



# Genetic Algorithm for Program Synthesis

Yutaka Nagashima<sup>(✉)</sup> 

Cambridge, UK

`united.reasoning@gmail.com`

**Abstract.** A deductive program synthesis tool takes a specification as input and derives a program that satisfies the specification. The drawback of this approach is that search spaces for such correct programs tend to be enormous, making it difficult to derive correct programs within a realistic timeout. To speed up such program derivation, we improve the search strategy of a deductive program synthesis tool, SuSLik, using evolutionary computation. Our cross-validation shows that the improvement brought by evolutionary computation generalises to unforeseen problems.

## 1 Introduction

A far-fetched goal of artificial intelligence research is to build a system that writes computer programs for humans. To achieve this goal, researchers take two distinct approaches: deductive program synthesis and inductive program synthesis. Both approaches attempt to produce programs requested by human users. The difference lies how they produce programs: deductive synthesis tries to *deduce* programs that satisfy specifications, while inductive program synthesis tries to *induce* programs from examples.

While such inductive synthesis alleviates the burden of implementation by guessing programs from given input-output examples, in inductive synthesis resulting programs are not trustworthy. Deductive synthesis overcomes this limitation with formal specifications: it allows users to formalise *what* they want as specifications, whereas inductive synthesis tools guess *how* programs should behave from examples provided by users. Thus, in deductive synthesis providing formal specifications remains as users' responsibility. The upside of deductive synthesis is, however, users can obtain *correct* programs upon success.

SuSLik [19], for example, is one of such deductive synthesis tools. It takes a specification provided by humans and attempts to produce heap-manipulating programs satisfying the specification in a language that resembles the C language. Internally, this derivation process is formulated as proof search: SuSLik composes a heap-manipulating program by conducting a best-first search for a proof goal presented as specification. The drawback is that the search algorithm often fails to find a proof within a realistic timeout. That is, even we

---

Y. Nagashima—Independent.

We would like to thank Andreea Costea for preparing additional SuSLik problems for cross-validations.

© IFIP International Federation for Information Processing 2023

Published by Springer Nature Switzerland AG 2023

H. Hojjat and E. Ábrahám (Eds.): FSEN 2023, LNCS 14155, pp. 104–111, 2023.

[https://doi.org/10.1007/978-3-031-42441-0\\_8](https://doi.org/10.1007/978-3-031-42441-0_8)

pass a specification to SuSLik, SuSLik may not produce a program satisfying the specification. According to Itzhaky *et al.* [5], *different synthesis tasks benefit from different search parameters, and that we might need a mechanism to tune SuSLik’s search strategy for a given synthesis task.*

## 2 SuSLik’s Search Strategy

SuSLik synthesises a program by searching for a corresponding proof. We can see SuSLik’s proof search as an exploration of an OR-tree, nodes of the tree represent (intermediate) synthesis goals, while edges of the tree represent rule applications. The shape of such search tree is not known in advance, and the task of SuSLik is to identify a solved node, in which a proof is complete.

Since such OR-trees can be too large to find proofs within a realistic timeout, SuSLik narrows the search space using a *proof strategy*. Essentially, proof strategy in SuSLik is a function that takes a synthesis goal and returns an ordered list of rules to apply next. Itzhaky *et al.* developed the default strategy by manually encoding human expertise. For example, the default strategy precludes the application of a rule called **CALL** when another rule **CLOSE** has been applied before reaching the current node. This way, the SuSLik rules are grouped into 10 ordered lists, and the order of these rules define how SuSLik explores the corresponding OR-tree.

Another decision SuSLik has to make for an effective search is to select the next node to expand. The current version of SuSLik make this decision using a cost function, manually developed and tuned by Itzhaky *et al.* [5].

Both the weights of the cost function and orders of derivation rules are manually tuned for the benchmark used in their evaluation [5]; however, as we show in Sect. 4, our evolutionary framework finds better strategies through evolution.

## 3 Evolutionary Computation for SuSLik

The aim of our evolutionary computation is to optimise the order of each group of derivation rules and the weights of the cost function, which is used to implement best-first search.

Algorithm 1 summarises the genetic algorithm we used in our framework to improve the search strategy of SuSLik. Firstly, the algorithm takes a set of training problems an inputs, using which we evolve SuSLik instances over 40




Fig. 1. Mutation and Elitist Selection

generations. Line 1 defines the initial population. Each individual in a population is evaluated according to the fitness function described in Sect. 3.

For each generation, we copy individuals from the previous iteration (Line 6), mutate them (Line 7), evaluate individuals (Line 8). Then, we sort all individuals in the current generation based on their performance (Line 9–10). And we continue to the next generation using the best 20 individuals from the current pool. In the following, we explain the mutation algorithm, the fitness function, and our selection algorithm.

*Mutation.* As we explained in Sect. 2, by default a search strategy of SuSLik is defined by two factors: the order of rule application and weights of each node in the search tree. To determine an effective way to apply genetic algorithms to program synthesis in SuSLik, we implemented the following three different mutation algorithms:

- *Order-only mutation* changes only the order of rule application for each node.
- *General rule-weight mutation* changes the weights of each node based on what rules have been applied to reach that node.
- *Goal-specific rule-weight mutation* allows SuSLik to choose a weight for each rule based on properties of a node during a search.

*Fitness.* The fitness function measures the performance of SuSLik instances. More specifically, it measures how many derivation problems each SuSLik instance solves within the timeout of 2.5s for each problem. When multiple SuSLik instances solve the same number of derivation problems, the fitness function uses the numbers of rules fired by the instances as a tie-breaker: it considers that the instance that solves a certain number of problems with a smaller number of rule applications is better than another instance that solves the same number of problems with a larger number of rule applications.

---

**Algorithm 1** Evolutionary Computation for SuSLik

---

**Input:** synthesis problems for SuSLik  
**Output:** a SuSLik search strategy

---

```

1: Let old_pop be the initial population.
2: fitness(old_pop)
3: generation  $\leftarrow$  1
4: while generation  $\leq$  40 do
5:   generation  $\leftarrow$  generation + 1
6:   new_pop  $\leftarrow$  old_pop
7:   mutate(new_pop)
8:   fitness(new_pop)
9:   whole_pop  $\leftarrow$  old_pop + new_pop
10:  sort(whole_pop)
11:  old_pop  $\leftarrow$  take (whole_pop, 20)
12: end while

```

---

*Selection.* We adopt a version of *elitist selection* as our selection method: we pass individuals from the current generation to the next generation. By copying them and mutating them if they show better performance in the current generation. Figure 1 provides the schematic view of our elitist selection. Unlike the standard elitist selection algorithm, ours prioritizes the best individual in each generation to speed up the evolution: the best individual in each generation, called *champion*, is entitled with three children, one original copy without mutation and two

mutated children, whereas each of other 19 winners has one original copy and only one mutated child in the next generation.

Note that each individual has two kinds of properties to mutate: the order of derivation rules, and weights used in the cost function. While we represent the weights as floating point numbers, we adopt permutation encoding for the orders of derivation rules.

For each permutation encoding, each individual has the probability of 0.1 to be moved, while we change weights by multiplying a random number between 0.8 and 1.2. In our framework, we do not apply crossover to permutation encoding: since our sequences denoting rule orders tend to be short, we are not sure if crossovers would result in a better performance of evolution.

Our evolutionary computation for program synthesis differs from genetic programming [9] or evolutionary programming [1]: we did not directly apply simulated evolution to programs, but our framework improves the search mechanism for deriving correct programs through evolution. We take this approach to take the best of both worlds: the correctness of resulting programs guaranteed by the deductive synthesis and its certification tool, and the search heuristics enhanced through evolutionary computation.


## 4 Evaluation

We conducted cross-validations to evaluate what improvements our evolutionary computation framework brought to SuSLik. We measured how many synthesis problems SuSLik failed to solve with in 2.5s of timeout. For this evaluation, we used a consumer laptop running Ubuntu 20.04.3 LTS on a machine with 16 CPUs of AMD Ryzen 7 4,800H with Radeon Graphics and 15,854 MB of main memory.


As SuSLik is a new tool, we have only 65 problems available in our benchmark: problems from a preceding work on SuSLik [5] and new problems prepared for this project. These problems include tasks on various data-structures such as integers, singly linked lists, sorted lists, doubly linked lists, lists of lists, binary trees, and packed trees.

Firstly, we randomly split our benchmarks into two groups: the validation dataset and training dataset. Then, using the training dataset we apply our evolutionary computation described in Algorithm 1 to evolve SuSLik’s search strategy. As explained in Sect. 3, the output of our evolutionary computation is just one search strategy produced after 40 generations. However, in this experiment we conducted cross-validations using the best individual from the training set for each generation to see how our framework produces transferable improvement over generations.


To reduce the influence from a specific random split, we conducted this experiment four times, and the result of each experiment is illustrated from Fig. 2 to Fig. 5. In these figures, the horizontal axes represent the number of generations, while the vertical axes represent the number of synthesis problems SuSLik did *not* solve within the timeout.




**Fig. 2.** Cross-validation 1



**Fig. 3.** Cross-validation 2



**Fig. 4.** Cross-validation 3



**Fig. 5.** Cross-validation 4

These figures show that when adopting the general rule-weight mutation, our evolutionary framework managed to improve SuSLik’s capability to find solutions in validation sets, even though evolution is based on training sets. That is, somewhat contrarily to the prediction by Itzhaky introduced in Sect. 1, we found that there are strategies that tend to perform better for unforeseen problems, and we can find such strategies using evolutionary computation.

On the other hand, the order-only mutation and goal-specific rule-weight mutation resulted in less promising results. In particular, the goal-specific rule-weight mutation over-fitted to training data in Fig. 2 and Fig. 5, probably due to its capability to fine tune the strategy for our small dataset.

## 5 Discussion

The limited size of available dataset is the main challenge we faced in this project. This problem is partially unavoidable since program synthesis itself is still an emerging field in Computer Science. Other AI projects for interactive theorem

provers take advantage of large existing proof corpora for training. For example, Nagashima built a tactic prediction tool, PaMpeR [16], for Isabelle/HOL by extracting 425,334 data points [13] from the Archive of Formal Proofs (AFP) [8]. Li *et al.* also mined the AFP and produced 820K training examples for conjecturing. For Coq, Yang *et al.* constructed a dataset containing 71K proofs from 123 projects [21], whereas Huang *et al.* [4] extracted a dataset consisting of 1,602 lemmas from the Feit-Thompson formalization. For HOL Light [3], The HOL-Step [6] used 1,013,046 training examples and 196,030 testing examples extracted from 11,400 proofs, while the HOList project presented a benchmark based on 2,199 definitions and 29,462 theorems and lemmas. These projects managed to gather large data sets since their underlying theorem provers, Isabelle/HOL, Coq, and HOL Light, have a larger user base than SuSLik [19] does.

For the moment, our framework improves *static* parameters for SuSLik. That is, the resulting weights and rule orders are fixed for all intermediate synthesis problems. Our evaluation has shown that our static parameter optimisation (general rule-weight mutation) using evolutionary computation generalises well: a SuSLik instance that performs well for a training dataset tends to perform well for an evaluation dataset. We expected that we could achieve even better performance by producing *dynamic* parameters (goal-specific rule-weight) for SuSLik: functions that inspect a node at hand and decide on a promising rule order and weights for that node. Our efforts in this direction are, unfortunately, unsuccessful so far. We hope that a larger training dataset would allow for such optimisation in the future.

## 6 Related Work

Even though there was an attempt to use reinforcement learning [20] for a connection-style proof search [7]; we mindfully chose evolutionary computation over reinforcement learning: since we do not have a changing environment in our setting, it is unclear if we gain any benefits from having two metrics, reward function for the long term goal and value function for the short term benefit. Instead, we improved SuSLik’s default search strategy for randomly chosen fixed training problem sets and measured how the improvement generalizes to validation sets.

When implementing our framework for evolutionary computation, we took the advantage of a Python framework for evolutionary computation called DEAP [2], even though SuSLik itself is implemented in Scala.

Previously, we attempted to improve proof strategies [17] for Isabelle/HOL using evolutionary computing [11]. However, the focus of that project shifted to the prediction of induction arguments [14, 15] using meta-languages [10, 12].

Nawaz *et al.* used a genetic algorithm to evolve random proof sequences to target proofs. The drawback of their approach is that the fitness function used in the genetic algorithm relies on the existence of a proof for a given problem. Therefore, this framework is not applicable to open conjectures without existing proofs [18].

## References

1. Fogel, L., Owens, A.J., Walsh, M.J.: Artificial intelligence through simulated evolution (1966)
2. Fortin, F.A., De Rainville, F.M., Gardner, M.A., Parizeau, M., Gagné, C.: DEAP: evolutionary algorithms made easy. *J. Mach. Learn. Res.* **13**, 2171–2175 (2012)
3. Harrison, J.: HOL light: a tutorial introduction. In: Srivas, M., Camilleri, A. (eds.) *FMCAD 1996*. LNCS, vol. 1166, pp. 265–269. Springer, Heidelberg (1996). <https://doi.org/10.1007/BFb0031814>
4. Huang, D., Dhariwal, P., Song, D., Sutskever, I.: GamePad: a learning environment for theorem proving. In: 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, 6–9 May 2019. OpenReview.net (2019). <https://openreview.net/forum?id=r1xwKor9Y7>
5. Itzhaky, S., Peleg, H., Polikarpova, N., Rowe, R.N.S., Sergey, I.: Deductive synthesis of programs with pointers: techniques, challenges, opportunities. In: Silva, A., Leino, K.R.M. (eds.) *CAV 2021, Part I*. LNCS, vol. 12759, pp. 110–134. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-81685-8\\_5](https://doi.org/10.1007/978-3-030-81685-8_5)
6. Kaliszzyk, C., Chollet, F., Szegedy, C.: HolStep: a machine learning dataset for higher-order logic theorem proving. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 24–26 April 2017, Conference Track Proceedings. OpenReview.net (2017). <https://openreview.net/forum?id=ryuxYmvel>
7. Kaliszzyk, C., Urban, J., Michalewski, H., Olsák, M.: Reinforcement learning of theorem proving. In: Bengio, S., Wallach, H.M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, Montréal, Canada, 3–8 December 2018*, pp. 8836–8847 (2018). <https://proceedings.neurips.cc/paper/2018/hash/55acf8539596d25624059980986aaa78-Abstract.html>
8. Klein, G., Nipkow, T., Paulson, L., Thiemann, R.: The Archive of Formal Proofs (2004). <https://www.isa-afp.org/>
9. Koza, J.R.: Genetic Programming - On the Programming of Computers by Means of Natural Selection. Complex Adaptive Systems. MIT Press, Cambridge (1993)
10. Nagashima, Y.: LiFtEr: language to encode induction heuristics for Isabelle/HOL. In: Lin, A.W. (ed.) *APLAS 2019*. LNCS, vol. 11893, pp. 266–287. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-34175-6\\_14](https://doi.org/10.1007/978-3-030-34175-6_14)
11. Nagashima, Y.: Towards evolutionary theorem proving for Isabelle/HOL. In: López-Ibáñez, M., Auger, A., Stützle, T. (eds.) *Proceedings of the Genetic and Evolutionary Computation Conference Companion, GECCO 2019, Prague, Czech Republic, 13–17 July 2019*, pp. 419–420. ACM (2019). <https://doi.org/10.1145/3319619.3321921>
12. Nagashima, Y.: Definitional quantifiers realise semantic reasoning for proof by induction. *CoRR* **abs/2010.10296** (2020). <https://arxiv.org/abs/2010.10296>
13. Nagashima, Y.: Simple dataset for proof method recommendation in Isabelle/HOL. In: Benz Müller, C., Miller, B. (eds.) *CICM 2020*. LNCS (LNAI), vol. 12236, pp. 297–302. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-53518-6\\_21](https://doi.org/10.1007/978-3-030-53518-6_21)
14. Nagashima, Y.: Smart induction for Isabelle/HOL (tool paper). In: 2020 Formal Methods in Computer Aided Design, FMCAD 2020, Haifa, Israel, 21–24 September 2020, pp. 245–254. IEEE (2020). [https://doi.org/10.34727/2020/isbn.978-3-85448-042-6\\_32](https://doi.org/10.34727/2020/isbn.978-3-85448-042-6_32)

15. Nagashima, Y.: Faster smarter proof by induction in Isabelle/HOL. In: Zhou, Z. (ed.) Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021, pp. 1981–1988. ijcai.org (2021). <https://doi.org/10.24963/ijcai.2021/273>
16. Nagashima, Y., He, Y.: PaMpeR: proof method recommendation system for Isabelle/HOL. In: Huchard, M., Kästner, C., Fraser, G. (eds.) Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Montpellier, France, 3–7 September 2018, pp. 362–372. ACM (2018). <https://doi.org/10.1145/3238147.3238210>
17. Nagashima, Y., Kumar, R.: A proof strategy language and proof script generation for Isabelle/HOL. In: de Moura, L. (ed.) CADE 2017. LNCS (LNAI), vol. 10395, pp. 528–545. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-63046-5\\_32](https://doi.org/10.1007/978-3-319-63046-5_32)
18. Nawaz, M.Z., Hasan, O., Nawaz, M.S., Fournier-Viger, P., Sun, M.: Proof searching in HOL4 with genetic algorithm. In: Hung, C., Cerný, T., Shin, D., Bechini, A. (eds.) SAC 2020: The 35th ACM/SIGAPP Symposium on Applied Computing, online event, [Brno, Czech Republic], March 30–April 3 2020, pp. 513–520. ACM (2020). <https://doi.org/10.1145/3341105.3373917>
19. Polikarpova, N., Sergey, I.: Structuring the synthesis of heap-manipulating programs. Proc. ACM Program. Lang. **3**(POPL), 72:1–72:30 (2019). <https://doi.org/10.1145/3290385>
20. Sutton, R.S., Barto, A.G.: Reinforcement learning: an introduction. IEEE Trans. Neural Netw. **9**(5), 1054–1054 (1998). <https://doi.org/10.1109/TNN.1998.712192>
21. Yang, K., Deng, J.: Learning to prove theorems via interacting with proof assistants. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning, ICML 2019, Long Beach, California, USA, 9–15 June 2019. Proceedings of Machine Learning Research, vol. 97, pp. 6984–6994. PMLR (2019). <http://proceedings.mlr.press/v97/yang19a.html>