

# Exploring the Verifiability of Code Generated by GitHub Copilot

DAKOTA WONG, University of Waterloo, Canada

AUSTIN KOTHIG, University of Waterloo, Canada

PATRICK LAM, University of Waterloo, Canada

GitHub's Copilot generates code quickly. We investigate whether it generates good code. Our approach is to identify a set of problems, ask Copilot to generate solutions, and attempt to formally verify these solutions with Dafny. Our formal verification is with respect to hand-crafted specifications. We have carried out this process on 6 problems and succeeded in formally verifying 4 of the created solutions. We found evidence which corroborates the current consensus in the literature: Copilot is a powerful tool; however, it should not be "flying the plane" by itself.

CCS Concepts: • **Software and its engineering** → **General programming languages**; • **Social and professional topics** → *History of programming languages*;

Additional Key Words and Phrases: GitHub Copilot, Program Synthesis, Software Verification.

## ACM Reference Format:

Dakota Wong, Austin Kothig, and Patrick Lam. 2022. Exploring the Verifiability of Code Generated by GitHub Copilot. *Proc. ACM Program. Lang.* 1, HATRA, Article 1 (December 2022), 10 pages.

## 1 INTRODUCTION

In aviation, commercial airplanes have two yokes (controls) which allows for sharing of responsibilities and workload between a primary pilot and a copilot. Both operators are fully trained, informed, and capable of flying the aircraft; however, only one of them needs to be controlling the aircraft at any one moment. In software development, traditionally, code is written by an individual. In team environments, this code is typically scrutinized and assessed by a number of people, but originally it had to be developed by an individual. Similar to a pilot, software developers are ordinarily well trained, informed, and capable of programming; however, inefficient or incorrect code segments can potentially slip between the cracks and into code-bases when a developer does not fully understand their tools, the domain, or the required program specifications.

GitHub, with the use of OpenAI's Codex [Chen et al. 2021], has released an "AI pair programmer" to assist in software development [GitHub 2022]. This programming partner is aptly named Copilot. The aim of this tool is to shift the programmer's time and attention away from boilerplate and repetitive code patterns, and toward the critical design of the system. Copilot works by providing *implementation suggestions* based on natural language descriptions of the required code. Copilot is currently capable of generating code in eleven programming languages, including C, C++, C#, Go, Java, JavaScript, PHP, Python, Ruby, Scala, and TypeScript<sup>1</sup>.

The community is struggling to understand what Copilot can and cannot do, and proofs are one indicator of where Copilot might stand in terms of generating correct code. We know that Copilot definitely generates some code, but what can we conclude about the code that it generates? Thus, our goal is to gain some understanding of the quality of Copilot's generated code; our approach is

<sup>1</sup><https://docs.github.com/en/copilot/overview-of-github-copilot/about-github-copilot>

Authors' addresses: Dakota Wong, University of Waterloo, Canada, [d82wong@uwaterloo.ca](mailto:d82wong@uwaterloo.ca); Austin Kothig, University of Waterloo, Canada, [akothig@uwaterloo.ca](mailto:akothig@uwaterloo.ca); Patrick Lam, University of Waterloo, Canada, [patrick.lam@uwaterloo.ca](mailto:patrick.lam@uwaterloo.ca).

2022. 2475-1421/2022/12-ART1 \$15.00  
<https://doi.org/>

to attempt to verify selected examples of Copilot-generated code using the software verification language Dafny [Leino 2010].

## 2 RELATED WORK

There has been a great deal of scrutiny of Copilot and of the code it produces. Early studies which aim to empirically evaluate Copilot [Nguyen and Nadi 2022] have found that while this early version of the model is not perfect, its current results are impressive. In their work, Nguyen and Nadi tested Copilot against 33 LeetCode<sup>2</sup> questions (4 from the easy category, 17 from the medium category, and 12 from the hard category) using four programming languages (Python, JavaScript, Java, C), ran the generated code against the LeetCode test cases, and examined the percentage of passing cases. Java performed best, being able to solve all test cases for 19 of the 33 problems. The authors found that, often, the code that Copilot generated could be simplified, and that sometimes the generated code would require methods or packages that have not been defined or declared.

Copilot generates code differently from traditional genetic programming paradigms. Sobania et al. [2022] found that Copilot could generate correct solutions to a variety of problems, with comparable accuracy to genetic programming. The improvement that Copilot makes over traditional approaches is that it was trained on a 159 GB dataset of unlabelled data, whereas genetic programming requires the dataset to be rigorously labeled. Thus, Sobania et al. claimed that the code Copilot generates is often easier to read and understand compared to that generated by genetic programming.

It is clear from the current research that Copilot is not a magic wand solution to problems in software engineering. With its flaws, how are developers able to benefit from this tool effectively? Dakhel et al. [2022] compared Copilot's proposed solutions against human programmers on a set of fundamental problems. The authors found that Copilot was unable to understand certain limitations embedded in the problem description, and while understanding the fundamentals of what is being requested, it can sometimes create errors when it does not understand. They conclude that while the generated code from Copilot is not a perfect solution, it can still be incorporated and modified into a project for faster early-stage prototyping.

Continuing on the question of usability of Copilot-generated code by developers, Vaithilingam et al. [2022] performed a user study. They found that using Copilot did not improve the task completion time or success rate, but that the developers preferred to start from the Copilot-generated code rather than a blank slate. The researchers also pointed out that it was challenging for developers to work with long blocks of generated code (and indeed, some developers failed to complete on time due to difficulties debugging generated code)—attempting to verify the code would also be a challenge, though one that might provide more understanding. One difference between our study and theirs, though, is that we did not encounter generated solutions that were known incorrect, and their users did.

Any introduction of small errors can potentially compound over time, creating large vulnerabilities in a software stack. Pearce et al. [2022] investigated code generated by Copilot for known security vulnerabilities. The authors prompted Copilot to generate code for 89 cybersecurity scenarios. They found that approximately 40% of generated code snippets were vulnerable to some form of known attack. Also, Siddiq et al. [2022] found code smells (including security smells) in training sets used for transformer-based code generation models for Python as well as in suggestions produced by Copilot. If code generated by Copilot is to be used in production, it needs to be rigorously vetted for exploits (and keeping in mind the caveat from Vaithilingam et al. about understanding Copilot's generated code).

---

<sup>2</sup><https://leetcode.com/>

### 3 METHODOLOGY

Our goal is to apply verification techniques to code generated by GitHub Copilot. We came up with sample problems and tasked Copilot with generating solutions in Python. We then used Dafny to verify that Copilot’s generated code satisfies the requirements.

More specifically, our methodology is as follows. First, we chose 6 problems to give to Copilot (five that we chose ourselves and one from [Nguyen and Nadi \[2022\]](#)). We wanted to test whether Copilot could generate a function which would solve the problems; by “solve”, we mean that the generated code satisfies the problem’s formal specifications. Thus, we wrote formal specifications in Dafny for each problem, consisting of *requires* and *ensures* statements.

Having defined the problems, we then asked Copilot to generate solutions in Python. To give Copilot the best chance of generating a good solution, we gave it a full function signature, including a descriptive function name, parameter name(s), parameter type(s), and an output type. With this input, Copilot then generated the function implementation. Since our goal was to test the capabilities of Copilot, we always chose the first suggested function.

Next, we carefully translated each Copilot implementation from Python to Dafny and added the specifications that we had created. We took care to preserve the behaviour of the implementation. In all cases, the implementations contained a loop of some type. Hence, the provided *requires* and *ensures* statements were inadequate to generate a proof that verified the implementation. We thus attempted to manually derive loop invariants so that Dafny could verify the implementation. During the verification process, we added extra *requires* and *ensures* clauses, if they were needed to verify the program, along with additional helper functions and methods when necessary.

### 4 RESULTS

We have attempted to verify 6 Copilot implementations, one of which intersects the LeetCode problems in [\[Nguyen and Nadi 2022\]](#). Table 1 summarizes the algorithms, implementations, and verification attempts. Incidentally, during testing, consistent with [Nguyen and Nadi \[2022\]](#), we found that Copilot would periodically include packages or methods that had not been defined; it happened to not do so for any of the implementations we describe here.

We have archived the current version of the Copilot-generated implementations and our Dafny code at <https://zenodo.org/record/7040924>.

Table 1. Algorithm implementations generated by Copilot. Dafny verification attempt statistics.

Algorithm Name	Python LOC	Verification Difficulty	Verified	Specification LOC	# Invariants
Binary Search (4.1)	12	Easy	✓	5	3
Two Sums (4.2)	5	Medium	✓	6	7
Largest Sum	10	Medium	✓	2	3
Sort Array	6	Medium	✓	4	11
Prime Factors (4.3)	12	Hard		3	> 23
Max Heapify (4.4)	14	Hard		4	> 4

Figure 1 presents the Copilot prompts that we used to generate implementations.

We were able to create a valid proof in Dafny for 4 of the 6 solutions generated by Copilot; we next discuss 2 of the valid proofs and our 2 failed attempts.

```

def binary_search(arr: list, target: int) -> int:
def twoSum(self, nums, target): // LeetCode; types on omitted subsequent lines
def largest_sum(nums: list) -> int:
def sortArray(arr: list) -> list:
def primeFactors(val: int) -> list:
def maxHeapify(arr: list) -> list:

```

Fig. 1. Copilot prompts for our algorithms.

#### 4.1 Binary Search

“Binary Search” is an algorithm which searches a sorted array for the index of a target element. The algorithm is appealing as it has a worst-case linear time complexity  $O(n)$ . Implementations are recognizable and simple to understand, making this algorithm a good candidate for our first verification example.

This example was inspired by a Dafny guide written by Arie Gurfinkel<sup>3</sup>. Our first observation was that Copilot’s implementation of Binary Search is different from that in the Dafny guide, leading to a different verification process.

We constructed formal specifications for binary search. As input, the implementation will take an array of distinct sorted integers and an integer target. It should then search the array for the target, returning the index of the target if found. If the target is not found, the program should have a return value ( $r.v$ ) of  $-1$ . These specifications can be expressed as:

```

requires Sorted( $arr$ )
requires Distinct( $arr$ )
ensures  $r.v \neq -1 \implies found(arr, target, r.v)$ 
ensures  $r.v = -1 \implies not\_found(arr, target)$ 

```

In the specifications above, *Sorted*, *Distinct*, *found*, and *not\_found* are all predicates. *Sorted* checks whether an array is sorted in ascending order and *Distinct* checks whether all elements of an array are distinct. We require that valid inputs satisfy these predicates. Next, we have the predicate *found*, which checks to see if the target was found at the index returned by the algorithm. This predicate should be true when the returned index is not  $-1$ . To handle the case where the target is not found, we have predicate *not\_found* that checks that the target does not exist in the array. This predicate is implied to be true when the return value is  $-1$ .

In parallel with manually producing specifications, we used Copilot to generate an implementation, given the following signature:

```

1 def binary_search(arr: list, target: int) -> int:

```

<sup>3</sup><https://ece.uwaterloo.ca/~agurfink/stqam/rise4fun-Dafny>

We accepted the first suggested implementation by Copilot based on the provided input:

```

1 def binary_search(arr: list, target: int) -> int:
2     low = 0
3     high = len(arr) - 1
4     while low <= high:
5         mid = (low + high) // 2
6         if arr[mid] == target:
7             return mid
8         elif arr[mid] < target:
9             low = mid + 1
10        else:
11            high = mid - 1
12    return -1

```

As stated above, we manually translated the implementation to Dafny, and added the formal specifications that we came up with. This also involved defining the predicates *Sorted* and *Distinct*. Given a specification and implementation, we could start verification.

However, Dafny also requires loop invariants. To provide them, we needed to (manually) reason about what happens in the implementation’s while loop. From our familiarity with binary search, we know that it uses two indices, high and low. High starts at the last element of the array; low starts at the first. In each iteration, the indices move closer together, closing the window of elements where the target can exist. In particular, we know that the target cannot exist at indices less than or equal to that of the low pointer or at indices greater than equal to that of the high pointer. This logic was to implement the primary loop invariant of the algorithm. In first-order logic,

$$\forall i \cdot 0 \leq i \leq low \wedge high \leq i < arr.Length \implies arr[i] \neq target.$$

With the loop invariants, Dafny was able to verify the implementation. So, in this case, Copilot was able to successfully generate an implementation which satisfied our formal specifications for binary search. This implementation could be verified without adding any extra *requires* or *ensures* statements beyond what we had initially identified as the specification.

## 4.2 Two Sums

We next picked a problem from LeetCode that was used in the study by [Nguyen and Nadi \[2022\]](#). “Two sums” takes an array of integers and an integer target, and returns two distinct indices in the array where the elements at those indices sum to the target. The Copilot input was:

```

1 def twoSum(self, nums, target):
2     """
3     :type nums: List[int]
4     :type target: int
5     :rtype: List[int]
6     """

```

The solution that Copilot generated passed all of the LeetCode test cases. LeetCode provides a number of constraints on the inputs—notably that a valid answer always exists. We used the constraints to formulate the following *requires* and *ensures* clauses:

```

requires 2 ≤ nums.Length
requires ∃i, j · (0 ≤ i < j < nums.Length ∧ nums[i] + nums[j] == target)
ensures  index1 ≠ index2
ensures  0 ≤ index1 < nums.Length
ensures  0 ≤ index2 < nums.Length
ensures  nums[index1] + nums[index2] == target

```

We converted the Python to Dafny and, with moderate effort, we then managed to write invariants that enabled Dafny to verify the generated solution. The 5 lines of Python code resulted in 8 lines of Dafny. We added 7 invariants (3 in the outer loop and 4 in the inner loop) to enable the verification.

Hence, in this case, Copilot generated a solution that passes the LeetCode test cases and that we managed to verify.

### 4.3 Prime Factors

“Prime Factors” is an algorithm that takes an integer and returns a list of prime numbers, which, when multiplied together, reproduce the original number. For example, the number 18 decomposes into the list [2, 3, 3]. The algorithm is  $O(n)$ , but does contain a nested loop.

We formally specified this algorithm; it accepts an integer greater than or equal to 2, and returns a list of prime factors of the integer. Our (arithmetic-heavy) specifications were:

```

requires val ≥ 2
ensures  ∏i=0len(arr) arr[i] == val

```

With these simple formal specifications defined, we were ready to prompt Copilot to generate our algorithm. We provided Copilot with the following function signature:

```
1 def primeFactors(val: int) -> list:
```

which prompted Copilot to suggest the following Python implementation:

```

1 def primeFactors(val: int) -> list:
2     factors = []
3     while val % 2 == 0:
4         factors.append(2)
5         val = val // 2
6     for i in range(3, int(val ** 0.5) + 1, 2):
7         while val % i == 0:
8             factors.append(i)
9             val = val // i
10    if val > 2:
11        factors.append(val)
12    return factors

```

To convert to Dafny: 1) As with previous implementations, we adapted Python for loops into while loops. 2) We had to implement integer square root in Dafny. Verifying integer square root is nontrivial; however, we found a verified implementation<sup>4</sup> which we used as a library. 3) We had to choose a Dafny data structure equivalent to a dynamic array.

For dynamic arrays (the Python equivalent of `list`), we used the Dafny `seq<int>` (sequence of integers), as it had a built in concatenate method which acts by *adding* to sequences to produce a new sequence which can be assigned. However, initialization was a challenge: Dafny does not permit empty sequences. We thus initialized our sequence with multiplicative unit 1; though this is not identical to the Python solution, it is reasonably equivalent.

We also created a function for multiplying the elements of the sequence, adding invariants to check that the current product times the remaining value produces the original value.

Using our chosen invariants, we encountered Dafny timeout issues and failed to verify the implementation. It may be the case that our invariants are correct and Dafny is unable to verify them in a reasonable timeframe, but we don't know that. As a last resort, to try to help the proof resolve with "what we know to be true", we added *lemmas* which assumed false, allowing us to say two specific *ensures* clauses were true:

$$\mathbf{ensures} \left( \prod_{i=0}^{\text{len}(\text{arr})} \text{arr}[i] \right) \cdot v == \text{val}$$

$$\mathbf{ensures} (v == 1) \implies \left( \prod_{i=0}^{\text{len}(\text{arr})} \text{arr}[i] \right) == \text{val}$$

before and after each *append* of a new value to the list of factors. Even assuming these two clauses to be true, the Dafny solver continued to hit the time limit for finding a proof. Thus, although we believe that Copilot generated a correct implementation, we have not yet been able to verify it.

#### 4.4 Max Heapify

"Max Heapify" is an algorithm for turning an array into a max-heap data structure. The array is conceptually a complete binary tree, where each index holds the value of a particular node, and that node's left and right children's indices are found using the following mapping:

$$\begin{aligned} \text{left\_child} &= (\text{index} * 2) + 1 \\ \text{right\_child} &= (\text{index} * 2) + 2. \end{aligned}$$

A max-heap satisfies the property that each node is greater than or equal to its left and right child. Since this property is recursive, it implies that any particular node is (transitively) larger than every node in both its right and left children, so the largest element in the tree should be at the root.

With this in mind, we wrote formal specifications:

$$\begin{aligned} \mathbf{requires} \quad & \text{arr is a list of distinct integers} \\ \mathbf{ensures} \quad & \forall x \cdot (0 \leq x < \text{len}(\text{arr})) \implies \\ & ((2x + 1 < \text{len}(\text{arr})) \implies (\text{arr}[x] \geq \text{arr}[2x + 1])) \wedge \\ & ((2x + 2 < \text{len}(\text{arr})) \implies (\text{arr}[x] \geq \text{arr}[2x + 2])). \end{aligned}$$

That is, the input array's values are unique; and the resulting list has the properties of a max-heap. Simultaneously, we provided the following function signature to Copilot:

<sup>4</sup><https://homepage.cs.uiowa.edu/~tinelli/classes/181/Spring11/Tools/Dafny/Examples/square-root.dfy>

```
1 def maxHeapify(arr: list) -> list:
```

Copilot produced a 14-line Python implementation, which we translated into Dafny and attempted to verify. We designed a predicate based on the *ensures* clause, which checks that each node is greater than or equal to the node at its childrens' indices, if those indices are not out of bounds for the array.

We noticed that the Copilot implementation differs from more traditional implementations by sifting down rather than sifting up: the recursive call to `maxHeapify` is given the entire (mutated) array. Often, implementations of max-heap traverse all non-leaf nodes backwards to the root and sift-up (with a recursive call) starting at any nodes that are out of place.

Prior to formal Dafny verification, we tested the the generated algorithm on a series of random inputs to ensure it produced a max-heap each time. Inspecting the algorithm with Python debugging tools showed that the trace-tree for this implementation goes quite deep. The implementation appears to make recursive calls, starting from the root, enough times to ensure that eventually the max-heap property holds. This would make finding a verifiable proof difficult. Instead of the standard *heapify* algorithm running in  $O(n \log n)$ , the version Copilot generated for us appears to run in time  $O(n^2)$ .

We tried a number of different invariants to show that the array was *partially* heapified for each step in the loop. However, we were unable to prove termination of the loop to begin with. For this reason, we were unable to formally verify the generated algorithm.

## 5 DISCUSSION

Copilot is undoubtedly good at generating code. In this work, we took a formal verification approach to investigating the code that Copilot generates, using Dafny. Of the 6 examples we tried, we managed to verify the 4 simpler problems. One of the problems appeared in [Nguyen and Nadi 2022], which used LeetCode test suites to establish correctness of Copilot-generated code.

Our work produced corroborating evidence that Copilot's generated code might be dubious. It also explored the difficulty of using Dafny to verify implementations; although Dafny is powerful, it can be nontrivial to verify implementations using it.

We believe that our work reinforces the fact that to effectively deploy Copilot, one should know what is needed. Similar to the work by Dakhel et al. [2022], we found that Copilot appears to skip over certain constraints imposed on the input space. As a specific example, the "two sums" example assumes that a solution exists; we had to manually record that assumption in our specification. Humans are not good at reasoning about edge cases [Wheeler 2021], and Copilot reproduces that bias. Thus, to productively use Copilot, it remains incredibly important to reason carefully about what the problem specifications should be.

We also found that verifying the generated software with Dafny was somewhat difficult, both due to Dafny implementation language limitations, and—more importantly—verifiability reasons. Some Python idioms do not translate cleanly to Dafny; we had to reformulate Python generator objects (e.g., `range` in a `for` loop) as `while` loops. Other Python data types such as `list` don't have an obvious substitution in Dafny. Dafny also omits math operations such as floor and square root, which Copilot generated in its Python solutions.

In the LeetCode study [Nguyen and Nadi 2022], not all of the Copilot-generated solutions passed the LeetCode test cases. In this work, we did not attempt to verify any solutions that failed their test cases. Of course, we expect such solutions to be unverifiable. However, attempting verification

could provide another way to understand how the solutions are incorrect, beyond simply inspecting the generated code. As is true of verification in general, an attempted verification can provide insight into the code, complementing a test-based lens.

We believe that algorithms designed with verifiability in mind are more likely to be easier to prove. A proof of correctness of binary search halves the search space at each step. If an implementation follows suit, we can focus on solving a subset of the problem in its proof. In reasoning about a sorting algorithm, we would like to be able to show that each step of a loop results in the array being partially sorted.

When we were choosing algorithms to verify, we thought that max-heap would be simple, as it has very clear sub-problems which, when solved, result in a well-defined data structure. However, the Copilot-generated implementation does not obviously reduce to sub-problems. We believe that an iterative version of the max-heap algorithm which walks backwards through the tree would have been easier to verify than the recursive version which Copilot generated for us.

Of course, we have not established that proving the generated max-heap implementation is impossible. Per [Leino \[2010\]](#), Dafny verification often requires an expert, and it is not an easy task for the average programmer. [Leino and Wüstholtz \[2014\]](#) have also identified that, for some problems, the solver timing out can be a barrier to verifiability. There are some shortcuts which can be taken in the proof to avoid timing out; however, these shortcuts require a great amount of practice to identify. Max-heapify has been formally verified [[Tafat and Marché 2011](#)] using other software verification techniques, but the proof and its description are both reasonably lengthy.

*Generalizability.* In this work, we attempted to verify 6 specific implementations generated by Copilot against specifications that we manually constructed. We found that we were able to verify 4 of the 6 implementations. The remaining 2 generated implementations may be correct, and others may be able to verify them, but we were not. We've thus shown that it is possible to verify some Copilot-generated code, and that other code is at least nontrivial to verify. We also know from [[Nguyen and Nadi 2022](#)] that, provided with a textual prompt and test cases, Copilot can generate code that does not output the expected results.

Our Python-mediated translation process meant that we avoided Dafny's treatment of programs with pointers and pointer arithmetic. Because we asked Copilot for Python implementations, and these implementations do not include pointer and memory management, we could not explore how Dafny can be used to verify functional correctness for pointer-based programs [[Leino 2010](#)].

We speculate that, even if we had chosen different target languages and verification tools, the fundamental difficulty of finding invariants would still be an issue with verifying code from Copilot. The fact that Copilot generated a recursive implementation that we could not verify means that we cannot blame Copilot for generating difficult-to-verify loops. Dafny's lack of square root implementation is, of course, a Dafny-specific problem.

Of course, our results would be stronger with more verification attempts, and across different domains as well. LeetCode examples, in particular, are more likely than average to be represented in the Copilot training corpus. Going beyond algorithmic examples (e.g. reading a CSV file, connecting to a database) would also be interesting, though harder to formally specify.

The other source of variability is expertise in using verification tools. We are not experts in using these tools, and hence our experience may be representative of average users, but the sample size of 2 verification users is small. A study involving more developers attempting to verify code would provide stronger results.

*Concluding thoughts.* In our experience, Copilot is a powerful tool for quickly producing (prototype) implementations: it can generate code faster than one could possibly write it. However, Copilot is certainly not a magic wand. A programmer using Copilot would be well advised to

not only look at what Copilot has produced, but also ensure it is correct and appropriate for the intended use case. If not with formal verification, one ought to at least rigorously vet the generated code with a test suite, prior to integrating it into any system. One should also consider copyright issues [Howard 2021]. Of course, in principle, one should vet any code before integrating into a system, whether human produced or machine produced.

We observed that Copilot was quite capable when generating solutions to well known or established problems, likely related to the fact that Copilot’s training data came from user data on GitHub. Like Dakhel et al. [2022], we found that Copilot generated sub-optimal and possibly incorrect solutions to some of our problems, e.g., Max Heapify. Due to how Copilot is trained, we do not believe that it would work as well on esoteric problems; it would be interesting to explore that point.

## REFERENCES

- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- Arghavan Moradi Dakhel, Vahid Majdinasab, Amin Nikanjam, Foutse Khomh, Michel C Desmarais, Zhen Ming, et al. 2022. GitHub Copilot AI pair programmer: Asset or Liability? *arXiv preprint arXiv:2206.15331* (2022).
- GitHub. 2022. Your AI pair programmer. (2022). <https://github.com/features/copilot> Accessed: 2022-07-31.
- Gavin D Howard. 2021. GitHub Copilot: Copyright, Fair Use, Creativity, Transformativity, and Algorithms. (2021).
- K Rustan M Leino. 2010. Dafny: An automatic program verifier for functional correctness. In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*. Springer, 348–370.
- K Rustan M Leino and Valentin Wüstholtz. 2014. The Dafny integrated development environment. *arXiv preprint arXiv:1404.6602* (2014).
- Nhan Nguyen and Sarah Nadi. 2022. An Empirical Evaluation of GitHub Copilot’s Code Suggestions. In *2022 IEEE/ACM 19th International Conference on Mining Software Repositories (MSR)*. IEEE, 1–5.
- Hammond Pearce, Baleegh Ahmad, Benjamin Tan, Brendan Dolan-Gavitt, and Ramesh Karri. 2022. Asleep at the Keyboard? Assessing the Security of GitHub Copilot’s Code Contributions. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 754–768.
- Mohammed Latif Siddiq, Shafayat H. Majumder, Maisha R. Mim, Sourov Jajodia, and Joanna C. S. Santos. 2022. An Empirical Study of Code Smells in Transformer-based Code Generation Techniques. In *2022 IEEE 22nd International Working Conference on Source Code Analysis and Manipulation (SCAM)*.
- Dominik Sobania, Martin Briesch, and Franz Rothlauf. 2022. Choose your programming Copilot: a comparison of the program synthesis performance of GitHub copilot and genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 1019–1027.
- Asma Tafat and Claude Marché. 2011. *Binary heaps formally verified in Why3*. Technical Report [Research Report] RR-7780. INRIA.
- Priyan Vaithilingam, Tianyi Zhang, and Elena L Glassman. 2022. Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–7.
- David A. Wheeler. 2021. The Apple goto fail vulnerability: lessons learned. <https://dwheeler.com/essays/apple-goto-fail.html>. (Jan 2021). Accessed 1 September 2022.