Codex has a higher  $\text{pass}@1000$  than the maximum  $\text{pass}@100$  measured for SEIDR with GPT-3.5. SEIDR with Codex also solves 20 problems at least once, whereas SEIDR with GPT-3.5 solves 18. SEIDR with Codex also outperforms other LLMs on PSB2 in C++. No parent selection strategy or tree arity  $N^*$  is leading across all PSB2 experiments.

In the HumanEval-C++ experiments, SEIDR performs well at larger  $k$ , i.e., when debugging steps are made. Remarkably, the union of SEIDR experiments with both GPT-3.5 and a much smaller Llama 3 solve 163 problems (or 99.4% in Table 5) and 162 problems (or 98.8%) in HumanEval-C++, correspondingly. SEIDR does not outperform other LLMs on HumanEval-Python, which, as mentioned earlier, could be the effect of HumanEval-Python being a popular benchmark for testing LLMs and indirectly optimizing the models.

Direct comparison of SEIDR results with other iterative program synthesis frameworks is challenging because, in addition to differences in benchmarks used, Jiang et al. [2023] and [Chen et al. 2023] do not always report the number of generated programs, which is the primary metric that we use to compare iteration against a replace-only baseline.<sup>21</sup>

<sup>21</sup>[Chen et al. 2023] does have a sentence “Note that typically one debugging turn is sufficient”

Table 4. Average number of solved PSB2 problems using SEIDR with different LLMs and PushGP. The best results among SEIDR-only experiments at each  $k$  are highlighted in bold. If pass@100 is not available, we cite the best available results in brackets.  $N^*$  stands for the number of debug candidates and debug instructions to generate (tree arity).

Language	Model in SEIDR	$N^*$	Parent selection	pass@1	pass@10	pass@100
Python	GPT-3.5	1	tournament	10.5	12.0	12.0
		2	tournament	9.3	11.5	12.0
		4	tournament	9.8	12.2	13.3
		10	tournament	10.7	<b>12.7</b>	<b>14.5</b>
		16	tournament	10.0	11.3	13.3
		16	lexicase	10.2	12.3	14.2
		100	tournament	10.8	12.2	13.7
	solved at least once by SEIDR with GPT-3.5			13	17	18
	Codex	10	tournament	5	10	14 (pass@1000=19)
	solved at least once by SEIDR with Codex			8	13	17 (pass@1000=20)
	PushGP (no SEIDR)	-	-	-	-	(17)
	Llama 3	1	tournament	1.2	1.3	2.3
		2	tournament	1.0	1.5	1.5
		4	tournament	1.0	1.5	2.3
		10	tournament	1.5	2.0	2.2
		16	tournament	1.7	<b>2.8</b>	2.8
		16	lexicase	1.0	1.7	2.0
		100	tournament	1.0	1.0	<b>3.8</b>
	solved at least once by SEIDR with Llama 3			4	5	11
C++	GPT-3.5	1	tournament	0.0	5.6	5.5
		2	tournament	1.0	5.0	6.0
		2	lexicase	1.3	<b>9.0</b>	<b>10.0</b>
		4	tournament	1.0	4.6	6.8
		10	tournament	1.0	1.3	5.6
		16	tournament	1.0	1.2	8.3
		100	tournament	1.0	1.2	2.3
	solved at least once by SEIDR with GPT-3.5			2	13	13
	Codex	10	tournament	3	12	14 (pass@1000=17)
	solved at least once by SEIDR with Codex			10	12	15 (pass@1000=18)
	PushGP (no SEIDR)	-	-	-	-	(17)
	Llama 3	1	tournament	0.0	1.0	2.0
		2	tournament	0.0	<b>2.0</b>	2.8
		4	tournament	1.0	<b>2.0</b>	2.6
		10	tournament	1.0	1.4	<b>4.0</b>
		16	tournament	0.0	1.8	3.8
		16	lexicase	1.0	1.4	3.5
		100	tournament	0.0	0.0	3.7
	solved at least once by SEIDR with Llama 3			4	5	11

Table 5. Percentage of solved tasks in HumanEval-X using SEIDR with different LLMs. The best results among SEIDR-only experiments at each  $k$  are highlighted in bold.  $N^*$  stands for the number of debug candidates and debug instructions to generate (tree arity).

Language	Model in SEIDR	$N^*$	Ranking	pass@1	pass@10	pass@100
Python	GPT-3.5	1	tournament	53.4	<b>60.4</b>	63.3
		2	tournament	51.3	58.2	63.7
		4	tournament	52.1	58.4	64.8
		4	lexicase	52.5	59.6	<b>66.1</b>
		10	tournament	51.1	58.3	61.7
		16	tournament	53.0	59.6	63.5
		100	tournament	54.0	56.7	62.5
	solved at least once by SEIDR with GPT-3.5			70.7	84.1	87.8
	Llama 3	1	tournament	22.0	44.0	47.1
		2	tournament	24.0	<b>48.5</b>	51.8
		4	tournament	22.5	43.1	51.8
		10	tournament	24.0	42.9	53.3
		10	lexicase	26.2	44.1	<b>54.1</b>
		16	tournament	23.5	41.6	51.5
		100	tournament	21.4	26.3	48.0
	solved at least once by SEIDR with Llama 3			56.7	73.8	79.3
	<b>LLM results without SEIDR</b>					
	GPT-3.5 (ChatGPT)	-	-	48.1	-	-
	GPT-4	-	-	67.0	-	-
	Code Llama 34B	-	-	48.8	76.8	93.0
	Unnatural Code Llama 34B	-	-	62.2	85.2	95.4
	CodeGeeX	-	-	22.89	39.57	60.92
C++	GPT-3.5	1	tournament	3.4	<b>55.9</b>	<b>78.5</b>
		2	tournament	4.7	36.6	69.6
		2	lexicase	4.3	35.9	73.6
		4	tournament	4.5	27.7	62.9
		10	tournament	3.6	13.7	49.7
		16	tournament	4.5	13.2	51.1
		100	tournament	3.6	7.1	31.3
	solved at least once by SEIDR with GPT-3.5			20.1	89.0	99.4
	Llama 3	1	tournament	29.6	59.5	79.5
		2	tournament	18.6	51.6	71.4
		4	tournament	20.0	50.6	73.9
		10	tournament	25.6	56.4	80.0
		10	lexicase	25.9	<b>60.9</b>	<b>84.2</b>
		16	tournament	20.9	53.0	75.5
		100	tournament	27.1	39.7	79.2
	solved at least once by SEIDR with Llama 3			72.0	97.6	98.8
	<b>LLM results without SEIDR</b>					
	Code Llama 34B	-	-	47.8	-	-
	CodeGeeX	-	-	17.06	32.21	51.00

**Summary of SEIDR results:** SEIDR with Codex benefits from the repair-replace trade-off when building a search tree of PSB2 solutions. The best results on PSB2 are achieved with Codex, tree arity of 10 and a maximum of 1000 programs generated. A trend of having the best tree arity set to 10 does not fully hold for other tested models, GPT-3.5 and Llama 3, where different tree arities performed better in different settings, e.g., depending on the parent selection algorithm, dataset, and programming language. Due to increasing costs of newer models time required for testing, we have run GPT-3.5 and Llama 3 six times with each set of hyperparameters and stopped building the search tree at 100 programs. In HumanEval-C++, the union of these runs has solved 163 problems with GPT-3.5, 162 problems with a much smaller Llama 3-8B model. SEIDR runs have also solved 18 PSB2 problems in C++ and 20 in Python with Codex at least once in all the experiments. The numbers for SEIDR with GPT-3.5 are lower than for Codex, possibly due to the focus of GPT-3.5 on general reasoning and of Codex — on coding. A smaller Llama 3 model performs poorly on PSB2, which we appoint to the difficulty of the benchmark and the popularity of HumanEval and to the fact that the models can be indirectly optimized for higher performance on HumanEval than on other programming benchmarks.

### 5.3 Threats to Validity

External threats to validity concern SEIDR performance with benchmarks and language models different from those tested. Specifically, PSB2 and HumanEval-X contain programming tasks that require smaller functions to be generated than production-scale software. Although some canonical solutions in HumanEval-Python have been criticised [Liu et al. 2023b], we primarily use unit tests to evaluate the output of SEIDR and do not compute the exact match. Therefore, these weaknesses do not impact our results.

Internal threats relate to the implementation. We use PSB2, which has both corner cases and regular tests available. To ensure a fair comparison with other studies on PSB2, we evaluate and report results on the provided test set of PSB2, of which we randomly pick 2000 tests. A risk that a program does not pass tests other than the ones picked persists but is assumed to be low given the large enough number of 2000 tests.

The large language models for code editing and text completion used in this study are non-deterministic, which can affect the results. Due to prohibitive model inference costs and discontinuation of support for earlier models, each experiment with Codex and GPT-3 is run only once. Experiments with Llama 3 and GPT-3.5 are restarted six times. We acknowledge that more restarts can provide a more complete picture of the approach performance, although this is not the standard practice in the LLM domain [Ouyang et al. 2023] and can come at considerable financial and environmental costs. Our temperature-based sampling procedure described in Section 2.1.1 reduces this stochasticity significantly, especially for low-EPG results, because earlier solutions in the search tree are obtained with lower temperatures, and lower temperature limits the stochasticity. Codex, GPT-3, and GPT-3.5 are black-box models and can generate malicious code [Pearce et al. 2022]. Therefore, when coding in real life, the LLM output should be carefully reviewed before running it.

The results can be skewed towards high performance in the programming languages prevailing in the pre-training dataset used by the authors of the tested LLMs – an issue also known as *data contamination*. However, if the models were able to directly reproduce the solutions seen in the training data, these problems would have been solved from the first attempt. Since we observe  $EPG > 1$  values, data contamination is likely to have only a partial effect on the results. Moreover, the results can be affected by popular evaluation benchmarks: even though the benchmarks themselves

are usually not parts of the pre-training dataset, the published LLMs are likely to be optimized for high performance on these benchmarks.

## 6 Conclusion

In this study, we propose SEIDR, a multi-agent framework to solve the challenge of fully autonomous programming. In SEIDR, the program synthesis procedure is augmented from the direct generation of code with large language models instructions to iterative calls to a DEBUG agent followed by the RANK agent. The DEBUG agent performs a tree search across program candidates generated by a large language model for code. The LLM used for code repair takes imperfect program candidates and instructions for their improvement as prompts. The instructions are obtained from both static templates with failing test case descriptions and templates with auto-generated bug summaries by a text completion language model.

In addition to the initial exploration of hyperparameters that influence the population size at each generation and the number of children for each parent in SEIDR iterations, we extend the framework to test two different ranking strategies (tournament selection and lexicase selection), three models in the coding part of SEIDR (Codex, GPT-3.5 and Llama 3), two datasets (PSB2 and HumanEval-X), two programming languages (Python and C++), and various branching factors. With the update to newer models, the prompts and parameters of the tree search are updated accordingly.

**Contributions:** We run one set of initial exploration experiments and two sets of generalizability experiments. In the initial exploration, we test SEIDR with Codex-edit as the model for draft program synthesis and debugging in Python and C++ on the PSB2 benchmark. In our experiments, SEIDR outperforms the PushGP baseline and achieves the state-of-the-art result with 19 solved problems out of 25. It requires under 1000 program executions to solve them, in stark contrast to billions<sup>22</sup> of executions in PushGP, making it feasible in the areas with costly testing, such as robotics. Investigation of the repair-replace trade-off shows that SEIDR with tree arity of 10 outperforms both the replace-only strategy and the repair-only approach.

To study the generalizability of SEIDR, we experiment with GPT-3.5 and Llama 3-8B as the models for draft program synthesis, explaining errors in synthesized programs and their debugging in Python and C++ on the PSB2 and HumanEval-X benchmarks. Our experiments with lexicase and tournament parent selection in the RANK agent do not show consistent improvement with either policy or any tree arity. One observation is that program candidates in generations prior to the final solution do not pass any test in the test suite in the vast majority of cases. The test pass rate for the final solution abruptly increases in one generation from 0 to all tests passed.

In our generalizability experiments, SEIDR shows better performance than using LLMs without SEIDR with the same prompts and the same budget in terms of programs to generate. The method achieves high results of 78.5% average pass@100 on HumanEval-C++ with GPT-3.5 and 84.2% with Llama 3. Remarkably, in HumanEval-C++, the union of SEIDR restarts with different hyperparameters has solved 163 problems with GPT-3.5, 162 problems with a much smaller Llama 3-8B model. The union of SEIDR runs has solved 18 PSB2 problems in C++ and 20 in Python with Codex.

**Future work:** Further investigation of SEIDR generalizability and ranking strategies are some of the areas for future work. Benchmarks with more tests than in HumanEval-X may shed more light on the most effective choice of the number of programs generated from each bug explanation, as well as the framework’s comparison on large-context projects. As an agent-based framework,

<sup>22</sup> A problem is considered “solved” by PushGP if at least 1 of 100 runs, each with a limit of 60 million programs, was successful.



SEIDR shows how LLM-based agents and other non-LLM agents or components can collaboratively interact and solve software engineering tasks.

## Data Availability

To support open science and allow for replication and verification of our work, a replication package containing the code and results is publicly available in Zenodo.<sup>23</sup>

## Acknowledgments

The work presented in this paper was supported by the European Commission through Horizon 2020 grant 812882, and by the Research Council of Norway through the secureIT project (#288787). The empirical evaluation made use of the Experimental Infrastructure for Exploration of Exascale Computing (eX3), supported by the Research Council of Norway through grant #270053.

## References

- Milad Taleby Ahvanooy, Qianmu Li, Ming Wu, and Shuo Wang. 2019. A Survey of Genetic Programming and Its Applications. *KSII Transactions on Internet and Information Systems (TIIS)* 13, 4 (2019), 1765–1794.
- Miltiadis Allamanis, Earl T. Barr, Premkumar Devanbu, and Charles Sutton. 2018. A Survey of Machine Learning for Big Code and Naturalness. *Comput. Surveys* 51, 4 (July 2018), 81:1–81:37. doi:10.1145/3212695
- Miltiadis Allamanis and Charles Sutton. 2014. Mining Idioms from Source Code. In *ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 472–483. doi:10.1145/2635868.2635901 arXiv:1404.0417 [cs]
- Rajeev Alur, Rastislav Bodik, Eric Dallal, Dana Fisman, Pranav Garg, Garvit Juniwal, Hadas Kress-Gazit, P. Madhusudan, Milo M. K. Martin, Mukund Raghothaman, Shamwadiya Saha, Sanjit A. Seshia, Rishabh Singh, Armando Solar-Lezama, Emina Torlak, and Abhishek Udupa. 2015. Syntax-Guided Synthesis. *Dependable Software Systems Engineering* (2015), 1–25. doi:10.3233/978-1-61499-495-4-1
- Alberto Bartoli, Mauro Castelli, and Eric Medvet. 2020. Weighted Hierarchical Grammatical Evolution. *IEEE Transactions on Cybernetics* 50, 2 (Feb. 2020), 476–488. doi:10.1109/TCYB.2018.2876563
- Osbert Bastani, Jeevana Priya Inala, and Armando Solar-Lezama. 2022. Interpretable, Verifiable, and Robust Reinforcement Learning via Program Synthesis. In *xxAI - Beyond Explainable AI: International Workshop in Conjunction with ICML - Revised and Extended Papers*, Andreas Holzinger, Randy Goebel, Ruth Fong, Taesup Moon, Klaus-Robert Müller, and Wojciech Samek (Eds.). Springer, 207–228. doi:10.1007/978-1-4939-9850-1\_12
- Rohan Bavishi, Harshit Joshi, José Cambronero, Anna Fariha, Sumit Gulwani, Vu Le, Ivan Radicek, and Ashish Tiwari. 2022. Neurosymbolic Repair for Low-Code Formula Languages. In *OOPSLA*. ACM, 1–30.
- Jenny Blain. 2002. *Nine Worlds of Seid-magic: Ecstasy and Neo-shamanism in North European Paganism*. Routledge.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models Are Few-Shot Learners. In *International Conference on Neural Information Processing Systems (NeurIPS) (NIPS'20)*. Curran Associates, Inc., Red Hook, NY, USA, 1877–1901.
- Saikat Chakraborty, Rahul Krishna, Yangruibo Ding, and Baishakhi Ray. 2022. Deep Learning Based Vulnerability Detection: Are We There Yet? *IEEE Transactions on Software Engineering* 48, 9 (Sept. 2022), 3280–3296. doi:10.1109/tse.2021.3087402
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebggen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021b. Evaluating Large Language Models Trained on Code. doi:10.48550/arXiv.2107.03374 arXiv:2107.03374 [cs]
- Xinyun Chen, Maxwell Lin, Nathanael Schärli, and Denny Zhou. 2023. Teaching Large Language Models to Self-Debug. doi:10.48550/arXiv.2304.05128 arXiv:2304.05128 [cs]

<sup>23</sup> Zenodo DOI: 10.5281/zenodo.13754705.

- Xinyun Chen, Dawn Song, and Yuandong Tian. 2021a. Latent Execution for Neural Program Synthesis Beyond Domain-Specific Languages. In *Advances in Neural Information Processing Systems*, Vol. 34. Curran Associates, Inc., 22196–22208.
- Thomas M. Connolly, Mario Soflano, and Petros Papadopoulos. 2023. Systematic Literature Review: XAI and Clinical Decision Support. In *Diverse Perspectives and State-of-the-Art Approaches to the Utilization of Data-Driven Clinical Decision Support Systems*. IGI Global, 161–188.
- Leonardo Lucio Custode, Chiara Camilla Migliore Rambaldi, Marco Roveri, and Giovanni Iacca. 2024. Comparing Large Language Models and Grammatical Evolution for Code Generation. In *Genetic and Evolutionary Computation Conference (GECCO) - Companion (GECCO '24 Companion)*. ACM, New York, NY, USA, 1830–1837. doi:10.1145/3638530.3664162
- Sander de Bruin, Vadim Liventsev, and Milan Petković. 2021. Autoencoders as Tools for Program Synthesis. doi:10.48550/arXiv.2108.07129 arXiv:2108.07129 [cs]
- Sauptik Dhar, Junyao Guo, Jiayi Liu, Samarth Tripathi, Unmesh Kurup, and Mohak Shah. 2021. A Survey of On-Device Machine Learning: An Algorithms and Learning Theory Perspective. *ACM Transactions on Internet of Things* 2, 3 (2021), 1–49. doi:10.1145/3450494
- Luca Di Grazia and Michael Pradel. 2023. Code Search: A Survey of Techniques for Finding Code. *ACM Comput. Surv.* 55, 11 (Feb. 2023), 220:1–220:31. doi:10.1145/3565971
- Yihong Dong, Xue Jiang, Zhi Jin, and Ge Li. 2024. Self-Collaboration Code Generation via ChatGPT. doi:10.48550/arXiv.2304.07590 arXiv:2304.07590 [cs]
- Zhiyu Fan, Xiang Gao, Martin Mirchev, Abhik Roychoudhury, and Shin Hwei Tan. 2023. Automated Repair of Programs from Large Language Models. doi:10.48550/arXiv.2205.10583 arXiv:2205.10583 [cs]
- Eduardo Fernandes, Johnatan Oliveira, Gustavo Vale, Thanis Paiva, and Eduardo Figueiredo. 2016. A Review-Based Comparative Study of Bad Smell Detection Tools. In *International Conference on Evaluation and Assessment in Software Engineering (EASE '16)*. ACM, New York, NY, USA, 1–12. doi:10.1145/2915970.2915984
- Sorin Grigorescu, Bogdan Trasnea, Tiberiu Cocias, and Gigel Macesanu. 2020. A Survey of Deep Learning Techniques for Autonomous Driving. *Journal of Field Robotics* 37, 3 (2020), 362–386. doi:10.1002/rob.21918
- Sumit Gulwani. 2016. *Programming by Examples (and Its Applications in Data Wrangling)*. Technical Report. Microsoft Corporation, Redmond, WA, USA. 22 pages.
- Sumit Gulwani, Oleksandr Polozov, and Rishabh Singh. 2017. Program Synthesis. *Foundations and Trends® in Programming Languages* 4, 1-2 (2017), 1–119. doi:10.1561/25000000010
- Kavi Gupta, Peter Ebert Christensen, Xinyun Chen, and Dawn Song. 2020. Synthesize, Execute and Debug: Learning to Repair for Neural Program Synthesis. In *Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates, Inc., 17685–17695.
- Daniel Conrad Halbert. 1984. *Programming by Example*. Ph.D. Dissertation. University of California, Berkeley.
- David Haussler. 1988. Quantifying Inductive Bias: AI Learning Algorithms and Valiant's Learning Framework. *Artificial Intelligence* 36, 2 (Sept. 1988), 177–221. doi:10.1016/0004-3702(88)90002-1
- Thomas Helmuth and Peter Kelly. 2021. PSB2: The Second Program Synthesis Benchmark Suite. doi:10.48550/arXiv.2106.06086 arXiv:2106.06086 [cs]
- Thomas Helmuth and Peter Kelly. 2022. Applying Genetic Programming to PSB2: The next Generation Program Synthesis Benchmark Suite. *Genetic Programming and Evolvable Machines* 23, 3 (Sept. 2022), 375–404. doi:10.1007/s10710-022-09434-y
- Thomas Helmuth and Lee Spector. 2015. General Program Synthesis Benchmark Suite. In *Genetic and Evolutionary Computation Conference (GECCO)*. ACM, Madrid Spain, 1039–1046. doi:10.1145/2739480.2754769
- Thomas Helmuth and Lee Spector. 2022. Problem-Solving Benefits of Down-Sampled Lexicase Selection. *Artificial Life* 27, 3–4 (March 2022), 183–203. doi:10.1162/artl\_a\_00341
- Thomas Helmuth, Lee Spector, and James Matheson. 2015. Solving Uncompromising Problems With Lexicase Selection. *IEEE Transactions on Evolutionary Computation* 19, 5 (Oct. 2015), 630–643. doi:10.1109/tevc.2014.2362729
- Srinivasan Iyer, Ioannis Konstas, Alvin Cheung, and Luke Zettlemoyer. 2018. Mapping Language to Code in Programmatic Context. doi:10.48550/arXiv.1808.09588 arXiv:1808.09588 [cs]
- Nat Jack and Frank Van der Duyn Schouten. 2000. Optimal Repair–Replace Strategies for a Warranted Product. *International Journal of Production Economics* 67, 1 (Aug. 2000), 95–100. doi:10.1016/s0925-5273(00)00012-8
- Yan Jia, John McDermid, Tom Lawton, and Ibrahim Habli. 2022. The Role of Explainability in Assuring Safety of Machine Learning in Healthcare. *IEEE Transactions on Emerging Topics in Computing* 10, 4 (Oct. 2022), 1746–1760. doi:10.1109/tetc.2022.3171314
- Shuyang Jiang, Yuhao Wang, and Yu Wang. 2023. SelfEvolve: A Code Evolution Framework via Large Language Models. doi:10.48550/arXiv.2306.02907 arXiv:2306.02907 [cs]
- Yanjie Jiang, Hui Liu, Yuxia Zhang, Weixing Ji, Hao Zhong, and Lu Zhang. 2022. Do Bugs Lead to Unnaturalness of Source Code?. In *ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE) (ESEC/FSE 2022)*. ACM, New York, NY, USA, 1085–1096. doi:10.1145/3540250.3549149

- Harshit Joshi, José Cambronero, Sumit Gulwani, Vu Le, Ivan Radicek, and Gust Verbruggen. 2022. Repair Is Nearly Generation: Multilingual Program Repair with LLMs. doi:10.48550/arXiv.2208.11640 arXiv:2208.11640 [cs]
- René Just, Darioush Jalali, and Michael D. Ernst. 2014. Defects4J: A Database of Existing Faults to Enable Controlled Testing Studies for Java Programs. In *International Symposium on Software Testing and Analysis (ISSTA)*. ACM, San Jose, CA, USA, 437–440. doi:10.1145/2610384.2628055
- Kisub Kim, Sankalp Ghatpande, Dongsun Kim, Xin Zhou, Kui Liu, Tegawendé F. Bissyandé, Jacques Klein, and Yves Le Traon. 2023. Big Code Search: A Bibliography. *ACM Comput. Surv.* 56, 1 (Aug. 2023), 25:1–25:49. doi:10.1145/3604905
- John R Koza. 1994. *Genetic Programming II*. Vol. 17. MIT.
- Sumith Kulal, Panupong Pasupat, Kartik Chandra, Mina Lee, Oded Padon, Alex Aiken, and Percy S Liang. 2019. SPoC: Search-based Pseudocode to Code. In *Advances in Neural Information Processing Systems*, Vol. 32. Curran Associates, Inc., 1–12.
- Claire Le Goues, Michael Dewey-Vogt, Stephanie Forrest, and Westley Weimer. 2012. A Systematic Study of Automated Program Repair: Fixing 55 out of 105 Bugs for \$8 Each. In *International Conference on Software Engineering (ICSE)*. IEEE, Zurich, 3–13. doi:10.1109/icse.2012.6227211
- Claire Le Goues, Michael Pradel, and Abhik Roychoudhury. 2019. Automated Program Repair. *Commun. ACM* 62, 12 (Nov. 2019), 56–65. doi:10.1145/3318162
- Joel Lehman, Jonathan Gordon, Shawn Jain, Kamal Ndousse, Cathy Yeh, and Kenneth O. Stanley. 2022. Evolution through Large Models. doi:10.48550/arXiv.2206.08896 arXiv:2206.08896 [cs]
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. 2022. Competition-Level Code Generation with AlphaCode. *Science* 378, 6624 (Dec. 2022), 1092–1097. doi:10.1126/science.abq1158
- Fei Liu, Xialiang Tong, Mingxuan Yuan, and Qingfu Zhang. 2023a. Algorithm Evolution Using Large Language Model. doi:10.48550/arXiv.2311.15249 arXiv:2311.15249 [cs]
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023b. Is Your Code Generated by ChatGPT Really Correct? Rigorous Evaluation of Large Language Models for Code Generation. In *Advances in Neural Information Processing Systems*, Vol. 36. Curran Associates, Inc., 21558–21572.
- Vadim Liventsev, Anastasiia Grishina, Aki Härmä, and Leon Moonen. 2023. Fully Autonomous Programming with Large Language Models. In *Genetic and Evolutionary Computation Conference (GECCO)*. ACM, 1146–1155. doi:10.1145/3583131.3590481
- Vadim Liventsev, Aki Härmä, and Milan Petković. 2021. BF++: A Language for General-Purpose Program Synthesis. doi:10.48550/arXiv.2101.09571 arXiv:2101.09571 [cs]
- Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, Dawn Drain, Daxin Jiang, Duyu Tang, Ge Li, Lidong Zhou, Linjun Shou, Long Zhou, Michele Tufano, Ming Gong, Ming Zhou, Nan Duan, Neel Sundaresan, Shao Kun Deng, Shengyu Fu, and Shujie Liu. 2021. CodeXGLUE: A Machine Learning Benchmark Dataset for Code Understanding and Generation. In *Neural Information Processing Systems Track on Datasets and Benchmarks*. Curran Associates, Inc., 1–16. arXiv:2102.04664 [cs]
- Z. Manna and R. Waldinger. 1992. Fundamentals of Deductive Program Synthesis. *IEEE Transactions on Software Engineering* 18, 8 (1992), 674–704. doi:10.1109/32.153379
- Zohar Manna and Richard J. Waldinger. 1971. Toward Automatic Program Synthesis. *Commun. ACM* 14, 3 (1971), 151–165. doi:10.1145/362566.362568
- Mauricio Marcano, Sergio Díaz, Joshué Pérez, and Eloy Irigoyen. 2020. A Review of Shared Control for Automated Vehicles: Theory and Applications. *IEEE Transactions on Human-Machine Systems* 50, 6 (2020), 475–491. doi:10.1109/thms.2020.3017748
- Changan Niu, Chuanyi Li, Vincent Ng, and Bin Luo. 2023. CrossCodeBench: Benchmarking Cross-Task Generalization of Source Code Models. doi:10.48550/arXiv.2302.04030 arXiv:2302.04030 [cs]
- Michael O’Neill, Anthony Brabazon, Miguel Nicolau, Sean Mc Garraghy, and Peter Keenan. 2004.  $\pi$ Grammatical Evolution. In *Genetic and Evolutionary Computation Conference (GECCO)*, Kalyanmoy Deb (Ed.). Springer, Berlin, Heidelberg, 617–629. doi:10.1007/978-3-540-24855-2\_70
- OpenAI. 2023. GPT-4 Technical Report. doi:10.48550/arXiv.2303.08774 arXiv:2303.08774 [cs]
- Dmitry Orlov. 2020. Finding Idioms in Source Code Using Subtree Counting Techniques. In *Leveraging Applications of Formal Methods, Verification and Validation: Engineering Principles*, Tiziana Margaria and Bernhard Steffen (Eds.). Springer, Cham, 44–54. doi:10.1007/978-3-030-61470-6\_4
- Shuyin Ouyang, Jie M. Zhang, Mark Harman, and Meng Wang. 2023. LLM Is Like a Box of Chocolates: The Non-determinism of ChatGPT in Code Generation. doi:10.48550/arXiv.2308.02828 arXiv:2308.02828 [cs]

- Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri. 2022. Asleep at the Keyboard? Assessing the Security of GitHub Copilot's Code Contributions. In *IEEE Symposium on Security and Privacy (SP)*. IEEE, 754–768. doi:10.1109/sp46214.2022.9833571
- Justyna Petke, Saemundur O. Haraldsson, Mark Harman, William B. Langdon, David R. White, and John R. Woodward. 2018. Genetic Improvement of Software: A Comprehensive Survey. *IEEE Transactions on Evolutionary Computation* 22, 3 (June 2018), 415–432. doi:10.1109/tevc.2017.2693219
- Oleksandr Polozov and Sumit Gulwani. 2015. FlashMeta: A Framework for Inductive Program Synthesis. In *OOPSLA*. ACM, 107–126. doi:10.1145/2814270.2814310
- Shuo Ren, Daya Guo, Shuai Lu, Long Zhou, Shujie Liu, Duyu Tang, Neel Sundaresan, Ming Zhou, Ambrosio Blanco, and Shuai Ma. 2020. CodeBLEU: A Method for Automatic Evaluation of Code Synthesis. doi:10.48550/arXiv.2009.10297 arXiv:2009.10297
- Julia Reuter, Hani Elmestikawy, Fabien Evrard, Sanaz Mostaghim, and Berend van Wachem. 2023. Graph Networks as Inductive Bias for Genetic Programming: Symbolic Models for Particle-Laden Flows. In *Genetic Programming*, Gisele Pappa, Mario Giacobini, and Zdenek Vasecek (Eds.). Springer, Cham, 36–51. doi:10.1007/978-3-031-29573-7\_3
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. 2023. Code Llama: Open Foundation Models for Code. doi:10.48550/arXiv.2308.12950 arXiv:2308.12950 [cs]
- Stuart J Russell. 2010. *Artificial Intelligence a Modern Approach*. Pearson Education, Inc.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. Reflexion: Language Agents with Verbal Reinforcement Learning. *Advances in Neural Information Processing Systems* 36 (Dec. 2023), 8634–8652.
- Atsushi Shirafuji, Yutaka Watanobe, Takumi Ito, Makoto Morishita, Yuki Nakamura, Yusuke Oda, and Jun Suzuki. 2023. Exploring the Robustness of Large Language Models for Solving Programming Problems. doi:10.48550/arXiv.2306.14583 arXiv:2306.14583 [cs]
- Yoav Shoham. 1993. Agent-Oriented Programming. *Artificial Intelligence* 60, 1 (March 1993), 51–92. doi:10.1016/0004-3702(93)90034-9
- Dominik Sobania, Martin Briesch, and Franz Rothlauf. 2022. Choose Your Programming Copilot: A Comparison of the Program Synthesis Performance of Github Copilot and Genetic Programming. In *Genetic and Evolutionary Computation Conference (GECCO)*. ACM, Boston Massachusetts, 1019–1027. doi:10.1145/3512290.3528700
- Dominik Sobania, Dirk Schweim, and Franz Rothlauf. 2021. Recent Developments in Program Synthesis with Evolutionary Algorithms. doi:10.48550/arXiv.2108.12227 arXiv:2108.12227 [cs]
- Kenneth O. Stanley and Joel Lehman. 2015. *Why Greatness Cannot Be Planned: The Myth of the Objective*. Springer, Cham Heidelberg New York Dordrecht London.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. doi:10.48550/arXiv.2302.13971 arXiv:2302.13971 [cs]
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *International Conference on Neural Information Processing Systems (NeurIPS)*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.). Curran Associates, Inc., 5998–6008. arXiv:1706.03762 [cs]
- Peter A. Whigham. 1996. Search Bias, Language Bias, and Genetic Programming. *Genetic Programming* 1996 (1996), 230–237.
- Chunqiu Steven Xia and Lingming Zhang. 2023. Conversational Automated Program Repair. doi:10.48550/arXiv.2301.13246 arXiv:2301.13246 [cs]
- Feng Xia, Jiaying Liu, Hansong Nie, Yonghao Fu, Liangtian Wan, and Xiangjie Kong. 2020. Random Walks: A Review of Algorithms and Applications. *IEEE Transactions on Emerging Topics in Computational Intelligence* 4, 2 (April 2020), 95–107. doi:10.1109/TETCI.2019.2952908
- Frank F. Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. 2022. A Systematic Evaluation of Large Language Models of Code. In *International Symposium on Machine Programming (MAPS 2022)*. ACM, New York, NY, USA, 1–10. doi:10.1145/3520312.3534862
- Weixiang Yan, Yuchen Tian, Yunzhe Li, Qian Chen, and Wen Wang. 2023. CodeTransOcean: A Comprehensive Multilingual Benchmark for Code Translation. In *Findings of the Association for Computational Linguistics (EMNLP)*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 5067–5089. doi:10.18653/v1/2023.findings-emnlp.337
- Zihan Yu, Liang He, Zhen Wu, Xinyu Dai, and Jiajun Chen. 2023. Towards Better Chain-of-Thought Prompting Strategies: A Survey. doi:10.48550/arXiv.2310.04959 arXiv:2310.04959 [cs]

- Daoguang Zan, Bei Chen, Dejian Yang, Zeqi Lin, Minsu Kim, Bei Guan, Yongji Wang, Weizhu Chen, and Jian-Guang Lou. 2022. CERT: Continual Pre-training on Sketches for Library-oriented Code Generation. In *International Joint Conference on Artificial Intelligence*, Vol. 3. International Joint Conferences on Artificial Intelligence, 2369–2375. doi:10.24963/ijcai.2022/329
- Daoguang Zan, Bei Chen, Fengji Zhang, Dianjie Lu, Bingchao Wu, Bei Guan, Wang Yongji, and Jian-Guang Lou. 2023. Large Language Models Meet NL2Code: A Survey. In *Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 7443–7464. doi:10.18653/v1/2023.acl-long.411
- Maksym Zavershynskiy, Alex Skidanov, and Illia Polosukhin. 2018. NAPS: Natural Program Synthesis Dataset. doi:10.48550/arXiv.1807.0316 arXiv:1807.0316
- Eric Zelikman, Eliana Lorch, Lester Mackey, and Adam Tauman Kalai. 2023. Self-Taught Optimizer (STOP): Recursively Self-Improving Code Generation. In *Annual Workshop on Optimization for Machine Learning (OPT)*. openreview.net, 1–44.
- Kechi Zhang, Zhuo Li, Jia Li, Ge Li, and Zhi Jin. 2023b. Self-Edit: Fault-Aware Code Editor for Code Generation. In *Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 769–787. doi:10.18653/v1/2023.acl-long.45
- Quanjun Zhang, Chunrong Fang, Yang Xie, YuXiang Ma, Weisong Sun, Yun Yang, and Zhenyu Chen. 2024b. A Systematic Literature Review on Large Language Models for Automated Program Repair. doi:10.48550/arXiv.2405.01466 arXiv:2405.01466 [cs]
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, and Guoyin Wang. 2024a. Instruction Tuning for Large Language Models: A Survey. arXiv:2308.10792 [cs]
- Yuwei Zhang, Zhi Jin, Ying Xing, and Ge Li. 2023a. STEAM: Simulating the InTeractive BEhavior of ProgrAMmers for Automatic Bug Fixing. doi:10.48550/arXiv.2308.14460 arXiv:2308.14460 [cs]
- Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Lei Shen, Zihan Wang, Andi Wang, Yang Li, Teng Su, Zhilin Yang, and Jie Tang. 2023. CodeGeeX: A Pre-Trained Model for Code Generation with Multilingual Benchmarking on HumanEval-X. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*. ACM, Long Beach CA USA, 5673–5684. doi:10.1145/3580305.3599790

accepted 19 February 2025