Reinforcement Learning for Mutation Operator Selection in Automated Program Repair

Carol Hanna[1], Aymeric Blot[2], Justyna Petke[1]

[1]University College London, UK, [2]Université du Littoral Côte d'Opale, France

Abstract—**Automated program repair techniques aim to aid software developers with the challenging task of fixing bugs. In heuristic-based program repair, a search space of program variants is created by applying mutation operations on the source code to find potential patches for bugs. Most commonly, every selection of a mutation operator during search is performed uniformly at random. The inefficiency of this critical step in the search creates many variants that do not compile or break intended functionality, wasting considerable resources as a result.**

In this paper, we address this issue and propose a reinforcement learning-based approach to optimise the selection of mutation operators in heuristic-based program repair. Our solution is programming language, granularity-level, and search strategy agnostic and allows for easy augmentation into existing heuristic-based repair tools. We conduct extensive experimentation on four operator selection techniques, two reward types, two credit assignment strategies, two integration methods, and three sets of mutation operators using 22,300 independent repair attempts. We evaluate our approach on 353 real-world bugs from the Defects4J benchmark. Results show that the epsilon-greedy multi-armed bandit algorithm with average credit assignment is best for mutation operator selection. Our approach exhibits a 17.3% improvement upon the baseline, by generating patches for 9 additional bugs for a total of 61 patched bugs in the Defects4J benchmark.

Index Terms—**Reinforcement Learning, Program Repair**

I. INTRODUCTION

FIXING bugs remains a largely manual and tedious process that often requires more time than is available to software developers [1], [2]. With the fast-evolving industry, this leads to the deployment of bug-prone products in an attempt to meet release deadlines [3], [4]. Research in the area of automated program repair (APR) aims to address this issue by automating the process of finding suitable patches for bugs in software systems.

The classification of APR approaches differs in the literature. According to a recent survey [5], the proposed approaches in the field of APR span across three main areas: constraint-based, learning-based, and heuristic-based techniques. Constraint-based approaches use the semantics of the buggy program to produce a constraint, then synthesise repairs that satisfy it [6]. Learning-based end-to-end repair techniques predict patches for buggy programs by learning features of the faulty code sections as well as their correct (developer-written) fixes [7]. Among APR approaches, heuristic-based APR has seen the earliest and thus far most industrial uptake: Heuristic-based APR was first applied to software during the development of a management system for a medical application [8]. The earliest deployment for the purpose of bug fix automation at scale was Meta's SapFix [9], which combines various approaches including a single generation search over the space of higher order mutants. More recently, Bloomberg applied heuristic-based APR within their infrastructure to target simple and frequently occurring bugs [10].

Heuristic-based APR uses search strategies, such as genetic programming [11] or local search [12], to navigate the space of software variants. These approaches require a test oracle for the buggy program to assess the correctness of the software variants. The oracle contains passing test cases and at least one failing test case that demonstrates the bug. As a first step, fault localisation techniques are employed to pinpoint the suspicious code sections. A program variant is created through the selection of a suspicious code location as well as a mutation operator to be applied at the location. The most common mutation operators are deletion, insertion, and replacement of fragments of code. A large population of candidate variants is produced through this process and validated against the test oracle to determine their fitness. The process is repeated until a suitable patch that passes all of the test cases is found.

In heuristic-based APR, the choice between the mutation operators in most approaches is random, causing the search to over-explore some operators that empirically lead to more failing software variants. As a result, more incorrect patches are produced and more resources are wasted which hinders the efficacy and efficiency of the technique. Smigielska et al. [13] improve on this by implementing a uniform selection strategy where the probability for each operator to get selected is proportional to the size of its respective search space. Alternatively, Soto et al. [14] use developer bug fix history to determine a fairer distribution of operators.

A variety of mutation operators is necessary for a diversified search space. However, if the application of the mutation operators throughout the search process is not optimised, this can create many incompatible variants which will delay or even prevent finding a correct patch. While previous work [13], [14] does address this concern, the probabilities for each operator remain predetermined and fixed throughout the search. Therefore, the online feedback of the search is not used to further tune the probabilities for the selection of each operator. Moreover, the optimal probabilities for the different operators are likely different for every software. We aim to tune the operators accordingly so that the search behaves at its best for the software that is currently being processed. Given that evolutionary computation is often used in APR as the search strategy, we drew inspiration from there. We found that augmenting a reinforcement learning agent to optimise the selection process of the mutation operators has been im-

plemented for evolutionary algorithms [15]–[21], differential evolution [22], and evolutionary programming [23]. To the best of our knowledge, this is yet to be applied for APR.

Our paper aims to tackle this gap in the literature by proposing a reinforcement learning based technique to inform the selection of mutation operators in heuristic-based APR to improve both the efficiency and efficacy of the process. We implement two credit assignment techniques: average and exponential recency-weighted average credit assignment and four reinforcement learning-based mutation operator selection strategies: probability matching, adaptive pursuit, epsilon-greedy multi-armed bandit, and upper confidence bound. We also experiment with two types of rewards: raw fitness values and fitness values relative to the parent. The activation point of the technique during the search process is also configured to best fit the context. We analyse the efficacy and efficiency of the different techniques to find the best strategy for guiding the search process in heuristic-based APR.

We find that epsilon-greedy mutation operator selection strategy is the most effective at generating patches for bugs while probability matching is the most efficient. Our results show that the RL-guided mutation operator selection approach improves on the standard random selection strategy by patching 9 additional bugs from a popular APR benchmark. Our contributions include:

1) Novel reinforcement learning-based approach for mutation selection in heuristic-based program repair that outperforms the standard random selection techniques;
2) Extensible implementation of multiple selection techniques, credit assignment policies, reward calculations, and integration strategies in a state-of-the-art heuristic-based APR tool;
3) Evaluation on 353 bugs from a popular APR benchmark.

II. BACKGROUND

Next, we discuss the key ideas in heuristic-based automated program repair (APR) and reinforcement learning (RL) that are relevant as background for our proposed approach.

A. Heuristic-based Program Repair

Heuristic-based program repair navigates the search space of software variants in order to find software that fixes a given bug [5]. The method is usually composed of three main stages. Given a buggy input program and its test oracle, the suspicious statements in the code are detected through the first fault localisation step. In the second stage, program variants are created by selecting mutation operators and applying them to the target locations to mutate them. Finally, the variants are all verified against the provided test oracle. A fitness function measures the viability of each variant based on the number of passing test cases and is used to guide the search strategy.

The two most commonly employed search strategies are local search [24] and genetic programming [25]. Local search is a heuristic approach to optimisation problems in which small and random local modifications are iteratively applied to find better solutions [12]. As for genetic programming, a population of candidate programs is created through evolving the source code using steps inspired by biological evolution [11].

Our approach is fully tool-agnostic and can be augmented in any heuristic-based tool regardless of the search strategy that it implements. However, our review of APR tools in the literature (additional details in Section V-A) revealed that genetic programming is the more common strategy in state-of-the-art tools. Thus, we provide more details on this search strategy in this section.

Traditionally, the genetic programming search strategy is compromised of 3 evolutionary steps: *selection*, *crossover*, and *mutation*. In the *selection* stage, the variants with the highest fitness in the population are selected. From there, the *crossover* transformation is applied to the population of variants in pairs. The result of crossing over two variants is a new offspring variant. In the *mutation* step, individual variants are mutated to create a second set of offspring variants.

Most commonly, program statement is mutated through a random choice between one of the following three options: deleting the statement entirely, inserting another statement chosen from elsewhere in the code after the current statement, or replacing the current statement with a different one.

More recent heuristic-based APR research improve upon this by using developer behaviour information during the search [26], improving the fitness function [27], addressing issues of overfitting [28], and proposing different mutation operators [29]. However, the choice between the selection of the mutation options remains predominantly random with very few papers addressing improvements on this choice [13], [14]. We propose the augmentation of reinforcement learning strategies to optimise the selection of mutation operators during the search process.

B. Reinforcement Learning

In reinforcement learning (RL) [30], sequential decision making is applied to problems where resources must be allocated between a finite number of conflicting choices. Given N different action choices where only partial historical information is known about each, the idea is to predict the best possible series of actions in order to maximise the cumulative outcome. Ideally, over time as more choices are made, the overall gain is optimised.

Balancing the exploration-exploitation tradeoff is at the core of RL techniques. *Exploitation* of actions that the algorithm has sufficient historical information about allows for a reliable estimation of the expected rewards. However, *exploration* is necessary in order to collect this information and optimise the selection of future actions. Choosing to exclusively exploit or explore actions will yield sub-optimal results.

RL-based approaches have been used to improve the mutation operator selection stage in evolutionary computation in the past [15], [16], [18], [31]. In these work, the integration of such techniques is shown to outperform standard methods and optimise the selection of operators throughout the search. In this section, we detail the state-of-the-art RL algorithms that have been applied in this context as well as various variants for reward calculation and quality estimation from the literature.

1) Rewards: Multiple variations have been suggested for the calculation of the reward values that the actions receive

in the RL algorithm. This can be the raw fitness itself, but usually it is the fitness improvement in reference to another individual. The reference individual can either be the offspring's ancestor [32], the currently most fit individual in the population [33], or even the median fitness of the population [34]. With this approach, an offspring that does not show an improvement is disregarded. In addition to just reflecting the fitness, the compass [35] technique also takes into account the population's diversity.

2) Estimating Action Qualities: In this section, we discuss two techniques for estimating the qualities of actions [30]. In the first technique, the estimated quality for an action A at time step $t+1$ is the average of the rewards that the action has received until time step t (Equation 1). This approach is appropriate for stationary environments where the distribution of the rewards is not altered over time. Alternatively, exponential recency-weighted average accounts for dynamic environments where the reward probabilities are non-stationary. With this approach, the estimated quality for an action A at time step $t+1$ is calculated according to Equation 2 such that $Q_A(0) = 1$ and $R_A(t)$ and $Q_A(t)$ are the reward and predicted quality for action A at time step t respectively. This equation introduces hyper-parameter α, which controls the learning rate.

$$Q_A(t+1) = \frac{\sum_{i=1}^{i=t} R_A(i)}{t} \qquad (1)$$

$$Q_A(t+1) = Q_A(t) + \alpha[R_A(t) - Q_A(t)] \qquad (2)$$

3) Reinforcement learning algorithms: We present the state-of-the-art reinforcement learning techniques used for evolutionary strategies in the literature: probability matching, adaptive pursuit, and bandit-based algorithms.

Probability matching (sometimes also referred to as Thompson Sampling or Posterior Sampling) is a simple technique that is widely observed in biological processes. With this approach, the probability of taking an action mirrors the probability of getting rewarded for it. Specific implementation details differ in the literature, however, in this paper we base the probability matching technique on the work by Thierens et al. [18] as their work applies to operator selection in genetic algorithms, often used in heuristic-based APR. To ensure that all actions continue to be explored throughout the search, they introduce a minimum probability value, $P_{\min}$. Given $P_{\min}$ and N actions, the probability for an action A at time t with the estimated quality $Q_A(t)$ gets calculated according to Equation 3.

$$P_A(t) = P_{\min} + (1 - N * P_{\min}) \frac{Q_A(t)}{\sum_{i=1}^{i=N} Q_i(t)} \qquad (3)$$

Adaptive pursuit (AP) [18] aims to solve the probability matching slow convergence issue by introducing a "winner takes all" strategy. As with PM, the actions are selected according to their probabilities. However, the probability of an action A at time step $t+1$ is calculated according to Equation 4. The action M is the action with highest quality value. This algorithm introduces a new hyper-parameter β and a maximum probability value $P_{\max}$.

$$P_A(t+1) = \begin{cases} P_A(t) + \beta[P_{\max} - P_A(t)] & \text{if A = M} \\ P_A(t) + \beta[P_{\min} - P_A(t)] & \text{otherwise} \end{cases} \qquad (4)$$

Bandit algorithms are the state-of-the-art approach to mutation operator selection in genetic algorithms. In this paper, we experiment with two different variations: epsilon-greedy and upper confidence bound which have been shown successful in related work [15], [16], [31].

The greedy bandit algorithm [30] tends to favour exploitation of the best observed action thus far. Often as a result, the current best observed action keeps getting re-selected which prevents the algorithm from converging towards the best action. The epsilon-greedy [30] improves on this by adding the option of exploration with a constant probability. With the epsilon-greedy algorithm, a random number between 0 and 1 is generated. If the generated number is less than ε, then an action is selected at random. Otherwise, the action with the highest quality (predicted reward) is greedily selected. Within bandit-based approaches the different action choices are referred to as *arms*.

The upper confidence bound (UCB) [36] has been proven to optimise the cumulative gain and convergence rate. Given N actions, Equation 5 presents the UCB action selection at time t for action A where $Q_A(t)$ is the action's estimated quality, E is a constant balancing exploration-exploitation tradeoff, and n_A is the number of times that arm A has been played. The chosen arm at time t is the one that maximises the result of Formula 5.

$$Q_A(t) + E * \frac{\sqrt{\log \sum_{j=1}^{j=N} n_j}}{n_A} \qquad (5)$$

III. APPROACH

In this work, we use reinforcement learning at the mutation operator selection stage of heuristic-based APR.

A. RL-Guided Mutation Operator Selection

Our approach uses reinforcement learning to optimise the mutation operator selection process in heuristic-based APR. For each mutation operator, we track the fitness of the variants that it creates throughout the repair attempt. The fitness that we use to guide the reward update in our approach is the same fitness that guides the given APR search process. We use this to inform the probability of selection for each operator while balancing the exploration-exploitation tradeoff.

We implement the four most commonly used strategies in heuristic numerical optimisation problems: selection based on probability matching, adaptive pursuit, epsilon-greedy bandits, and upper confidence bound bandits. We chose these operator selection strategies, in particular, as these are simple, yet shown successful in the context of genetic algorithms [15], [18] — a search strategy frequently used in heuristic-based APR. We are also the first to apply these in this context, thus it is yet unclear which ones, or variants thereof, would be successful. For each of the above techniques, we experiment with an operator's reward as the raw fitness value or the fitness value relative to the parent and its credit being either based on exponential recency-weighted average or the average of the rewards (see Sections II-B1 and II-B2). The variation in the credit assignment technique will inform whether the search process is stationary or dynamic.

This technique can be augmented into any heuristic-based APR tool that uses mutation operators. This is because we do not limit the number of actions (mutation operator options) to any specific value. We do, however, explore tuning the approach depending on the number of operators. With this, we are able to inform the community of how to best apply the technique depending on the number of mutation operators that the underlying tool implements.

B. Credit Assignment

When thinking about credit assignment in adaptive operator selection there are four main choices that must be made [16].

Firstly, **which type of reward will be used?** As this is the first paper to attempt reinforcement learning-based mutation operator selection, we experimented with the two most simple approaches: the raw fitness value of the individual and the relative fitness of the individual in reference to its direct parent. For an action A at time step t, we calculate the relative fitness $R'_A(t)$ according to Equation 6. This representation allows us to avoid negative numbers while still accounting for the exact level of improvement or deterioration in the offspring.

$$R'_A(t) = \begin{cases} \frac{R_A(t)}{R_A(t-1)} & \text{if } R_A(t-1) \neq 0, \\ R_A(t) & \text{otherwise} \end{cases} \quad (6)$$

Secondly, **which operators to reward during the search?** Most commonly, only the operator that was used to create the offspring gets a reward. However, alternatives suggest rewarding the older ancestors of the offspring as well since they also had a part in its creation [16]. Later work suggested that rewarding older individual's operators in that way is less effective [37]. Therefore, we only award the direct operator that was applied to create the offspring.

Thirdly, **how will the rewards accumulate throughout the search?** For each operator type, we can simply assign it its most recently received reward. We can also average all of the rewards that it has ever received. Since older rewards might be less relevant, Thierens et al. [18] assign a window with a fixed size N where only the last N rewards contribute to the average. An improvement was made on this by Fialho et al. [16] to use an extreme-based credit assignment which is based on the assumption that infrequent large improvements in the fitness are more significant that frequent smaller ones [38]. We experiment with two simple strategies in this paper: using the average reward or the exponential recency-weighted average (see Section II-B2). Using these strategies, we are able to calculate the estimated quality of each mutation operator. Experimenting with both of these strategies will help us better understand how stationary/dynamic the search process is in APR.

Finally, **at which points in the search should the credit assignment occur?** An option would be to do so immediately after each operator is applied. This would mean that every time an operator is selected to create a variant, the variant's fitness would be evaluated and used to update the credit of the operator. Another approach would be to do this in batches, i.e., after a constant or variable number of variants. We experiment with both options.

IV. RESEARCH QUESTIONS

In heuristic-based APR the variety of mutation operators increases the chances of patching bugs [29]. However, this is at the cost of a much larger search space, and a slower search as a result. We propose an RL-based mutation operator selection solution to better guide the search and mitigate this slowdown.

RQ1: Which credit assignment technique is suitable for mutation selection in heuristic-based APR?

The choice of RL credit assignment strategy depends on whether the expected reward values can change over the course of the search. We compare the performance of two main stationary and dynamic methods to determine the environment type in heuristic-based APR.

RQ2: Which mutation operator selection strategy is best in heuristic-based APR?

RL has a long history of various selection strategies. We investigate the efficacy and efficiency of four different strategies that have been shown to be effective in similar contexts and compare with the standard operator selection technique in heuristic-based APR.

RQ3: How does efficacy/efficiency of RL-based mutation selection in heuristic-based APR change with increase in the number of mutation operators?

The addition of more fine-grained mutation operators that target specific bugs could improve efficacy but impede efficiency of traditional heuristic-based APR by creating a larger search space. We investigate whether this issue can be fixed by our RL-based approach.

RQ4: To what extent does RL-based mutation operator selection improve the bug fixing ability of heuristic-based APR?

Through online feedback of the search process, we hypothesise that the probability distribution of the mutation operators will be tuned to favour the more effective ones for the given buggy program, thus improving performance.

V. METHODOLOGY

Next we describe our methodology for answering RQs. First, we conduct a pre-study, where we fine-tune learning rate α (see Equation 2) values for each of the four selection strategies. We use 10% of the dataset for these preliminary experiments. While the optimal dataset size for parameter tuning remains an open problem, running these experiments on the full dataset would be too costly and would also risk overfitting to the benchmark.

To answer RQ1, we run each RL algorithm with both of the credit assignment techniques on a variety of real-world bugs from an APR benchmark. For each operator selection strategy, we analyse its efficiency and efficacy when combined with the various credit assignment techniques. We then compare each of the optimal combinations of operator selection and credit assignment identified from RQ1 with the baseline, i.e., without using RL. To answer RQ3, we explore the effect of changing the number of arms on the efficacy and efficiency rates of the proposed approach.

For the preliminary experiments and the first three research questions, we use the raw fitness value as the reward as

it is the most simple strategy. We chose to assign credit to the operators at the end of each generation in batch at this stage to avoid too frequent updates that might overfit. Furthermore, most heuristic-based APR work relies on genetic programming which evolves populations of variants, already providing a natural division into batches.

For RQ4, we use the optimal settings of credit assignment and operator selection. In answering this question we add experimentation for the reward types and integration strategy and compare our proposed approach with the baseline. We assess the quality of the patches by running the patched code on a second set of evaluation held-out test suites, which is a common strategy for patch evaluation in heuristic-based APR [14], [26], [39] (further details on patch correctness evaluation in Section V-C).

We use the same measures for efficacy and efficiency to answer all of the research questions. Efficacy is measured by assessing the number of bugs for which a patch was generated, the frequency of such successful repair attempts for each bug, as well as patch quality, measured on held-out evaluation test suites. As for the efficiency measure we use the median and average numbers of individuals evaluated until a test-suite adequate patch is found.

A. Tool

The field of APR has seen increasing growth in the last few years, with now tens of tools available[1] to automate the task of bug fixing [40]. After a thorough review of the heuristic-based APR literature, we found that JaRFly [39] would be the best fit for our implementation. JaRFly is a novel open-source framework for search-based APR. It implements all three statement-level mutations that are used in GenProg [41] and TrpAutoRepair [42]: append, delete, and replace. It also implements 18 PAR templates [29] such as null checker and object initializer. JaRFly is well documented and very modular, making it an excellent candidate for extensibility.

When choosing a mutation operator to apply, JaRFly currently implements two options. The default is a uniform probability across all available mutations. However, they also allow for a probability distribution based on a probabilistic model. The model was created by mining open source repositories and analysing the frequency in which developers apply each of the available mutations [14]. We chose to use the uniform distribution approach as the baseline because the results that were reported show that operator selection informed by the probabilistic model generated a smaller number of patches than the uniform distribution [26].

JaRFly Modifications: We implement a third reinforcement based option in JaRFly that selects the mutation operator based on their current saved probabilities. These probabilities can be calculated according to any combination between the four operator selection strategies, the two credit assignment strategies, the two reward calculation techniques, and the two activation point variations. We implement each algorithm separately and depending on the mode with which the repair

---

[1]Our full tool comparison report will be made publicly available upon acceptance of the paper.

attempt was launched, the correct algorithm gets activated to set the probability values in the search process. We used the JaRFly fitness value as is in our implementation. The fitness value in JaRFly is calculated on the basis of the number of passing tests in the provided test suite.

Running the experiments revealed some bugs in the underlying code. We were able to locate five bugs that were causing uncaught exceptions in the execution of the code. These bugs are likely due to the changing version of external libraries that JaRFly depends on, e.g., in one of the cases this caused a change in the treatment of two-dimensional array identifiers. Patches for these bugs were added and can be found in our artefact. While JaRFly does implement 18 PAR templates, currently the most recent version has a bug that throws a NullPointerException within some scenarios of the application when 3 of the PAR templates generate variants, namely Parameter Replacer, Adder, and Remover. Therefore, we exclude these 3 PAR templates from all of our experiments.

B. Benchmark

We considered various benchmarks for evaluation through the overview that [43] provides. However, we chose Defects4J as it is currently the most comprehensive and popular dataset for evaluating Java APR tools [25], [44] which allows us to be able to compare our results with the state-of-the-art. Additionally, Defects4J was used to evaluate the tool that we chose to implement our approach in, JaRFly [39], which allows for a direct comparison.

Our paper uses the most recent (2.0.0) version of Defects4J [45] for evaluation. Defects4J is a collection of 835 real-word Java bugs from open-source repositories. Each defect in the dataset includes the defective version of code, its developer-fixed version, as well as test oracles. There are two types of tests that accompany every Defects4J bug: developer written tests as well as the infrastructure for generating automated tests using EvoSuite [46] or Randoop [47].

To fairly evaluate our approach, we used the exact Defects4J bugs that were used for the baseline JaRFly program. At the time that the JaRFly framework was evaluated, version 1.1.0 was used which consisted of 395 bugs. Given that the Mockito project was excluded in the JaRFly paper, we did not consider it. Finally, we omitted the 4 bugs that have since been deprecated. We evaluated our approach on the 353 remaining bugs. The JaRFly paper reports to be able to successfully repair 49 out of these bugs using the GenProg algorithm and an additional 15 active bugs using PAR templates. The breakdown of the bugs is presented in Table I.

We do not extend to the full dataset of 835 bugs from Defects4J. Doing so would have been an enormous waste of computational resources as we would have had to run the baseline for comparison as well. We already conduct a very large-scale evaluation in this paper as we launch over 22,300 repair attempts using the latest version of this state-of-the-art benchmark. However, this extension is indeed possible as we intend to make all of our scripts and replication instructions public upon acceptance.

TABLE I
BREAKDOWN OF THE DEFECTS4J BUGS THAT WERE USED IN THE JARFLY
STUDY [39] AND PATCHED USING THEIR TOOL.

Project	Currently Active Bugs	Patched with GenProg	Patched with PAR
JFreeChart	26	6	7
Closure Compiler	131	15	20
Apache Commons Lang	64	9	12
Apache Commons Math	106	18	23
Joda-Time	26	1	2
Total	353	49	64

TABLE II
HYPER-PARAMETER VALUES FOR THE PM, AP, EPSILON-GREEDY, AND
UCB ALGORITHMS; N: NO. OF MUTATIONS.

Name	Associated Algorithms	Value
P_{min}	PM, AP	$\frac{1}{2N}$
P_{max}	PM, AP	$1 - (N - 1) * P_{min}$
β	AP	0.8
ε	epsilon-greedy	0.2
E	UCB	10

C. Experimental Set Up

To answer our research questions, we ran the latest version of JaRFly. We used the default uniform selection strategy to get a baseline, before enabling the implemented RL-approaches.

General search settings: For each bug, 20 repair attempts were launched independently, to account for the heuristic nature of the underlying genetic algorithm. The search parameters were set to the same values as specified in the original paper of JaRFly [39]. As such, each repair attempt was bound to 10 generations with a population size of 40. Since we were not able to run our experiments in the identical environment of JaRFly in terms of the hardware specifications, we eliminated the 4-hour timeout per repair attempt and ran all experiments to completion of the 10 generations instead. JaRFly provides a replication package [48] which includes the bugs that were successfully repaired as well as scripts for launching the tool to repair a specific bug. These scripts were used both for replicating their results as well as for conducting our experiments.

Hyper-parameter settings: The various reinforcement learning algorithms that we experiment with in this paper introduce hyper-parameters. Table II depicts the different parameters that are associated with the different algorithms (see Section II-B for more details on the parameters and formulas). The values of the parameters in the table are those used in the literature [16], [18].

As for the learning rate α (from the exponential recency-weighted average formula — see Section II-B2), we conducted preliminary experiments to tune its value for each experiment type on a subset of bugs from the Defects4J dataset. The subset includes 5 diverse bugs which constitute just over 10% of the 49 bugs that the baseline in JaRFly successfully repaired. Each of these five bugs was the first in lexicographical order from those that the baseline in JaRFly correctly patched from each of the five projects (Table I). The five bugs we used in the experiment were: Chart 1, Closure 102, Lang 10, Math 18, Time 19. For each of the four operator selection strategies, we conducted four experiments using 4 learning rate values for a total of 16 preliminary experiments. The value that we found in the literature for the learning rate in this context is 0.8 [18]. We experimented further with the learning rates 0.2, 0.4, and 0.6. Each repair attempt was repeated 20 times, to account for the heuristic

nature of JaRFly's genetic programming algorithm. Overall, in this preliminary study, we conducted 1600 repair runs.

RQ1 and RQ2: To answer the first 2 research questions, we activated our modified version of JaRFly with the 3 simple mutation operators from GenProg: insert, delete, replace (see Section II-A) as these are the most commonly used and would provide a foundation to build on for evaluation. The evaluation for these 2 RQs was on the 49 bugs that the original JaRFly paper reports to repair using the GenProg algorithm. As explained in Section V, at this stage we also used the raw fitness value as the reward and assigned credits to the operators every generation.

RQ3: The experimentation for RQ3 was extended to include the 15 PAR templates that are implemented in JaRFly for a total of 18 mutation operators. Therefore, evaluation for RQ3 was conducted on all of the bugs that were repaired with either GenProg or PAR in JaRFly for a total of 64 bugs.

RQ4: Given that we update weights after each generation, we note that it might not be enough time to evaluate all 18 mutation operators and learn optimal rewards. We also note that several of the 18 operators naturally belong into groups, e.g., some manipulate parameter values, while others are concerned with bound checks. Therefore in RQ4, we activate the RL algorithm in the search using 3 different sets of arms for further evaluation:

1) 3 arms: 3 GenProg mutations
2) 18 arms: 3 GenProg mutations + 15 PAR templates
3) 7 arms: 3 GenProg mutations + 15 PAR templates aggregated into four groups.

The third set of arms groups together PAR templates that pertain to bugs within a similar category. We divide the PAR templates into four groups: functions and expressions (FunRep, ExpRep, ExpAdd, ExpRem), bound and null checks (NullCheck, RangeCheck, SizeCheck, LBoundSet, UBoundSet, OffByOne), casting and initialisation (CastCheck, ObjInit, CasterMut, CasteeMut), and multi-line edits (SeqExch). The mutations are chosen and applied as normal in the search algorithm. However, each operator's credit gets assigned to its respective group. From there, an operator's probability for selection is what was assigned to the arm of the group that the operator belongs to.

We experimented with two reward types: raw fitness and fitness relative to parent as well as two modes for assigning credit: credit assignment after every generation and after every mutation. We extend the evaluation to 353 bugs from the Defects4J benchmark.

TABLE III
PRELIMINARY EXPERIMENT RESULTS. THE TABLE PRESENTS THE
EXPERIMENT TYPE, LEARNING RATE (α), PERCENTAGE OF REPAIR
ATTEMPTS THAT PRODUCED A TEST-SUITE ADEQUATE PATCH, NUMBER OF
BUGS PATCHED, AVERAGE/MEDIAN NUMBER OF VARIANTS THAT WERE
EVALUATED UNTIL A PATCH WAS FOUND.

Repair Type	α	Success Rate	Bugs Patched (/5)	Avg. Variant	Med. Variant
Baseline	-	30%	4	177	171.5
PM	0.2	33%	3	154	131
	0.4	27%	3	141	80
	0.6	28%	3	178	93
	0.8	**34%**	**4**	**136**	**138**
UCB	0.2	36%	4	154	136
	0.4	34%	4	145	103.5
	0.6	36%	4	139	119.5
	0.8	**38%**	**4**	**155**	**117**
AP	**0.2**	**35%**	**4**	**127**	**79**
	0.4	33%	3	130	78
	0.6	28%	3	141	77
	0.8	31%	3	142	79
epsilon-greedy	0.2	24%	4	136	123
	0.4	**28%**	**4**	**133**	**97.5**
	0.6	30%	3	156	140
	0.8	18%	3	103	90

We evaluate the correctness of the patches produced by our approach according to the same methodology used in JaRFly [39] for a fair comparison. The authors of JaRFly use two versions of the automated test generation tool Evo-Suite [46] to create held-out evaluation test suites. They created evaluation test suites for 71 defects from Defect4J that passed their criteria for statement coverage which are publicly available [48]. For each of the defect patches that we generated, if the defect had an evaluation test suite, we applied the patch to the defect and executed both sets of tests (v1.0.3 and v1.0.6) on the patched code. From there, we compared the quality scores of patches. The quality score for a patch is: $\frac{T_{\text{Pass}}}{T_{\text{Total}}}$ where T_{Pass} is the number of tests that passed and T_{Total} is the total number of test cases.

VI. RESULTS

We ran the experiment on a cluster of 1028 compute nodes. The experiments were assigned to a subset of these nodes with varying specifications. RAM size ranged from 16GB to 375GB and local SSD disk space ranged from 80GB to 780GB. All CPUs were from the Intel Xeon processors range from a variety of generations with the number of cores ranging from 4 to 48. This did not affect the results as we ran all repair attempts to completion of the 10 generations. On average, the computation time for the experiments ranged from 45 minutes to 1.5 hours per repair attempt. However, for some it reached upwards of 12 hours. To run the preliminary experiments and answer all RQs, overall 22,300 repair attempts were run.

A. Preliminary Experiments

The results of the parameter tuning preliminary experiments described in Section V-C are depicted in Table III. We used these experiments to determine the optimal learning rate value for credit assignment using the exponential recency-weighted average for each operator selection strategy independently.

All 16 of the experiments that we set up at this stage had both higher successful repair attempts than the baseline and lower medians for the repair variant number. For each operator selection strategy, we chose the learning rate value that maximised the number of successful repair attempts. Given this criteria, for the probability matching and UCB algorithms we found that the optimal α value is 0.8. As for adaptive pursuit, it is 0.2 and for epsilon-greedy it is 0.4. These are the experimental set ups that we proceeded with.

B. RQ1: Best Credit Assignment Technique

The results of the experiments with average and exponential recency-weighted average credit assignment are presented in Table IV. Interestingly, we can see that in all of the experiments the success rate of the repair attempts is higher with the average credit assignment than with the exponential recency-weighted average credit assignment. Moreover, we can see that in all of the experiments, except for AP, the number of unique bugs that were fixed is higher as well. Therefore, average credit assignment is best suited for PM, UCB, and the epsilon-greedy algorithms whereas for AP, the exponential recency-weighted average is better.

We investigated the discrepancy for the unique number of bugs patched in the adaptive pursuit results and found that there were 3 bugs that were successfully patched in combination with the exponential recency-weighted average credit assignment that were not patched in combination with the average credit assignment. However, for these bugs, only 1 or 2 repair attempts successfully generated patches. i.e., their patches are harder to find in the search space. AP with the average credit assignment might have failed to produce these patches due to limiting the number of repair attempts to 20 in our experiments. We can thus conclude that the search process in heuristic-based APR is stationary.

RQ1: Average credit assignment is best suited for APR. Thus, the search process in heuristic-based APR is stationary.

C. RQ2: Best Mutation Operator Selection Technique

To fairly compare with the baseline in answering RQ2, we tried to replicate the baseline results from the JaRFly paper by running the 49 bugs (Section V-B) that the authors report to generate patches for using the GenProg setting. Our results are presented in Table V. When attempting to reproduce the results, we were only able to generate a patch for 41 out of the 49 bugs. Since we used the exact search parameter settings and seeds that were used in JaRFly, we think that the discrepancy is likely due to mismatched versions (see details in Section IX). The latest versions of JaRFly and Defects4J were used in our experiments.

To answer this question, we compared all of the RL algorithms with their optimal credit assignment technique that was identified in RQ1 (Table IV). For PM, we can see that 41

Repair Type	α	Success Rate	Bugs Patched (/49)	Avg. Variant	Med. Variant
Average Credit Assignment					
PM	-	43.3%	41	99	52.5
UCB	-	46.5%	40	125	88
AP	-	45.1%	39	96	54
epsilon-greedy	-	47.9%	43	112	81
Exponential Recency-Weighted Average Credit Assignment					
PM	0.8	43.2%	38	107	52
UCB	0.8	44.1%	39	121	80
AP	0.2	43.6%	41	101	62
epsilon-greedy	0.4	44.7%	43	109	73

Mutations	Success Rate	Bugs Patched	Avg. Variant	Med. Variant
GenProg	43.9%	41/49	105	57
GenProg and PAR	35.6%	52/64	119	74

Repair Type	Credit Assignment	Success Rate	Avg. Variant	Med. Variant
Baseline	-	58.6%	105	57.5
PM	A	57.9%	97	50
UCB	A	64%	125	85.5
AP	W	58.7%	99	61
epsilon-greedy	A	63.9%	109	76

unique bugs were patched which is the same as the baseline. However, the success rate of the overall repair attempts was 0.6% lower. As for UCB, only 40 unique bugs were patched, but the overall success rate was 2.6% higher. AP had the same number of unique bug patches as the baseline with a 0.3% decrease in overall success rate. Finally, with the epsilon-greedy experiment 43 unique bugs were patched (2 more than the baseline) with a 4% higher success rate for repair attempts. Given these results, it is evident that the epsilon-greedy algorithm with average credit assignment is the best reinforcement learning-based mutation selection strategy.

Since the bugs that were patched with each experiment varied, we had to look at the intersection of the 35 common bugs that were patched in all of the experiments including the baseline in order to conduct an efficiency analysis. The results of the overall comparison are presented in Table VI. We can see that PM with average credit assignment was the only experiment to achieve lower median and average numbers than the baseline for the variant numbers of successful patches. However, it is important to note that the success rate of the experiments varied. Therefore, while the epsilon-greedy experiment does have higher median and average values for the variant number of the patches, this might be due to the higher success rate which skewed the values.

RQ2: Epsilon-greedy is the most effective mutation operator selection strategy. Probability matching is the most efficient.

D. RQ3: Additional Mutation Operators

Given that the epsilon-greedy algorithm with the average credit assignment was the best strategy in answering RQ2, we proceed with this technique for experimentation with additional operators.

We run the experiment using 15 PAR templates (recall Section V-A where we detail that the remaining 3 templates were excluded in our study) and the 3 GenProg mutations for a total of 18 mutation operators. The experiment resulted in a success rate of 35.2% and patches for 51 unique bugs. As for efficiency, the average variant number for variants that resulted

in a patch was 131 and the median 87. When comparing with the JaRFly results with 18 arms in Table V, we can see that our approach did not improve the results of the standard selection strategy. Our approach with 18 arms was able to solve one bug fewer and the success rate for repair attempts was also 0.4% lower.

RQ3: With the addition of 15 arms, reinforcement learning-guided mutation operator selection does not patch more bugs than the standard operator selection strategy.

E. RQ4: RL-Aided Mutation Operator Selection Performance

The results of RQ3 evidence that a drastic increase in the number of mutation operators lowers the efficacy and efficiency of the approach. To control the number of arms in the RL algorithm, we decided to experiment with separating the PAR templates into groups (details in Section V-C). We additionally experiment with a second type of reward which is based on the fitness of the variant in reference to its direct parent in the mutation step (recall Section II-B1). Finally, we alter the rate of learning and experiment with two more aggressive approaches for the reinforcement learning. The first approach considers recalculating the mutation operator probabilities after every mutation instead of every generation. The second recalculates the probabilities every generation, but decreases the population within each one. Since our experiments set a bound of 40 on the population size and a bound of 10 on the number of generations, we flip these values and launch repair attempts bound to a population size of 10 over 40 generations.

Table VII presents the results of the mini-experiments that we ran to test the additional settings of grouped mutations, reward types, and aggressive learning. These experiments were on the same sample of five bugs that was used in the

TABLE VII
RQ4. RESULTS OF MINI-EXPERIMENTS ON THE SAMPLE OF 5 BUGS FOR 3, 7, AND 18 ARMS WITH 10 OR 40 GENERATIONS USING RAW AND RELATIVE FITNESS VALUES AS THE REWARDS.

Arms	Generations	Activation Point	Reward Type	Success Rate	Bugs Patched	Avg. Variant	Med. Variant
3	10	Every Generation	Raw	32%	4	176	123
	10	Every Generation	Relative	32%	4	218	152
	10	Every Mutation	Raw	0%	0	-	-
	40	Every Generation	Raw	16%	3	158	120
18	10	Every Generation	Raw	31%	3	116	92
	10	Every Generation	Relative	31%	4	219	133
	10	Every Mutation	Raw	0%	0	-	-
	40	Every Generation	Raw	15%	3	175	129
7	**10**	**Every Generation**	**Raw**	**43%**	**4**	**145**	**108.5**
	10	**Every Generation**	**Relative**	**43%**	**4**	**229**	**186**
	10	Every Mutation	Raw	0%	0	-	-
	40	Every Generation	Raw	18%	3	178	101.5

TABLE VIII
RQ4. RESULTS FOR THE EPSILON-GREEDY OPERATOR SELECTION ALGORITHM WITH 7 ARMS (GENPROG AND GROUPED PAR MUTATIONS) WITH RAW/RELATIVE FITNESS RUN ON THE 64 BUGS.

Reward Type	Success Rate	Bugs Patched	Avg. Variant	Med. Variant
Raw	34.6%	48	125	82
Relative	38.1%	51	159	113

TABLE IX
RQ4. COMPARISON OF QUALITY VALUES, AND PERCENTAGE OF PATCHES WITH 100% QUALITY BETWEEN OUR APPROACH AND THE QUALITY RESULTS PRESENTED IN THE JARFLY PAPER [39].

Repair	Min.	Mean	Median	Max.	100% Quality
JaRFly GenProg	64.8%	95.7%	98.4%	100%	24.3%
JaRFly PAR	64.8%	96.1%	98.5%	100%	13.8%
Our Paper	63.5%	95.8%	98.5%	100%	34%

preliminary experiments for tuning the learning rate (recall Section V-C). The results confirm that increasing the number of arms from 3 to 18 does hinder the performance of the approach. However, we can also see that keeping the 18 mutations but grouping them into 7 arms instead does greatly improve the results. The more aggressive learning achieved both through a smaller generation size or adjusting the operator probabilities after every mutation yielded worse results.

From here, we proceeded with additional experimentation on the two most successful experiments (emboldened in Table VII). Table VIII presents the results of these two experiments on the dataset of 64 bugs (that were reported as patched using GenProg and PAR in JaRFly). The experiment with the fitness relative to the parent as the reward achieved a 2.5% increase in the success rate of the repair attempts and patched 3 additional bugs. Therefore, we decided to test it on the additional 289 bugs that the JaRFly reports to not be able to generate patches for with GenProg and PAR (recall Table I). Our approach produced 80 patches for 10 bugs from the set

that JaRFly didn't patch. Thus, on the full dataset of 353 bugs, our approach generates patches for 61 bugs — 9 more than the baseline random selection strategy.

We follow the patch evaluation methodology of JaRFly [39] to assess the quality of the patches generated by our approach. We assess the quality of 487 patches that correspond to the 51 defects that the epsilon-greedy operator selection algorithm with average credit assignment, relative fitness values as the reward and 7 arms generated on the benchmark of 64 bugs (recall Table VIII). Table IX presents the results of the quality scores of the defects. We remove the 2 patches generated for the Time 19 bug from our analysis as they did not pass any test in the evaluation test suite. The results do not show improvement in the minimum, maximum, median, or average values. However, we can see that the percentage of patches that were evaluated to have 100% quality is 34% which is significantly higher than JaRFly.

> **RQ4**: Reinforcement learning-aided mutation operator selection patched 9 more bugs than the baseline which is a 17.3% improvement. Thus, we conclude that it outperforms than the standard uniform selection approach.

VII. DISCUSSION

We deduce that the environment in heuristic-based APR is stationary (recall Section VI-B), i.e., the optimal probabilities for the arms in the RL algorithms do not change over time. As a result, the longer the search progresses, the more likely that the probabilities will be tuned to these constant optimal values.

As expected, the probability matching algorithm had the lowest success rate in our experiments as it is the most simple approach (recall Section II-B). Adaptive pursuit performed slightly better which can be explained by the improvement that this algorithm accounts for in the convergence rate. The UCB algorithm improves the success rate even further since it has been proven to optimise the convergence rate in regards to the cumulative gain [36]. However, to do so, this algorithm continues to tune the exploration-exploitation tradeoff as the search progresses. Since we showed that the environment in heuristic-based APR is stationary, we hypothesize that this simply added

noise to the learning process. In addition, UCB restricts the sampling to the best actions in exploration as opposed to the epsilon-greedy algorithm which chooses randomly between all of the actions. We hypothesize that the combination of these two factors given the stationary environment quickly made the UCB over-exploit and thus resulted in the epsilon-greedy approach being the most fitting for the task.

The results demonstrate (Section VI-C) that the most effective mutation operator selection strategy (epsilon-greedy) is not the same as the most efficient (PM). This is likely due to the variance in the exploration/ exploitation tradeoff between the different RL algorithms. Algorithms that favour exploitation will be less effective at finding patches for defects that are more difficult to fix since they do not explore the search space as extensively. However, their more exploitative nature will allow them to find the patches for easier-to-patch defects more quickly. The hyper-parameter settings of the different algorithm also affects this tradeoff (recall Section V-C).

Increasing the number of mutation operators from three to eighteen greatly increased the search space and effectively slowed the rate of learning in our approach. Provided that we limit the search to 10 generations, efficacy and efficiency rates were reduced as well. However, grouping mutations that target similar defect types together into a single arm representation within the RL algorithm mitigated this problem and effectively improved results. This solution allowed us to maintain the power of the PAR fix templates within the search while controlling the noise that the large number of arms added as a result. Therefore, while our approach does not put a limit on the number of mutations operators, we do recommend the grouping method to control for the number of arms that the RL algorithm rewards. The success of the grouping method hints at the importance of defect type classification within APR.

Out of all APR approaches to-date, heuristic-based techniques have been the most adopted in industrial settings [8]–[10][2]. We pose this is largely due to heuristic-based techniques not requiring a training stage, pre-trained model, or constraint solvers. We recognise the potential in the continued scaling of heuristic-based APR while accounting for the high success rates of end-to-end learning-based strategies [49]–[51]. Therefore, our hybrid approach integrates a classical machine learning technique (RL) within heuristic-based automated program repair. Our results show that our technique surpasses all of the state-of-the-art tools in heuristic-based automated repair. Our approach generates patches for 61 bugs. Tools such as Arja [25] and JaRFly [39] generate patches for 59 bugs and 52 bugs[3] respectively. While we do conduct this comparison given the results of our approach when implemented in JaRFly, our approach can be deployed in any heuristic-based tool regardless of programming language, repair granularity, or search strategy.

VIII. Related Work

Our literature review revealed four popular approaches for RL-based mutation selection. None of these, however, have been applied to improve mutation selection strategies for improvement of existing software. We are the first to do so for the problem of APR.

Mutation Operators: Papers that investigate varying operator probabilities [19], [52] and tuning parameters generally [53], [54] in the field of genetic algorithms have been around for decades [55], [56]. These works generally focus on the adaptive operator selection based on a probability matching strategy that references the fitness improvements of the individuals [57], [58]. Thierens et al. propose an alternative to this through their adaptive pursuit strategy [18]. Other papers focused on the mutation operator specifically both in the field of genetic algorithms [59], [60] and genetic programming [61], [62]. The shortcomings of mutation operation, particularly in promoting diversity for GP [63] are discussed in the literature as well. Glickman et al. [64] and Friedrich et al. [65] address some of these issues like the premature convergence problem and escaping local optima.

Multiple studies alter operator probabilities using reinforcement learning. These studies address both credit assignment strategies as well as operator selection techniques, therefore they are particularly relevant to our paper [66]–[68]. More specifically, extensive research has been done on operator selection strategies that use multi-armed bandits [15]–[21], [31]. However, none of these papers were aimed at APR. They were intended for the purpose of evolution strategies that mainly focus on numerical optimisation problems.

In the field of APR specifically, we found some work by Le Goues et al. [69] on operator design choices. Smigielska et al. [13] propose a uniform strategy when choosing operators. Enhancements were suggested using probabilistic models [14], program context [70], [71], and bug fix history [72]. Soto et al. [26] further propose a set of techniques for enhancing the quality of patches. All of this work focuses on a fixed distribution of probabilities for operator selection. Our work takes this a step further by using reinforcement learning to modify this distribution during search.

Machine Learning for APR: The utility of the machine learning component within the approaches that we found varies from predicting patch correctness [73], [74], to predicting the type of fault [75], to predicting whether continuing the search process will result in a repair [27]. Existing work guides the search process by learning from code patterns and features [7], [75], bug reports [76], or context and statistics [77], [78]. An approach by Ji et al. [79] applies program synthesis to the problem of automated program repair. Conner et al. [80] use neural machine translation models that generate edit operations for patching bugs instead of translating from the buggy to fixed source code directly. Additional work focuses on more specific types of repair, e.g., for conditional statements [81], compilation errors [82], API misuses [83], or specific program languages [84]. Our approach does not aim to target specific bug classes or languages and can be augmented into any heuristic-based APR technique to tune the probabilities of selection for mutation operators during search.

---

[2]We acknowledge [9] uses multiple approaches.

[3]JaRFly reports to patch 64 distinct bugs with GenProg/PAR. We were only able to reproduce 52. ARJA evaluation does not include the Closure project.

IX. Threats to Validity

External threats: The Defects4J benchmark is composed of real-world bugs and has been widely used in the literature. However, there always remains a threat to generalisability of the results that we present in this paper. Moreover, our methodology in this paper uses Java specifically which may not generalise to other object-oriented programming languages. Our choices of benchmark and language enable direct comparison with state-of-the-art approaches. We do not limit our approach to any specific benchmark, programming language, or tool therefore these threats can be mitigated through additional implementations.

Internal threats: Our experiments were conducted on a cluster of computation nodes with varying specifications. We mitigated this threat by comparing the variant number of the successful patch instead of execution time to measure efficiency.

Threats to construct validity: We were not able to reproduce the results of the JaRFly baseline, despite using the same search parameters. However, our approach uses the most recent version of both JaRFly and Defects4J so we conclude that the discrepancy is likely a version issue. Patch quality assessment remains an open question in the field. We mitigate the threat of potential overfitting by evaluating the patches that were generated by our approach on held-out evaluation test suites.

X. Conclusions and Future Work

We introduce a reinforcement learning approach for mutation operator selection in heuristic-based automated program repair. We experiment with four RL algorithms: probability matching, adaptive pursuit, epsilon-greedy, and upper confidence bound as well as two credit assignment techniques: average and exponential recency-weighted average. Moreover, we assess the effect of reward types, number of mutation operators, and activation points on the performance of the technique. Evaluation is conducted on 353 bugs from the Defects4J benchmark. Our results show that the epsilon-greedy algorithm is most effective for guiding mutation operator selection and that the probability matching algorithm is the most efficient. Moreover, we find that the average credit assignment technique is best suited for this context and conclude that the search process is stationary. Additionally, we experiment with altering the number of mutation operators and investigate its effect on performance. Our results show that our RL-based approach to mutation selection improves upon the baseline by patching 9 additional bugs from the Defects4J benchmark, making it an excellent and easy to plug-in addition to existing heuristic-based automated repair tools.

Directions for future work include additional experimentation with integration strategies, i.e., integration at varying points in the search instead of at the fixed points of every mutation or generation as we implement in this paper. To do this, our approach can be coupled with statistical tests to identify the optimal activation points [15]. Advanced credit assignment such as extreme-based [16] or compass [31] credit assignment can also be further explored.

All source code, supplementary materials, and replication instructions will be made publicly available upon acceptance.

References

[1] C. Weiß, R. Premraj, T. Zimmermann, and A. Zeller, "How long will it take to fix this bug?," *MSR*, 2007.

[2] M. Böhme, E. O. Soremekun, S. Chattopadhyay, E. Ugherughe, and A. Zeller, "Where is the bug and how is it fixed? An experiment with practitioners," *ESEC/FSE*, pp. 117–128, 2017.

[3] C. Liu, J. Lu, G. Li, T. Yuan, L. Li, F. Tan, J. Yang, L. You, and J. Xue, "Detecting TensorFlow Program Bugs in Real-World Industrial Environment," *ASE*, pp. 55–66, 2021.

[4] G. Tassey, "The Economic Impacts of Inadequate Infrastructure for Software Testing," *NIST*, 2002.

[5] C. L. Goues, M. Pradel, and A. Roychoudhury, "Automated program repair," *Communications of the ACM*, 2019.

[6] J. Xuan, M. Martinez, F. Demarco, M. Clément, S. Lamelas, T. Durieux, D. L. Berre, and M. Monperrus, "Nopol: Automatic repair of conditional statement bugs in java programs," *IEEE TSE*, vol. 43, pp. 34–55, 2018.

[7] L. Chen, Y. Pei, M. Pan, T. Zhang, Q. Wang, and C. A. Furia, "Program Repair with Repeated Learning," *IEEE TSE*, 2022.

[8] S. O. Haraldsson, J. R. Woodward, A. E. I. Brownlee, and K. Siggeirsdottir, "Fixing bugs in your sleep: How genetic improvement became an overnight success," in *GECCO*, pp. 1513–1520, ACM, 2017.

[9] A. Marginean, J. Bader, S. Chandra, M. Harman, Y. Jia, K. Mao, A. Mols, and A. Scott, "Sapfix: Automated end-to-end repair at scale," *ICSE-SEIP*, 2019.

[10] S. Kirbas, E. Windels, O. Mcbello, K. Kells, M. Pagano, R. Szalanski, V. Nowack, E. Winter, S. Counsell, D. Bowes, T. Hall, S. Haraldsson, and J. Woodward, "On the introduction of automatic program repair in bloomberg," *IEEE Software*, 2020.

[11] J. R. Koza, "Genetic programming as a means for programming computers by natural selection," *Statistics and Computing*, vol. 4, no. 2, pp. 87–112, 1994.

[12] H. H. Hoos and T. Stützle, *Stochastic local search: Foundations and Applications.* Morgan Kaufmann Publishers Inc., 2004.

[13] M. Smigielska, A. Blot, and J. Petke, "Uniform Edit Selection for Genetic Improvement: Empirical Analysis of Mutation Operator Efficacy," *Int. Workshop on GI*, pp. 1–8, 2021.

[14] M. Soto and C. Le Goues, "Using a probabilistic model to predict bug fixes," *SANER*, vol. 2018-March, pp. 221–231, 2018.

[15] L. da Costa, Á. Fialho, M. Schoenauer, and M. Sebag, "Adaptive Operator Selection with Dynamic Multi-Armed Bandits," *GECCO*, pp. 913–920, 2008.

[16] Á. Fialho, L. Da Costa, M. Schoenauer, and M. Sebag, "Dynamic Multi-Armed Bandits and Extreme Value-Based Rewards for Adaptive Operator Selection in Evolutionary Algorithms," *LNTCS*, vol. 5851, pp. 176–190, 2009.

[17] Á. Fialho, M. Schoenauer, and M. Sebag, "Toward Comparison-based Adaptive Operator Selection," *GECCO*, 2010.

[18] D. Thierens, "An adaptive pursuit strategy for allocating operator probabilities," *GECCO*, pp. 1539–1546, 2005.

[19] T. Murata and H. Ishibuchi, "Positive and negative combination effects of crossover and mutation operators in sequencing problems," *IEEE CEC*, pp. 170–175, 1996.

[20] A. E. Eiben, M. Horvath, W. Kowalczyk, and M. C. Schut, "Reinforcement Learning for Online Control of Evolutionary Algorithms," *LNAI*, vol. 4335, pp. 151–160, 2007.

[21] R. E. J. E. Pettinger, "Controlling Genetic Algorithms With Reinforcement Learning," *GECCO*, 2002.

[22] M. Sharma, M. López-Ibáñez, A. Komninos, and D. Kazakov, "Deep reinforcement learning based parameter control in differential evolution," *GECCO*, pp. 709–717, 2019.

[23] H. Zhang and J. Lu, "Adaptive evolutionary programming based on reinforcement learning," *Information Sciences*, vol. 178, no. 4, pp. 971–984, 2008.

[24] Y. Qi, X. Mao, Y. Lei, Z. Dai, and C. Wang, "The strength of random search on automated program repair," *ICSE*, pp. 254–265, 2014.

[25] Y. Yuan and W. Banzhaf, "ARJA: Automated Repair of Java Programs via Multi-Objective Genetic Programming," *IEEE TSE*, 2018.

[26] M. Soto, "Improving patch quality by enhancing key components of automatic program repair," *ASE*, pp. 1230–1233, 2019.

[27] X.-B. D. Le, "Towards Efficient and Effective Automatic Program Repair," *ASE*, 2016.

[28] E. K. Smith, E. T. Barr, C. Le Goues, and Y. Brun, "Is the cure worse than the disease? Overfitting in automated program repair," *ESEC/FSE*, pp. 532–543, 2015.

[29] D. Kim, J. Nam, J. Song, and S. Kim, "Automatic Patch Generation Learned from Human-Written Patches," *ICSE*, 2013.

[30] R. S. Sutton and A. G. Barto, *Reinforcement Learning An Introduction*. MIT Press, 2018.

[31] J. Maturana, Á. Fialho, F. Saubion, M. Schoenauer, and M. Sebag, "Extreme compass and dynamic multi-armed bandits for adaptive operator selection," *GEC*, pp. 365–372, 2009.

[32] Á. Fialho, M. Schoenauer, and M. Sebag, "Analysis of Adaptive Operator Selection Techniques on the Royal Road and Long K-Path Problems," *GECCO*, 2009.

[33] L. Davis, "Adapting Operator Probabilities in Genetic Algorithms," in *Int. Conf. on GA*, pp. 61–69, 1989.

[34] B. A. Julstrum, "What Have You Done for Me Lately? Adapting Operator Probabilities in a Steady-State Genetic Algorithm," in *Int. Conf. on GA*, pp. 81–87, 1995.

[35] J. Maturana and F. Saubion, "A compass to guide genetic algorithms," *LNCS*, vol. 5199, pp. 256–265, 2008.

[36] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time Analysis of the Multiarmed Bandit Problem," *Machine Learning*, vol. 47, no. 2, pp. 235–256, 2002.

[37] H. J. C. Barbosa and A. Medeiros, "On Adaptive Operator Probabilities in Real Coded Genetic Algorithms," *Workshop on Advances and Trends in Artificial Intelligence for Problem Solving*, 2000.

[38] J. M. Whitacre, T. Q. Pham, and R. A. Sarker, "Use of Statistical Outlier Detection Method in Adaptive Evolutionary Algorithms," *GECCO*, 2006.

[39] M. Motwani, M. Soto, Y. Brun, R. Just, and C. Le Goues, "Quality of Automated Program Repair on Real-World Defects," *IEEE TSE*, vol. 48, no. 2, pp. 637–661, 2022.

[40] program repair.org, "Tools for automated program repair." https://program-repair.org/tools.html, 2023. Accessed: 2rd Feb 2023.

[41] C. Le Goues, T. V. Nguyen, S. Forrest, and W. Weimer, "GenProg: A generic method for automatic software repair," *IEEE TSE*, vol. 38, no. 1, pp. 54–72, 2012.

[42] Y. Qi, X. Mao, and Y. Lei, "Efficient automated program repair through fault-recorded testing prioritization," *IEEE ICSM*, pp. 180–189, 2013.

[43] program repair.org, "Benchmarks for automated program repair." https://program-repair.org/benchmarks.html, 2023. Accessed: 2rd Feb 2023.

[44] M. Martinez and M. Monperrus, "ASTOR: A program repair library for Java (Demo)," *ISSTA*, pp. 441–444, 2016.

[45] Defects4J, "Defects4J – version 2.0.0." https://github.com/rjust/defects4j/, 2023. Accessed: 2rd Feb 2023.

[46] EvoSuite, "Automatic test suite generation for java." https://www.evosuite.org/, 2023. Accessed: 2rd Feb 2023.

[47] Randoop, "Automatic unit test generation for java." https://randoop.github.io/randoop/, 2023. Accessed: 2rd Feb 2023.

[48] LASER-UMASS, "Artifact Replication Repository for the study on Quality of Automated Program Repair on Real-World Defects." https://github.com/LASER-UMASS/JavaRepair-replication-package, 2020. Accessed: 2rd Feb 2023.

[49] N. Jiang, T. Lutellier, and L. Tan, "Cure: Code-aware neural machine translation for automatic program repair," *ICSE*, pp. 1161–1173, 2 2021.

[50] Q. Zhu, Z. Sun, Y. A. Xiao, W. Zhang, K. Yuan, Y. Xiong, and L. Zhang, "A syntax-guided edit decoder for neural program repair," *ESEC/FSE*, pp. 341–353, 8 2021.

[51] C. S. Xia and L. Zhang, "Less training, more repairing please: Revisiting automated program repair via zero-shot learning," *ESEC/FSE*, pp. 959–971, 7 2022.

[52] A. Tuson and P. Ross, "Cost based operator rate adaptation: An investigation," *LNCS*, vol. 1141, pp. 461–469, 1996.

[53] A. Aleti and I. Moser, "A Systematic Literature Review of Adaptive Parameter Control Methods for Evolutionary Algorithms," *ACM CSUR*, vol. 49, no. 3, 2016.

[54] H. Moazen, S. Molaei, L. Farzinvash, and M. Sabaei, "Pso-elpm: Pso with elite learning, enhanced parameter updating, and exponential mutation operator," *Information Sciences*, vol. 628, pp. 70–91, 5 2023.

[55] J. T. Stanczak, J. J. Mulawka, and B. K. Verma, "Genetic algorithms with adaptive probabilities of operator selection," *ICCIMA*, pp. 464–468, 1999.

[56] B. A. Julstrom, "Adaptive operator probabilities in a genetic algorithm that applies three operators," *ACM SAC*, pp. 233–238, 1997.

[57] F. Vafaee, W. Xiao, P. C. Nelson, and C. Zhou, "Adaptively evolving probabilities of genetic operators," *ICMLA*, pp. 292–299, 2008.

[58] J. A. Soria-Alcaraz, G. Ochoa, M. Carpio, and H. Puga, "Evolvability metrics in adaptive operator selection," *GECCO*, pp. 1327–1334, 2014.

[59] J. Hesser and R. Mäinner, "Towards an optimal mutation probability for genetic algorithms," *LNCS*, vol. 496, pp. 23–32, 1991.

[60] K. I. Ali and K. Brohi, "An adaptive learning automata for genetic operators allocation probabilities," *FIT*, pp. 55–59, 2013.

[61] M. T. A. Anik and S. Ahmed, "A mixed mutation approach for evolutionary programming based on guided selection strategy," *ICIEV*, 2013.

[62] L. Hong, J. H. Drake, and E. Özcan, "A step size based self-adaptive mutation operator for Evolutionary Programming," *GECCO*, pp. 1381–1387, 2014.

[63] D. Jackson, "Mutation as a Diversity Enhancing Mechanism in Genetic Programming," *GECCO*, 2011.

[64] M. R. Glickman, M. R. Glickman, and K. Sycara, "Reasons for Premature Convergence of Self-Adapting Mutation Rates," *IEEE CEC*, pp. 62–69, 2000.

[65] T. Friedrich, F. Quinzan, and M. Wagner, "Escaping large deceptive basins of attraction with heavy-tailed mutation operators," *GECCO*, pp. 293–300, 2018.

[66] A. Awad, A. Hawash, and B. Abdalhaq, "A genetic algorithm (ga) and swarm based binary decision diagram (bdd) reordering optimizer reinforced with recent operators," *IEEE TEVC*, 2022.

[67] T. Li, Y. Meng, and L. Tang, "Scheduling of continuous annealing with a multi-objective differential evolution algorithm based on deep reinforcement learning," *ASE*, 2023.

[68] Z. Zhang, Q. Tang, M. Chica, and Z. Li, "Reinforcement learning-based multiobjective evolutionary algorithm for mixed-model multimanned assembly line balancing under uncertain demand," *IEEE Trans. on Cybernetics*, pp. 1–14, 1 2023.

[69] C. Le Goues, W. Wemer, and S. Forrest, "Representations and operators for improving evolutionary software repair," *GECCO*, pp. 959–966, 2012.

[70] M. Wen, J. Chen, R. Wu, D. Hao, and S. C. Cheung, "Context-aware patch generation for better automated program repair," *ICSE*, vol. 2018-Janua, pp. 1–11, 2018.

[71] M. R. Ullah, N. S. Chowdhury, and F. M. Tawsif, "Impact of combining syntactic and semantic similarities on patch prioritization while using the insertion mutation operators," *SEKE*, pp. 190–195, 2023.

[72] X. B. D. Le, D. Lo, and C. Le Goues, "History driven program repair," *SANER*, vol. 1, pp. 213–224, 2016.

[73] L. Schramm, "Improving Performance of Automatic Program Repair using Learned Heuristics," *ESEC/FSE*, 2017.

[74] H. Tian, K. Liu, Y. Li, A. Kader, A. Koyuncu, A. Habib, L. Li, J. Wen, J. Klein, and T. F. Bissyandé, "The best of both worlds: Combining learned embeddings with engineered features for accurate prediction of correct patches," *ACM TOSEM*, vol. 1, 2023.

[75] M. Valueian, M. Vahidi-Asl, and A. Khalilian, "SituRepair: Incorporating machine-learning fault class prediction to inform situational multiple fault automatic program repair," *IJCIP*, vol. 37, p. 100527, 2022.

[76] C. Liu, J. Yang, L. Tan, and M. Hafiz, "R2Fix: Automatically Generating Bug Fixes from Bug Reports," *ICST*, 2013.

[77] J. Jiang, L. Ren, Y. Xiong, and L. Zhang, "Inferring Program Transformations From Singular Examples via Big Code," *ASE*, 2019.

[78] Z. Yu, M. Martinez, Z. Chen, T. F. Bissyandé, and M. Monperrus, "Learning the relation between code features and code transforms with structured prediction," *IEEE TSE*, pp. 1–29, 2023.

[79] S. Ji, S.-M. Choi, S.-K. Ko, D. Kim, and H. Im, "RepCoder: An Automated Program Repair Framework for Probability-Based Program Synthesis," *ACM/SIGAPP SAC*, 2022.

[80] A. Connor, A. Harris, N. Cooper, and D. Poshyvanyk, "Can we automatically fix bugs by learning edit operations?," *SANER*, pp. 782–792, 2022.

[81] D. Gopinath, K. Wang, J. Hua, and S. Khurshid, "Repairing Intricate Faults in Code Using Machine Learning and Path Exploration," *ICSME*, 2016.

[82] T. Ahmed, N. R. Ledesma, and P. Devanbu, "Synshine: Improved fixing of syntax errors," *IEEE TSE*, 2022.

[83] D. Wu, B. Shen, Y. Chen, H. Jiang, and L. Qiao, "Automatically repairing tensor shape faults in deep learning programs," *Information and Software Technology*, vol. 151, p. 107027, 2022.

[84] M. Lajko, V. Csuvik, and L. Vidacs, "Towards javascript program repair with generative pre-trained transformer (gpt-2)," *Int. Workshop on APR*, pp. 61–68, 2022.