Expanding the Reach of Fuzz Testing

Caroline Lemieux

The University of British Columbia

CSER New Faculty Talk

June 7, 2023

Software Has Bugs

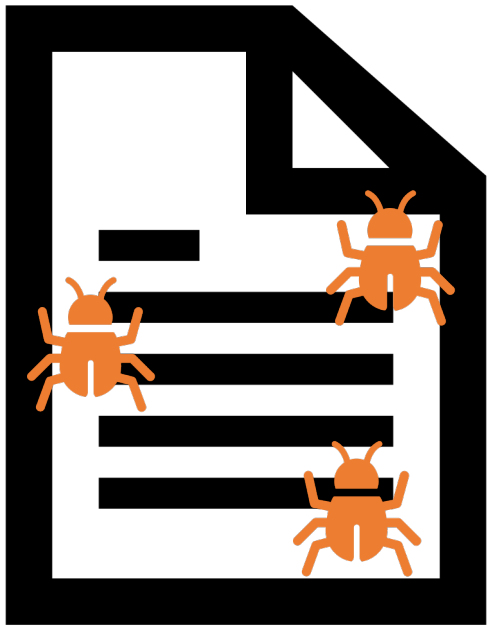

Bugs Have Increasing Consequences

Bugs Have Increasing Consequences

Badlock

Cloudbleed

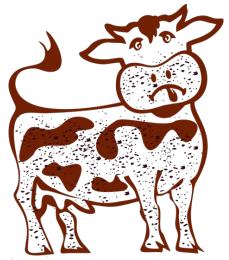

Dirty COW

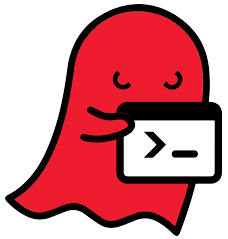

GHOST

Heartbleed

StageFright

ShellShock

Bugs Have Increasing Consequences

Badlock

Cloudbl

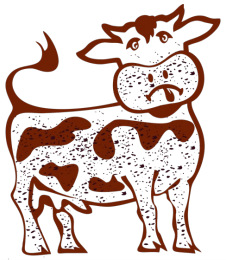

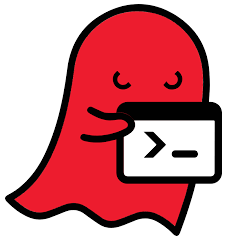

GHOST

Heartbleed

StageFright

ShellShock

the cost of
poor quality software is

$2.08 trillion

in the US in 2020

My Work:

build tools to help developers improve **correctness**, **security** and **performance** of software

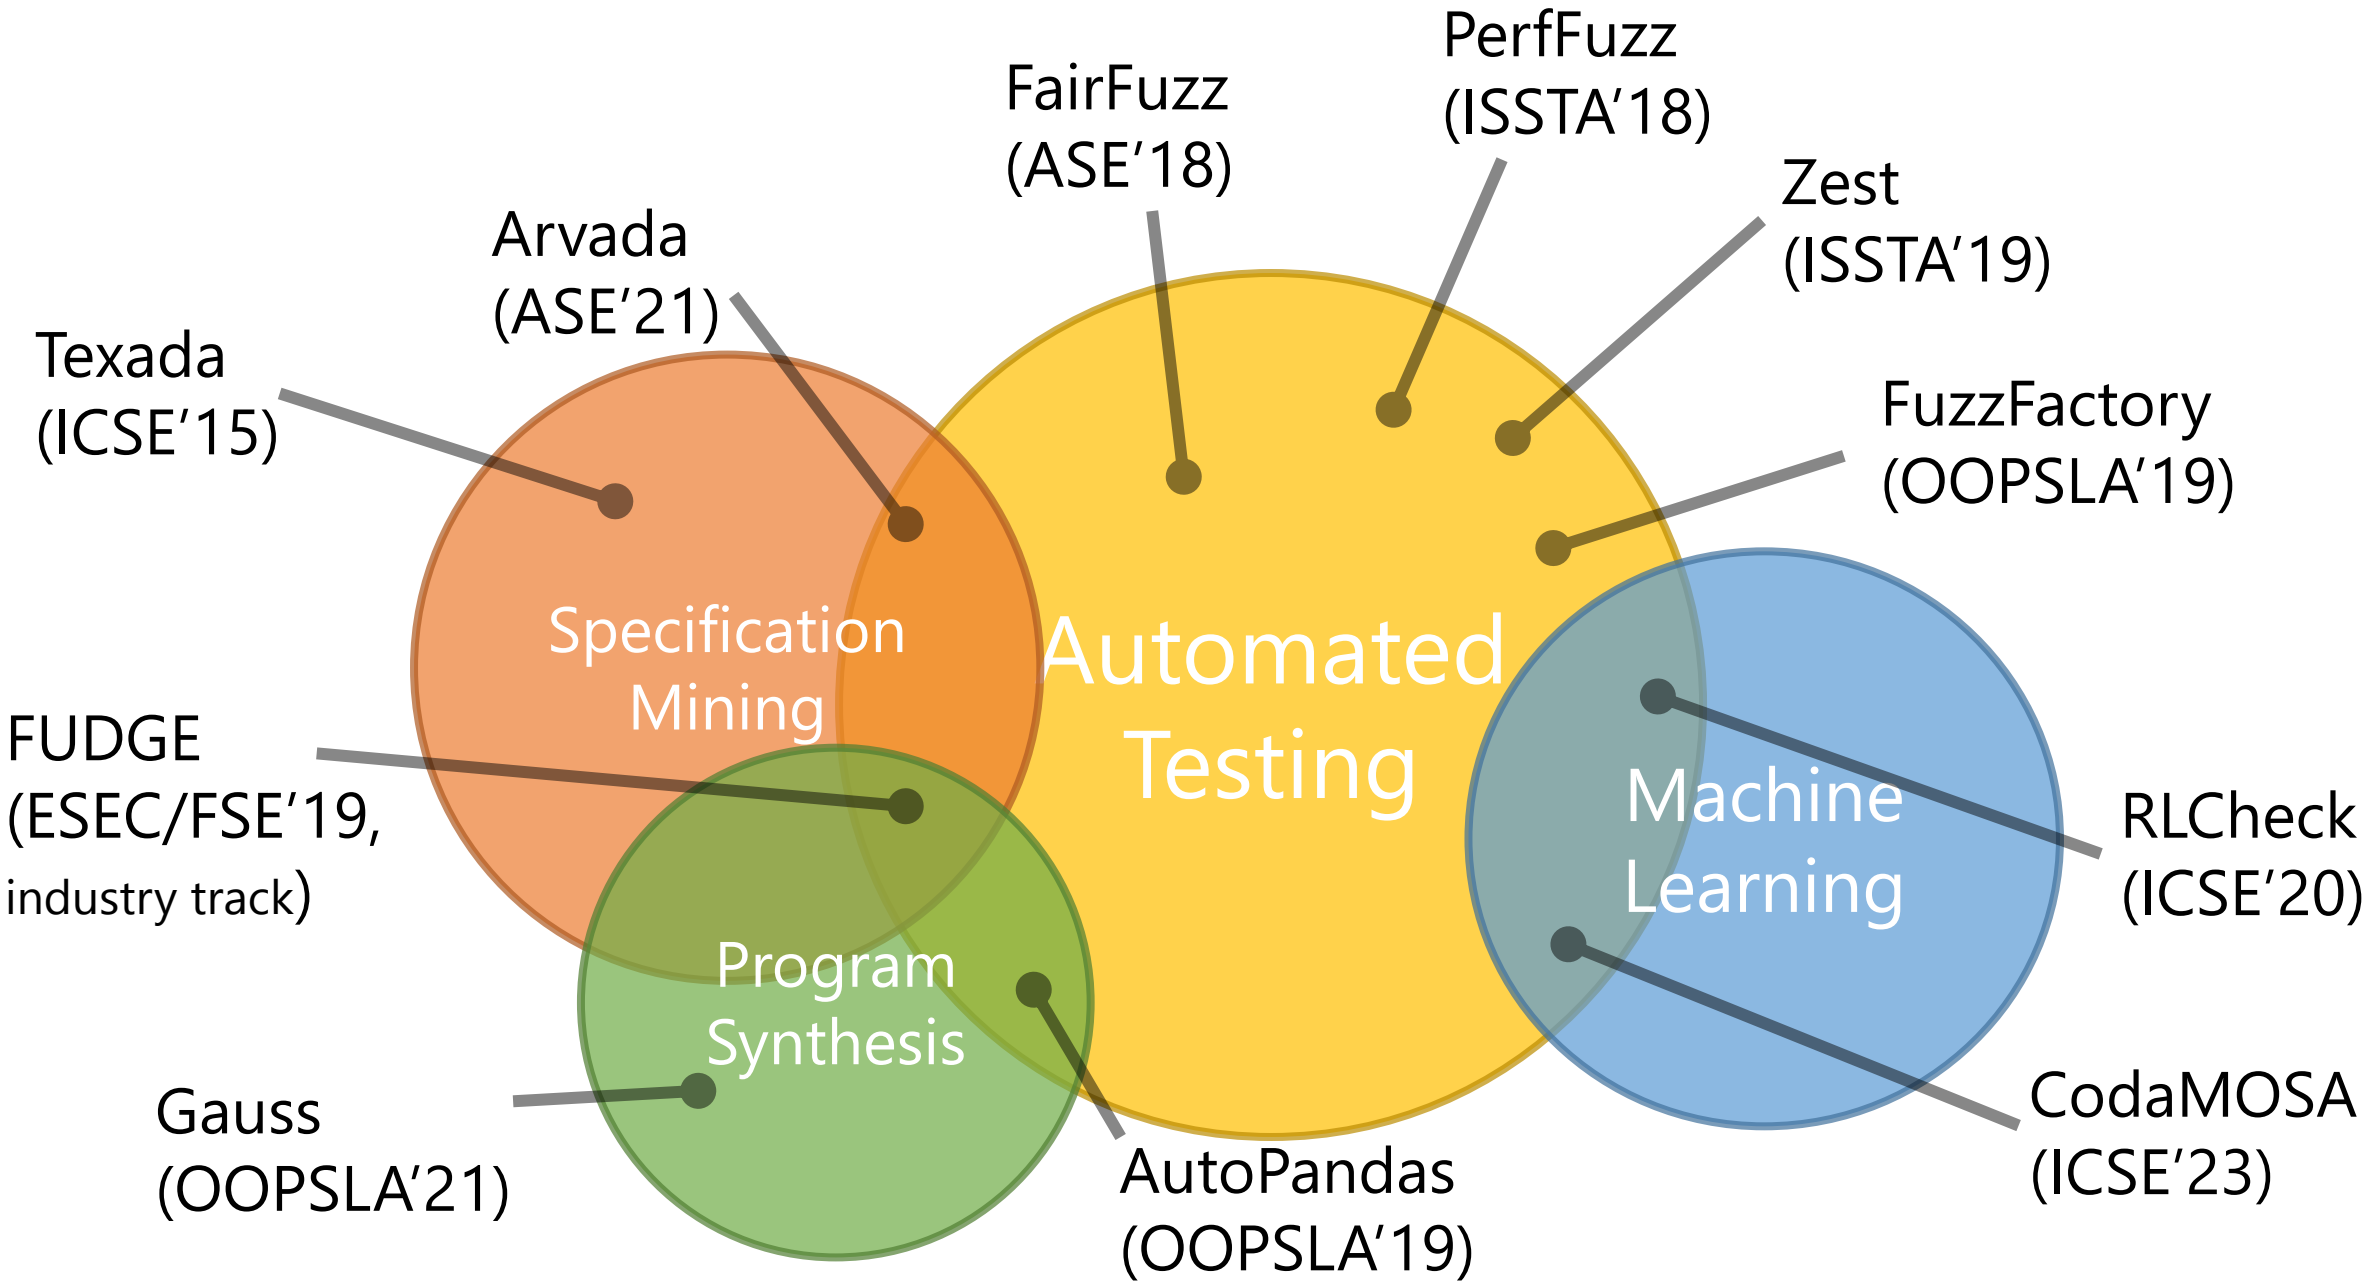
PerfFuzz
(ISSTA'18)
FairFuzz
(ASE'18)
Zest
(ISSTA'19)
Arvada
(ASE'21)
Texada
(ICSE'15)
FuzzFactory
(OOPSLA'19)
Specification
Mining
Automated
Testing
Machine
Learning
FUDGE
(ESEC/FSE'19,
industry track)
RLCheck
(ICSE'20)
Program
Synthesis
Gauss
(OOPSLA'21)
AutoPandas
(OOPSLA'19)
CodaMOSA
(ICSE'23)

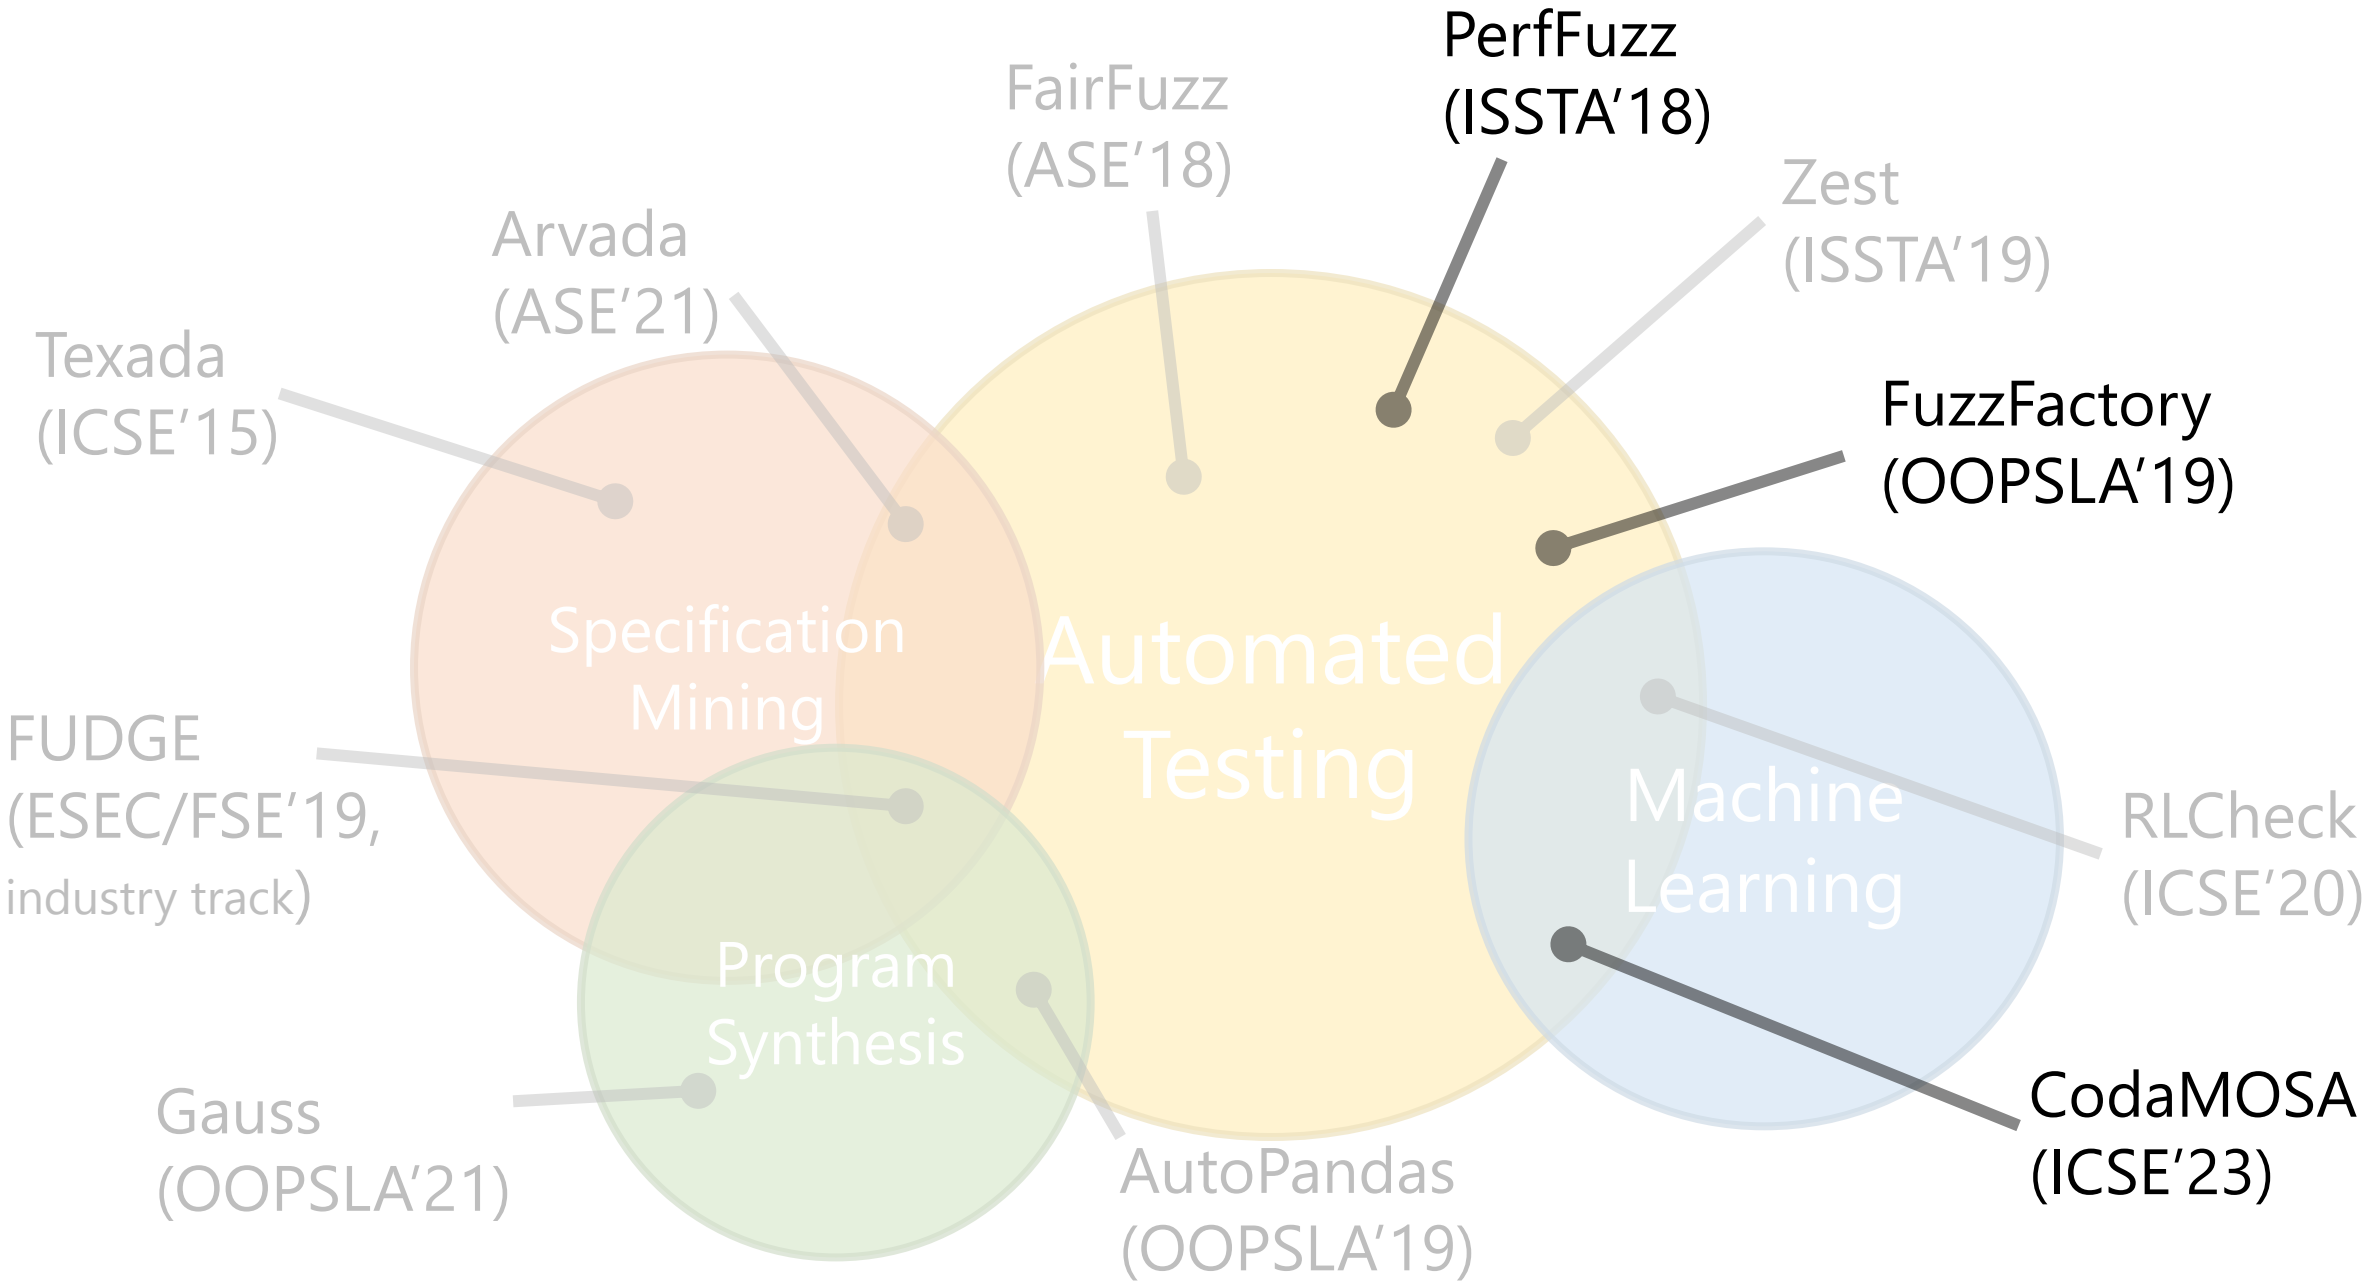
PerfFuzz
(ISSTA'18)
FairFuzz
(ASE'18)
Zest
(ISSTA'19)
Arvada
(ASE'21)
Texada
(ICSE'15)
FuzzFactory
(OOPSLA'19)
Specification
Mining
Automated
Testing
Machine
Learning
FUDGE
(ESEC/FSE'19,
industry track)
RLCheck
(ICSE'20)
Program
Synthesis
Gauss
(OOPSLA'21)
AutoPandas
(OOPSLA'19)
CodaMOSA
(ICSE'23)

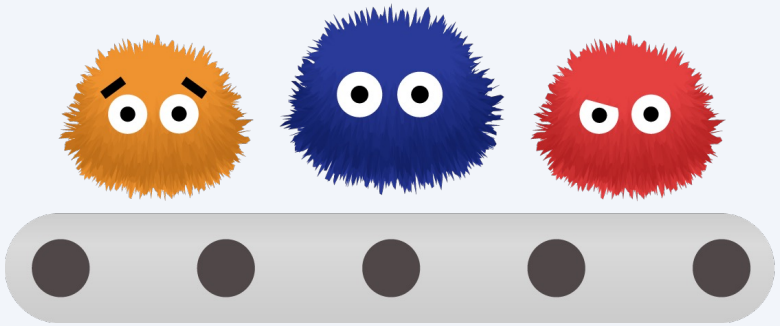

Using ***generalized feedback maps*** to expand ***bugs findable by fuzz testing***

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

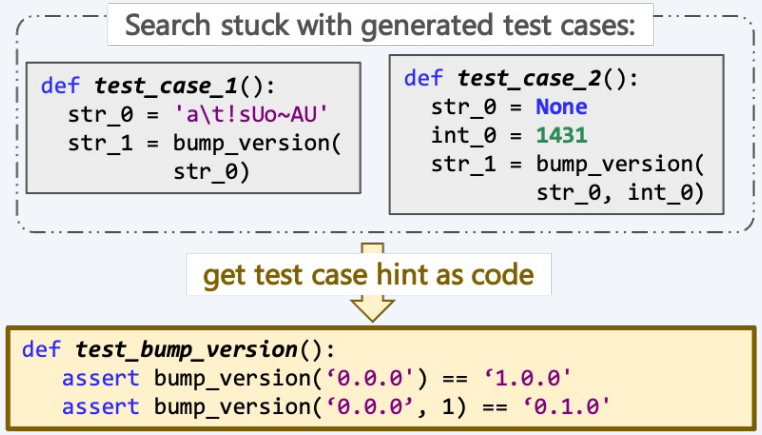

Using ***large language models*** to improve ***automated test suite generation***

CodaMOSA (ICSE'23)

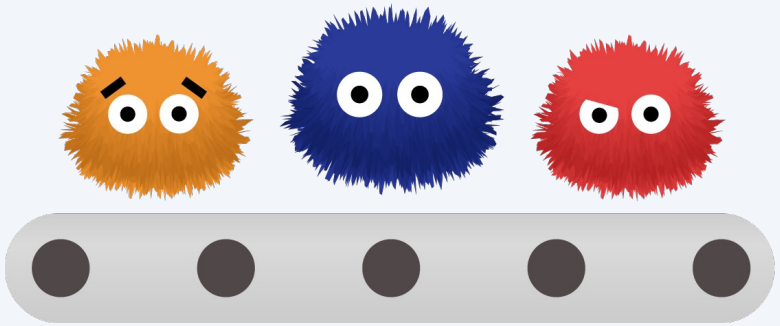

Using ***generalized feedback maps*** to expand ***bugs findable by fuzz testing***

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

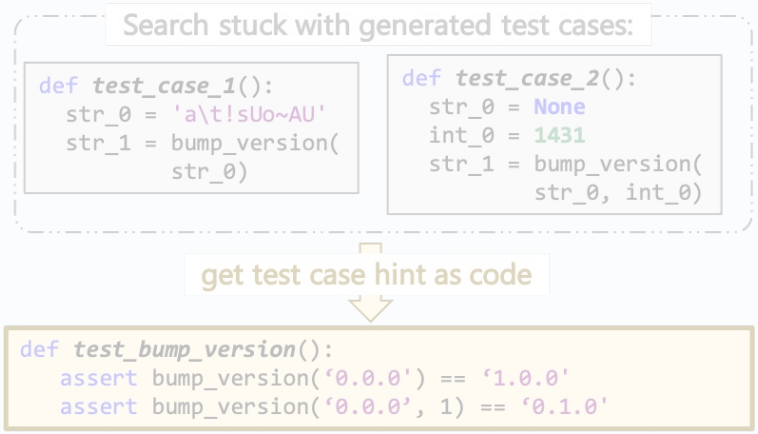

Using ***large language models*** to improve ***automated test suite generation***

CodaMOSA (ICSE'23)

Performance Bugs

Performance Bugs

Performance Bugs

Takes **1+ minute** to report syntax error:

```
((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((((e foo = 1; => 1;
```

Performance Bugs

CVE-2020-7212 Detail

Current Description

The _encode_invalid_chars function in util/url.py in the urllib3 library 1.25.2 through 1.25.7 for Python allows a denial of service (CPU consumption) because of an inefficient algorithm. The percent_encodings array contains all matches of percent encodings. It is not deduplicated. For a URL of length N, the size of percent_encodings may be up to O(N). The next step (normalize existing percent-encoded bytes) also takes up to O(N) for each step, so the total time is O(N^2). If percent_encodings were deduplicated, the time to compute _encode_invalid_chars would be O(kN), where k is at most 484 ((10+6*2)^2).

+View Analysis Description

Severity CVSS Version 3.x | CVSS Version 2.0

CVSS 3.x Severity and Metrics:

NVD **NIST:** NVD **Base Score:** 7.5 HIGH

Takes **1+ minute** to report syntax error:

```
((((((((((((((((((((((((((((((((((((((((((((((((((((((((e foo = 1; => 1;
```

Performance Bugs

CVE-2020-7212 Detail

Current Description

The _encode_invalid_chars function in util/url.p
through 1.25.7 for Python allows a denial of serv
because of an inefficient algorithm. The percent
all matches of percent encodings. It is not dedup
N, the size of percent_encodings may be up to O
(normalize existing percent-encoded bytes) also
step, so the total time is O(N^2). If percent_enco
the time to compute _encode_invalid_chars wo
most 484 ((10+6*2)^2).

View Analysis Description

Severity CVSS Version 3.x | CVSS Version 2.0

CVSS 3.x Severity and Metrics:

NVD **NIST:** NVD **Base Score:** 7.5 HIGH

Takes 1+ minute to report syntax error:

```
((((((((((((((e foo = 1; => 1;
```

Can we find inputs revealing these bugs automatically? 🤔

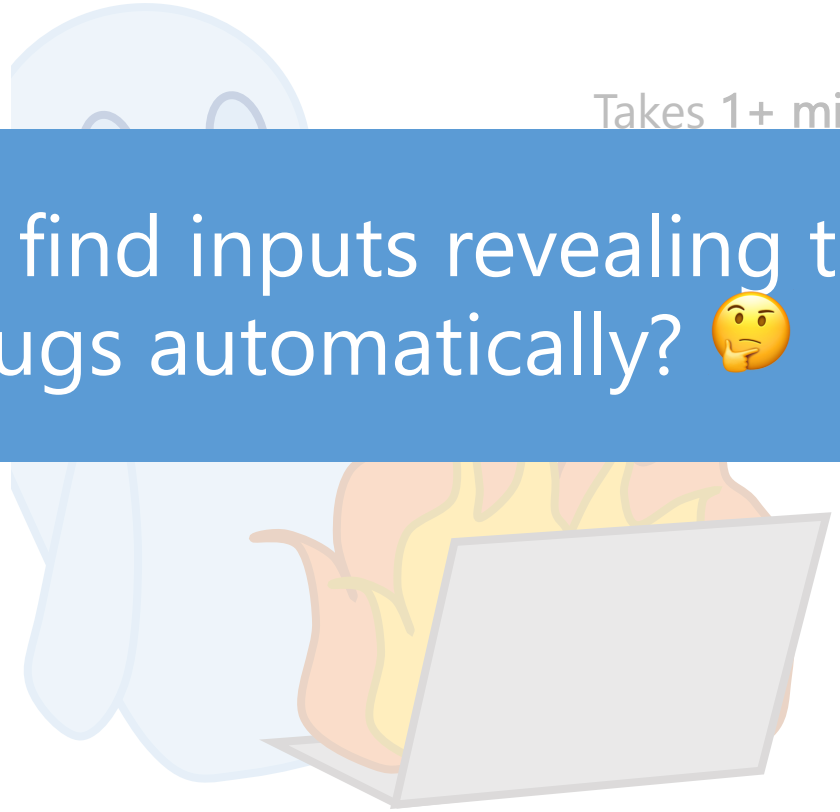

The bug-o-rama trophy case

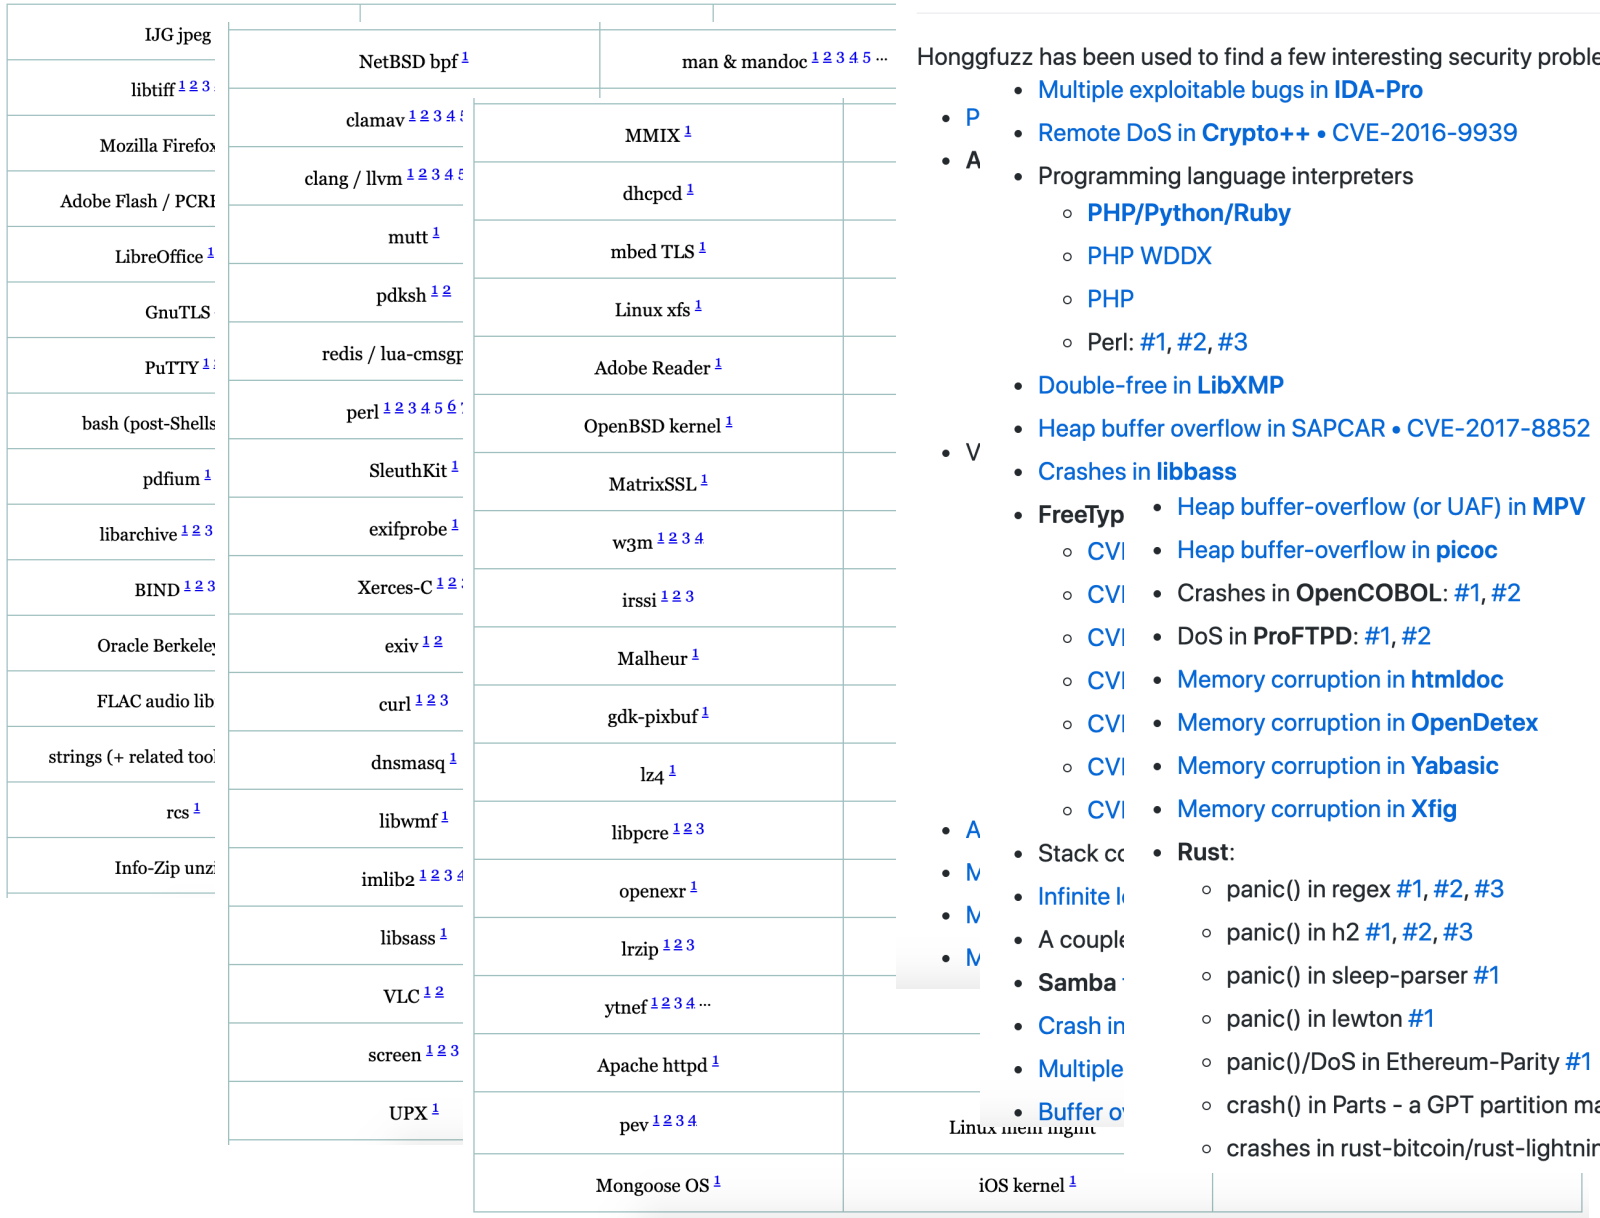

Trophies

Honggfuzz has been used to find a few interesting security problems in major software packages; An incomplete list:

- P
- A
- V
- A
- M
- M
- M
- Multiple exploitable bugs in **IDA-Pro**
- Remote DoS in **Crypto++** • CVE-2016-9939
- Programming language interpreters
  - **PHP/Python/Ruby**
  - PHP WDDX
  - PHP
  - Perl: #1, #2, #3
- Double-free in **LibXMP**
- Heap buffer overflow in SAPCAR • CVE-2017-8852
- Crashes in **libbass**
- **FreeTyp**
  - CVI
  - CVI
  - CVI
  - CVI
  - CVI
  - CVI
  - CVI
- Stack cc
- Infinite l
- A couple
- **Samba**
- Crash in
- Multiple
- Buffer o
  - Heap buffer-overflow (or UAF) in **MPV**
  - Heap buffer-overflow in **picoc**
  - Crashes in **OpenCOBOL**: #1, #2
  - DoS in **ProFTPD**: #1, #2
  - Memory corruption in **htmldoc**
  - Memory corruption in **OpenDetex**
  - Memory corruption in **Yabasic**
  - Memory corruption in **Xfig**
- **Rust**:
  - panic() in regex #1, #2, #3
  - panic() in h2 #1, #2, #3
  - panic() in sleep-parser #1
  - panic() in lewton #1
  - panic()/DoS in Ethereum-Parity #1
  - crash() in Parts - a GPT partition manager #1
  - crashes in rust-bitcoin/rust-lightning #1

The bug-o-rama trophy case

IJG jpeg	NetBSD bpf [1]	man & mandoc [1 2 3 4 5 ...]	
libtiff [1 2 3]	clamav [1 2 3 4 5]	MMIX [1]	
Mozilla Firefox	clang / llvm [1 2 3 4 5]	dhcpcd [1]	
Adobe Flash / PCRE	mutt [1]	mbed TLS [1]	
LibreOffice [1]	pdksh [1 2]	Linux xfs [1]	
GnuTLS	redis / lua-cmsgp	Adobe Reader [1]	
PuTTY [1]	perl [1 2 3 4 5 6]	OpenBSD kernel [1]	
bash (post-Shells	SleuthKit [1]		
pdfium [1]	exifprobe [1]		
libarchive [1 2 3]	Xerces-C [1 2]		
BIND [1 2 3]	exiv [1 2]		
Oracle Berkeley	curl [1 2 3]		
FLAC audio lib	dnsmasq [1]	lz4 [1]	
strings (+ related too	libwmf [1]	libpcre [1 2 3]	
rcs [1]	imlib2 [1 2 3 4]	openexr [1]	
Info-Zip unzi	libsass [1]	lrzip [1 2 3]	
	VLC [1 2]	ytnef [1 2 3 4 ...]	
	screen [1 2 3]	Apache httpd [1]	
	UPX [1]	pev [1 2 3 4]	Linux mem mgmt
		Mongoose OS [1]	iOS kernel [1]

Trophies

Honggfuzz has been used to find a few interesting security problems in major software packages; An incomplete list:

- Multiple exploitable bugs in **IDA-Pro**
- Remote DoS in **Crypto++** • CVE-2016-9939
- Programming language interpreters
  - **PHP/Python/Ruby**
  - PHP WDDX
  - PHP
  - Perl: #1, #2, #3
- Double-free in **LibXMP**
- Memory corruption in **Yabasic**
- Memory corruption in **Xfig**
- **Rust**:
  - panic() in regex #1, #2, #3
  - panic() in h2 #1, #2, #3
  - panic() in sleep-parser #1
  - panic() in lewton #1
  - panic()/DoS in Ethereum-Parity #1
  - crash() in Parts - a GPT partition manager #1
  - crashes in rust-bitcoin/rust-lightning #1

How does it work?

Coverage-Guided Fuzzing

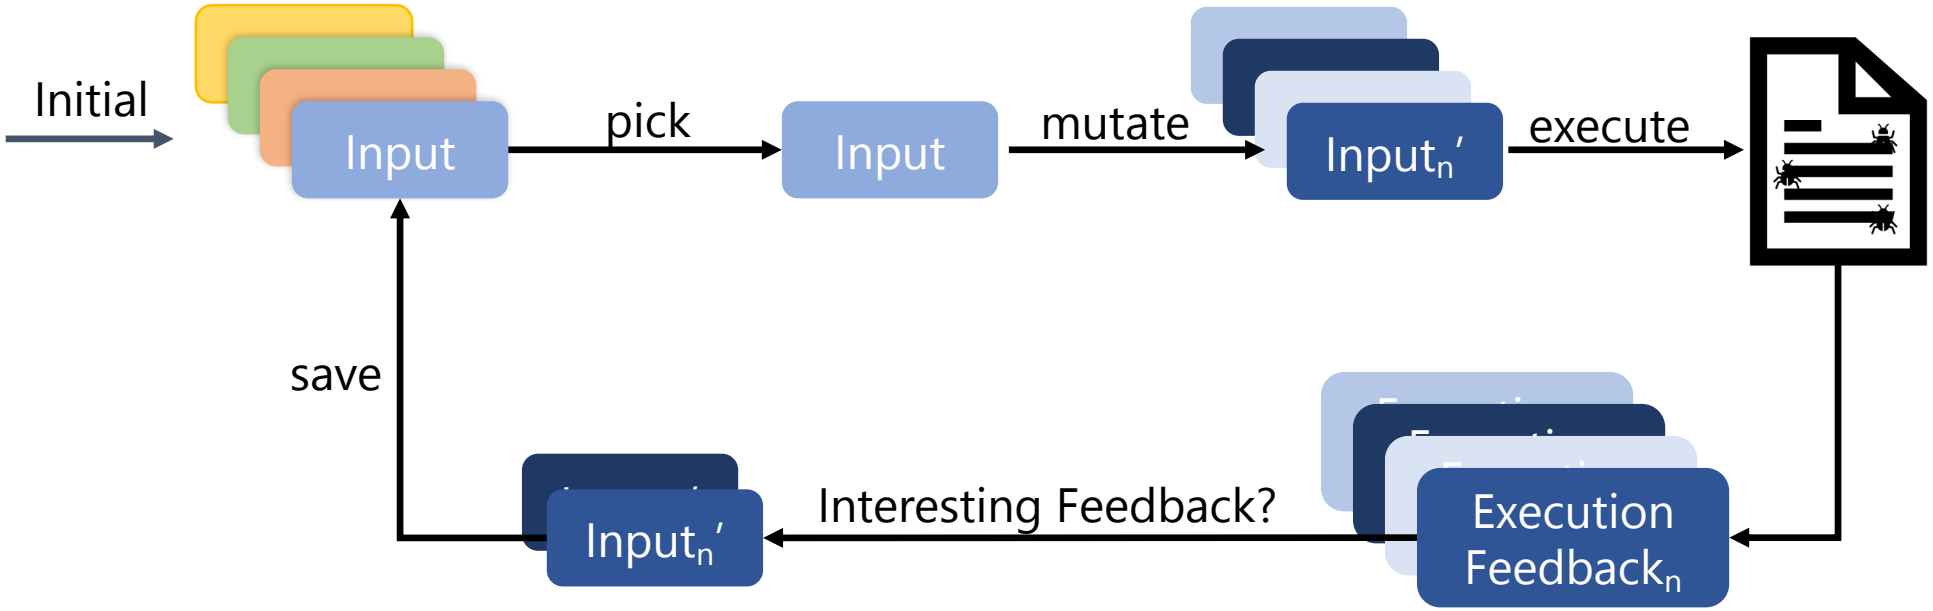

Coverage-Guided Fuzzing

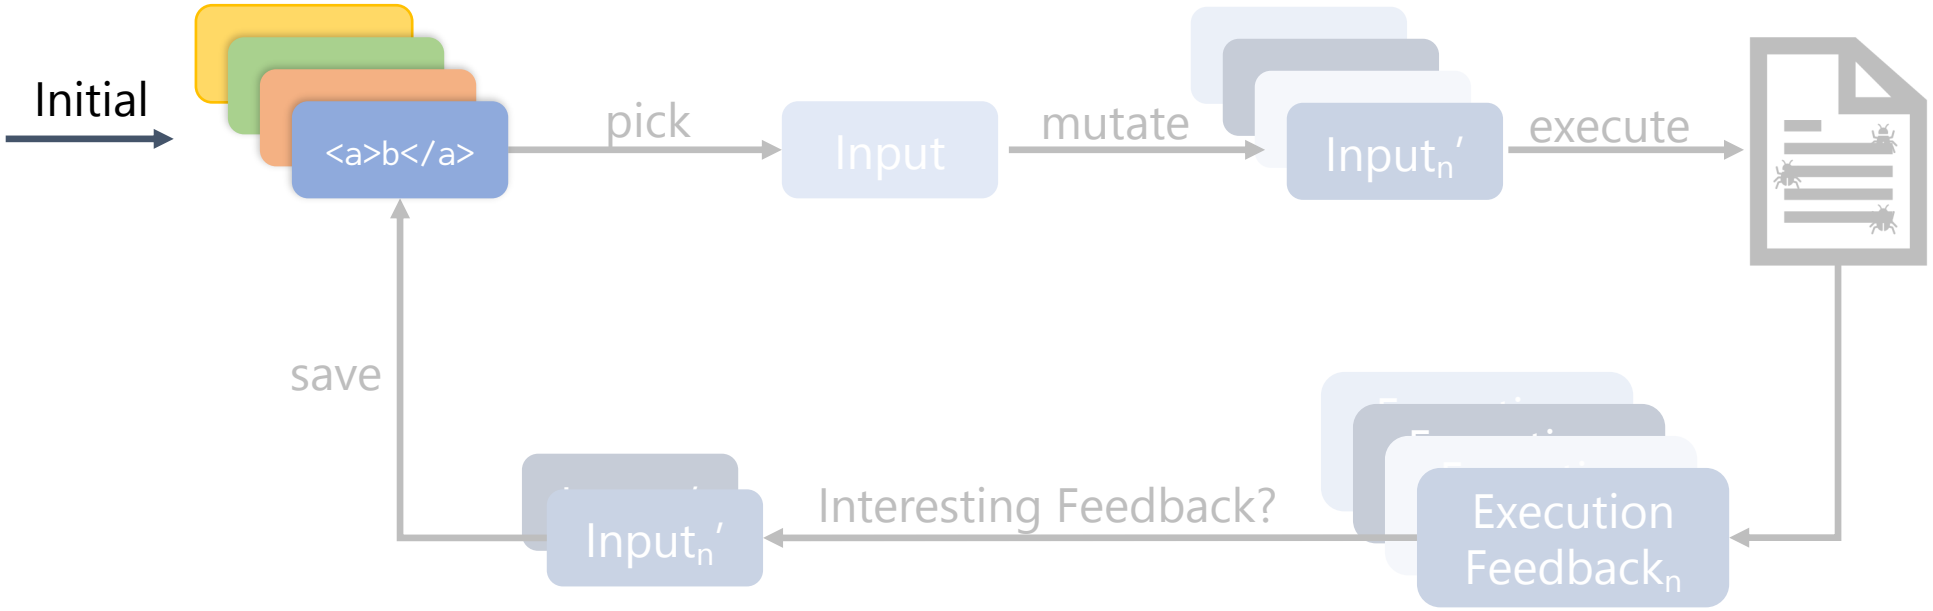

Coverage-Guided Fuzzing

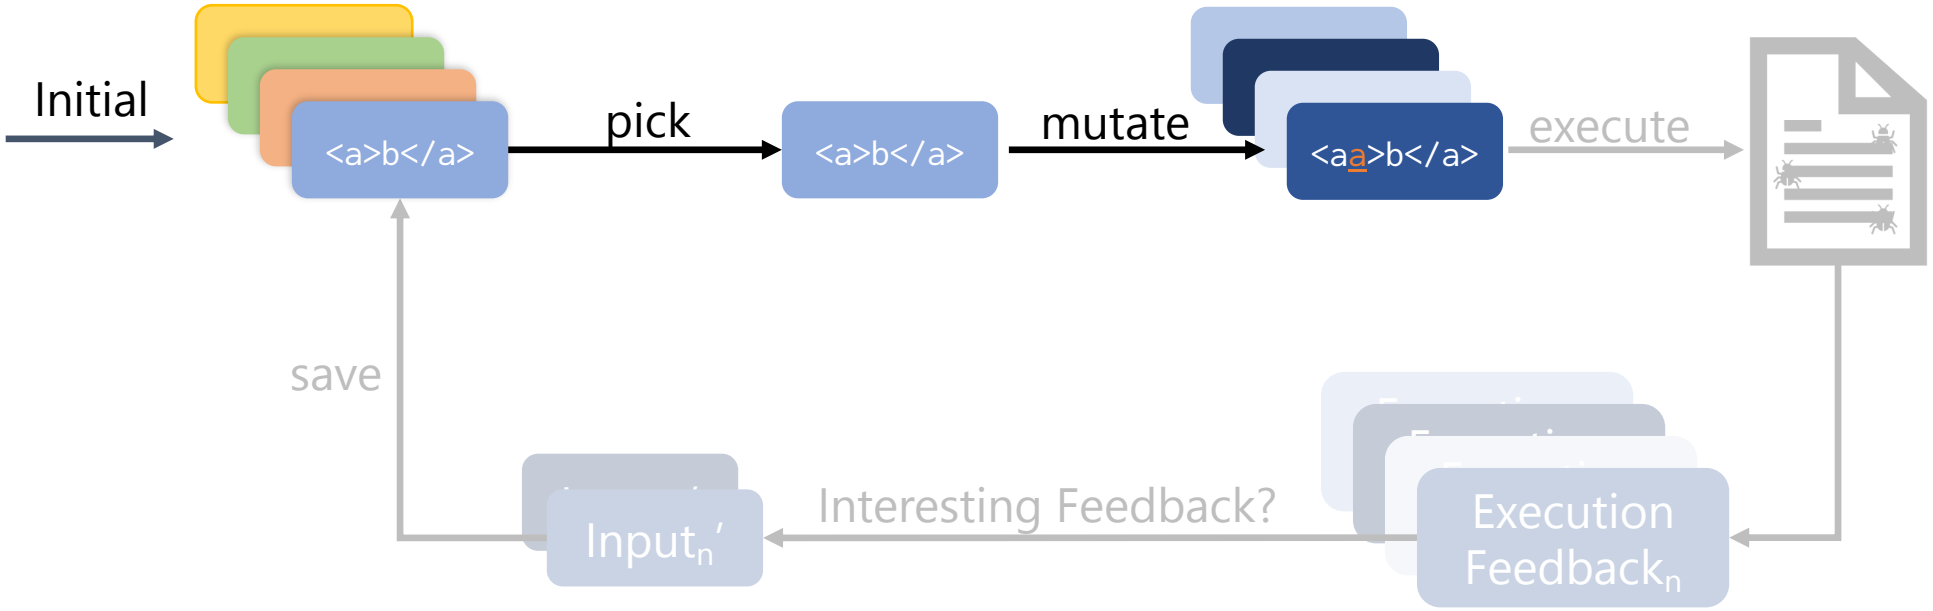

Coverage-Guided Fuzzing

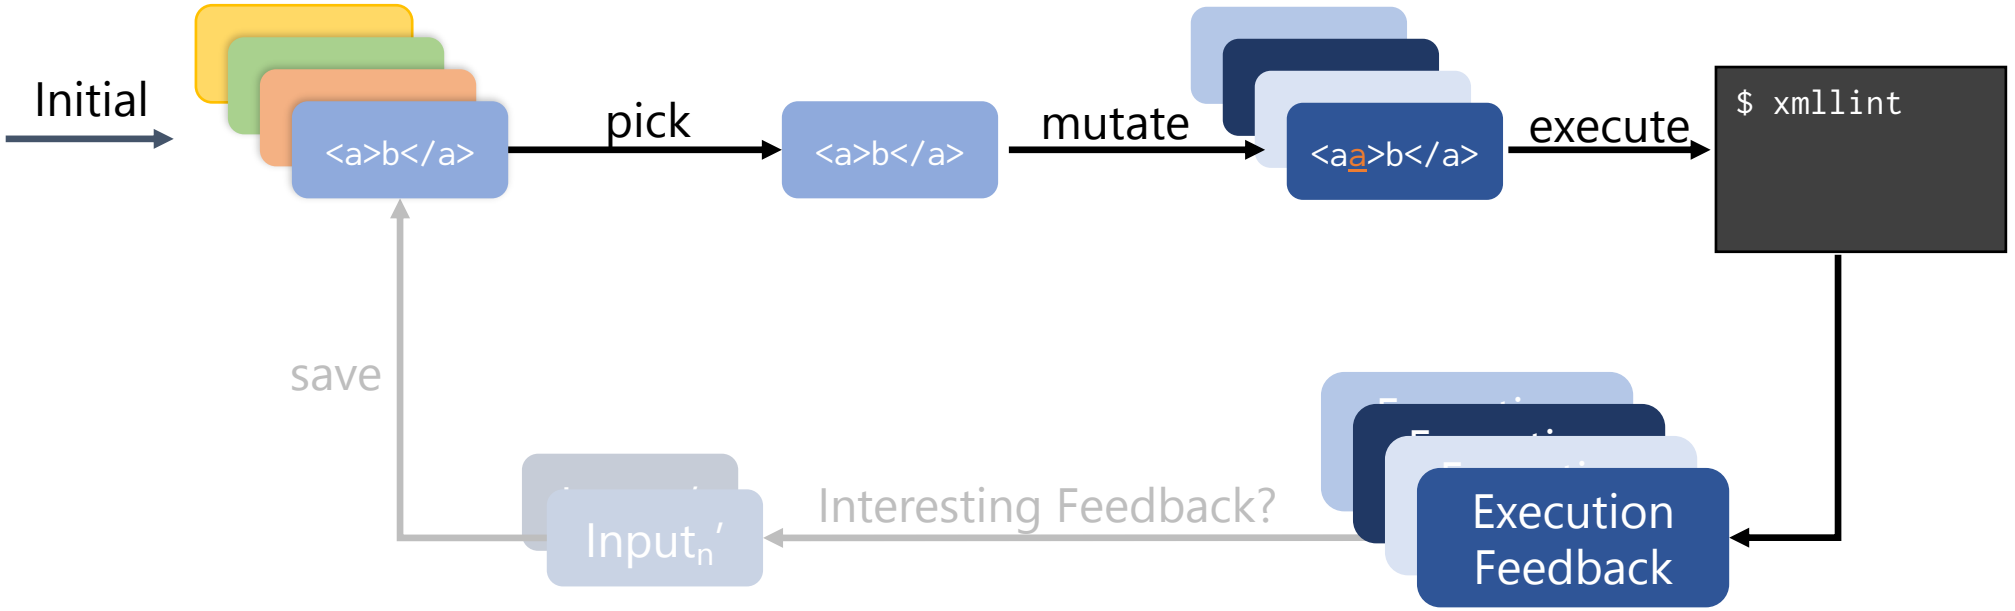

Coverage-Guided Fuzzing

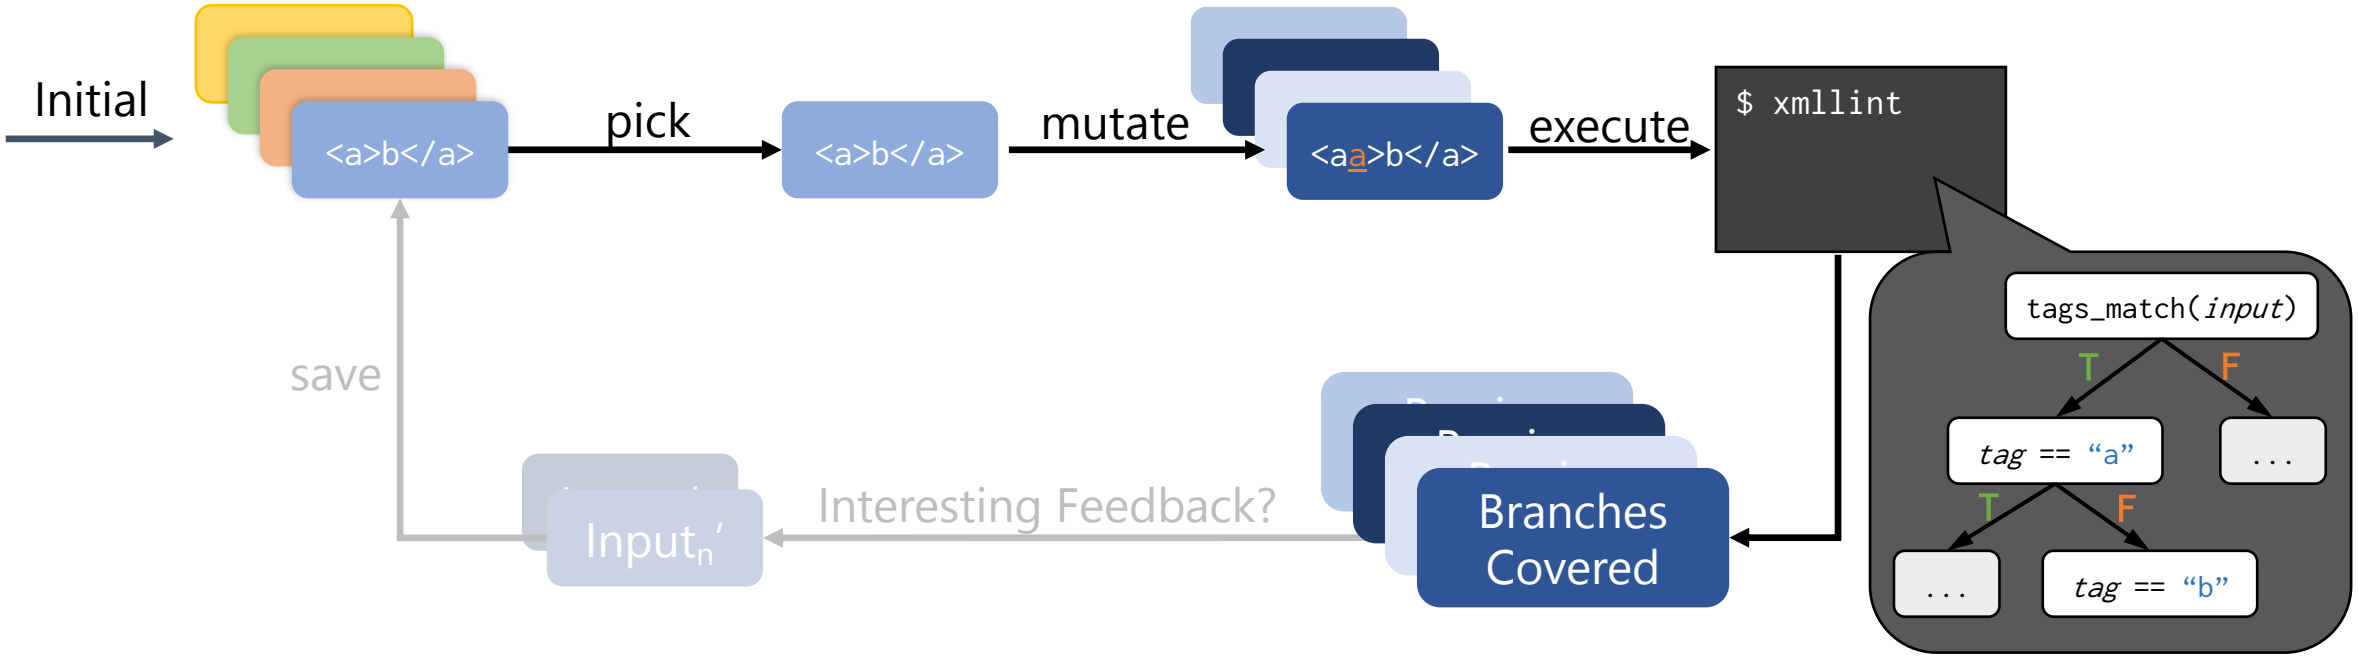

Coverage-Guided Fuzzing

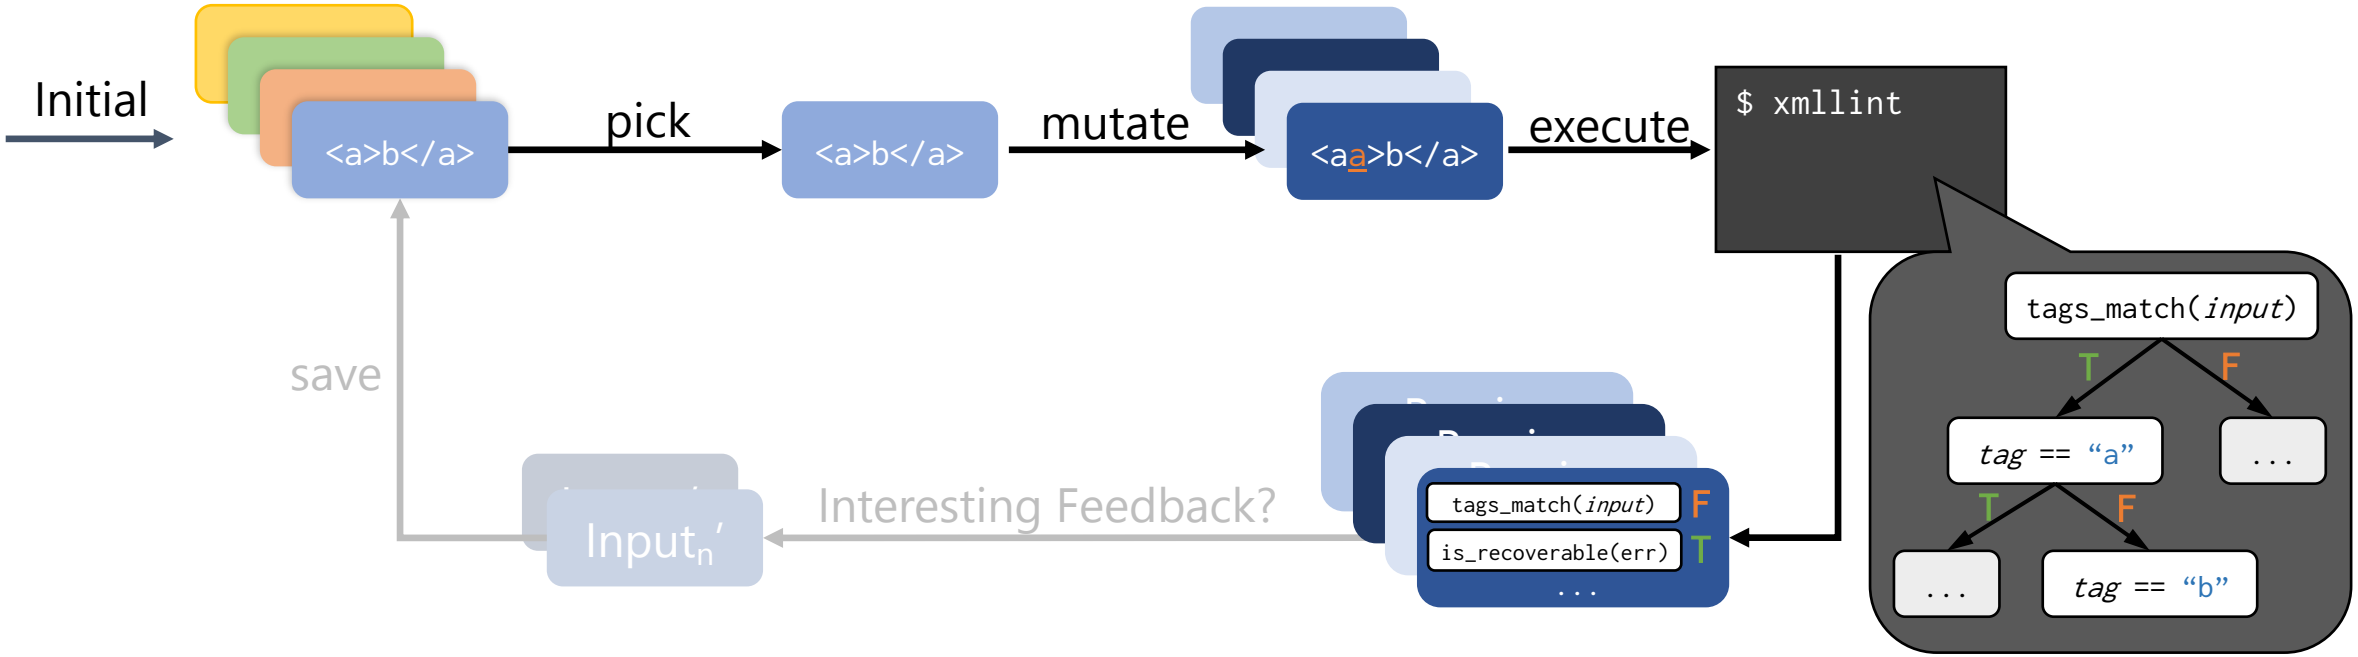

Coverage-Guided Fuzzing

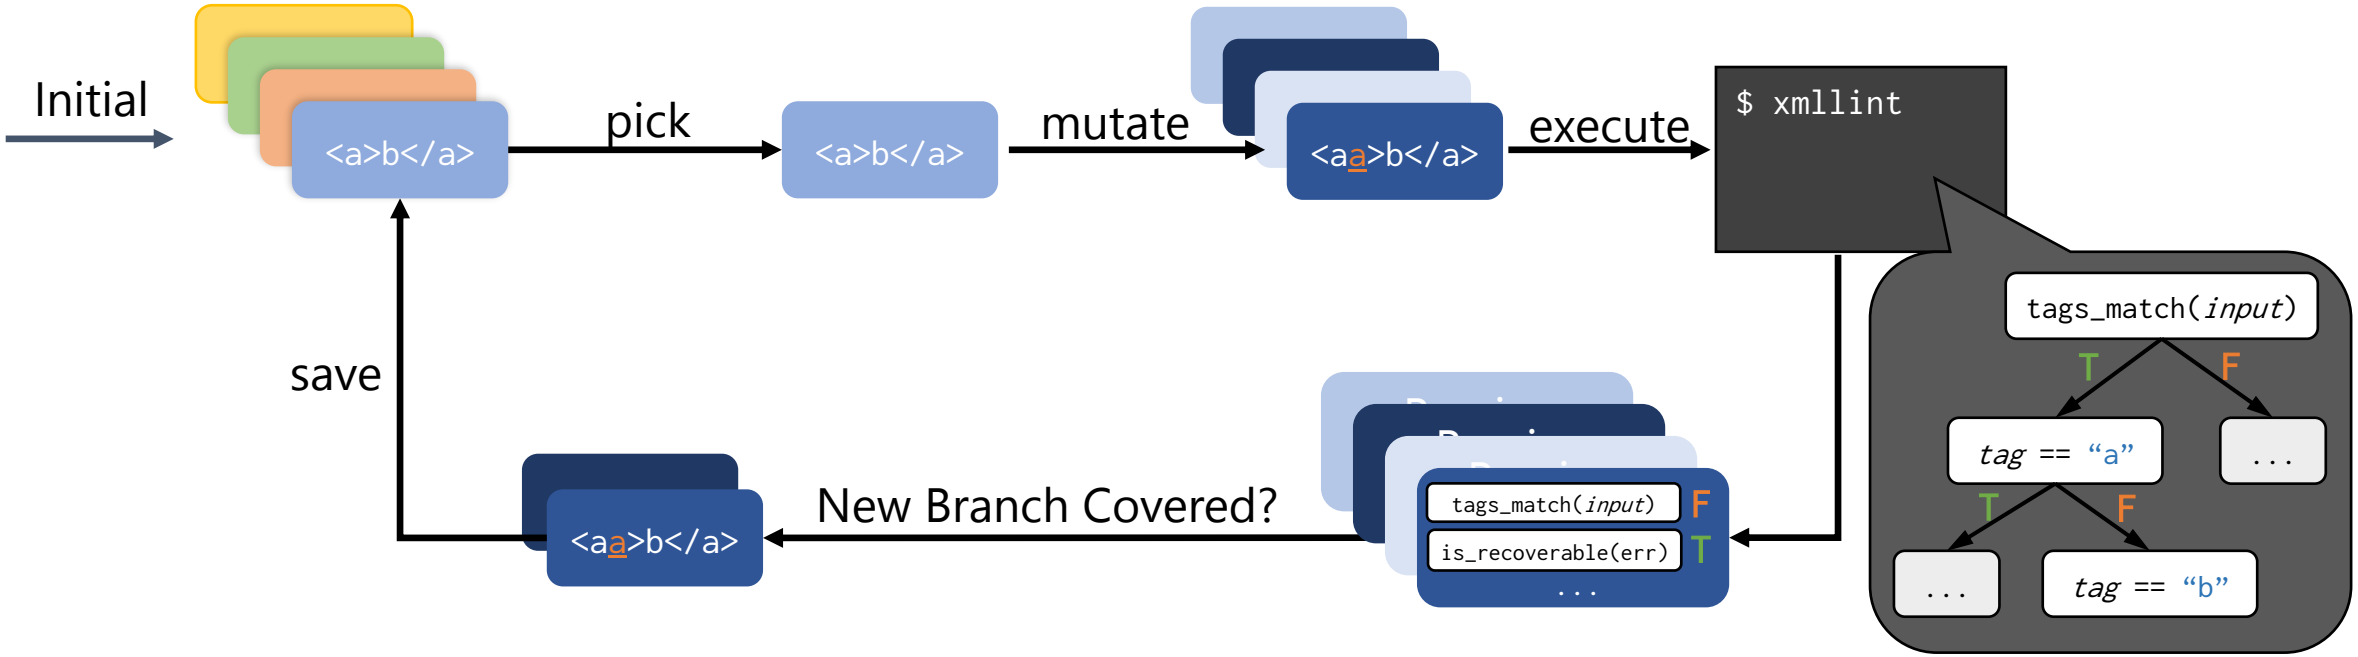

Coverage-Guided Fuzzing

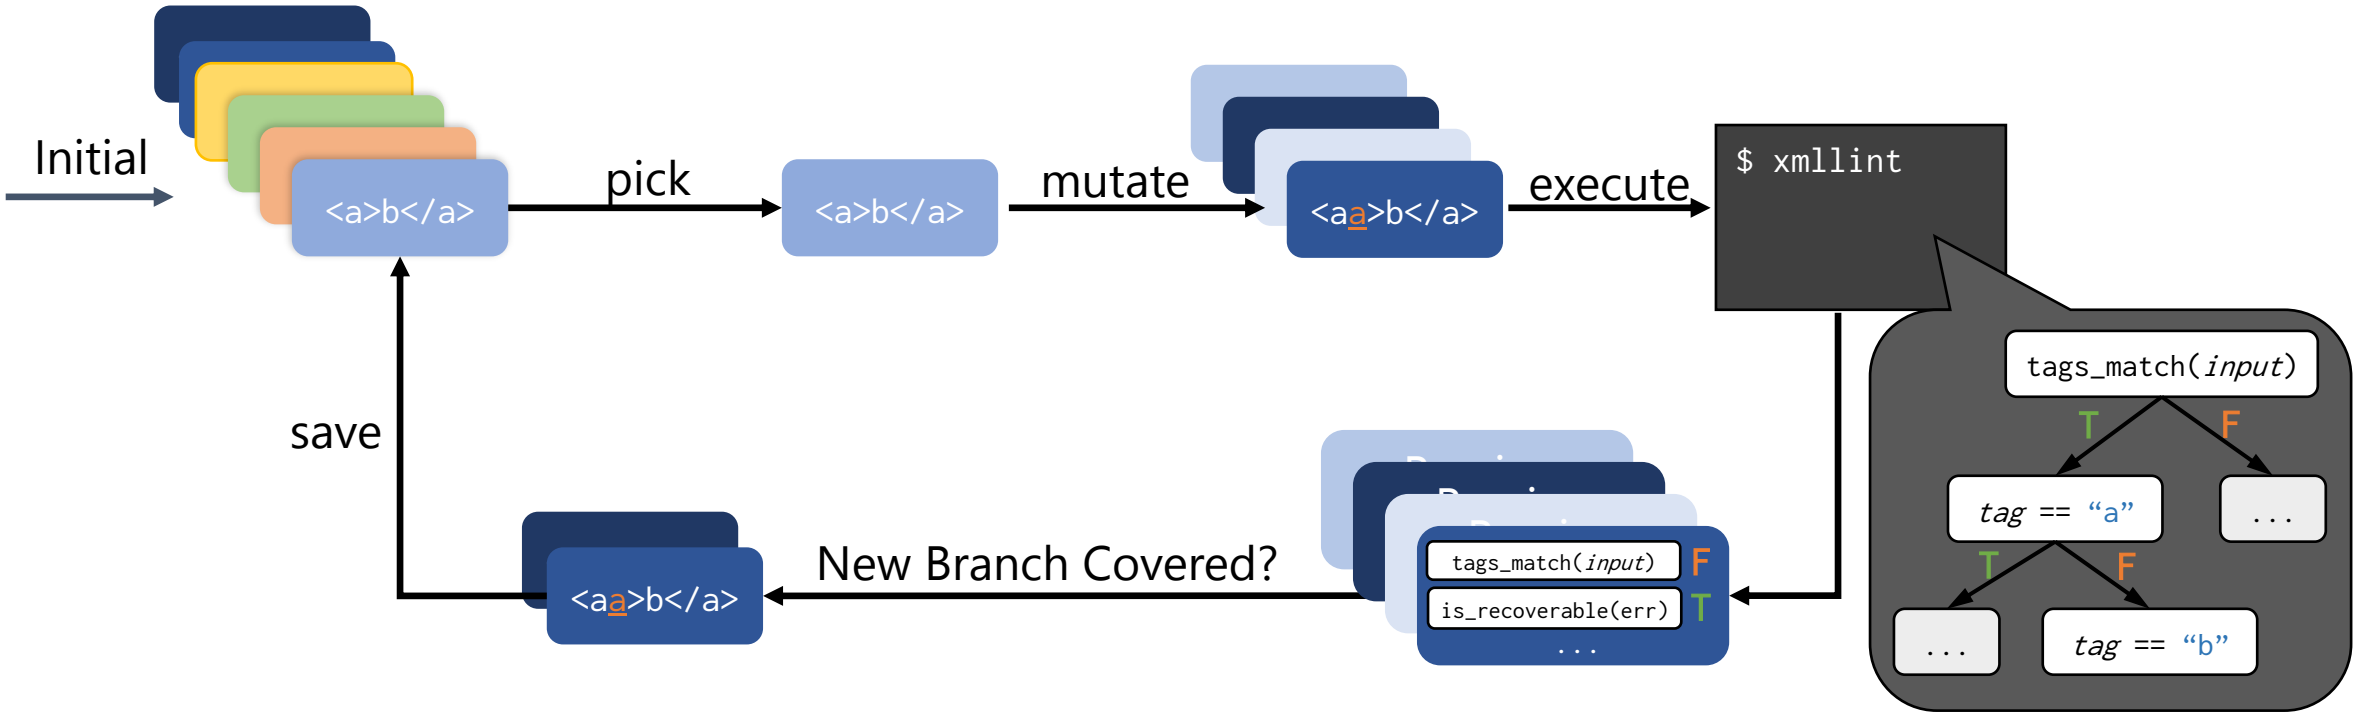

Coverage-Guided Fuzzing

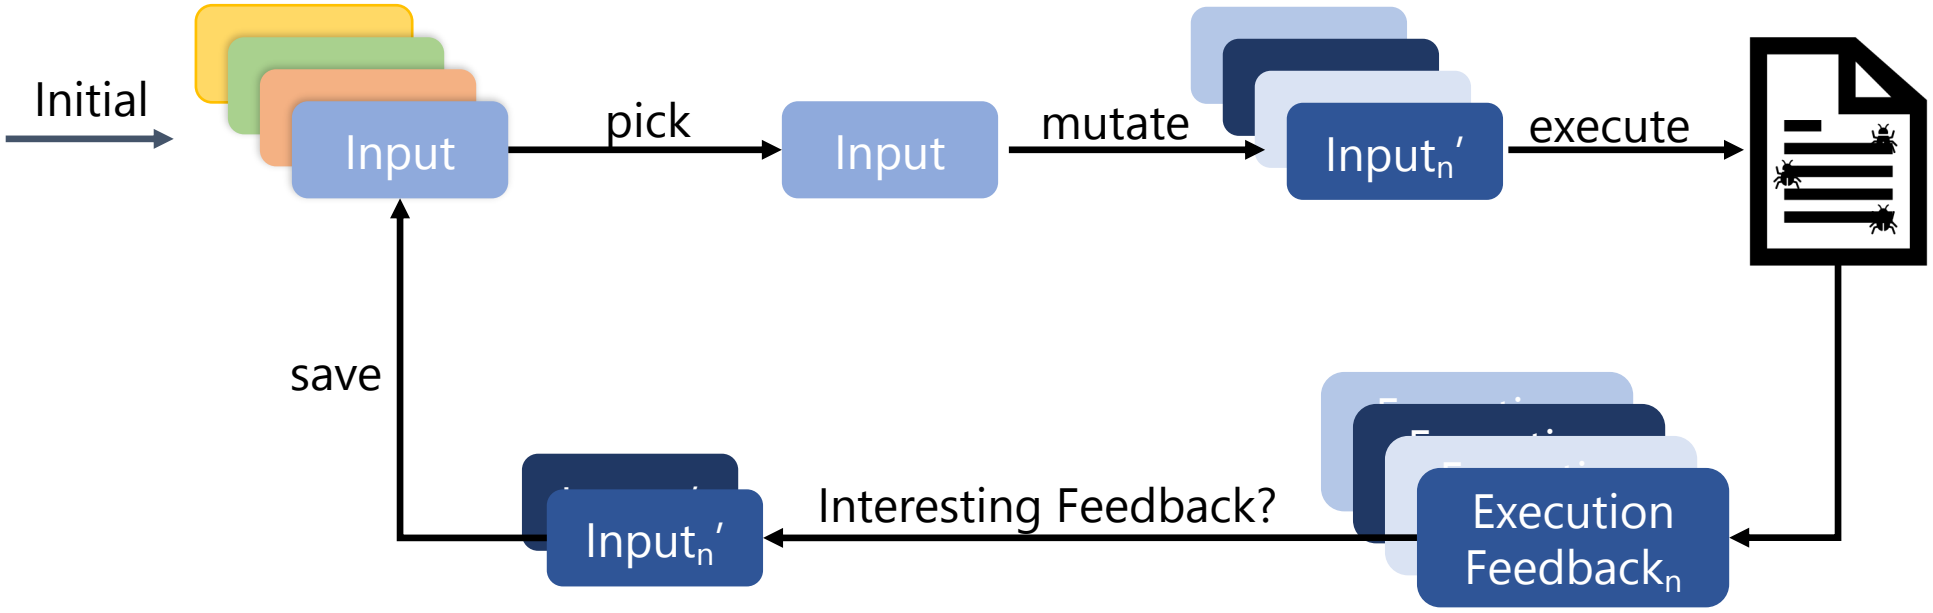

Fuzzing to Find Performance Bugs?

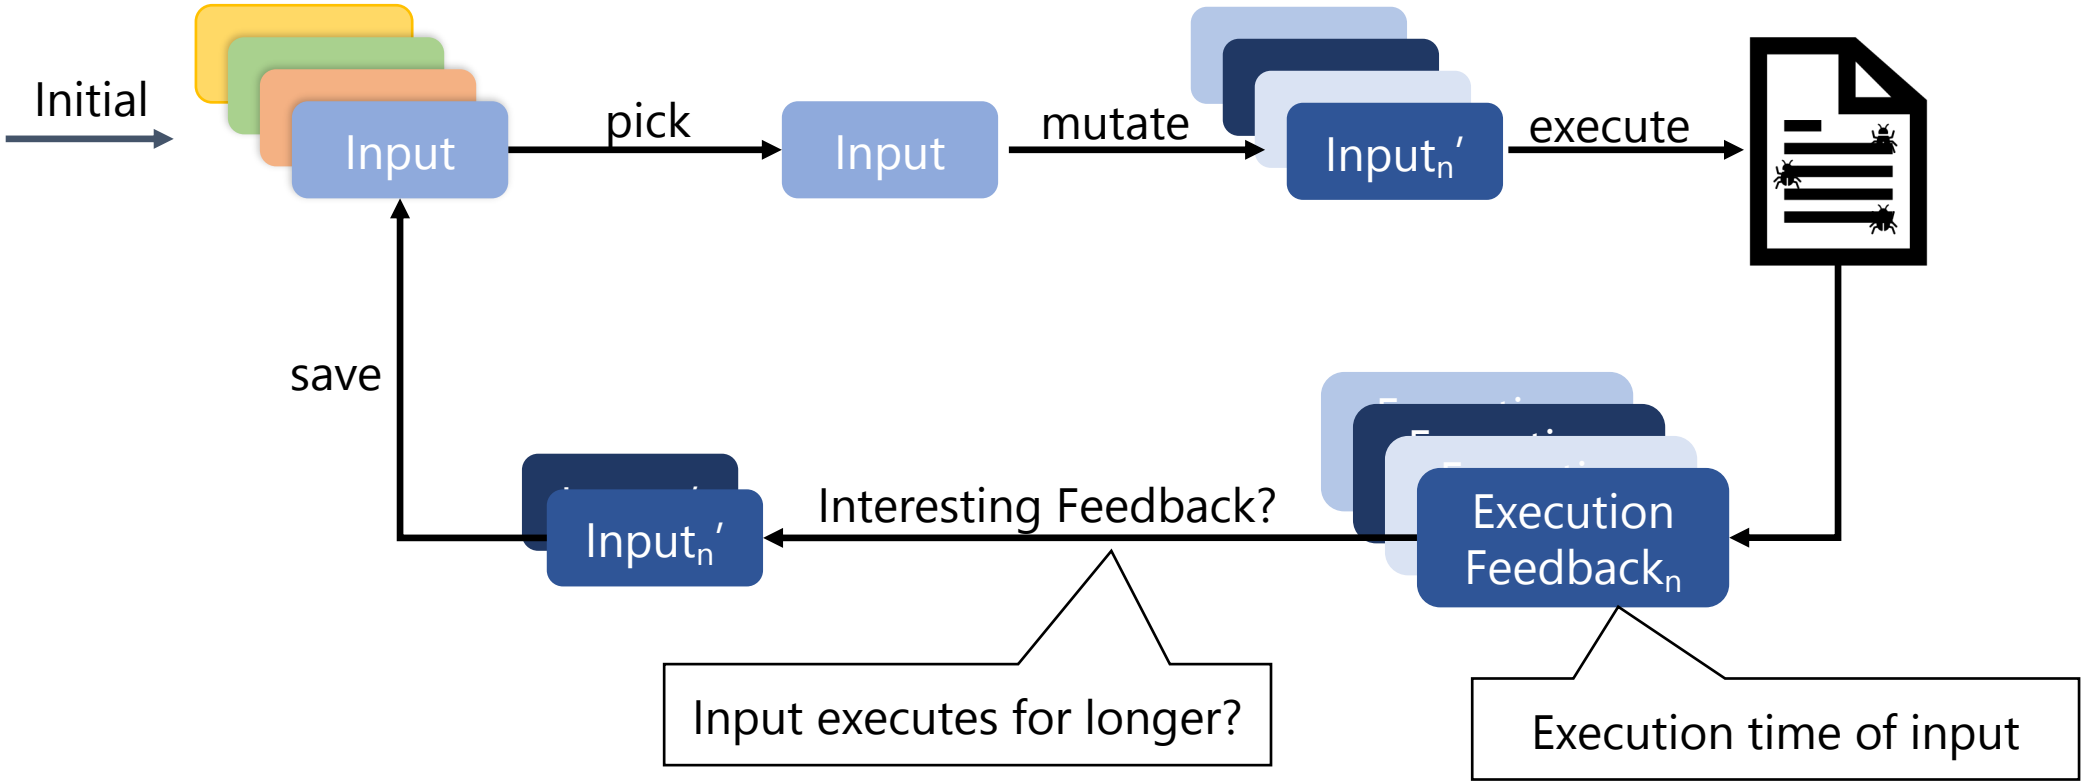

Fuzzing to Find Performance Bugs?

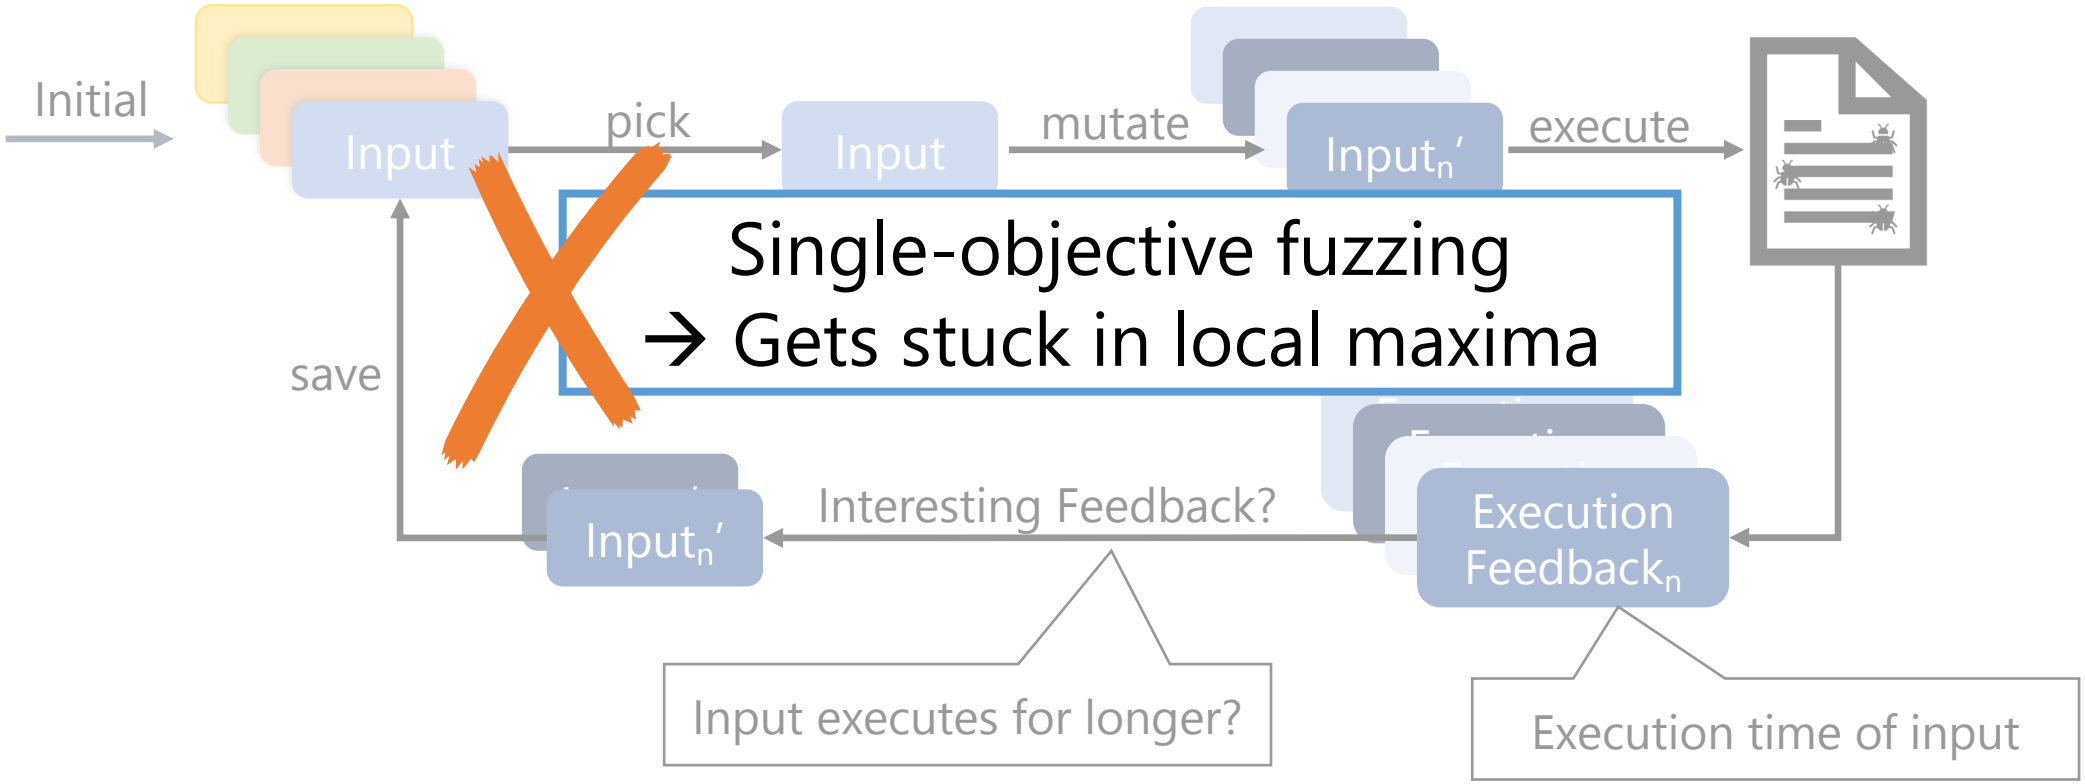

PerfFuzz

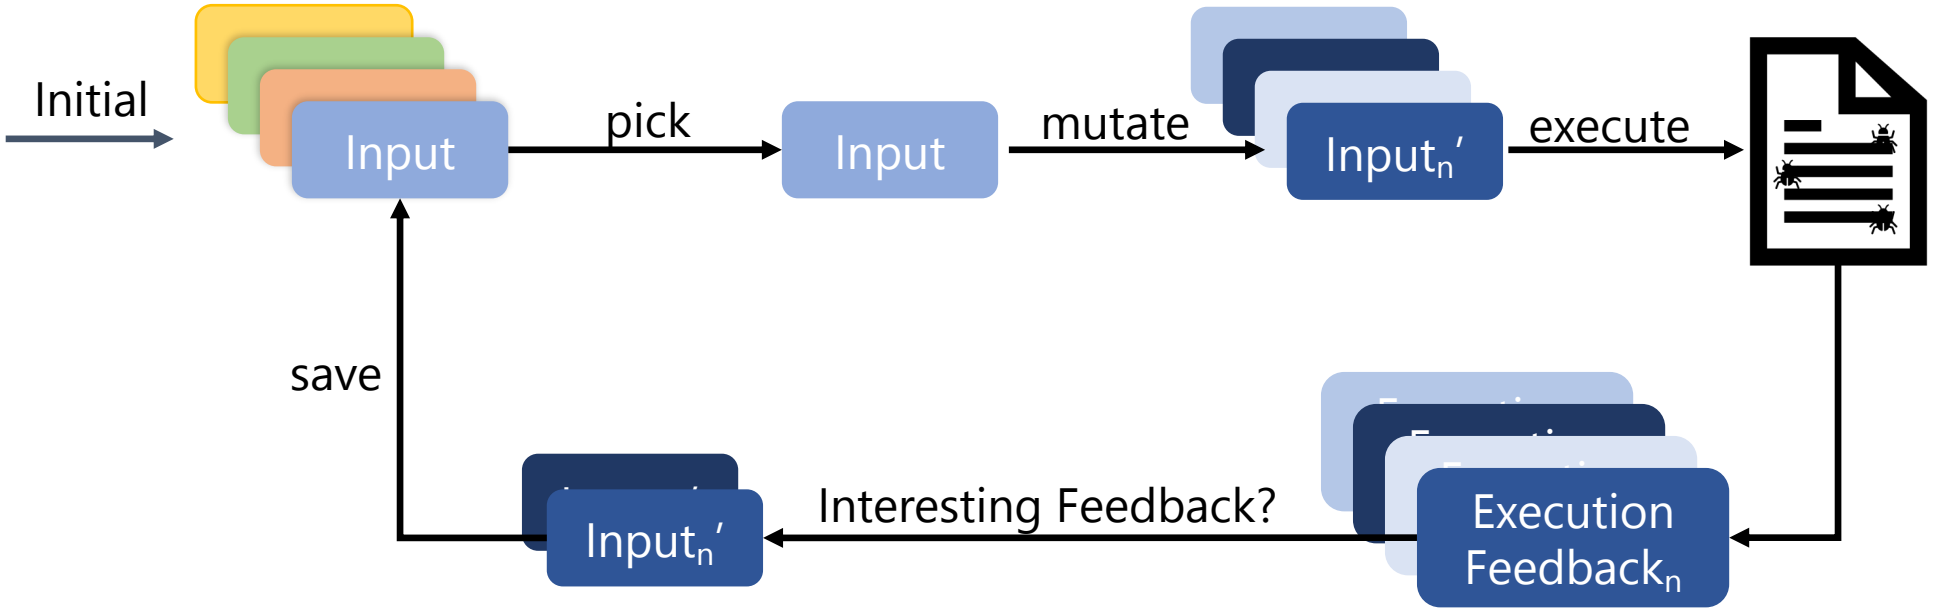
Initial
Input
pick
Input
mutate
Input$_n$'
execute
save
Input$_n$'
Interesting Feedback?
Execution Feedback$_n$

PerfFuzz

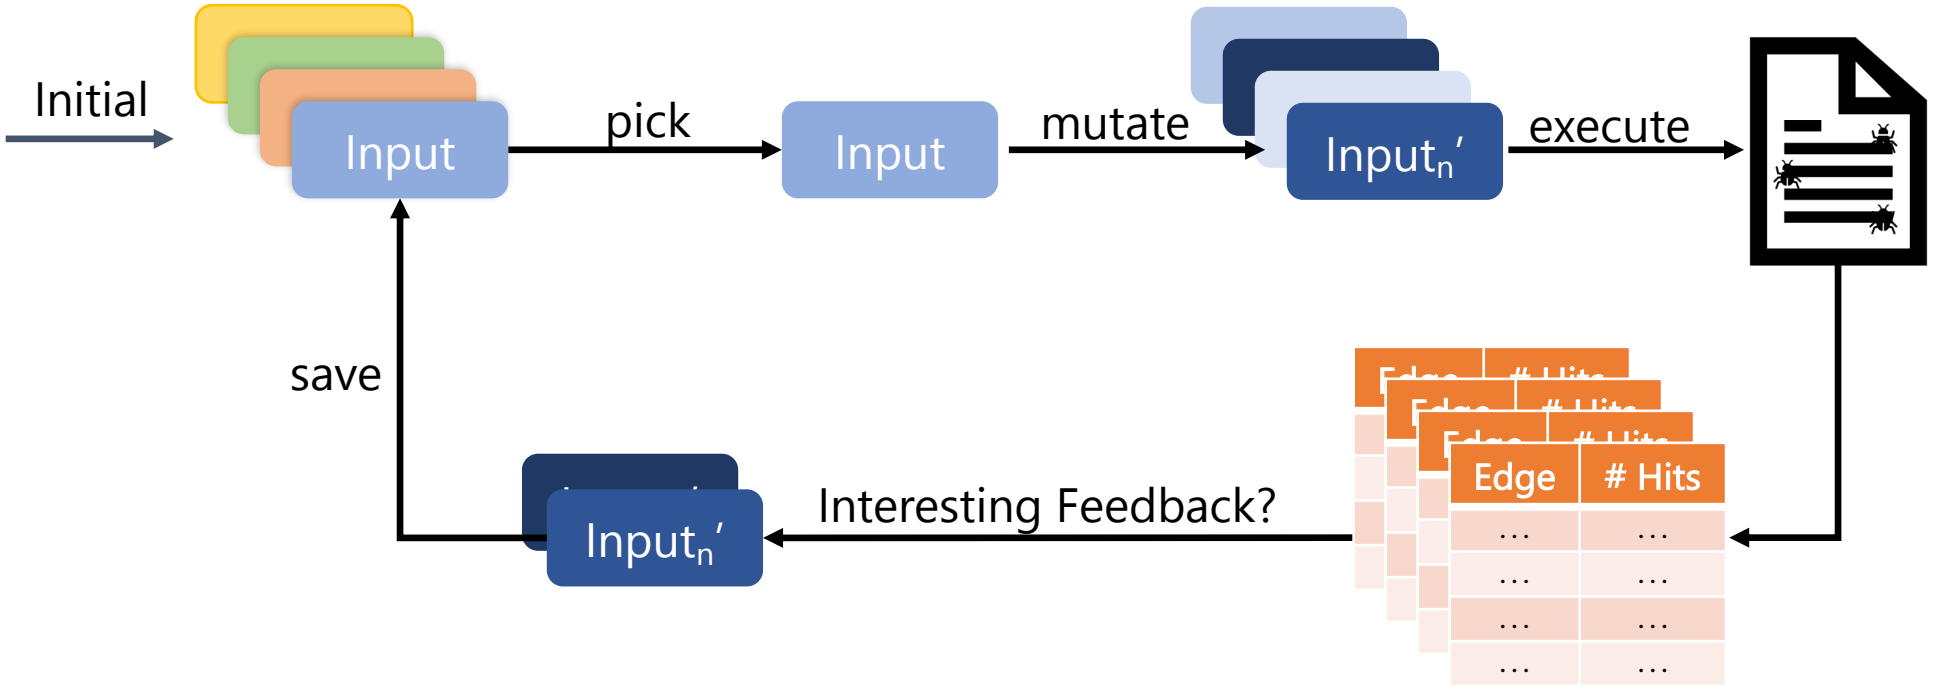
Initial
Input
pick
Input
mutate
Input$_n$'
execute
save
Input$_n$'
Interesting Feedback?
Edge
Hits

PerfFuzz

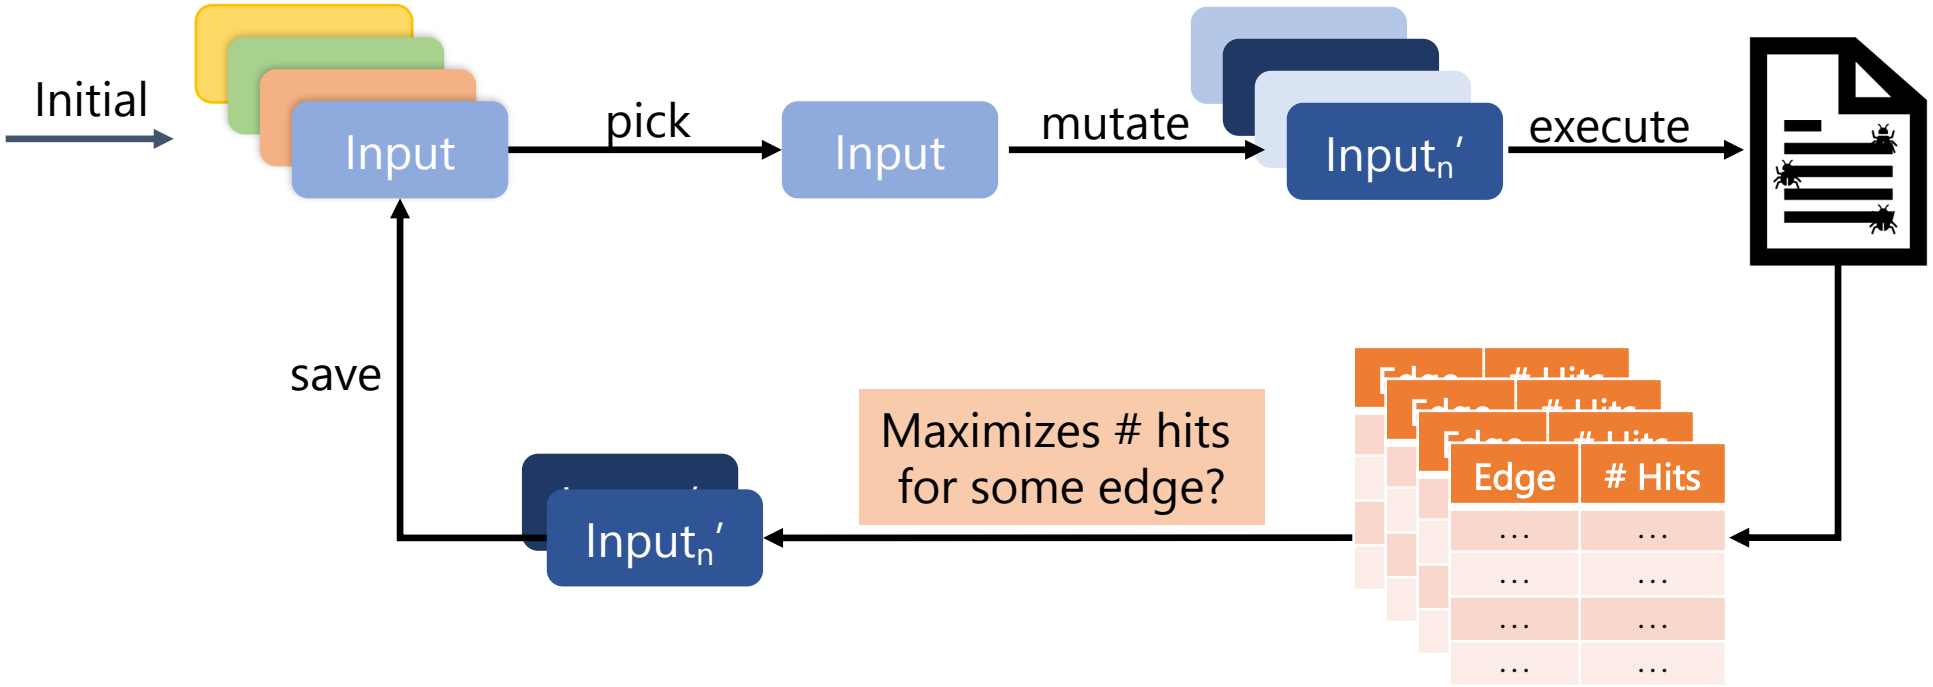
Initial
Input
pick
Input
mutate
Input$_n$'
execute
save
Input$_n$'
Maximizes # hits for some edge?
Edge
Hits
…
…
…
…
…
…
…
…

PerfFuzz

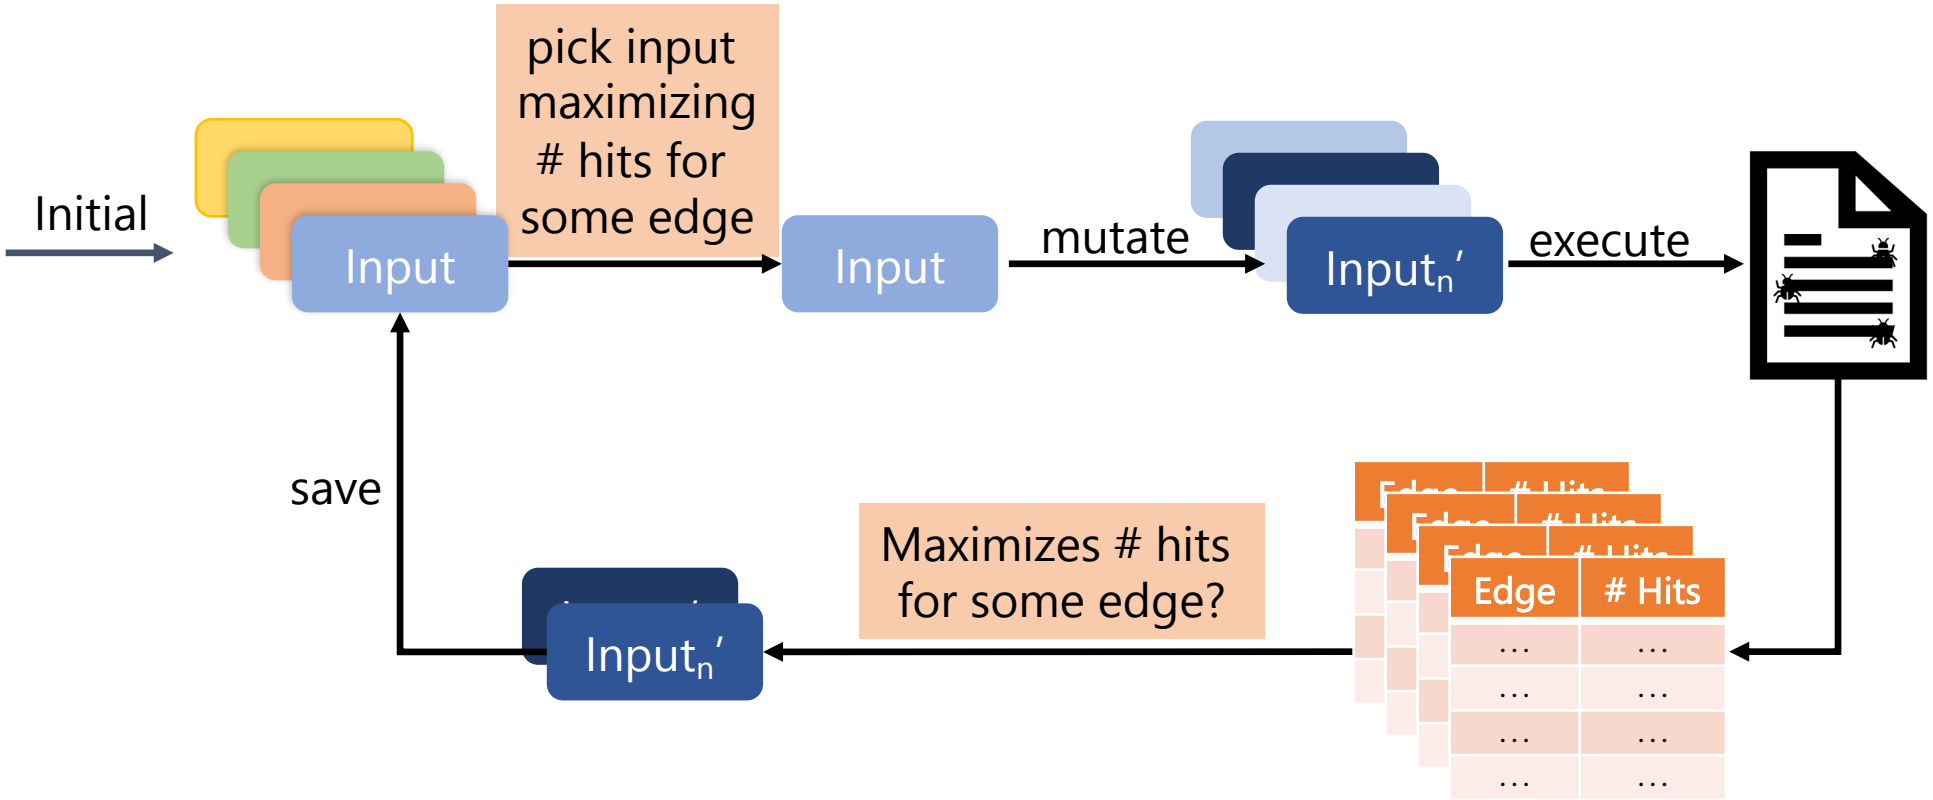
Initial
Input
pick input maximizing # hits for some edge
Input
mutate
Input$_n$'
execute
save
Input$_n$'
Maximizes # hits for some edge?
Edge
Hits

PerfFuzz

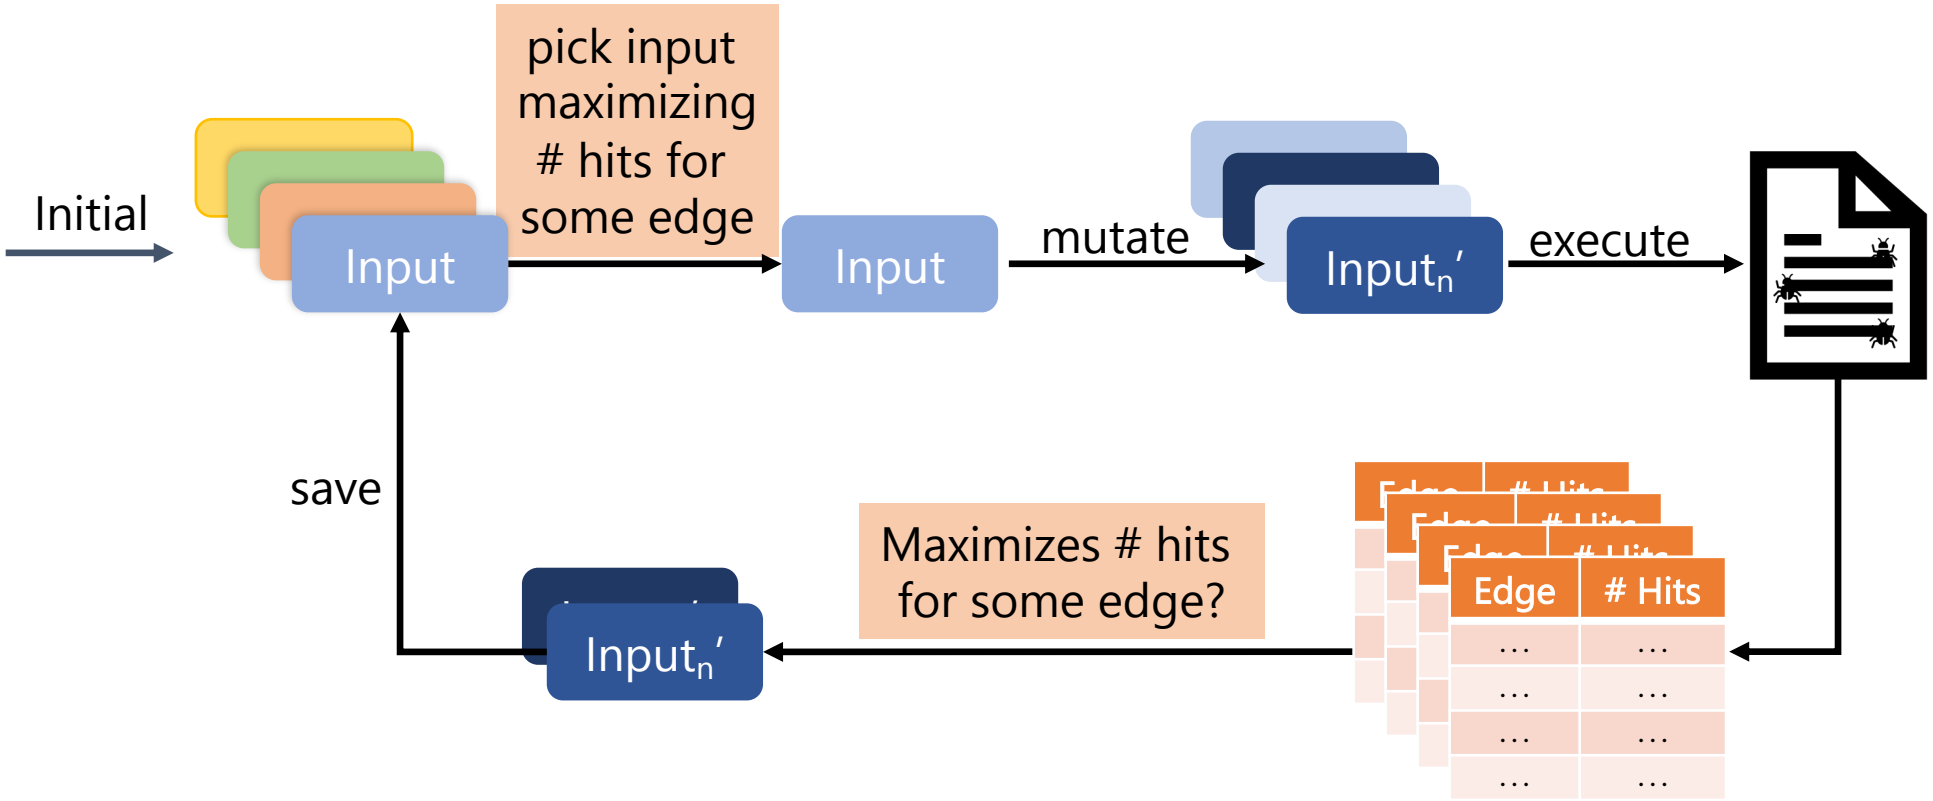
Initial
Input
pick input maximizing # hits for some edge
Input
mutate
Input$_n$'
execute
save
Input$_n$'
Maximizes # hits for some edge?
Edge
Hits

Example program: `wf`

- Count frequency of words in string

input:

```
the quick brown the dog
```

output:

```
brown: 1
dog: 1
quick: 1
the: 2
```

Example program: `wf`

- Count frequency of words in string

input:

```
the quick brown the dog
```

output:

```
brown: 1
dog: 1
quick: 1
the: 2
```

Fedora Linux implementation: linked list hash table.
Quadratic worst-case behavior ☹
(when all words hash collide)

`wf` results: PerfFuzz Finds True Worst Cases

`wf` results: PerfFuzz Finds True Worst Cases

SlowFuzz (single objective maximization) worst case:

```
t r t t s f o Öe r t s f o r t x x t s f o r t x x
```

`wf` results: PerfFuzz Finds True Worst Cases

SlowFuzz (single objective maximization) worst case:

```
t r t t s f o Öe r t s f o r t x x t s f o r t x x
```

PerfFuzz worst case:

```
t <81>v ^?@t <80>!^?@t <80>!t t^Rn t t t t t t t t t
```

wf results: PerfFuzz Finds True Worst Cases

SlowFuzz (single objective maximization) worst case:

```
t r t t s f o Öe r t s f o r t x x t s f o r t x x
```

PerfFuzz worst case:

```
t <81>v ^?@t <80>!^?@t <80>!t t^Rn t t t t t t t t t
```

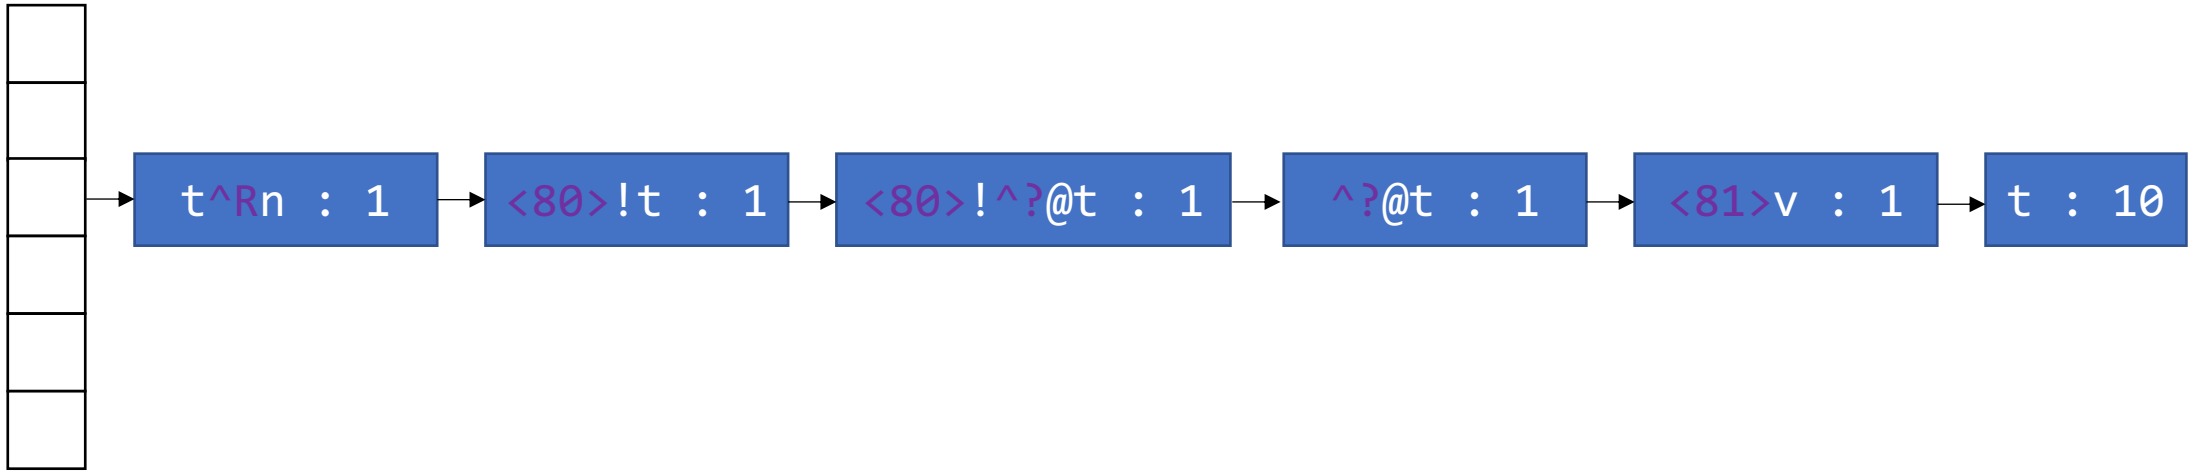

PerfFuzz

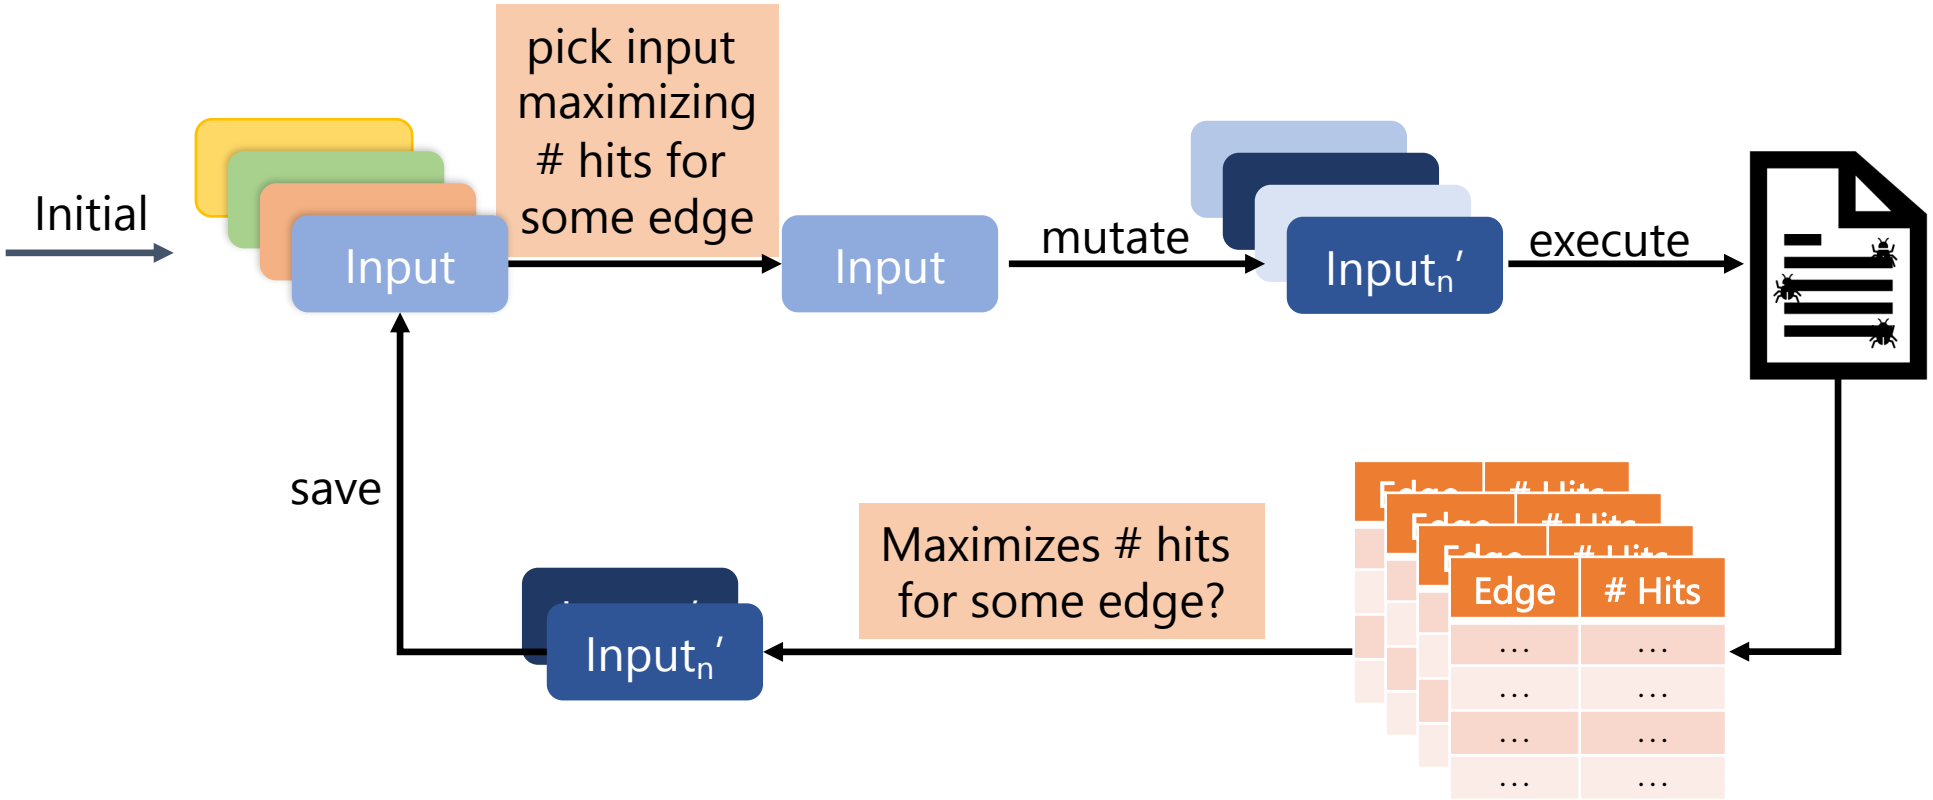
Initial
Input
pick input maximizing # hits for some edge
Input
mutate
Input$_n$'
execute
save
Input$_n$'
Maximizes # hits for some edge?
Edge
Hits

Observation: Algorithm is More General

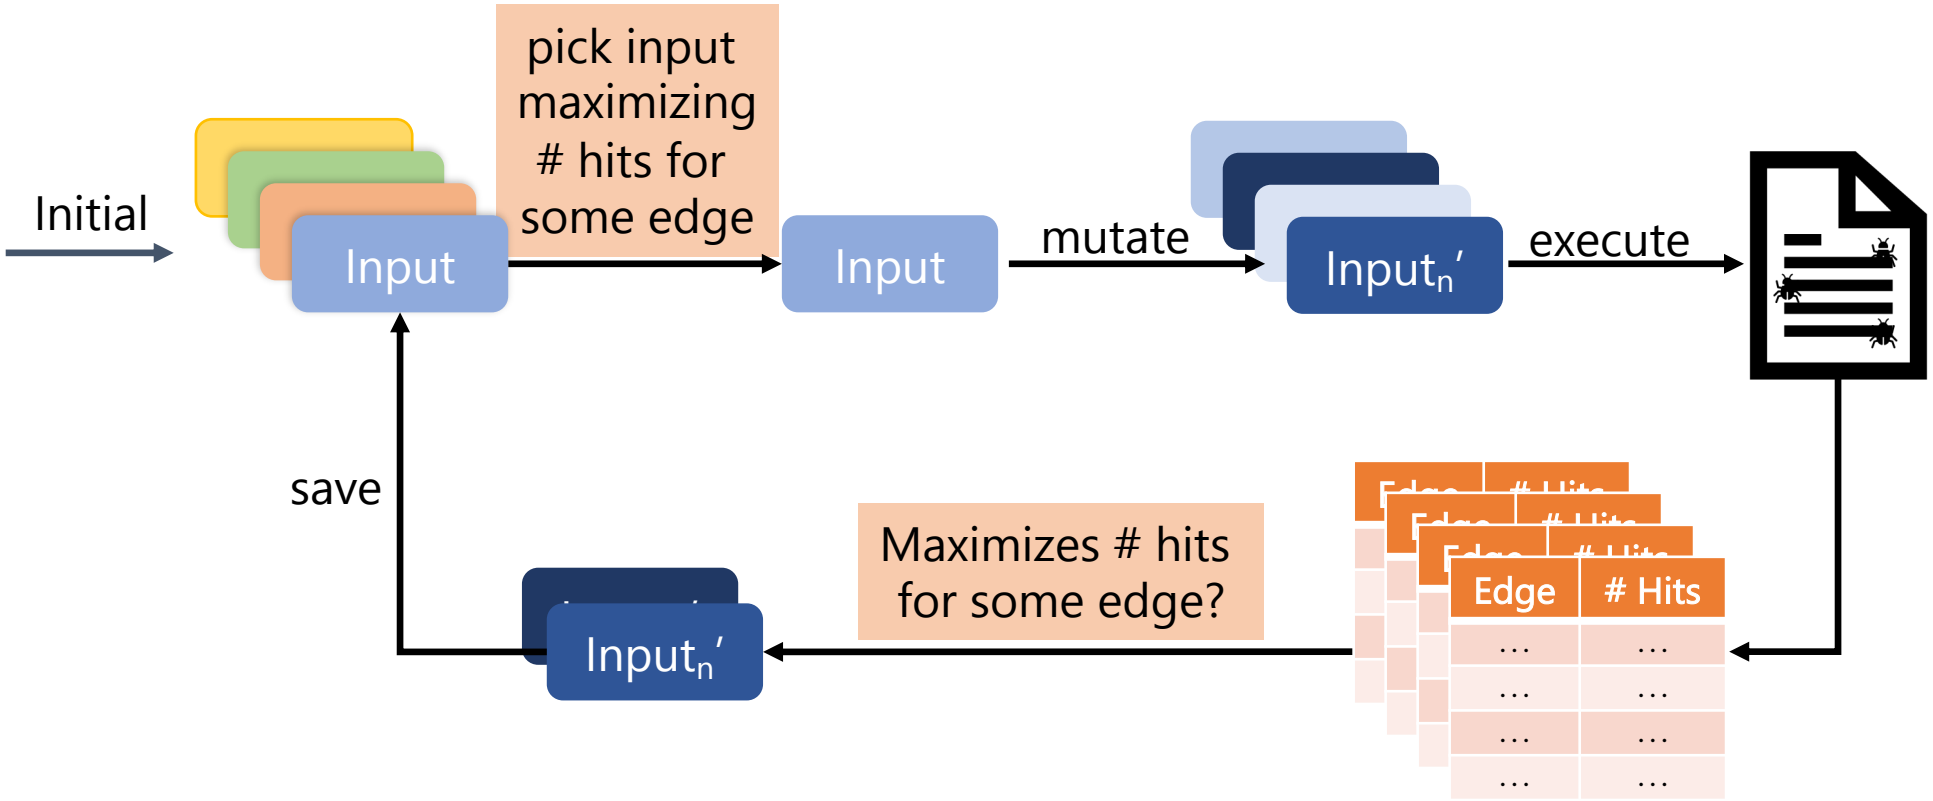

Observation: Algorithm is More General

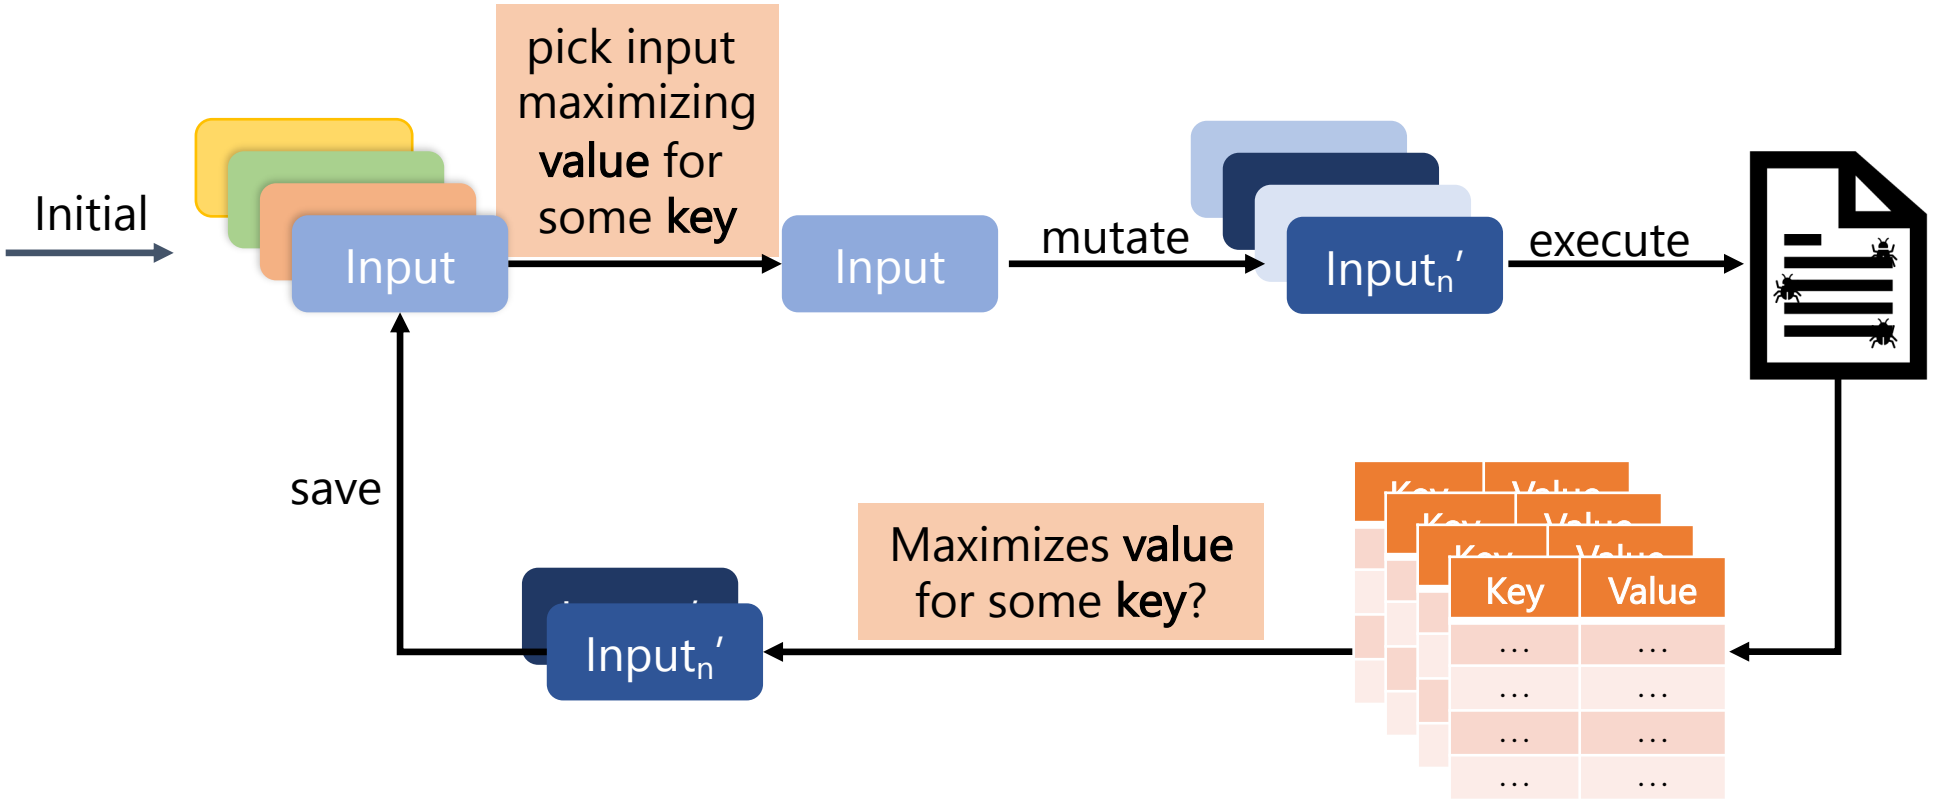

FuzzFactory

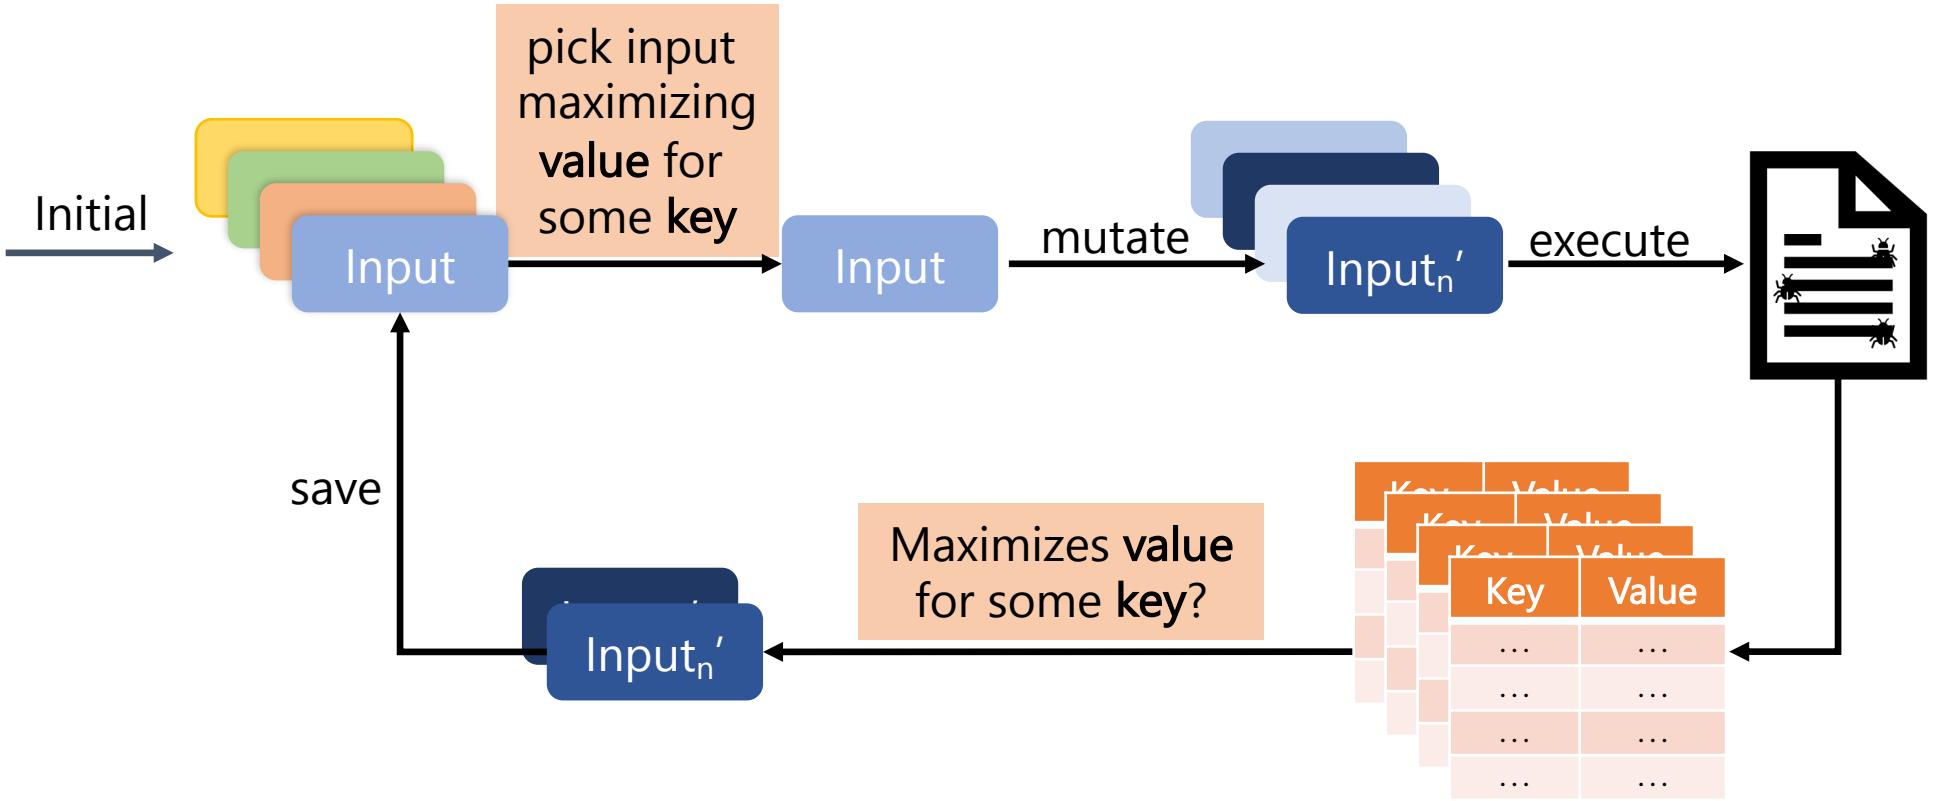

Super Fuzzer: Hard Comparisons + Memory Allocations

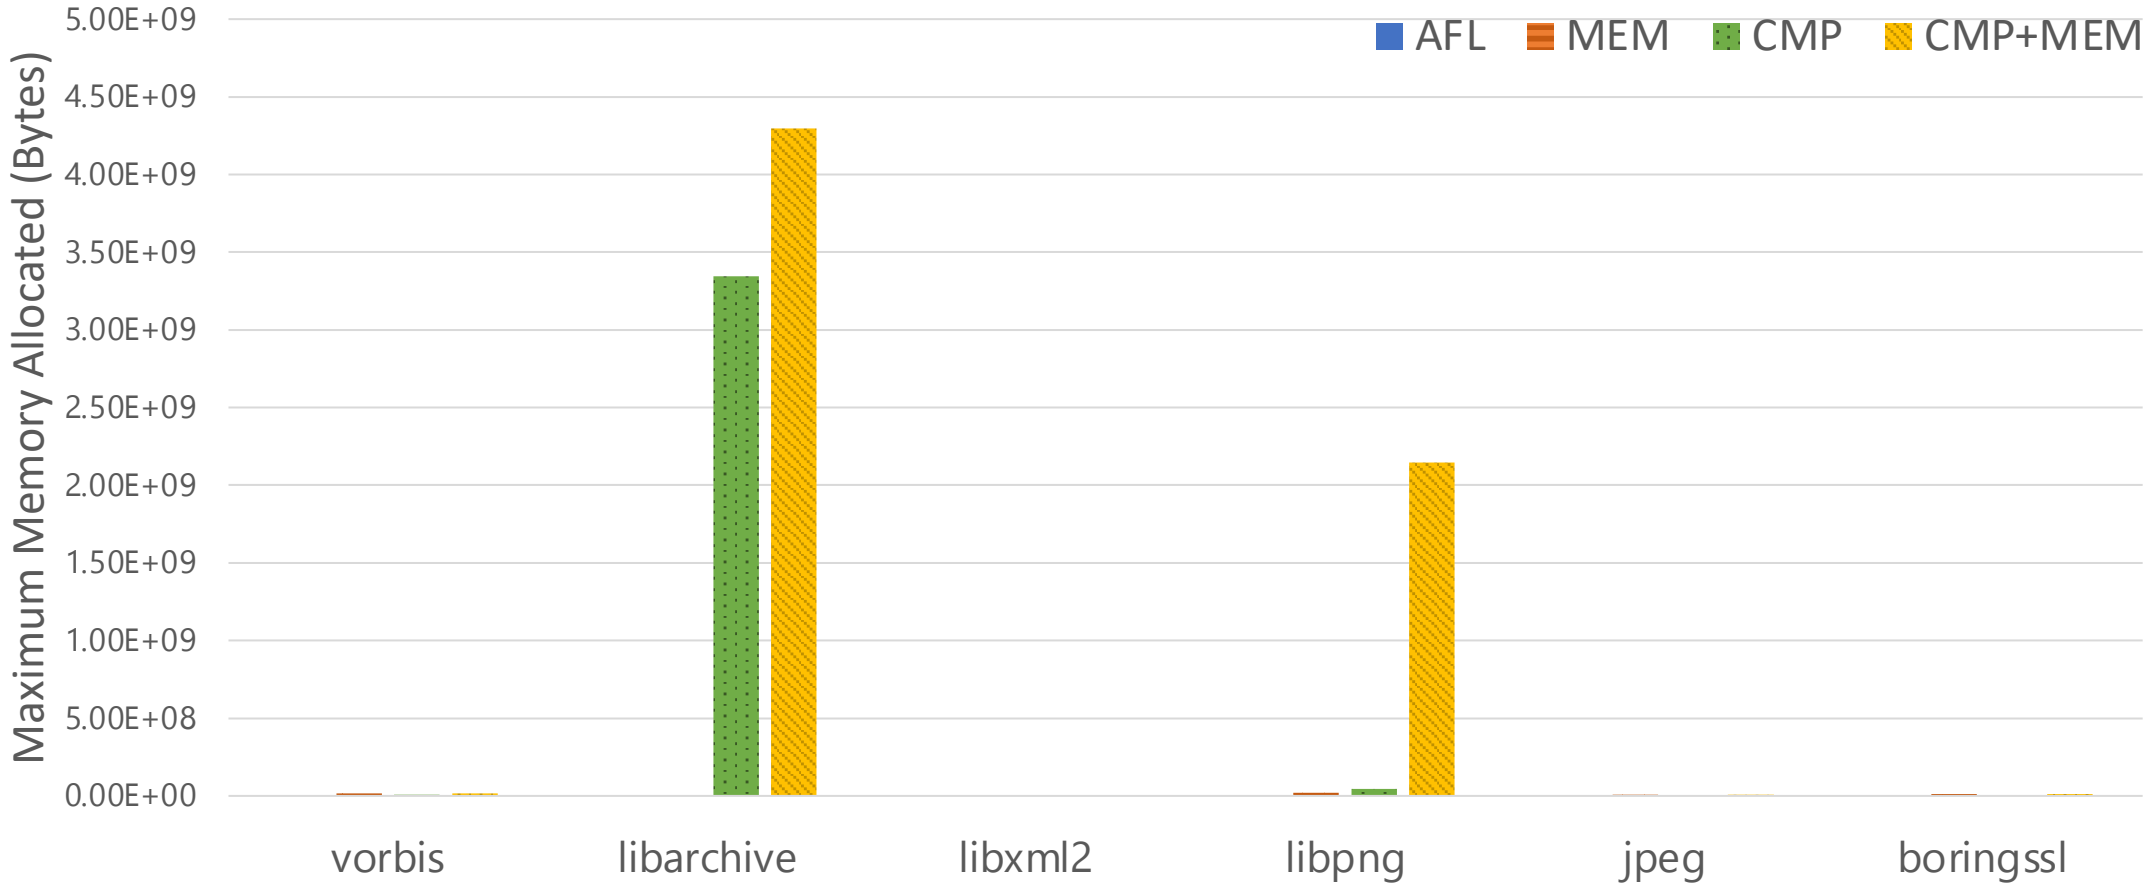

Super Fuzzer: Hard Comparisons + Memory Allocations

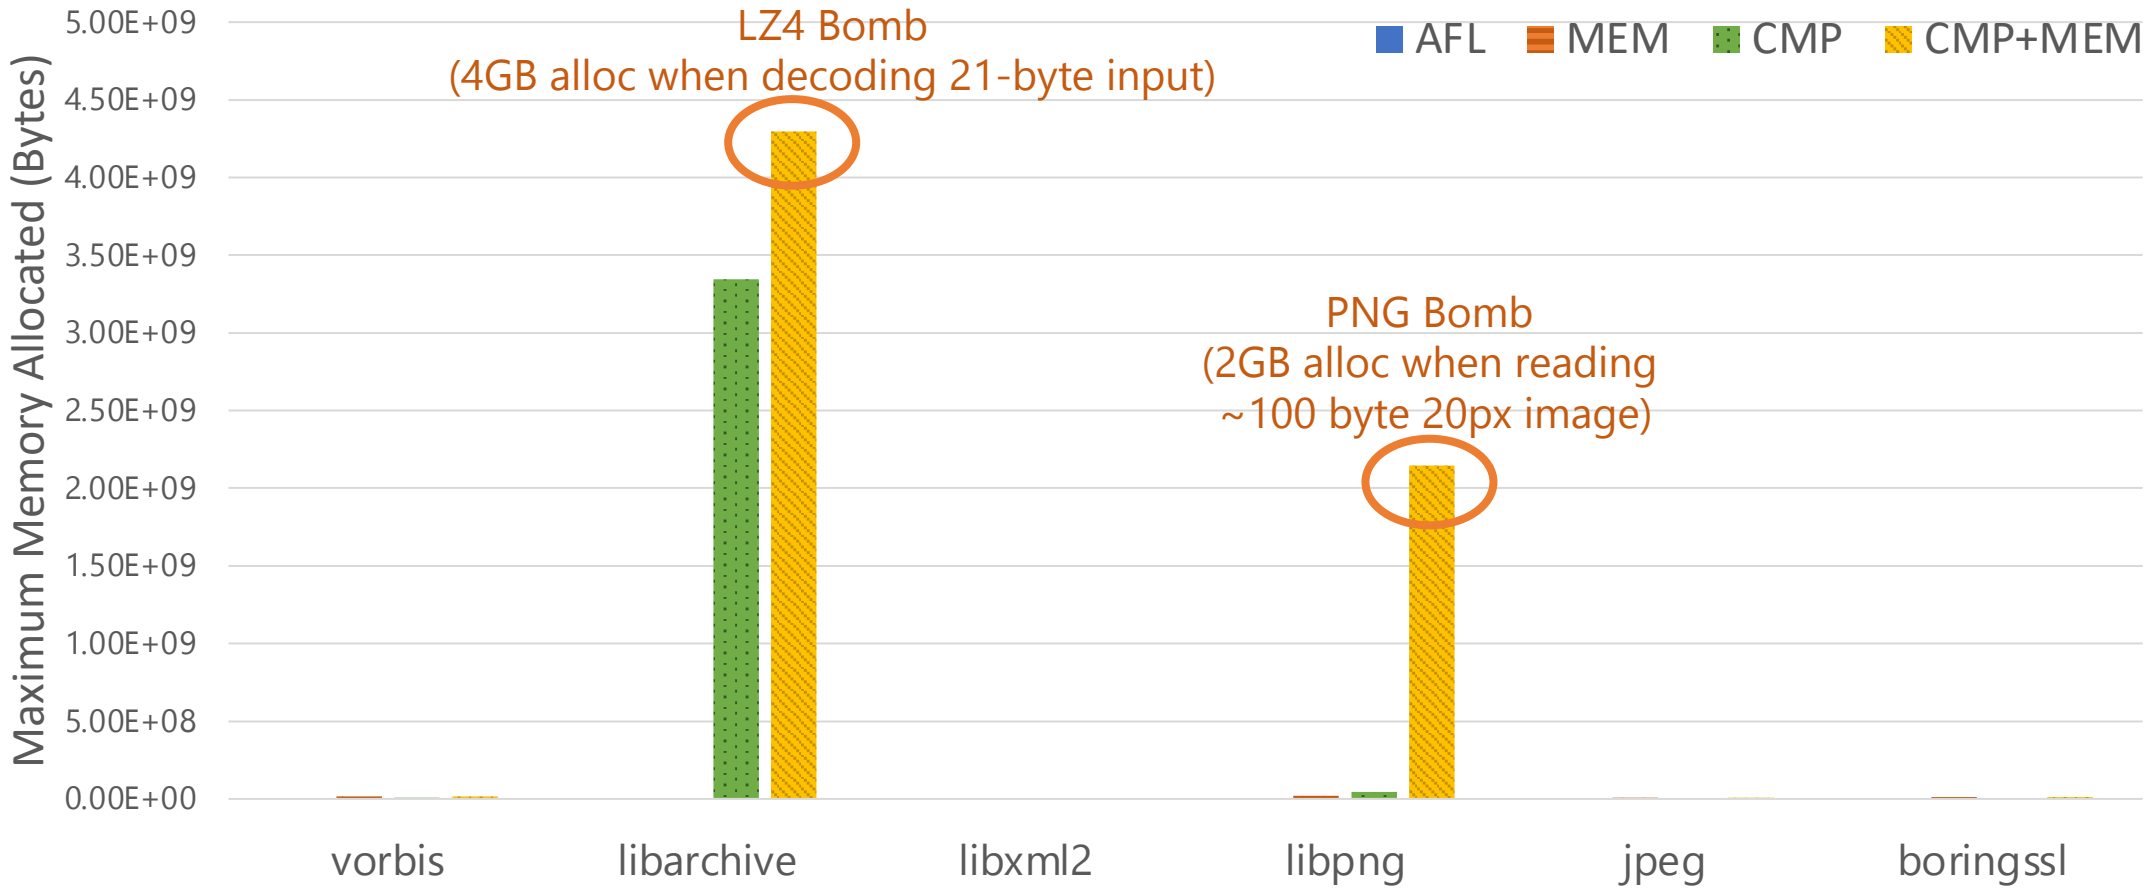

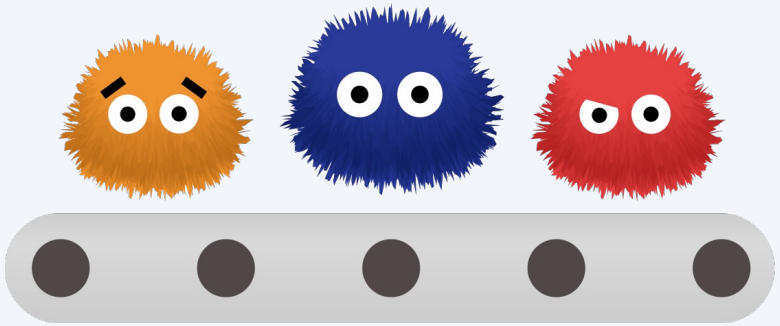

Using ***generalized feedback maps*** to expand ***bugs findable by fuzz testing***

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

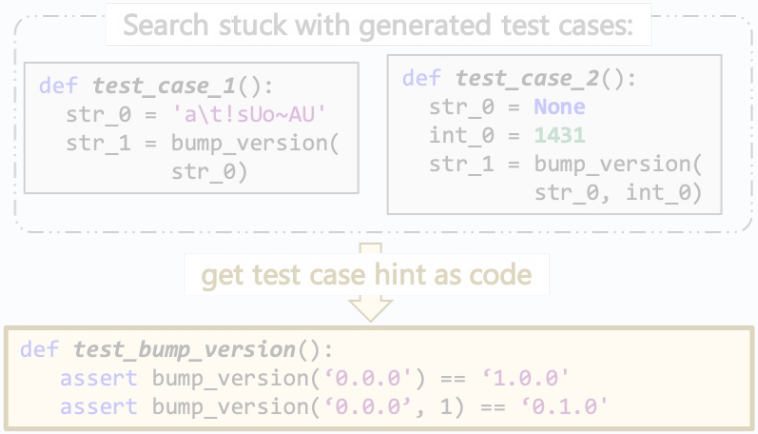

Using ***large language models*** to improve ***automated test suite generation***

CodaMOSA (ICSE'23)

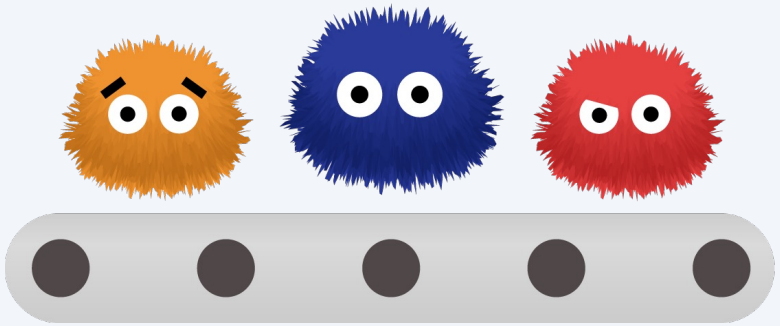

Using ***generalized feedback maps*** to expand ***bugs findable by fuzz testing***

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

→ Feedback-directed fuzzing is a flexible abstraction to explore different dimensions of program behavior

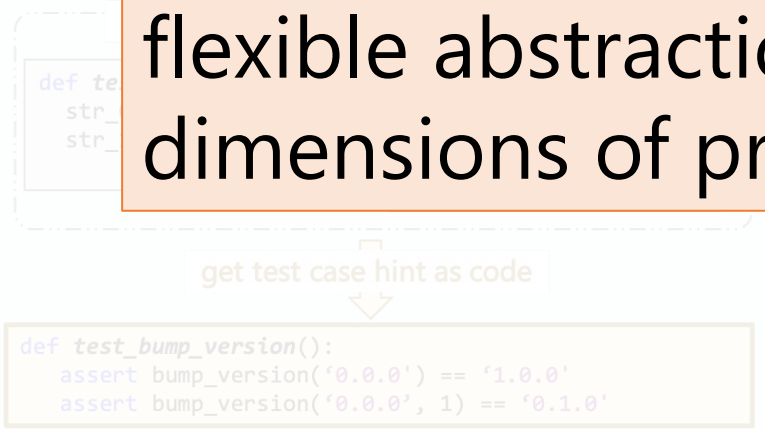

automated test suite generation

CodaMOSA (ICSE'23)

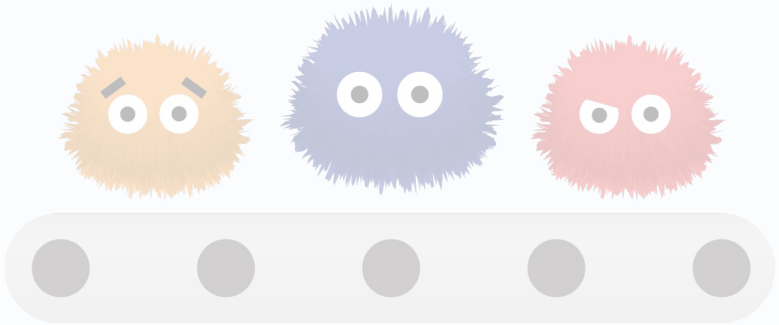

Using *generalized feedback maps* to expand *bugs findable by fuzz testing*

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

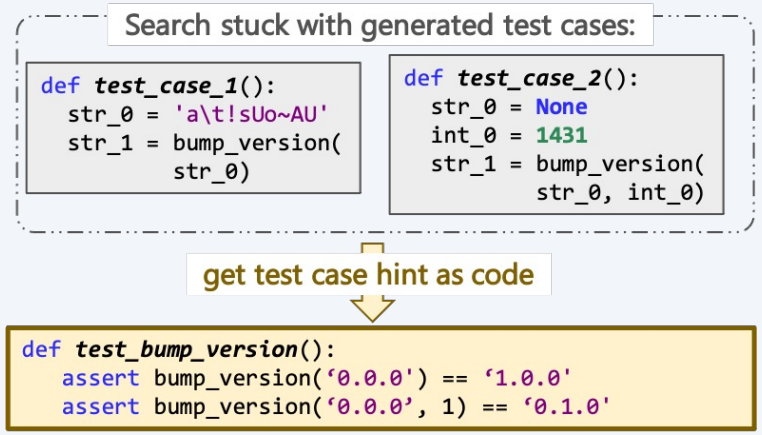

Using *large language models* to improve *automated test suite generation*

CodaMOSA (ICSE'23)

Test Input Generation (PerfFuzz, etc.)

Generate **inputs** to a **parameterized** test function

```
the quick brown the dog
```

```
t <81>v <80>!^?@t t t
```

```
t e q I k b o n t e d g
```

Parameterized Test Function

```
def test_wf(document):
    frequencies = wf(document)
    print(frequencies)
```

Test *Suite* Generation

Generate **test cases** for a file **(e.g., python module, java class)** under test

```
def test_BST_insert():
    tree = None
    tree = BST_insert(tree, 5)
    tree = BST_insert(tree, 3)
    tree = BST_insert(tree, 7)
```

```
def test_BST_search():
    tree = Node(5)
    tree.left = Node(3)
    tree.left.left = Node(1)
    res = BST_search(tree, 3)
```

```
def test_BST_delete():
    tree = Node(5)
    tree.left = Node(3)
    tree.left.left = Node(1)
    tree.left.right = Node(4)
    BST_delete(tree, 4)
```

Module Under test

```
def BST_insert(tree, to_add):
    # Insert to_add into tree
    <...>

def BST_search(tree, to_search):
    # Search tree for to_search
    <...>

def BST_delete(tree, to_delete):
    # Delete to_delete from tree
    <...>
```

Search-Based Test Suite Generation

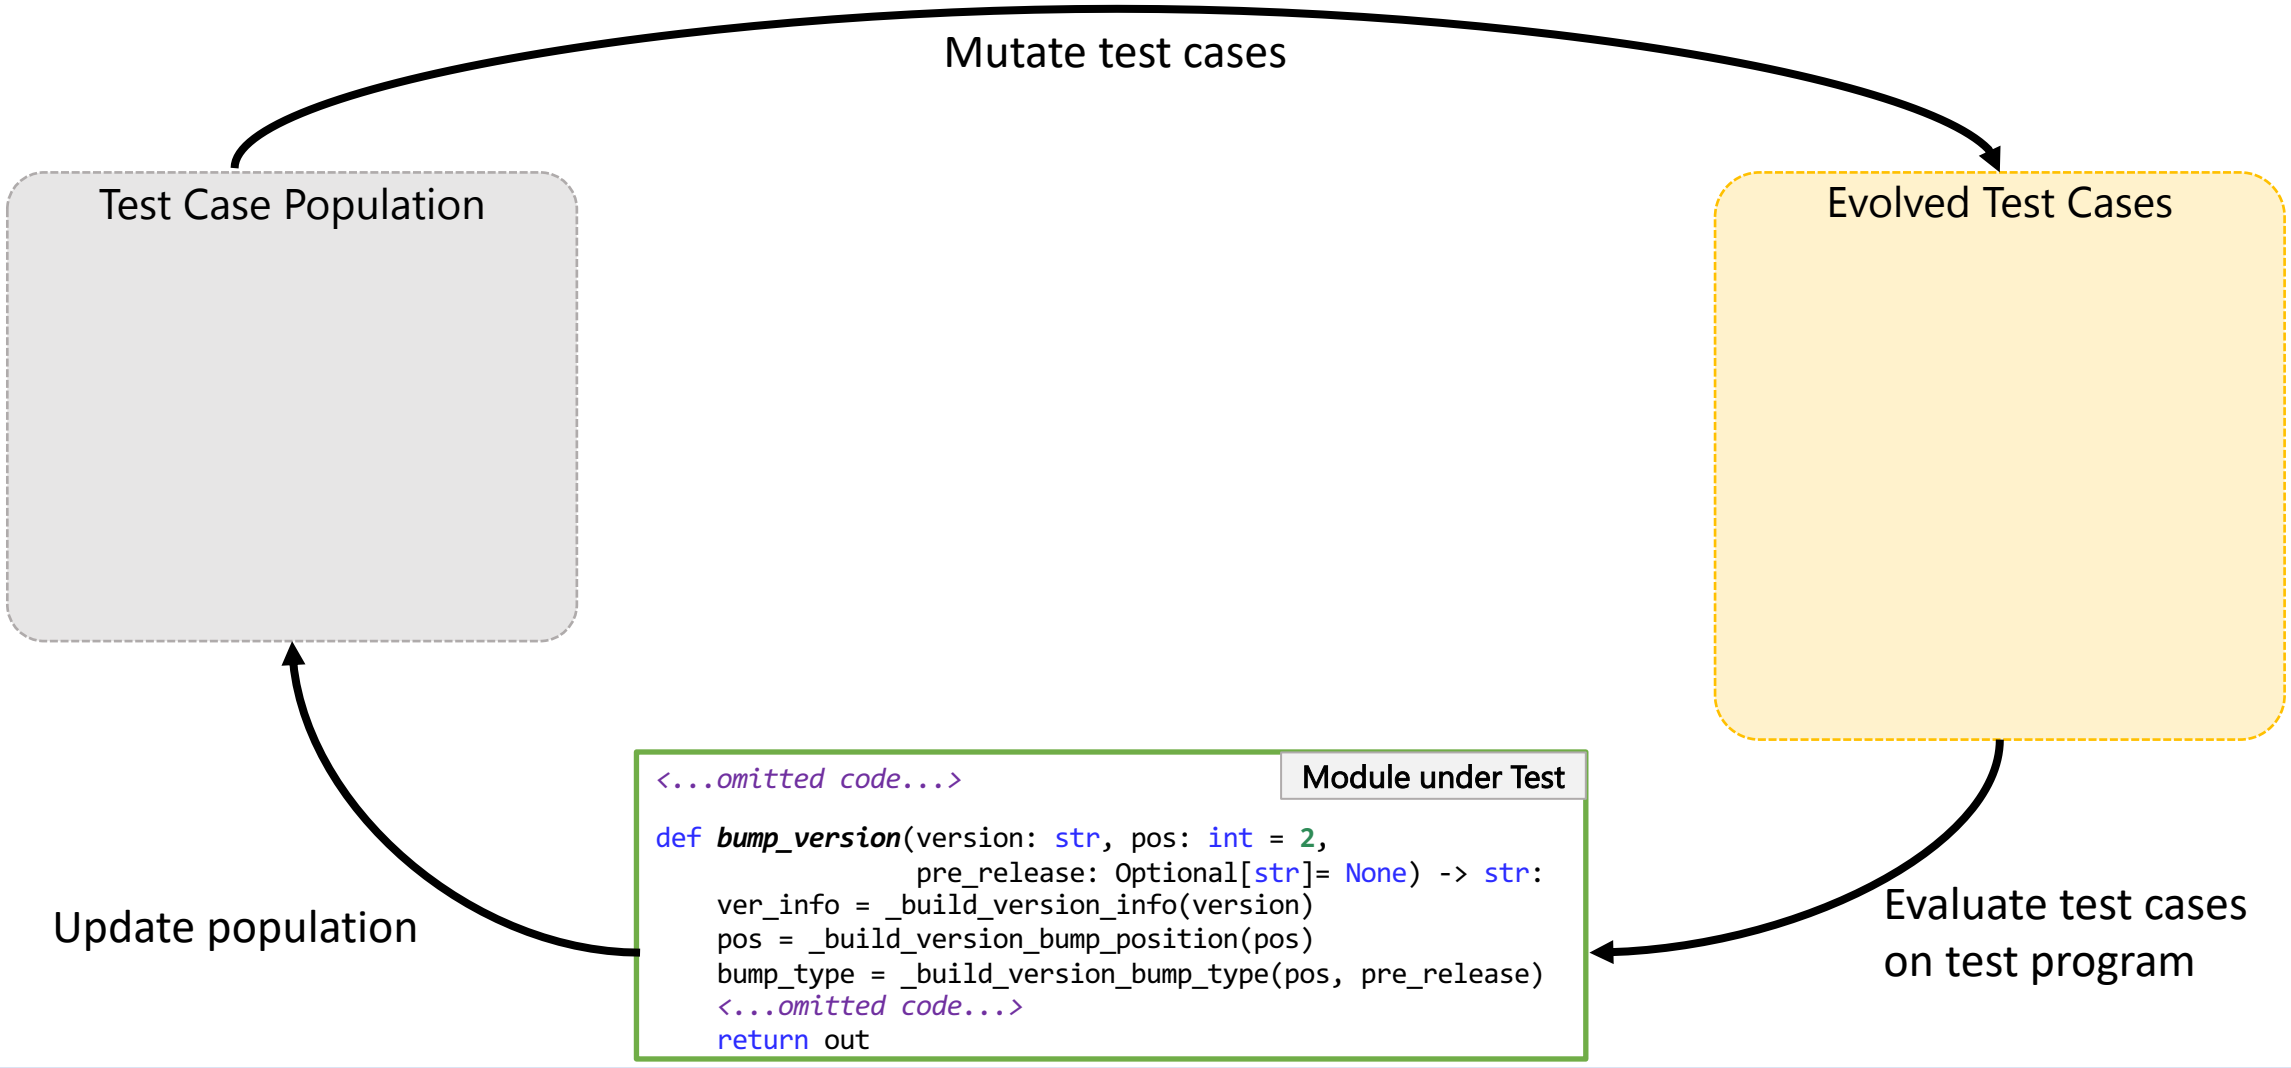

Example: Expected Behavior of Function

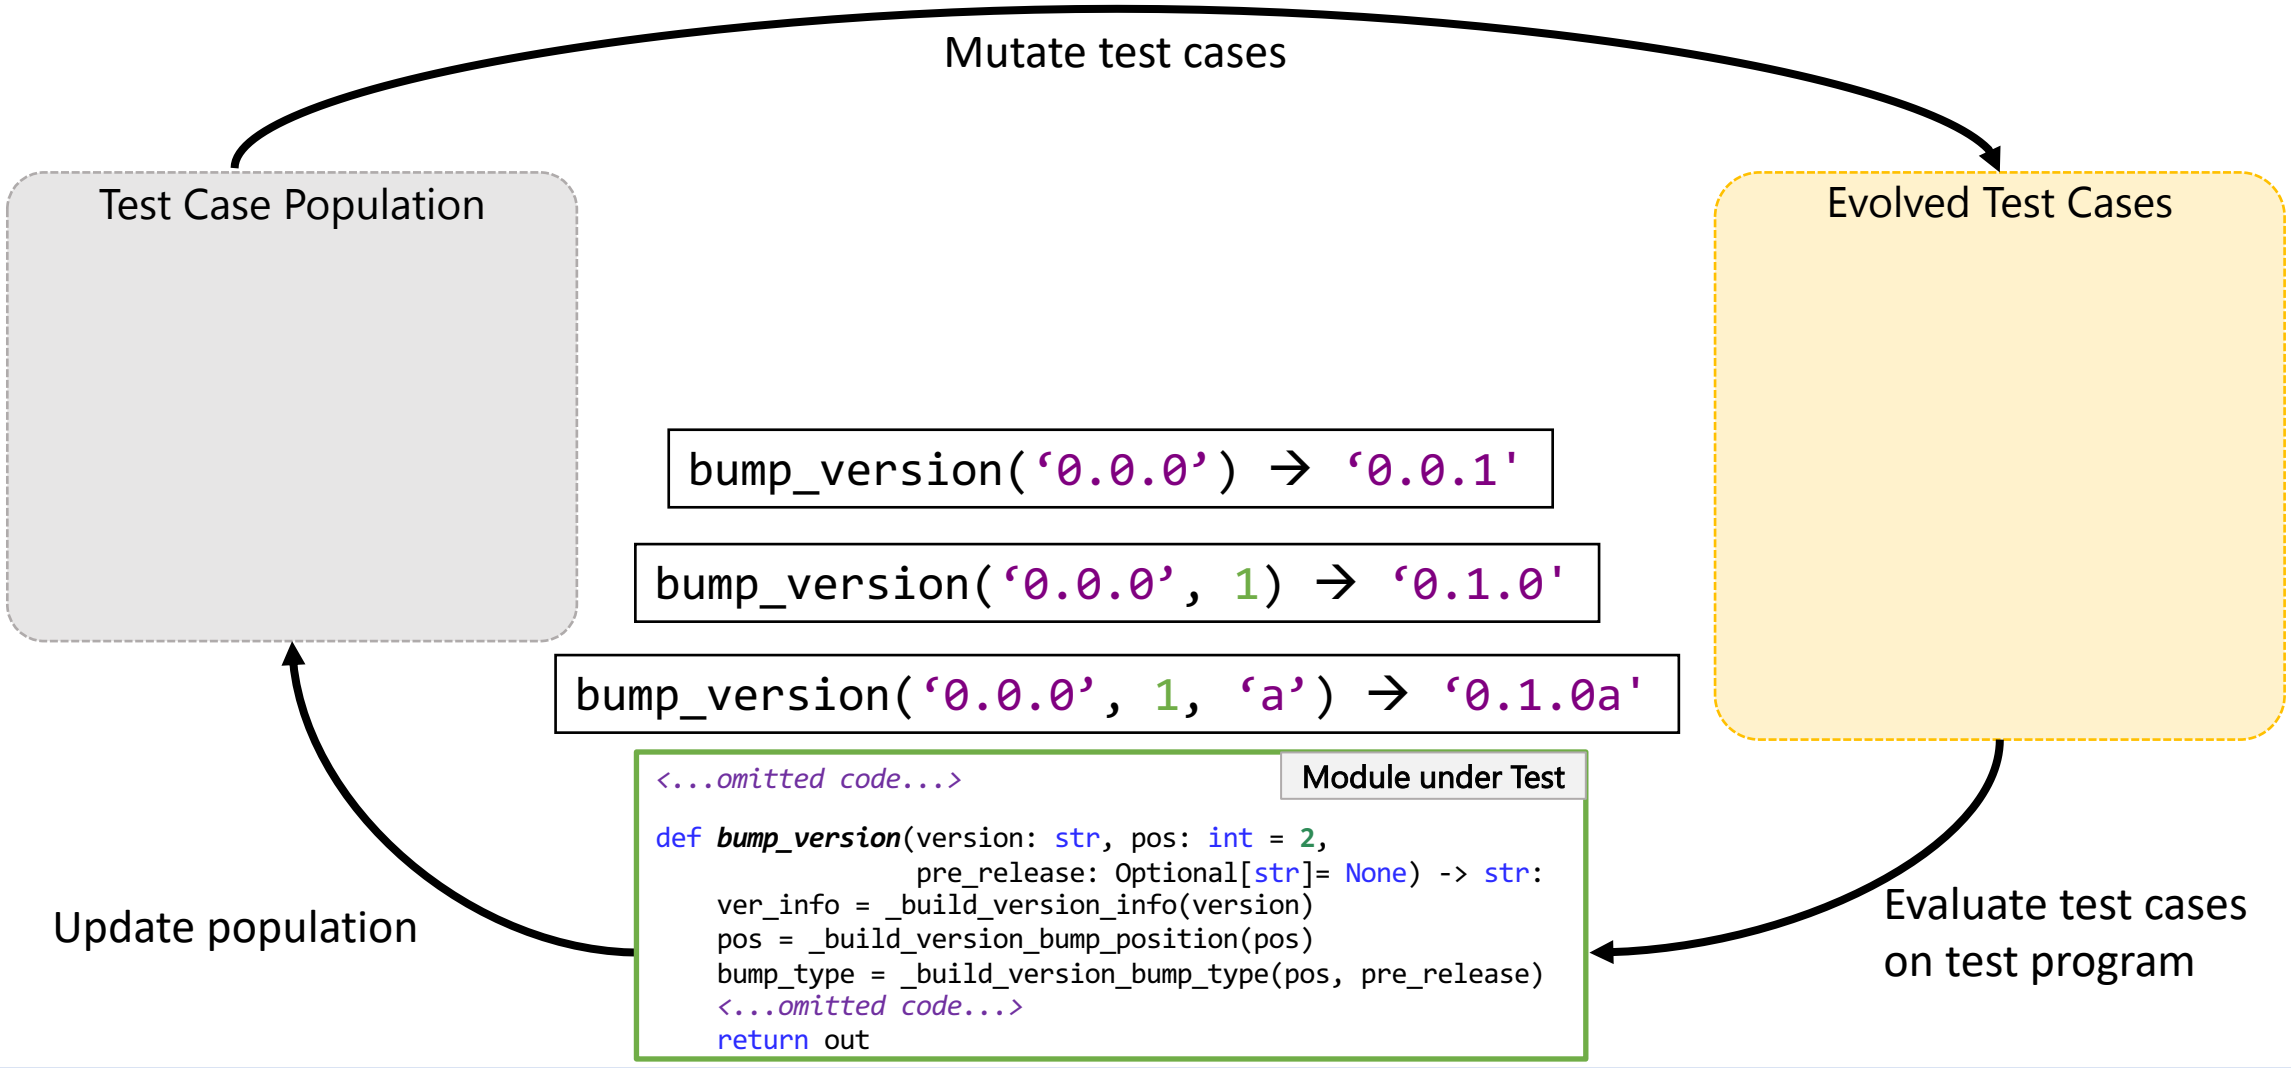

Initialize Test Population

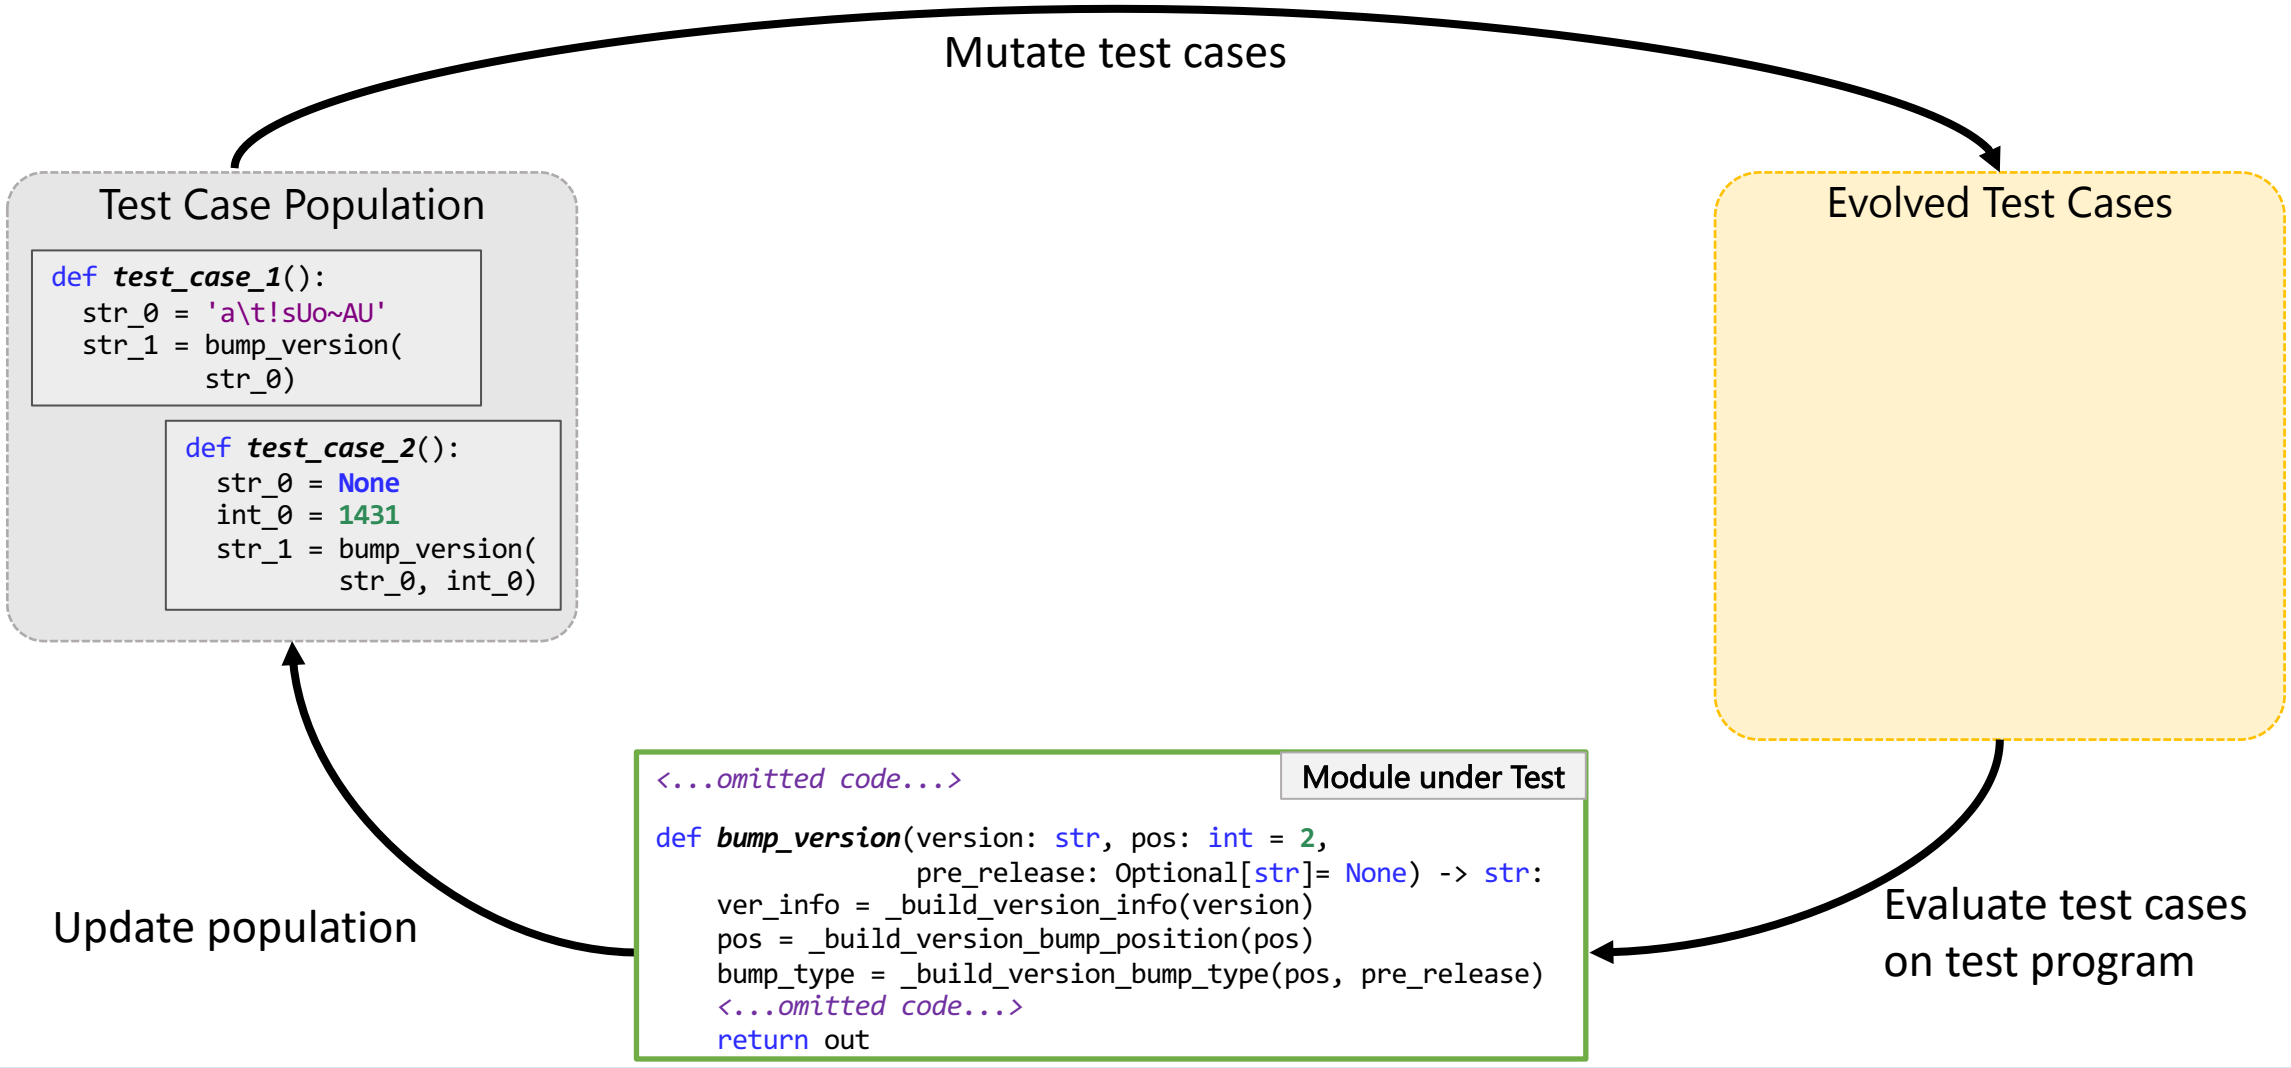

Current Tests have Low Coverage

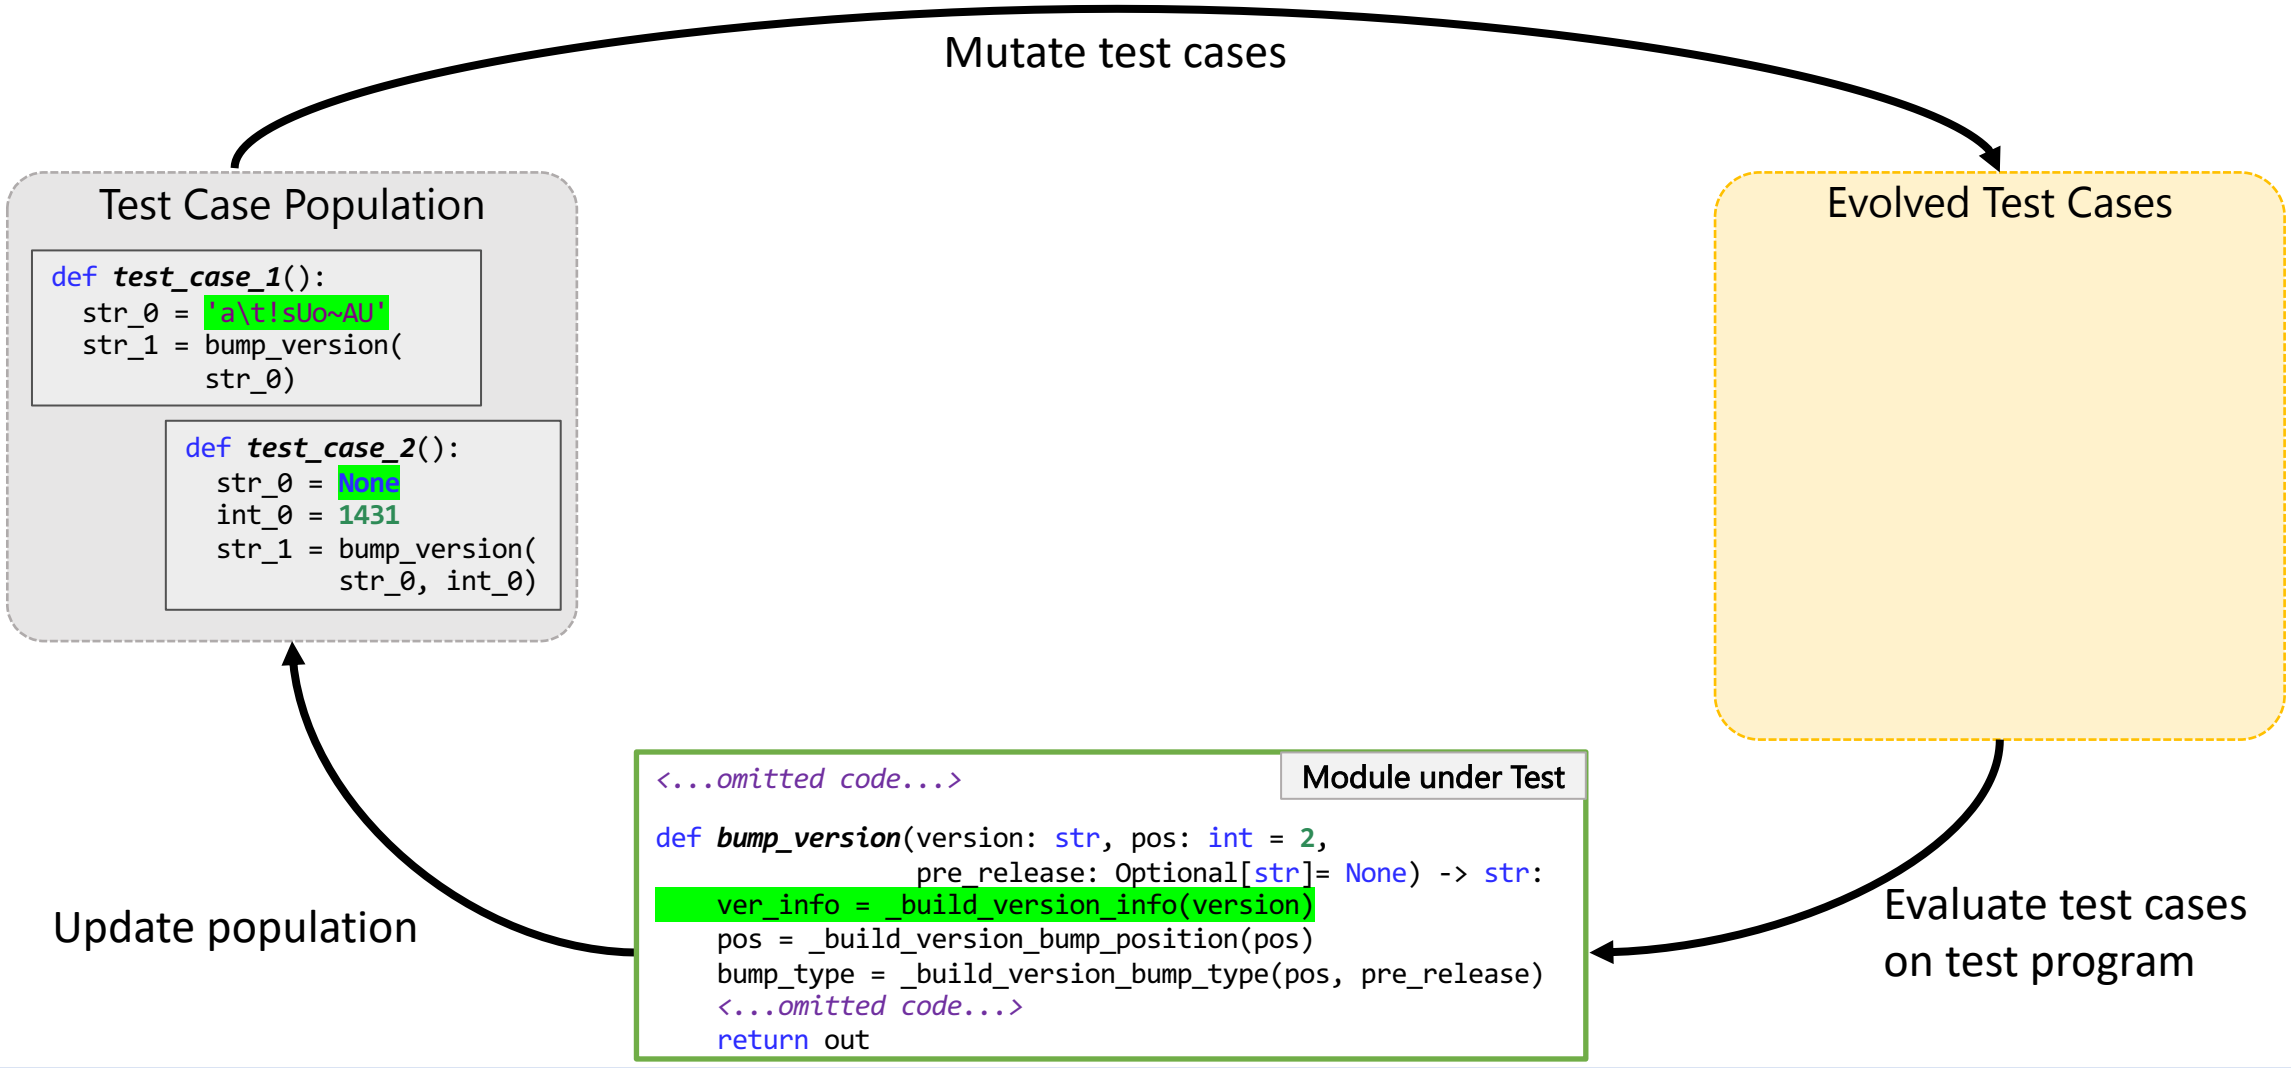

Create New Test Cases via Mutation

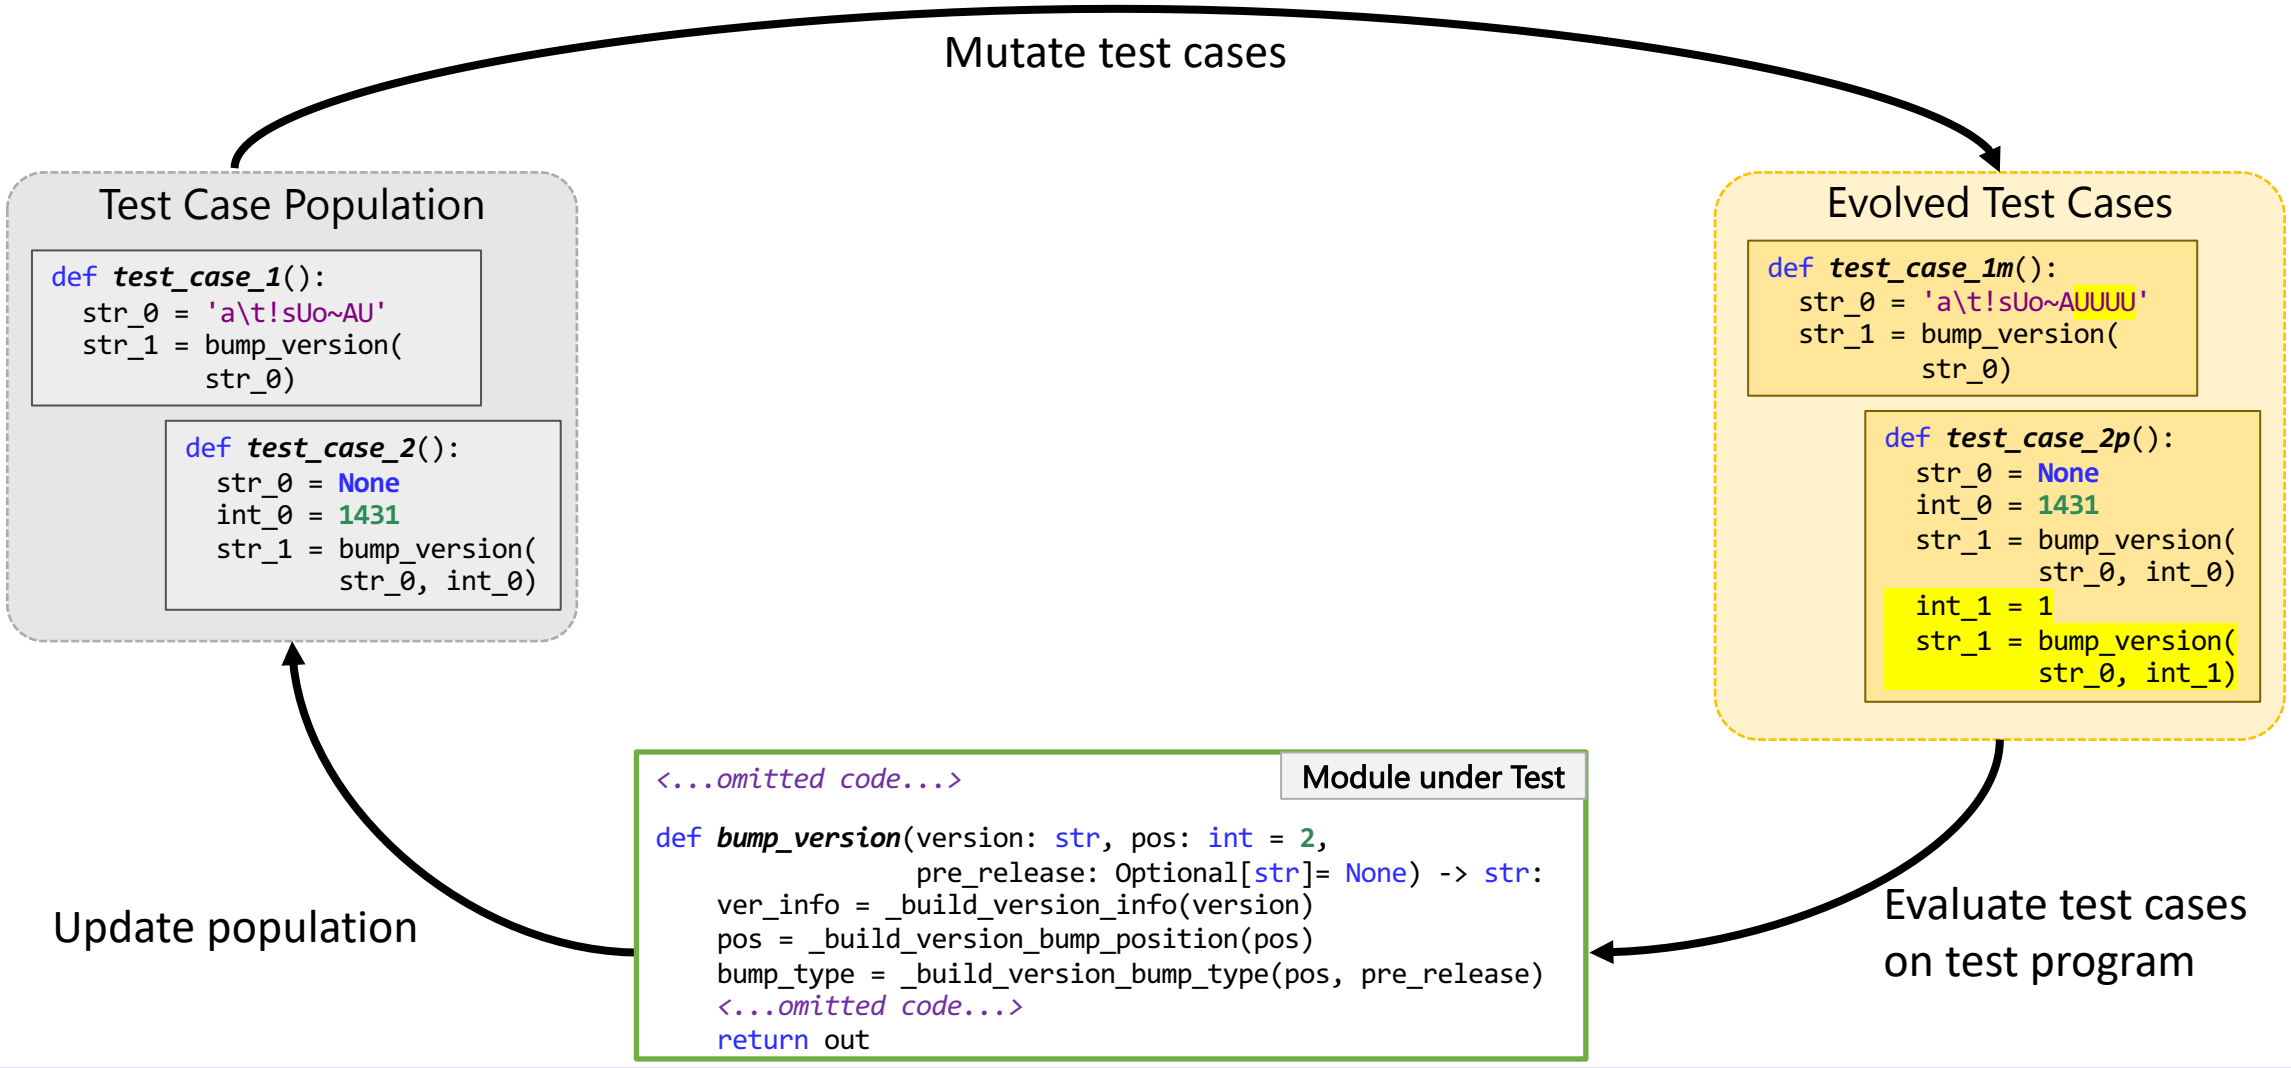

Mutation Unable to Increase Coverage

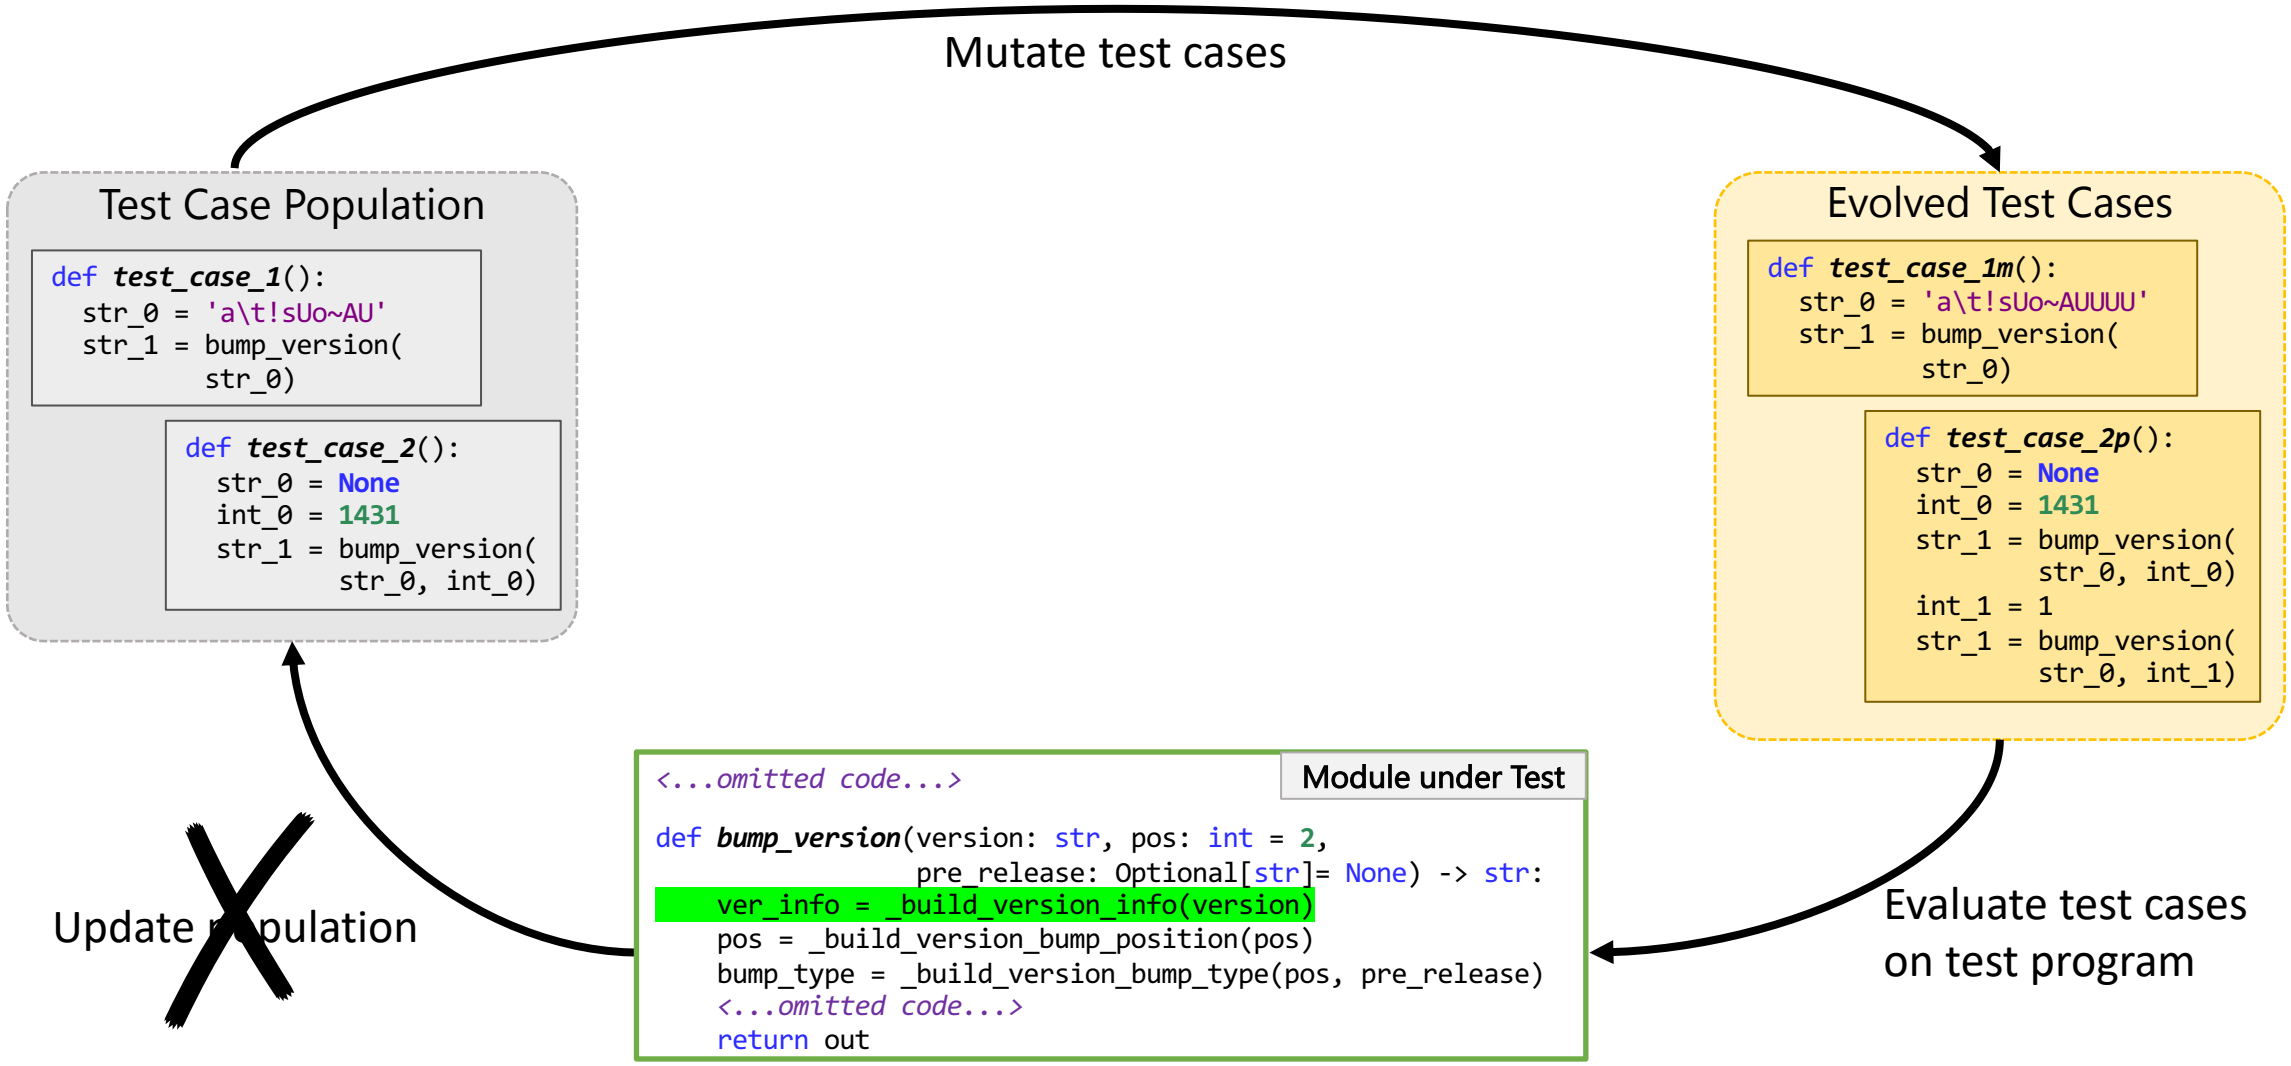

Search stalled. What to do now?

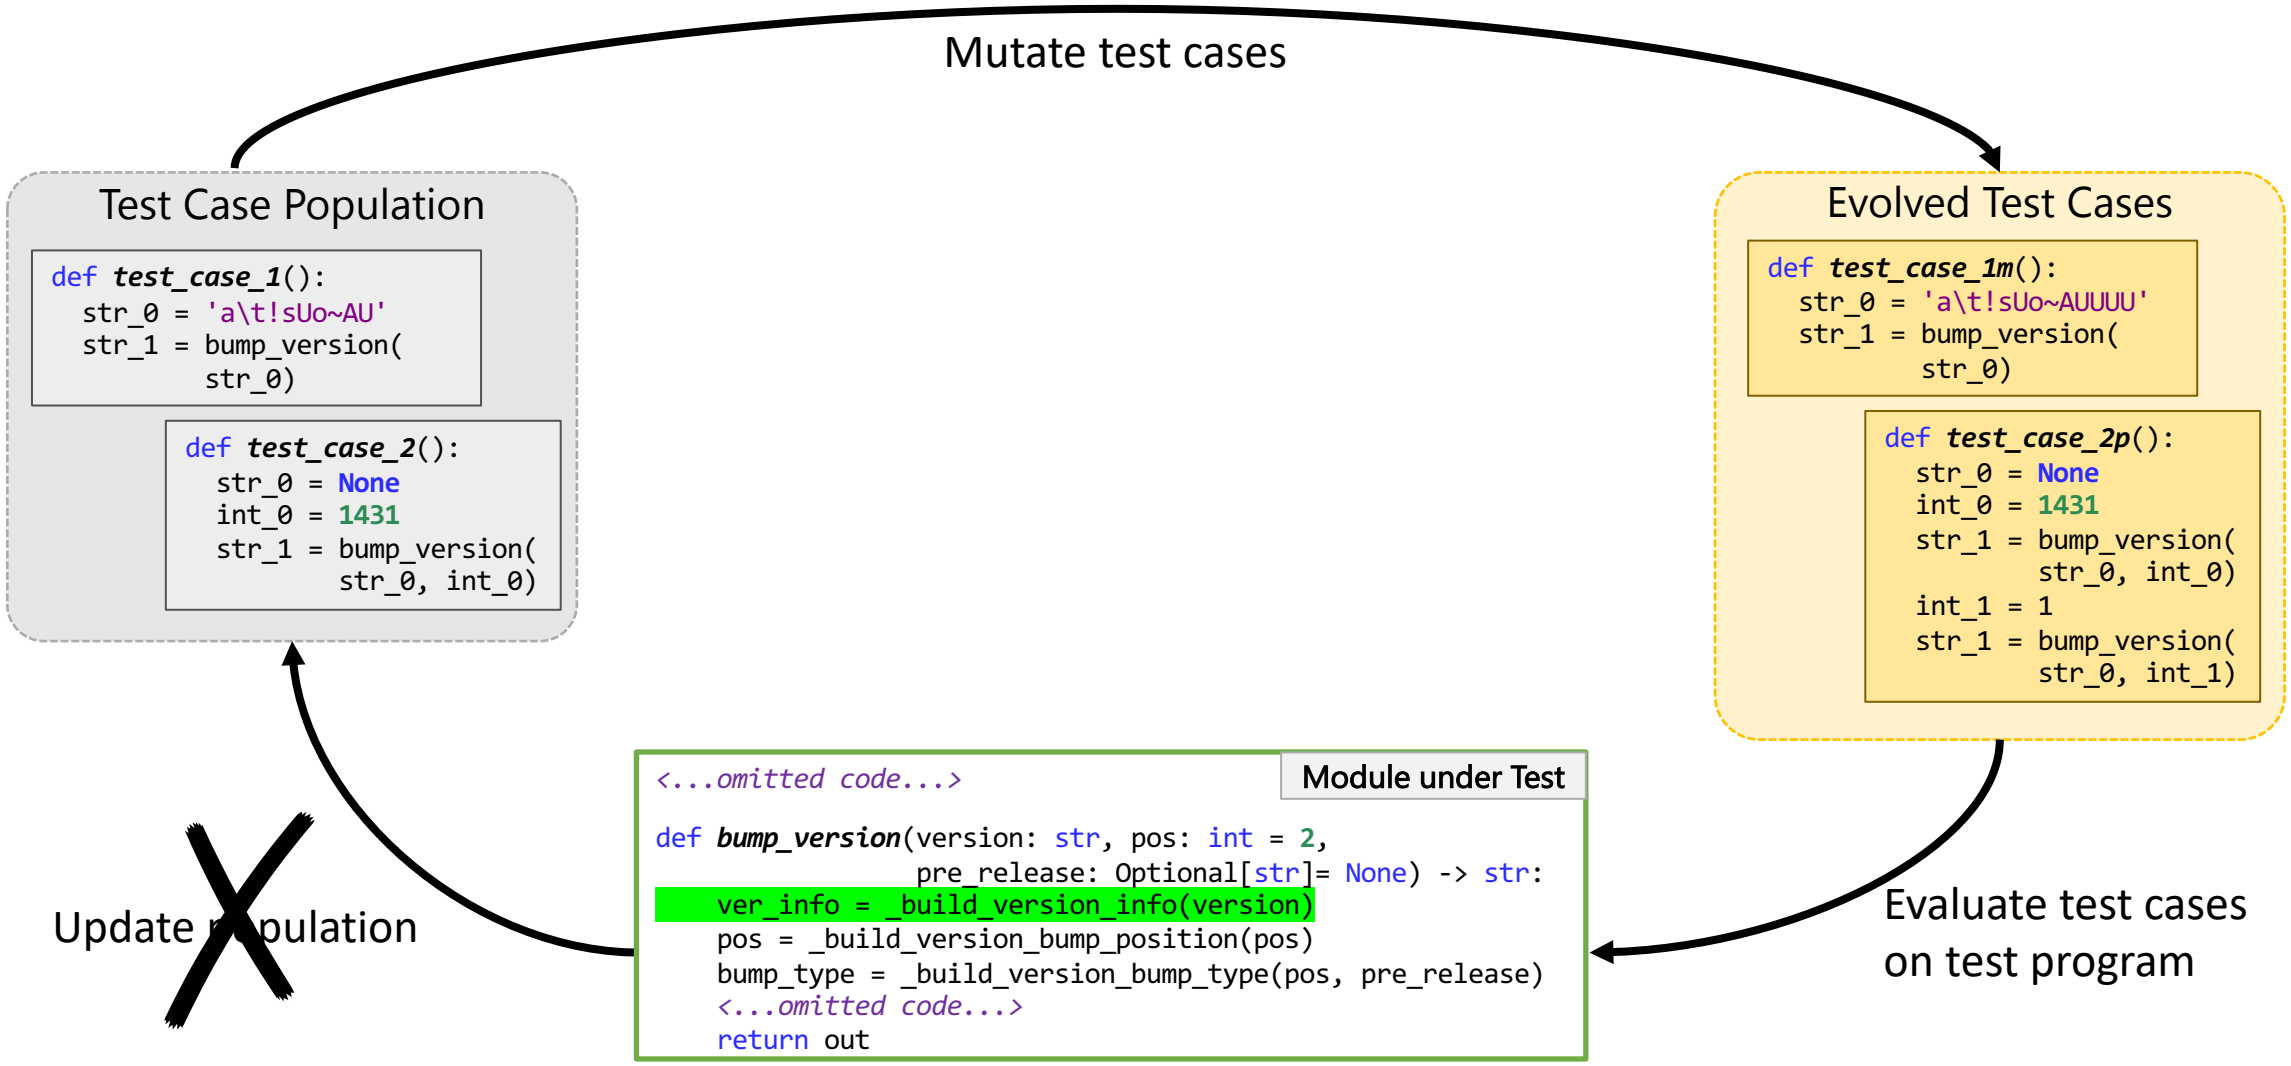

CodaMOSA: Asks for hints when stuck

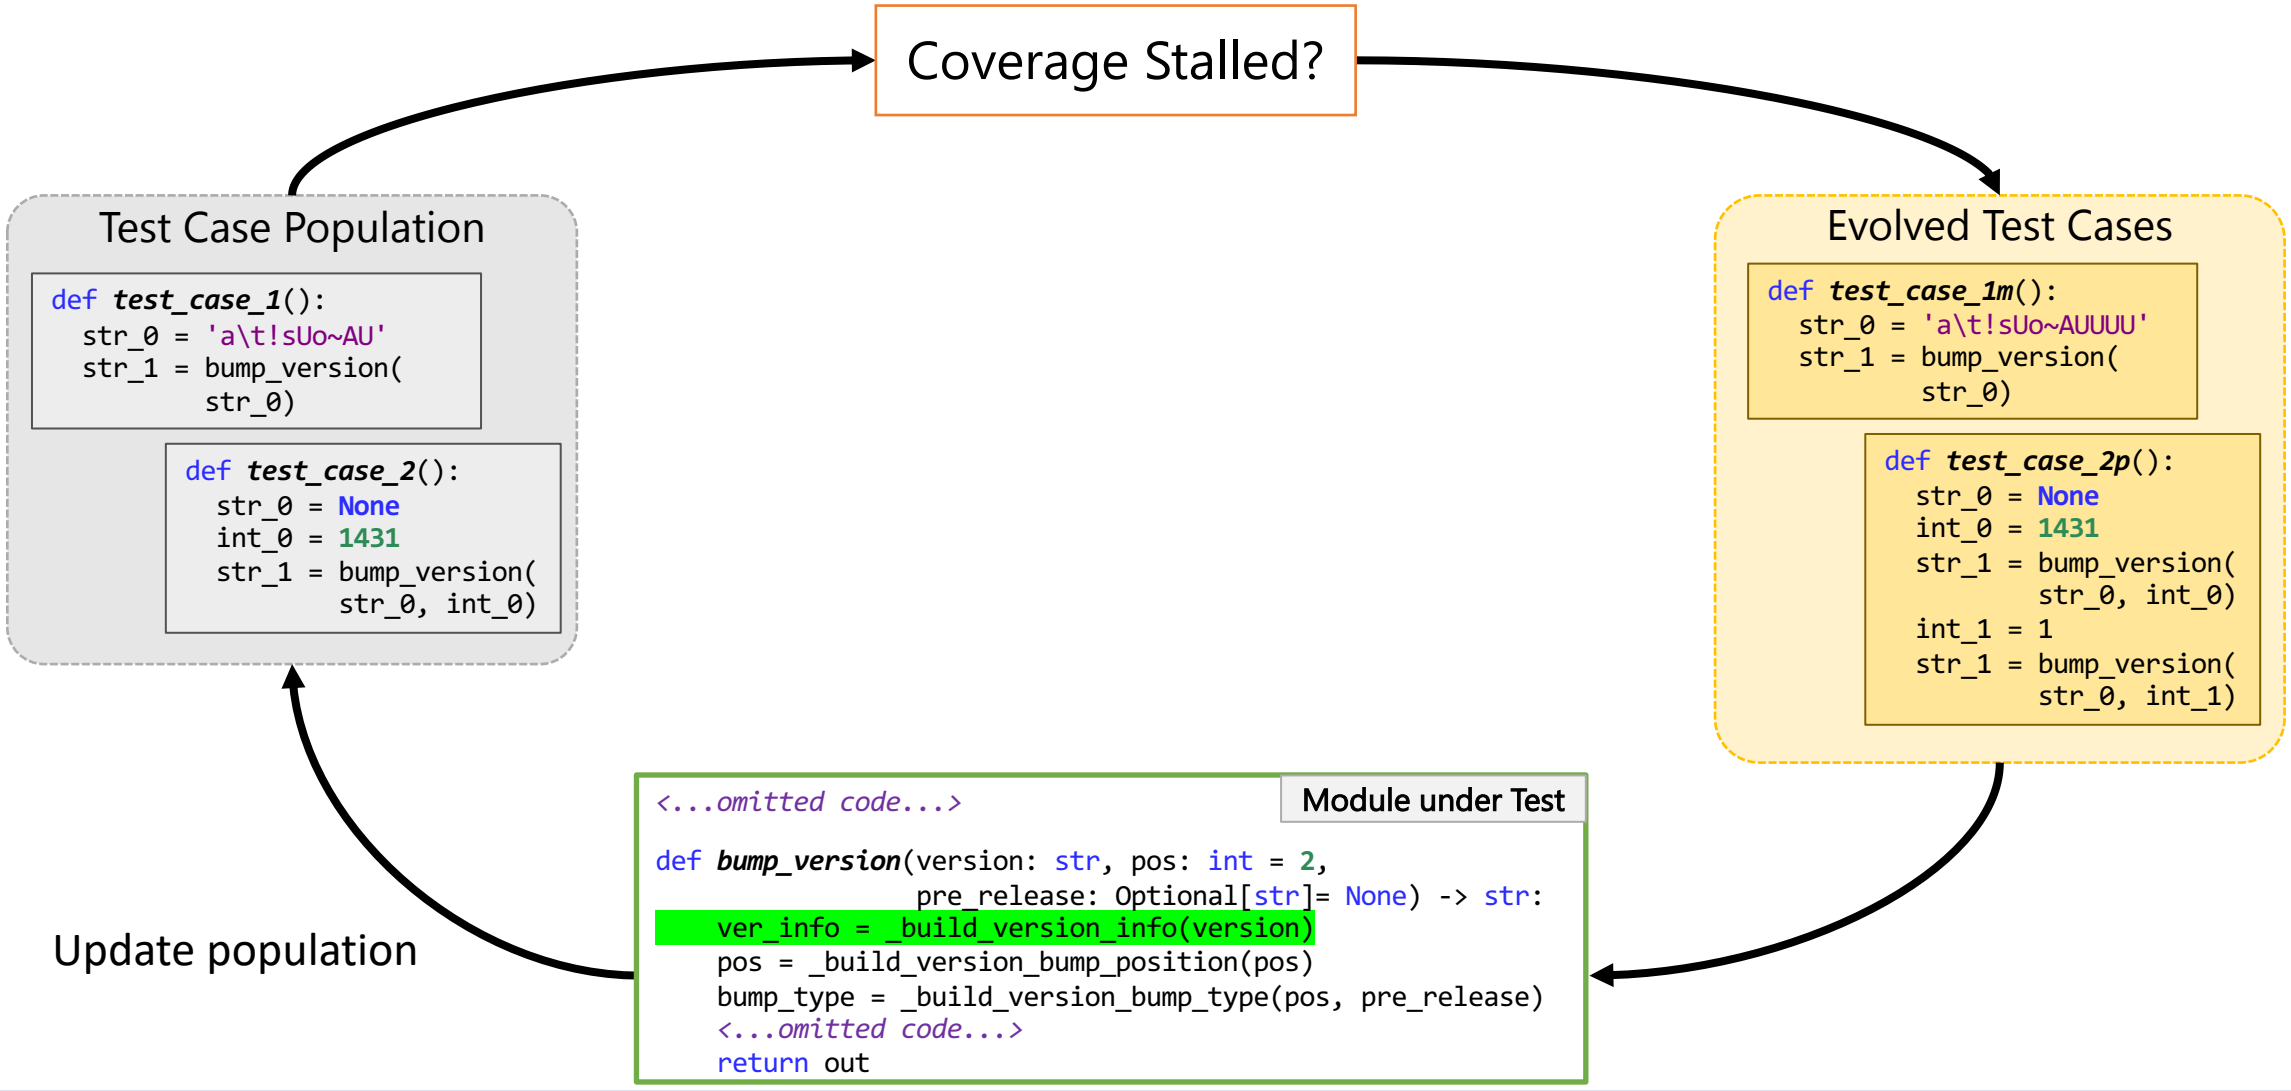

CodaMOSA: Asks for hints when stuck

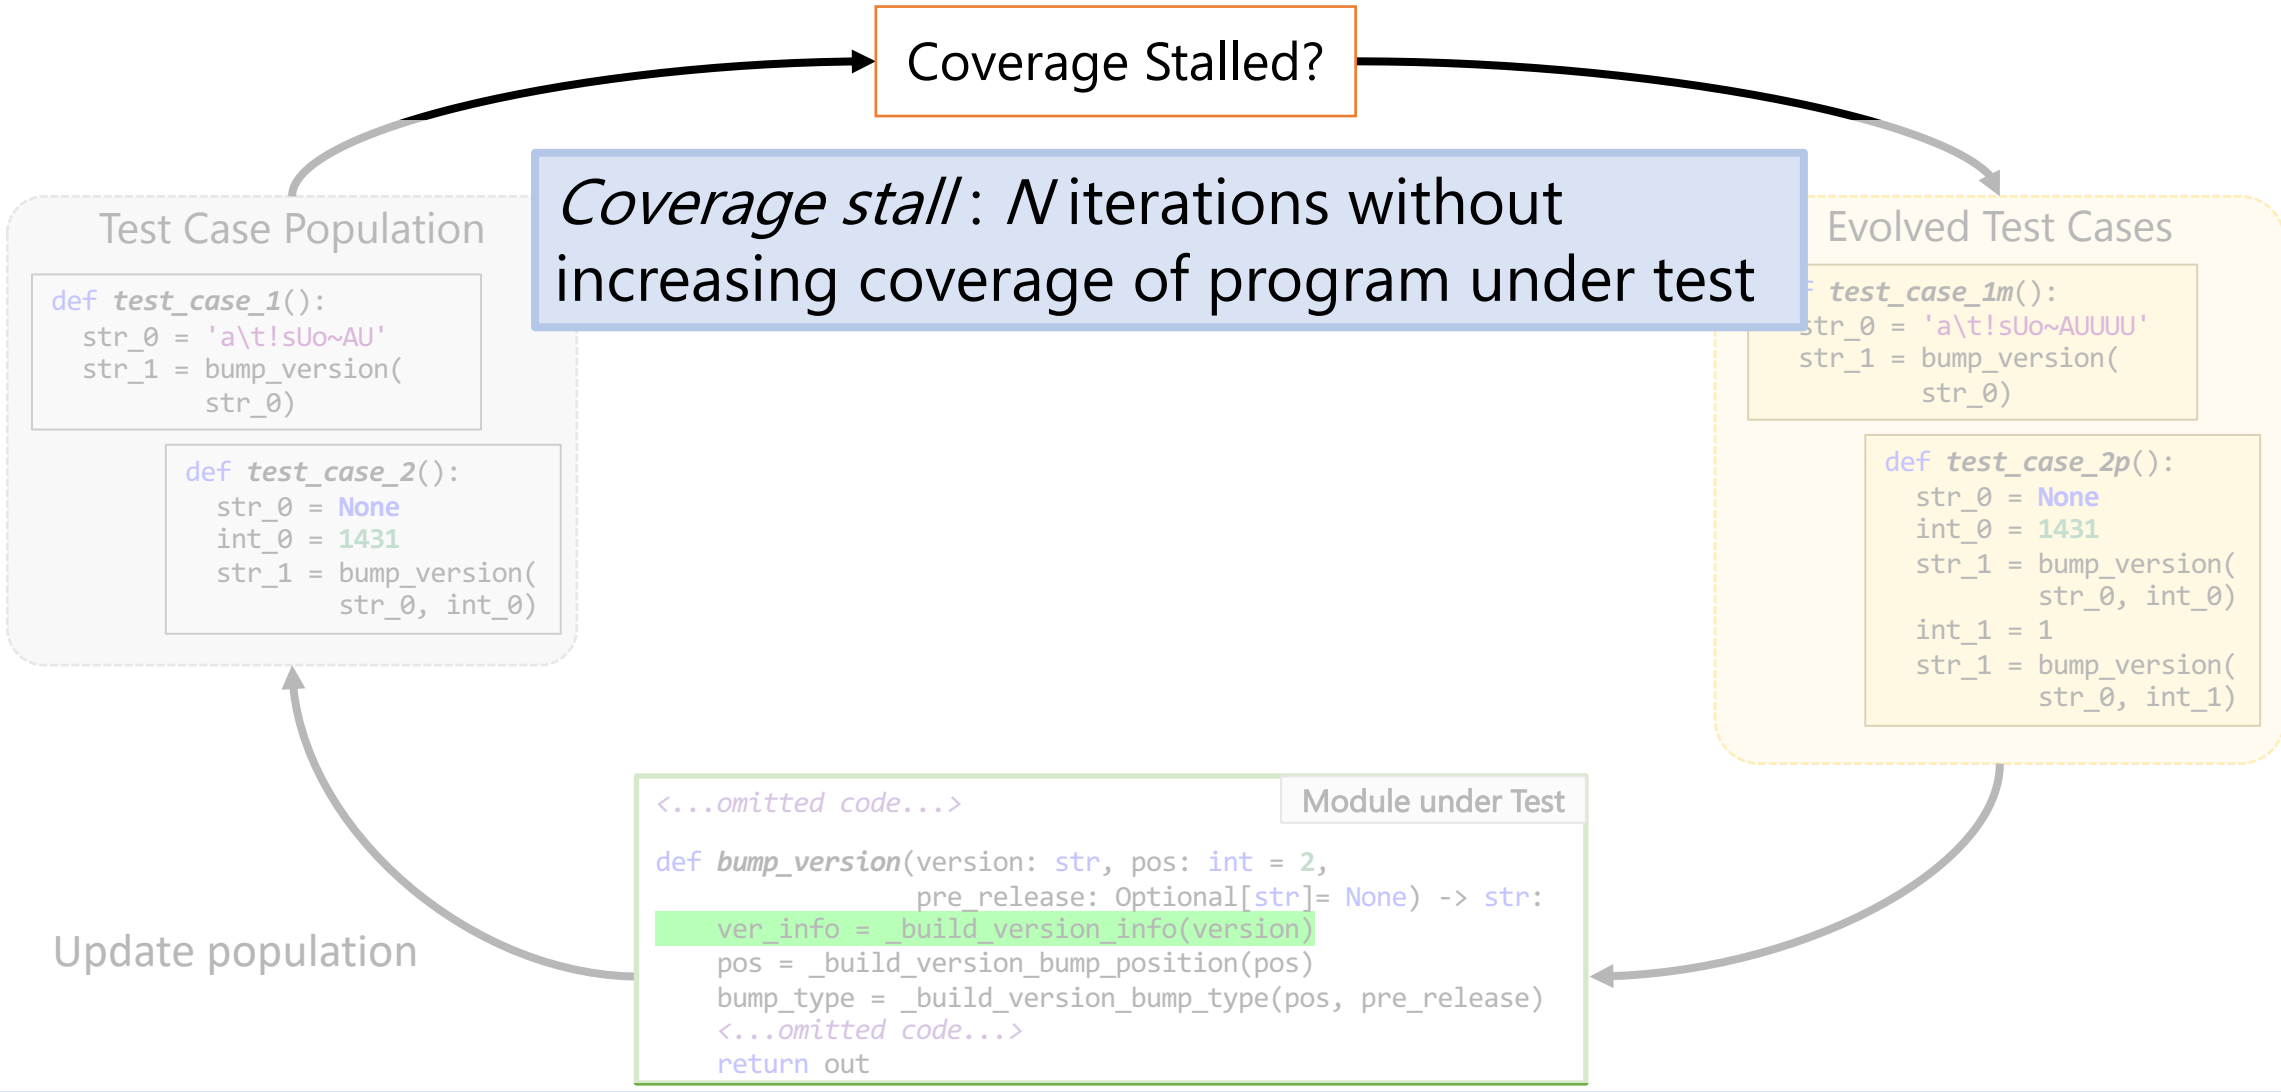

CodaMOSA: Asks for hints when stuck

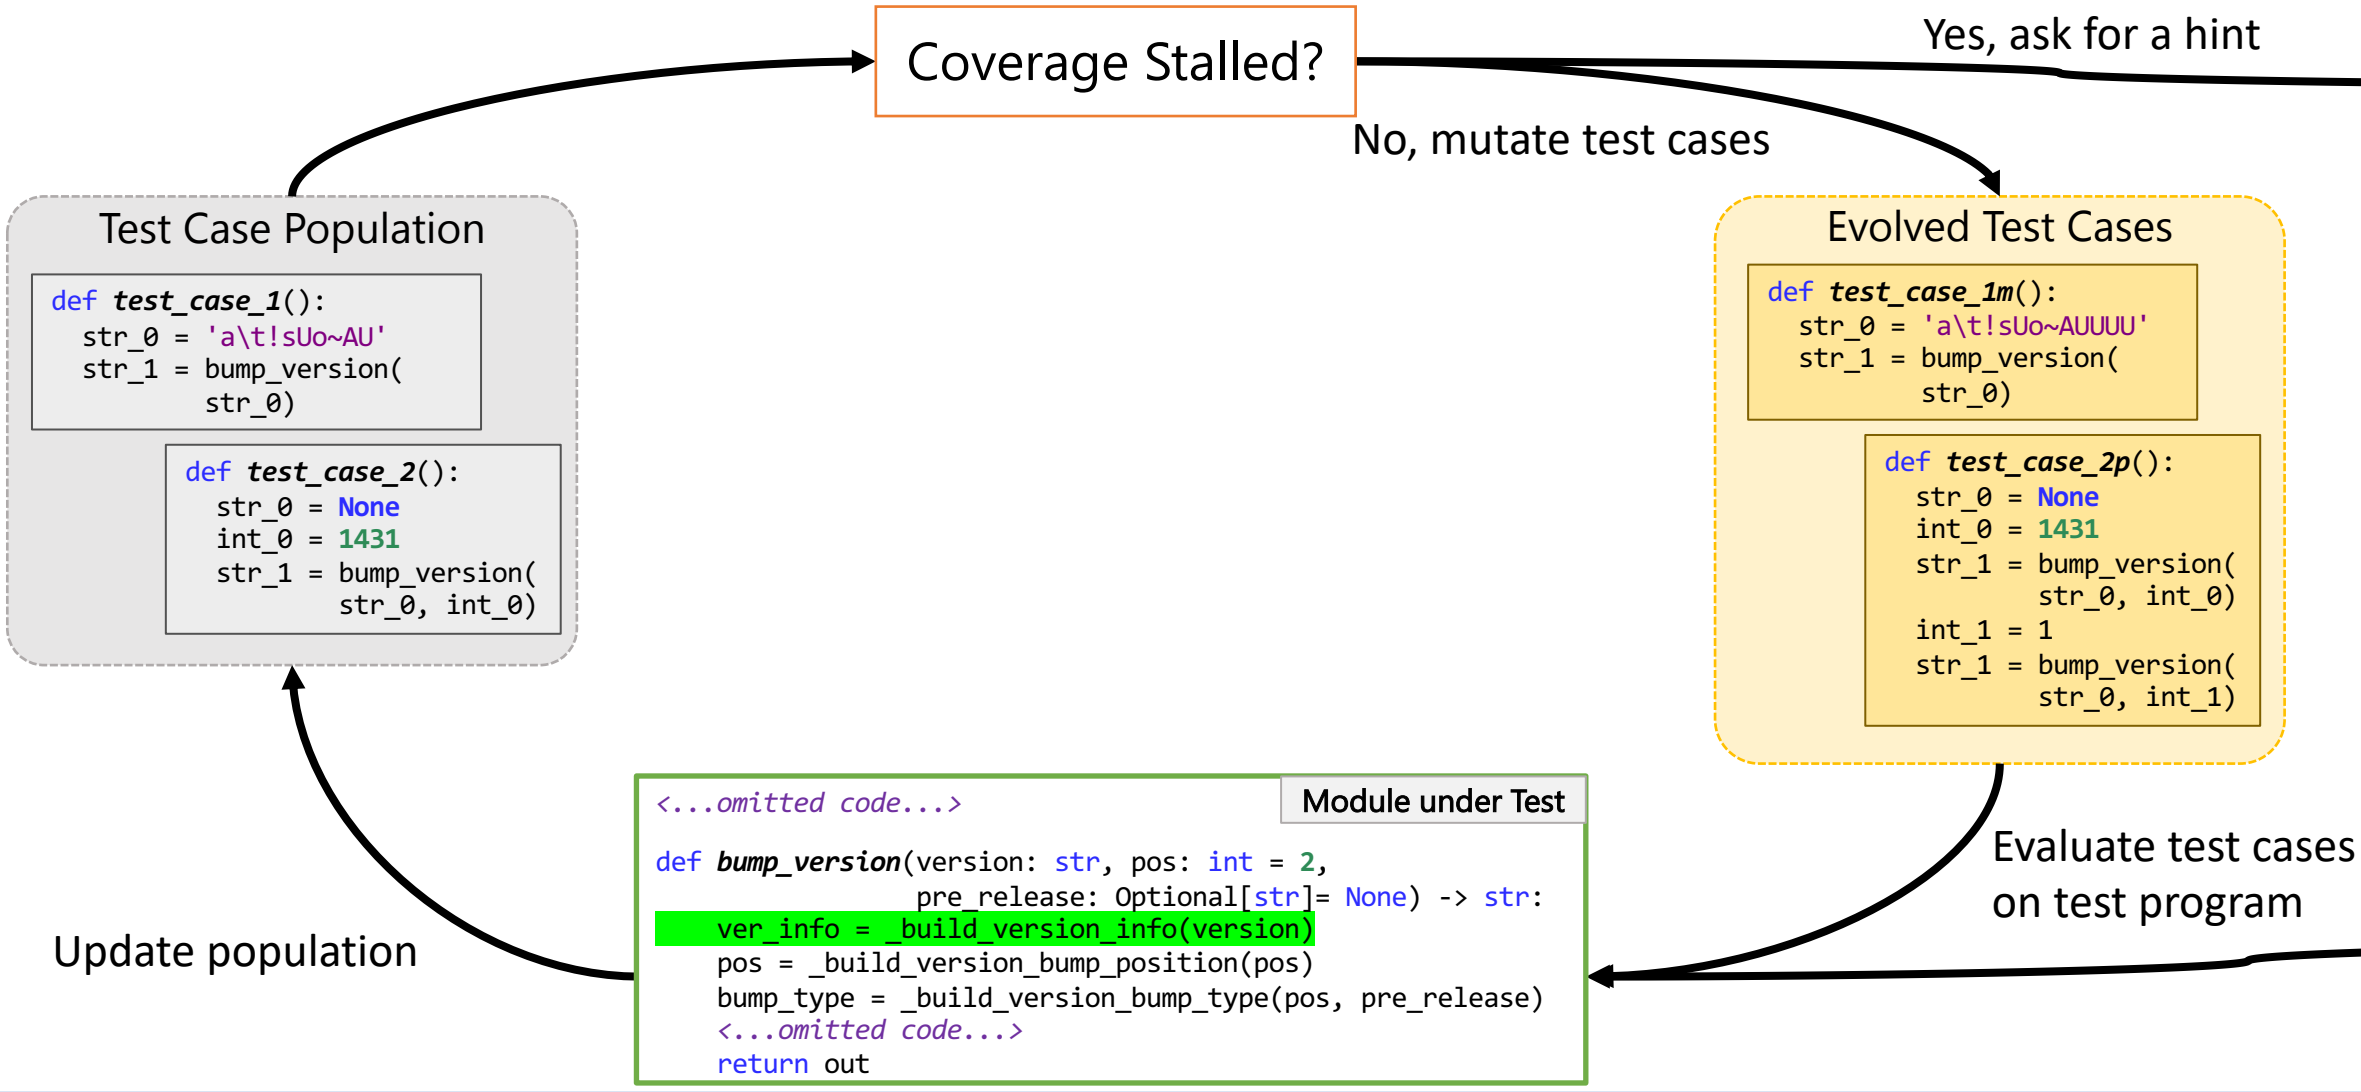

Search Stalled

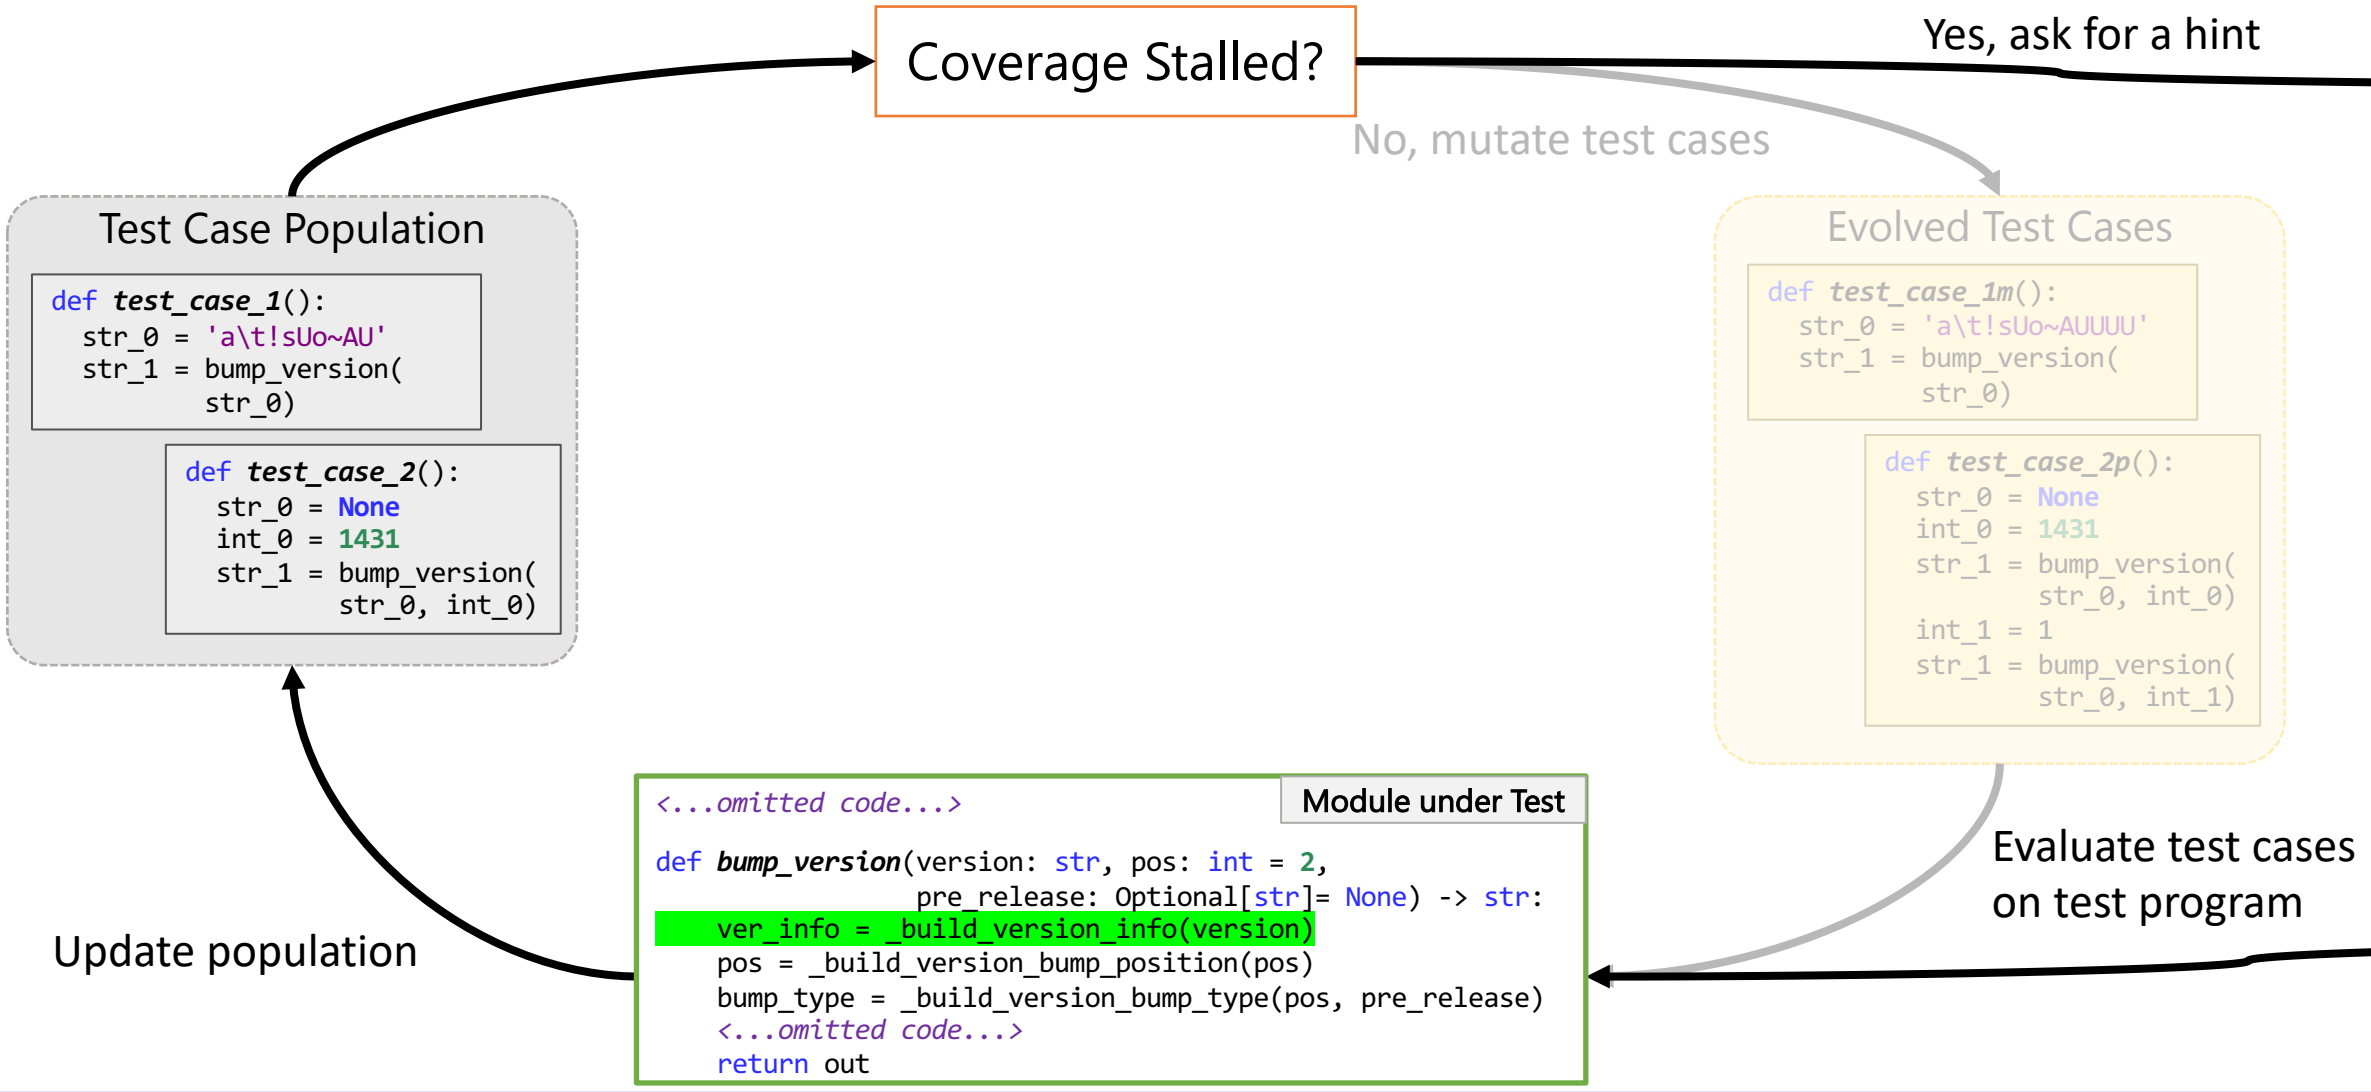

Time to Ask for a Hint

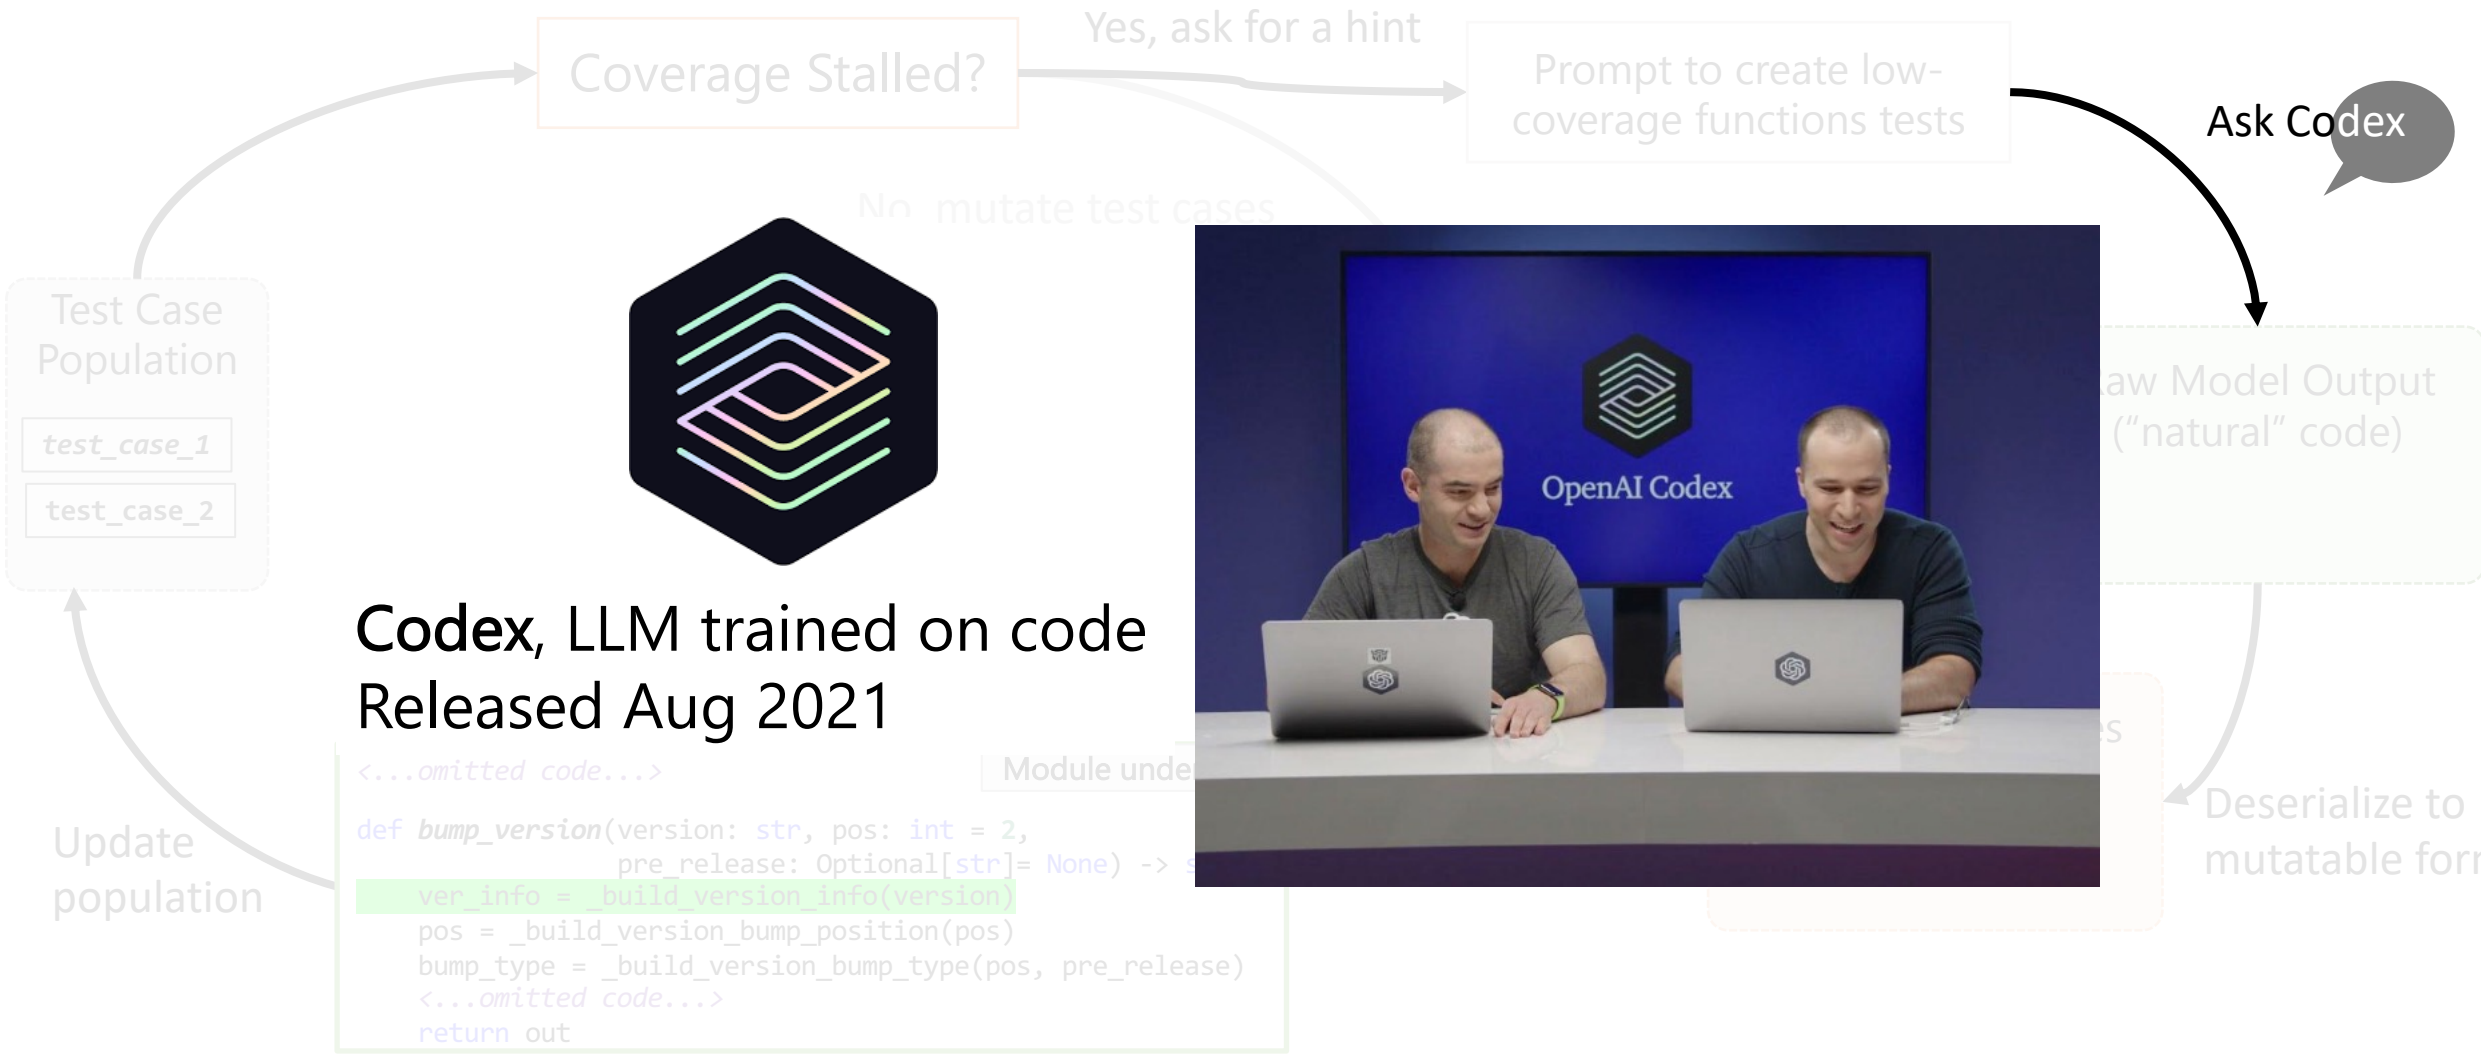

Codex, LLM trained on code
Released Aug 2021

Time to Ask for a Hint

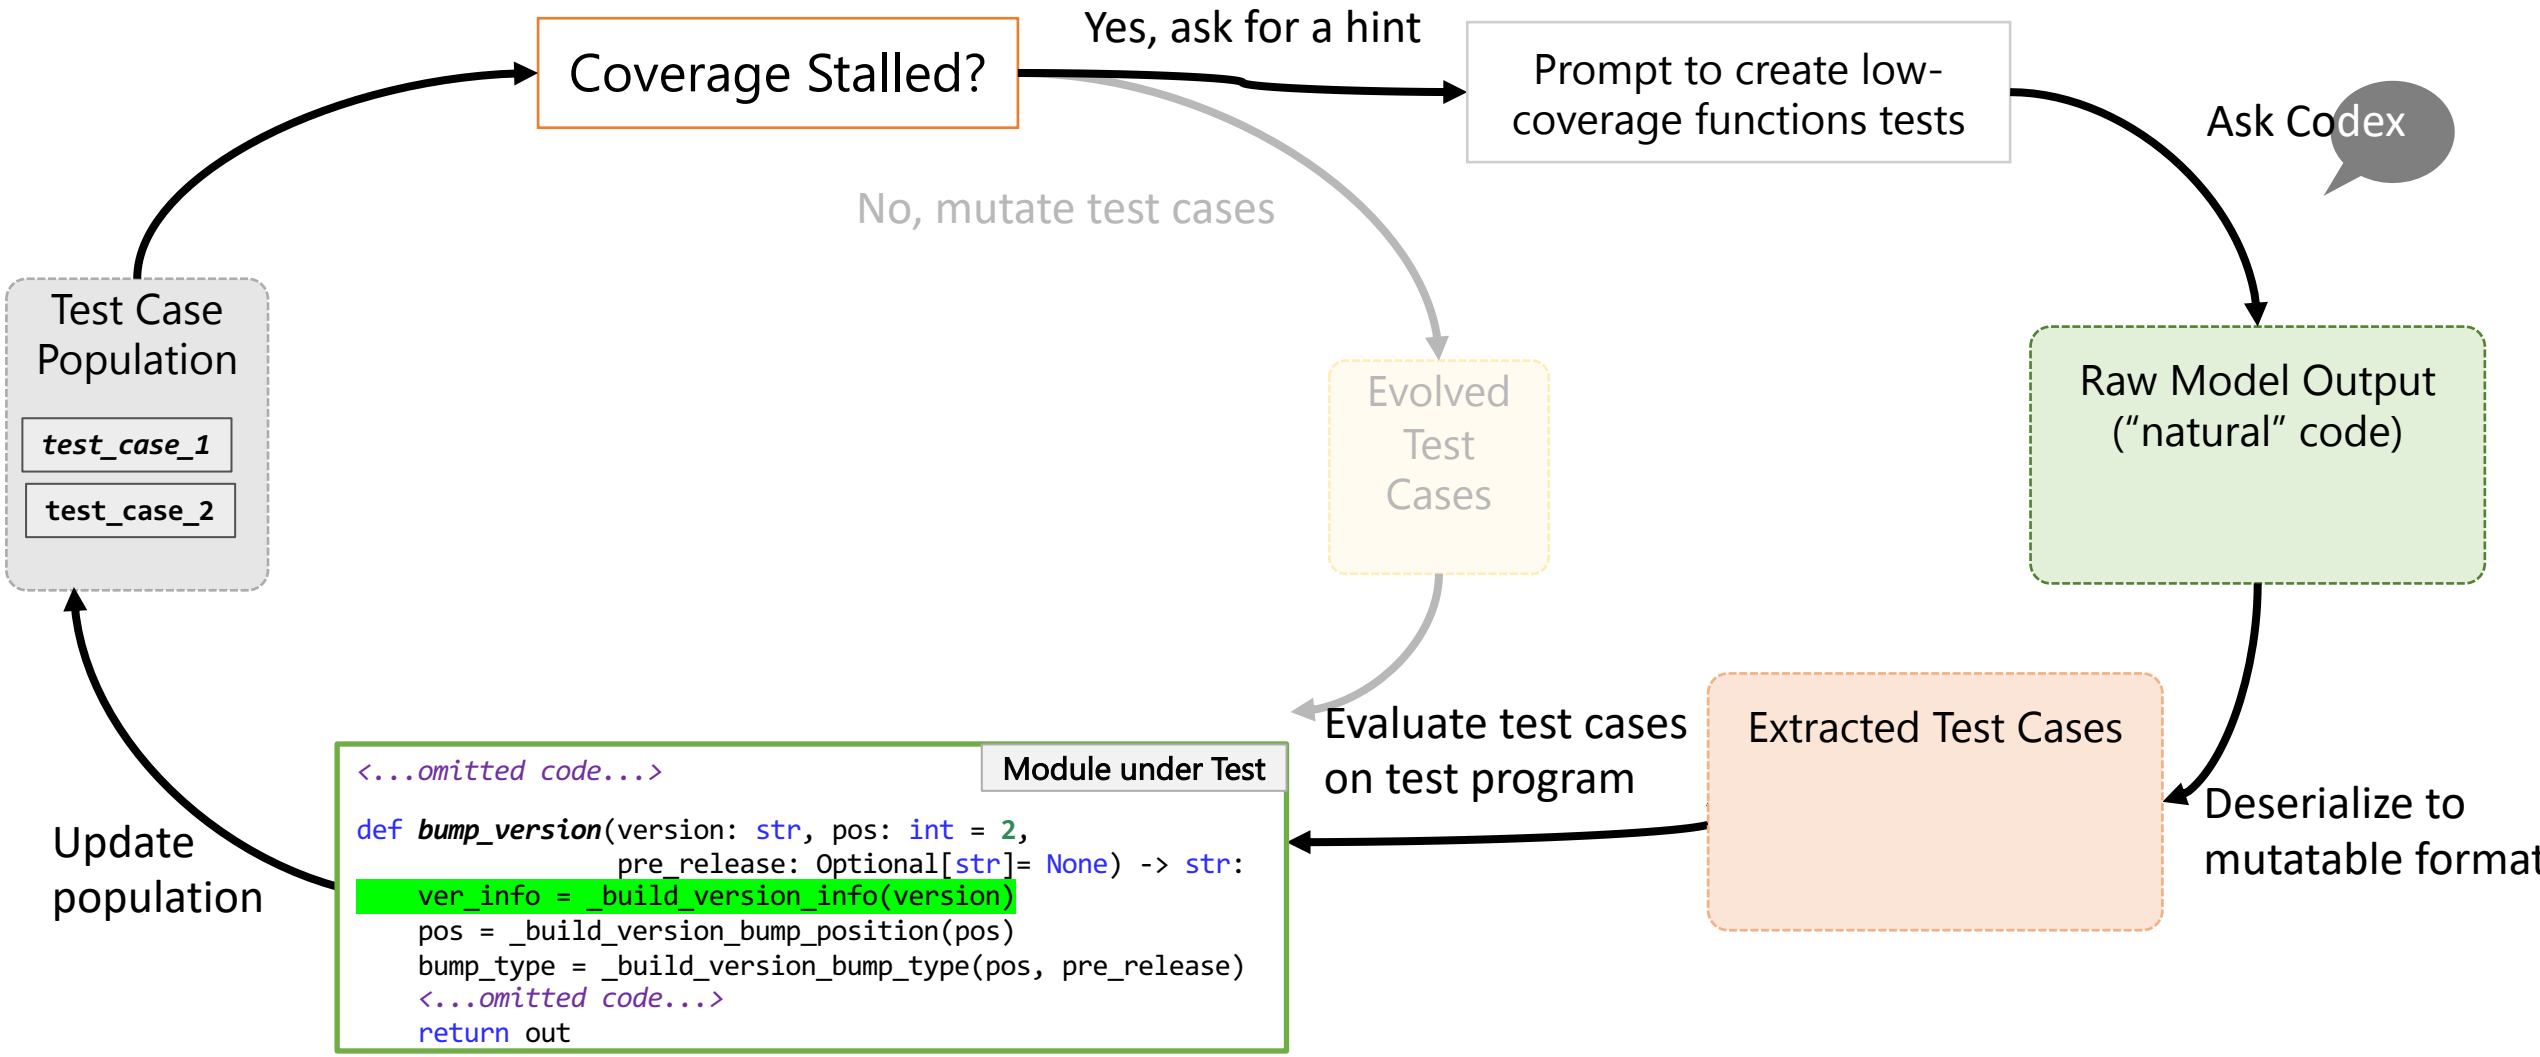

Time to Ask for a Hint

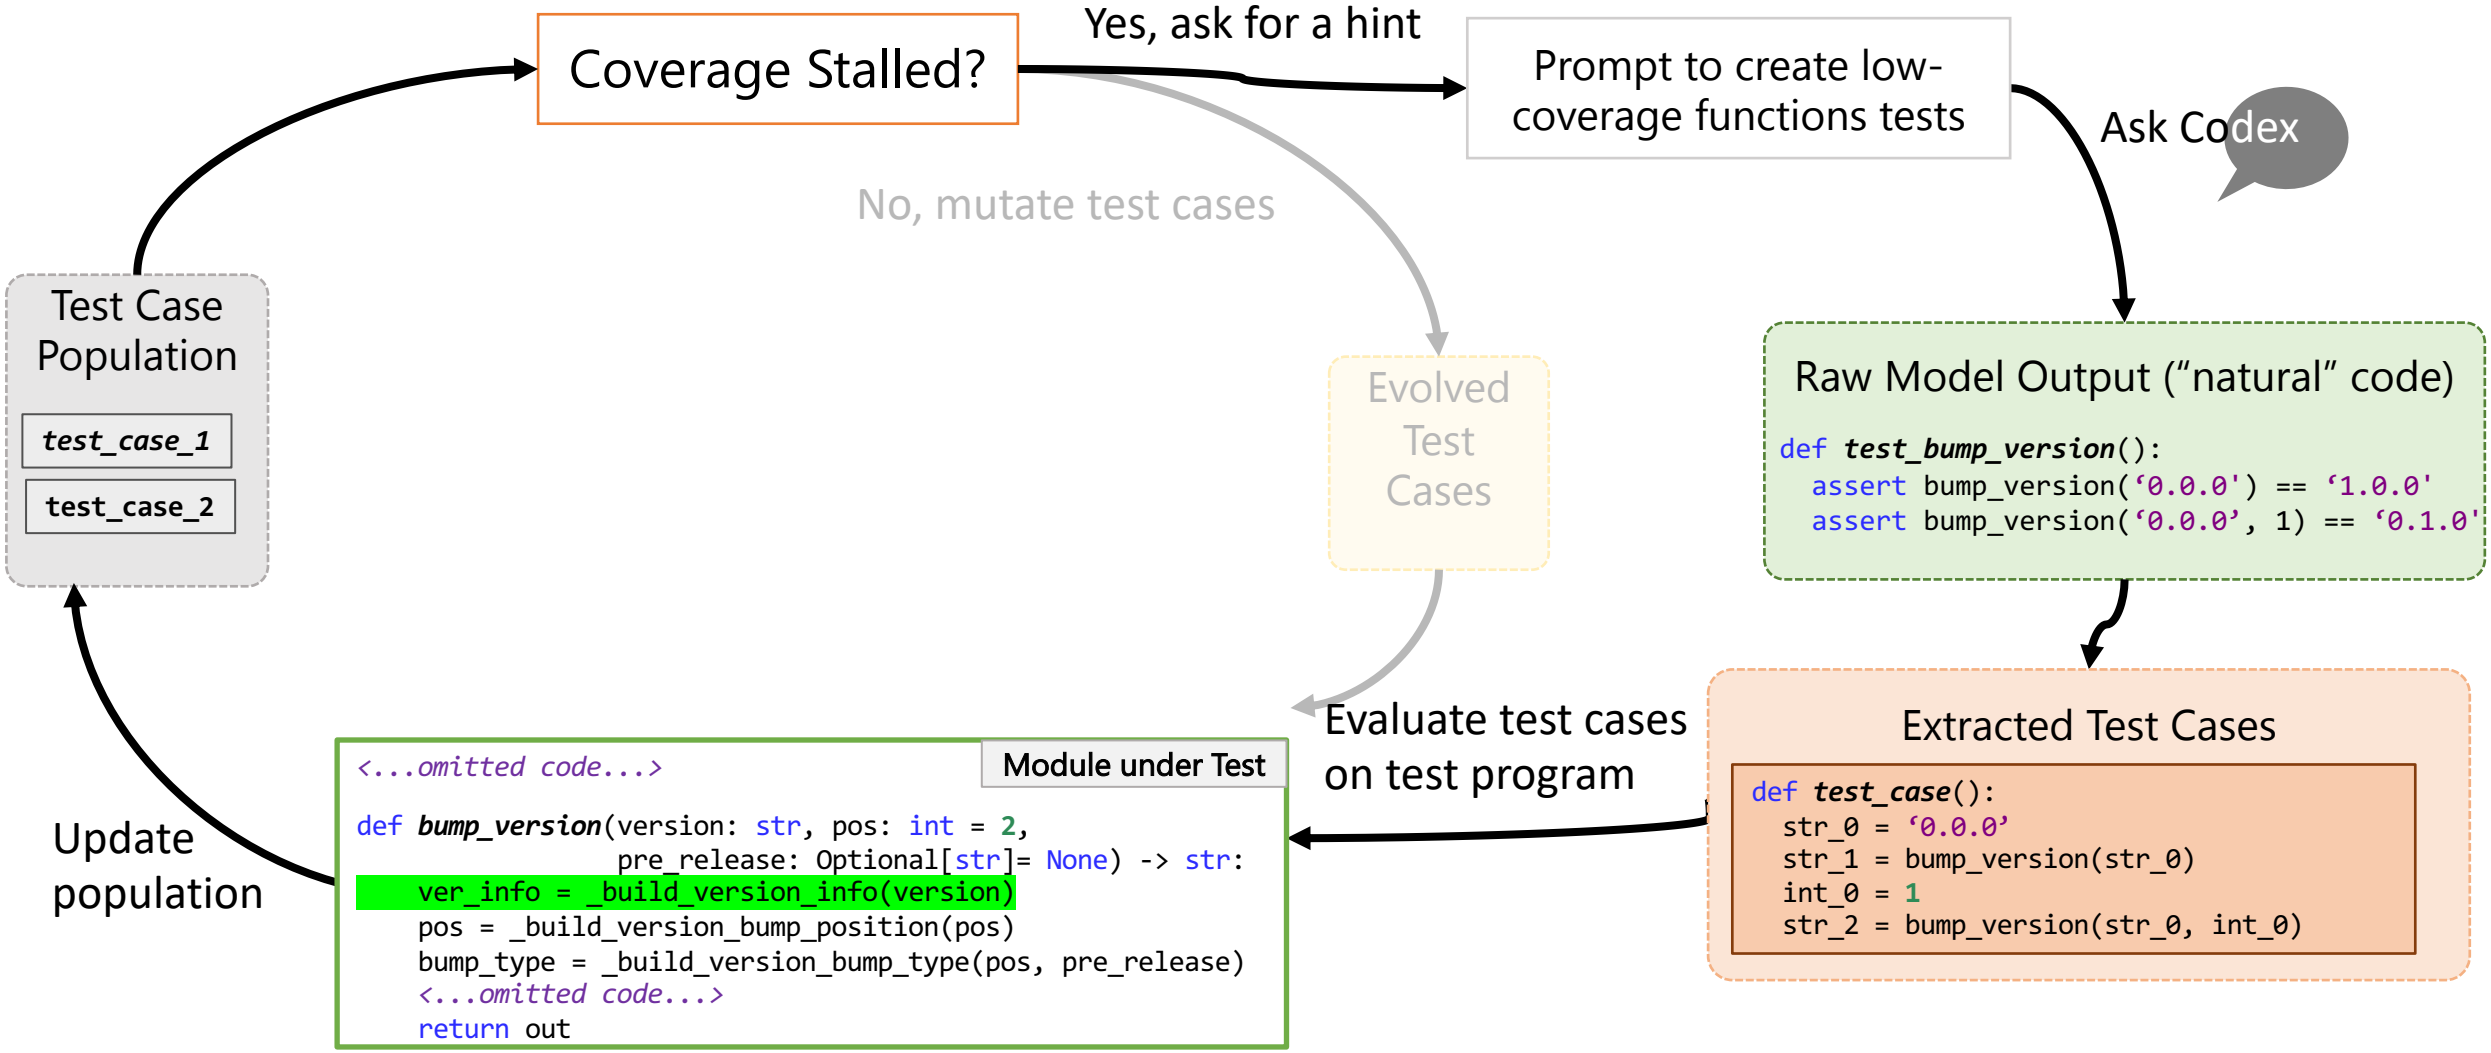

Suggested Test Case Increases Coverage

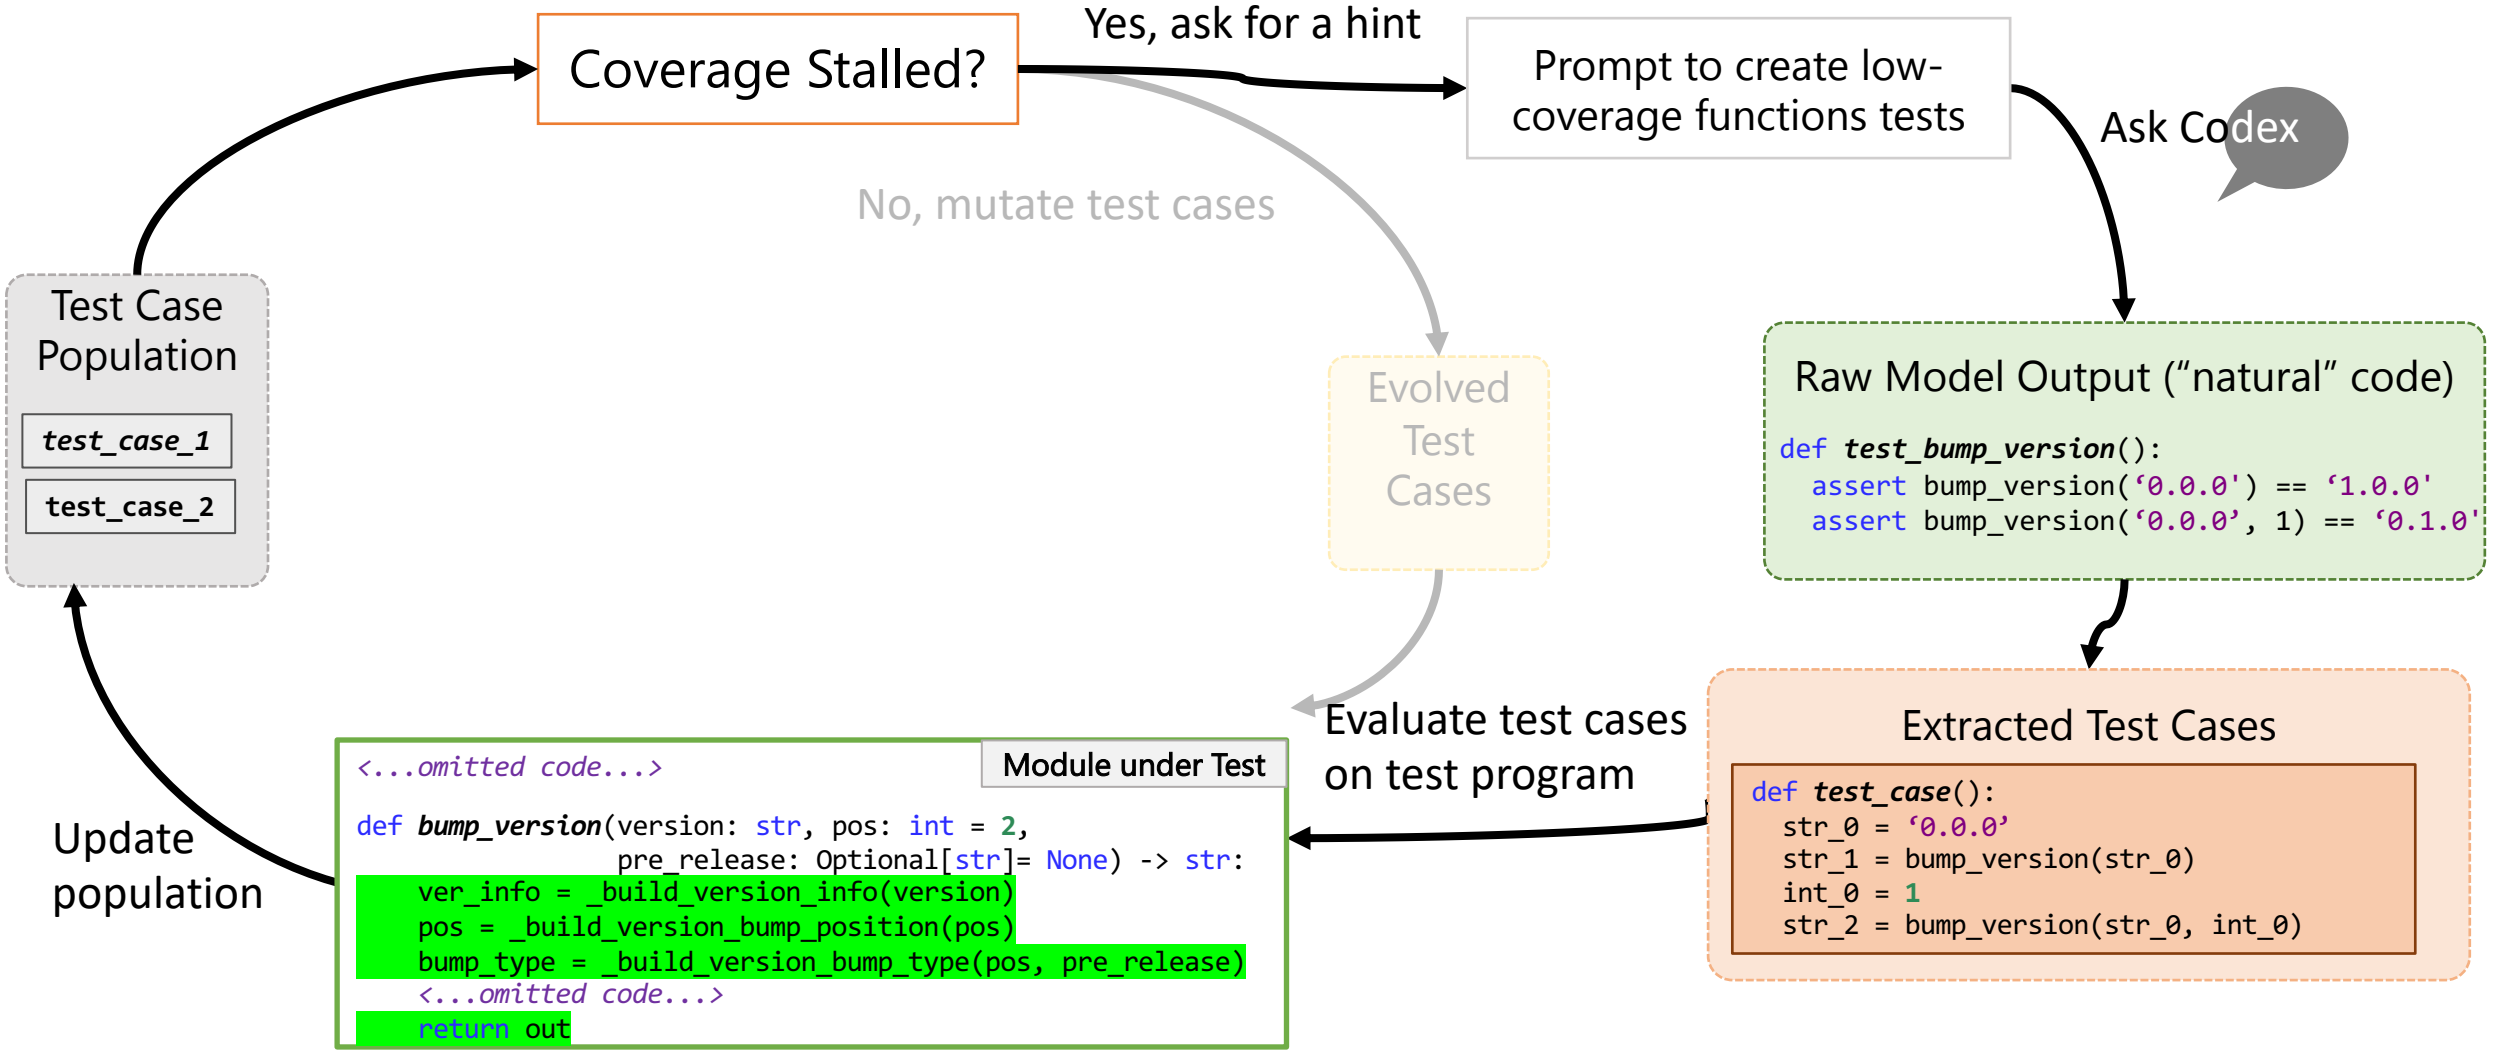

Suggested Test Case Increases Coverage

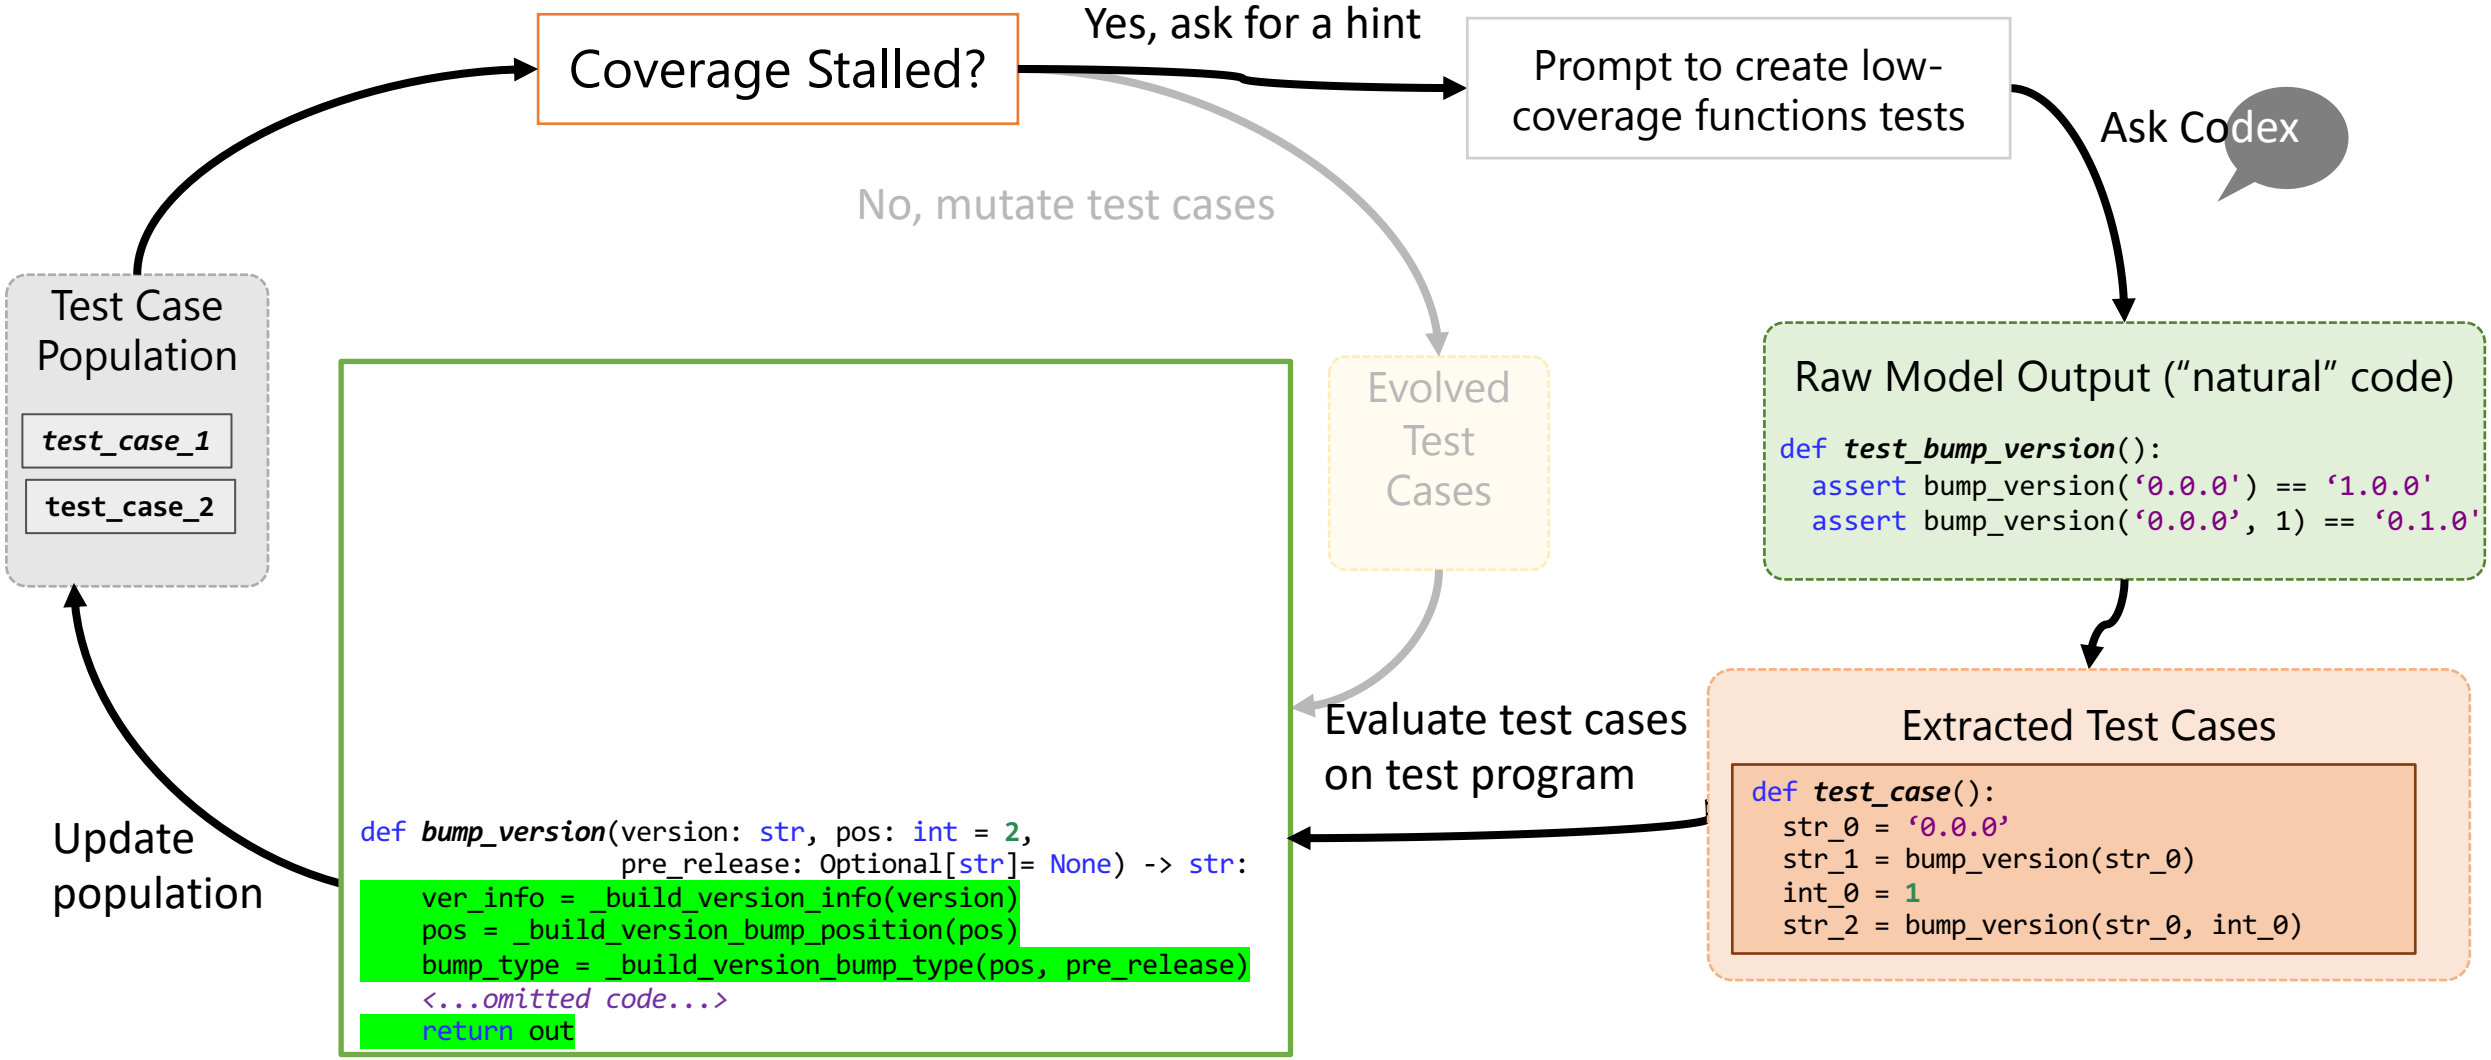

Suggested Test Case Increases Coverage

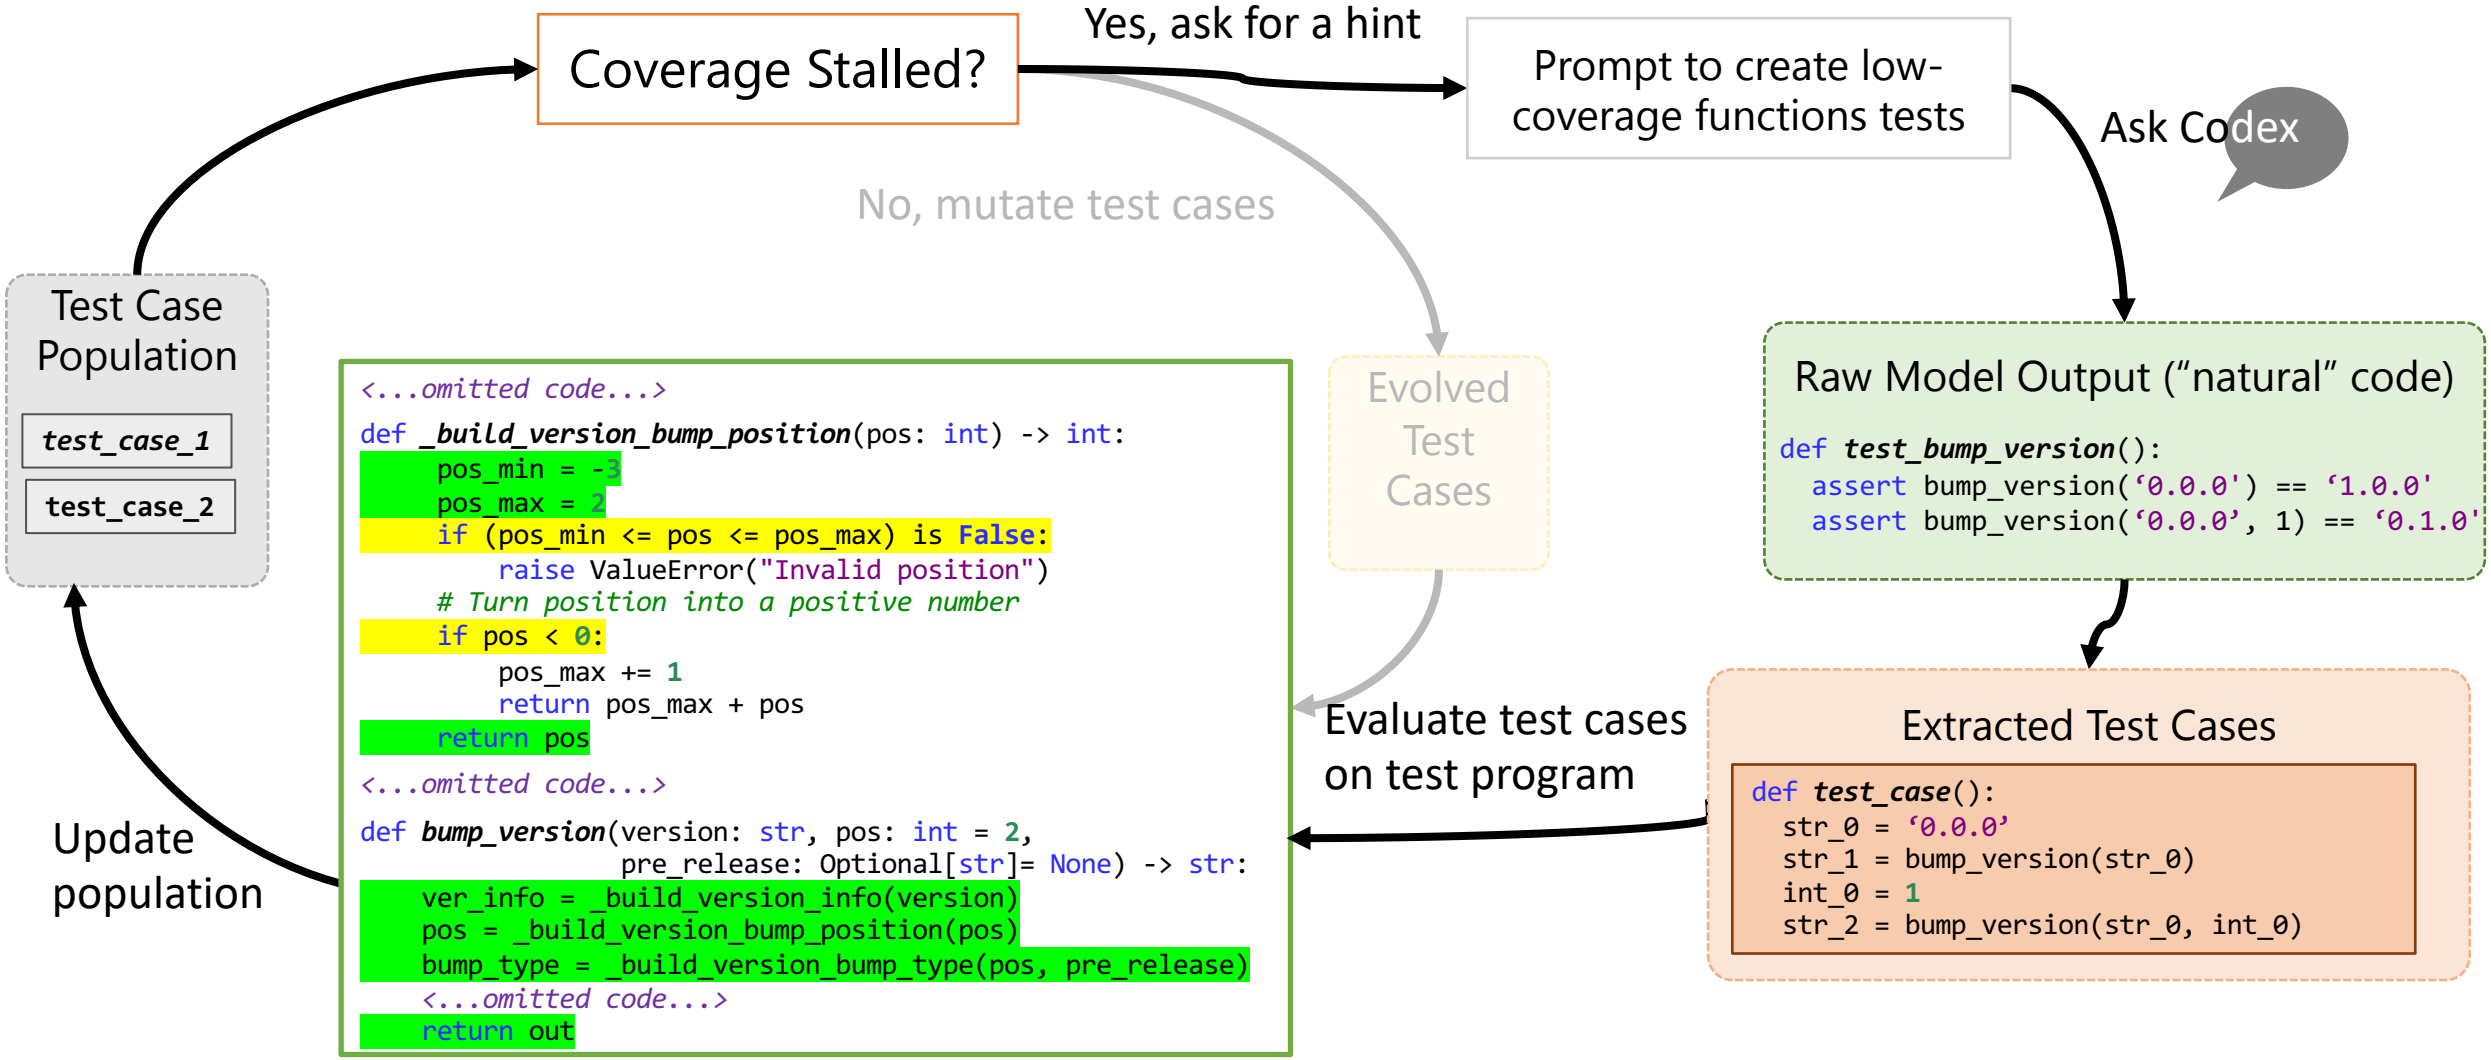

Update Population

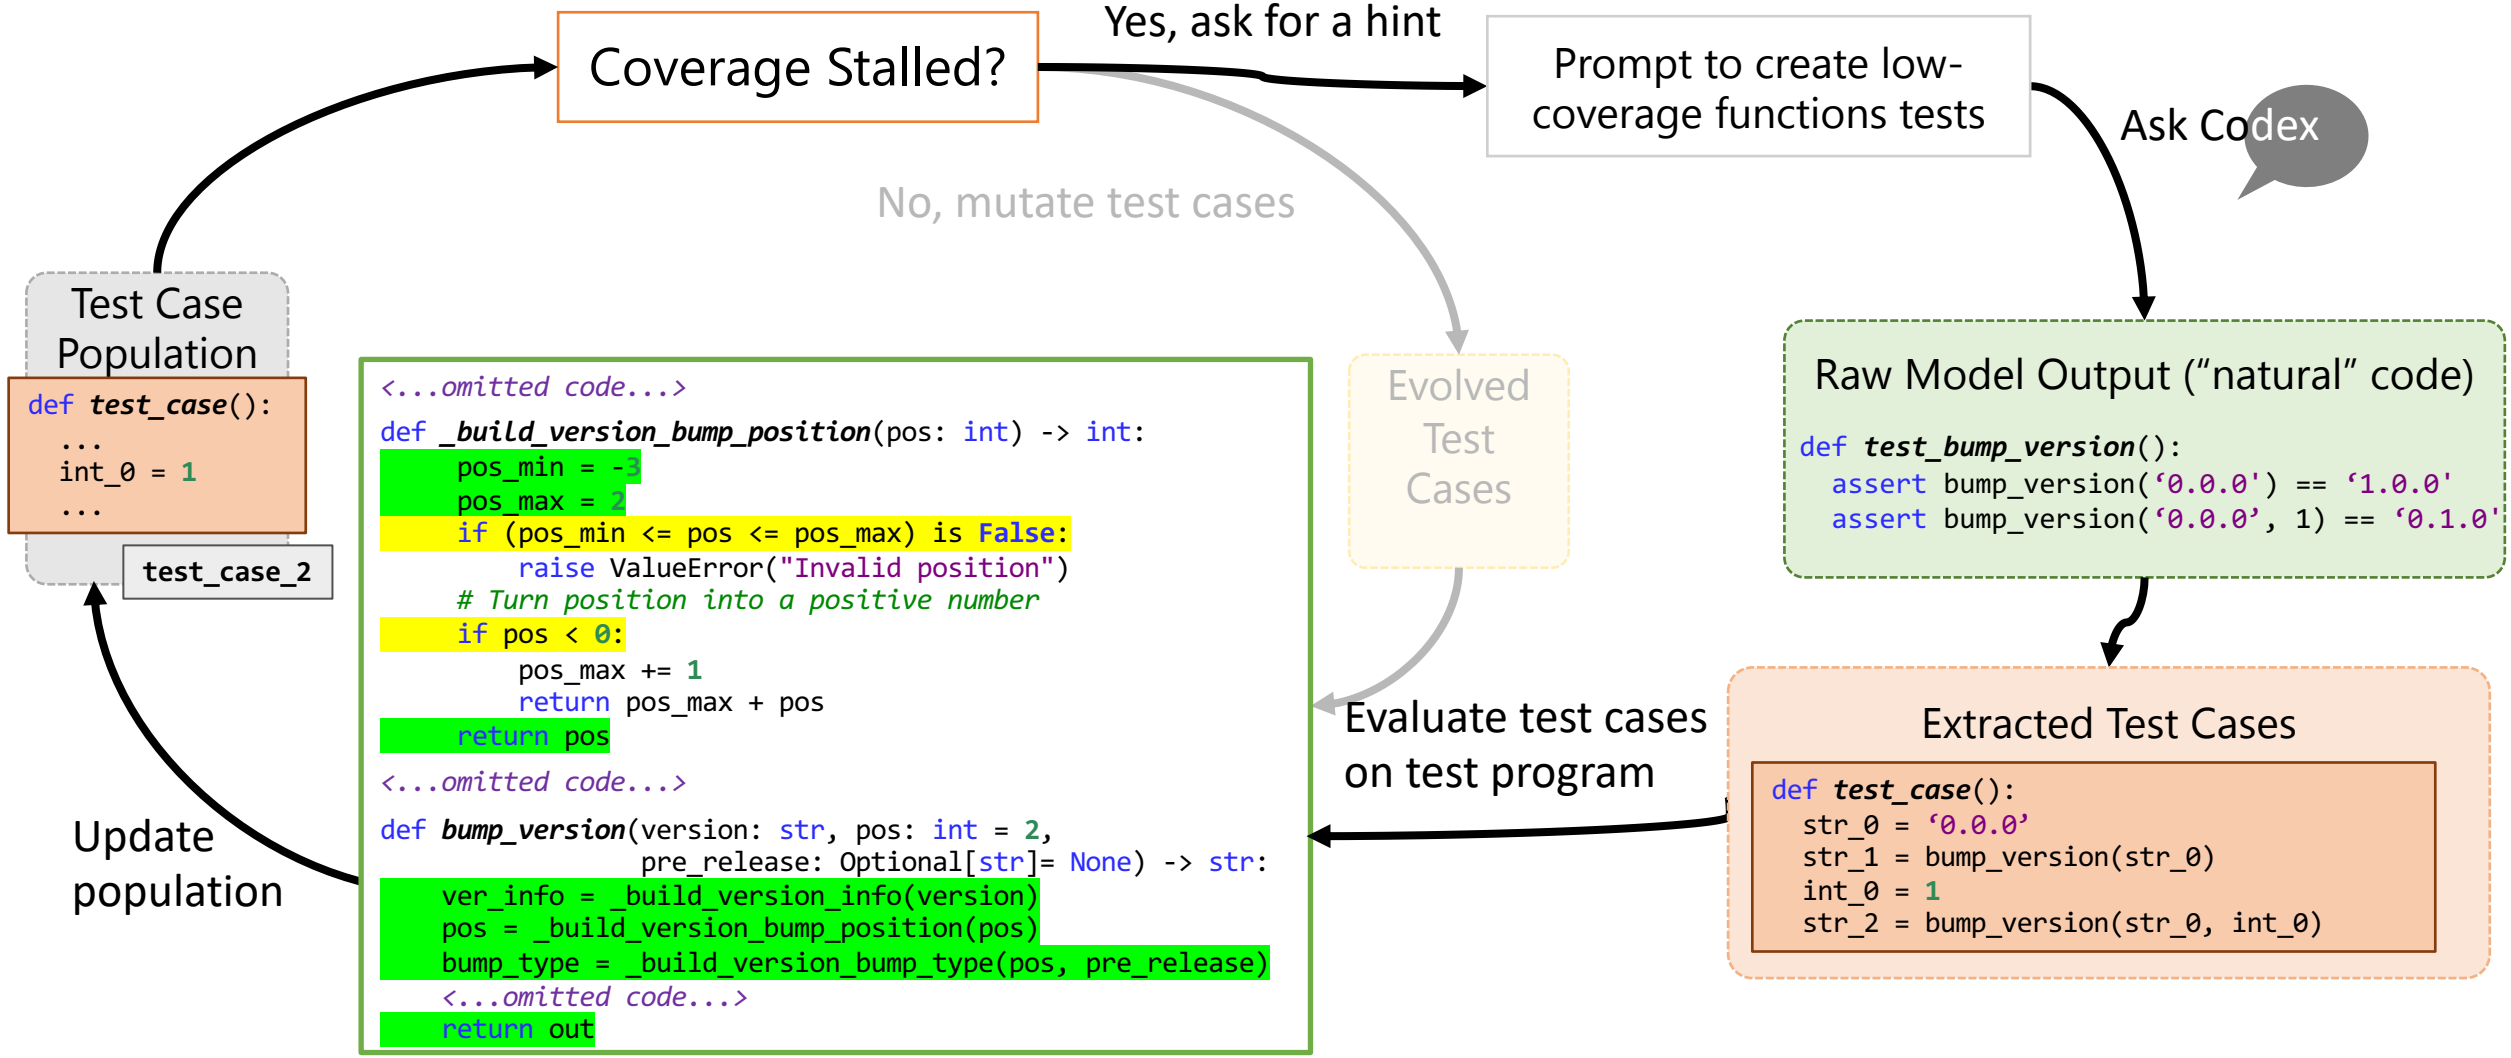

Search No Longer Stalled

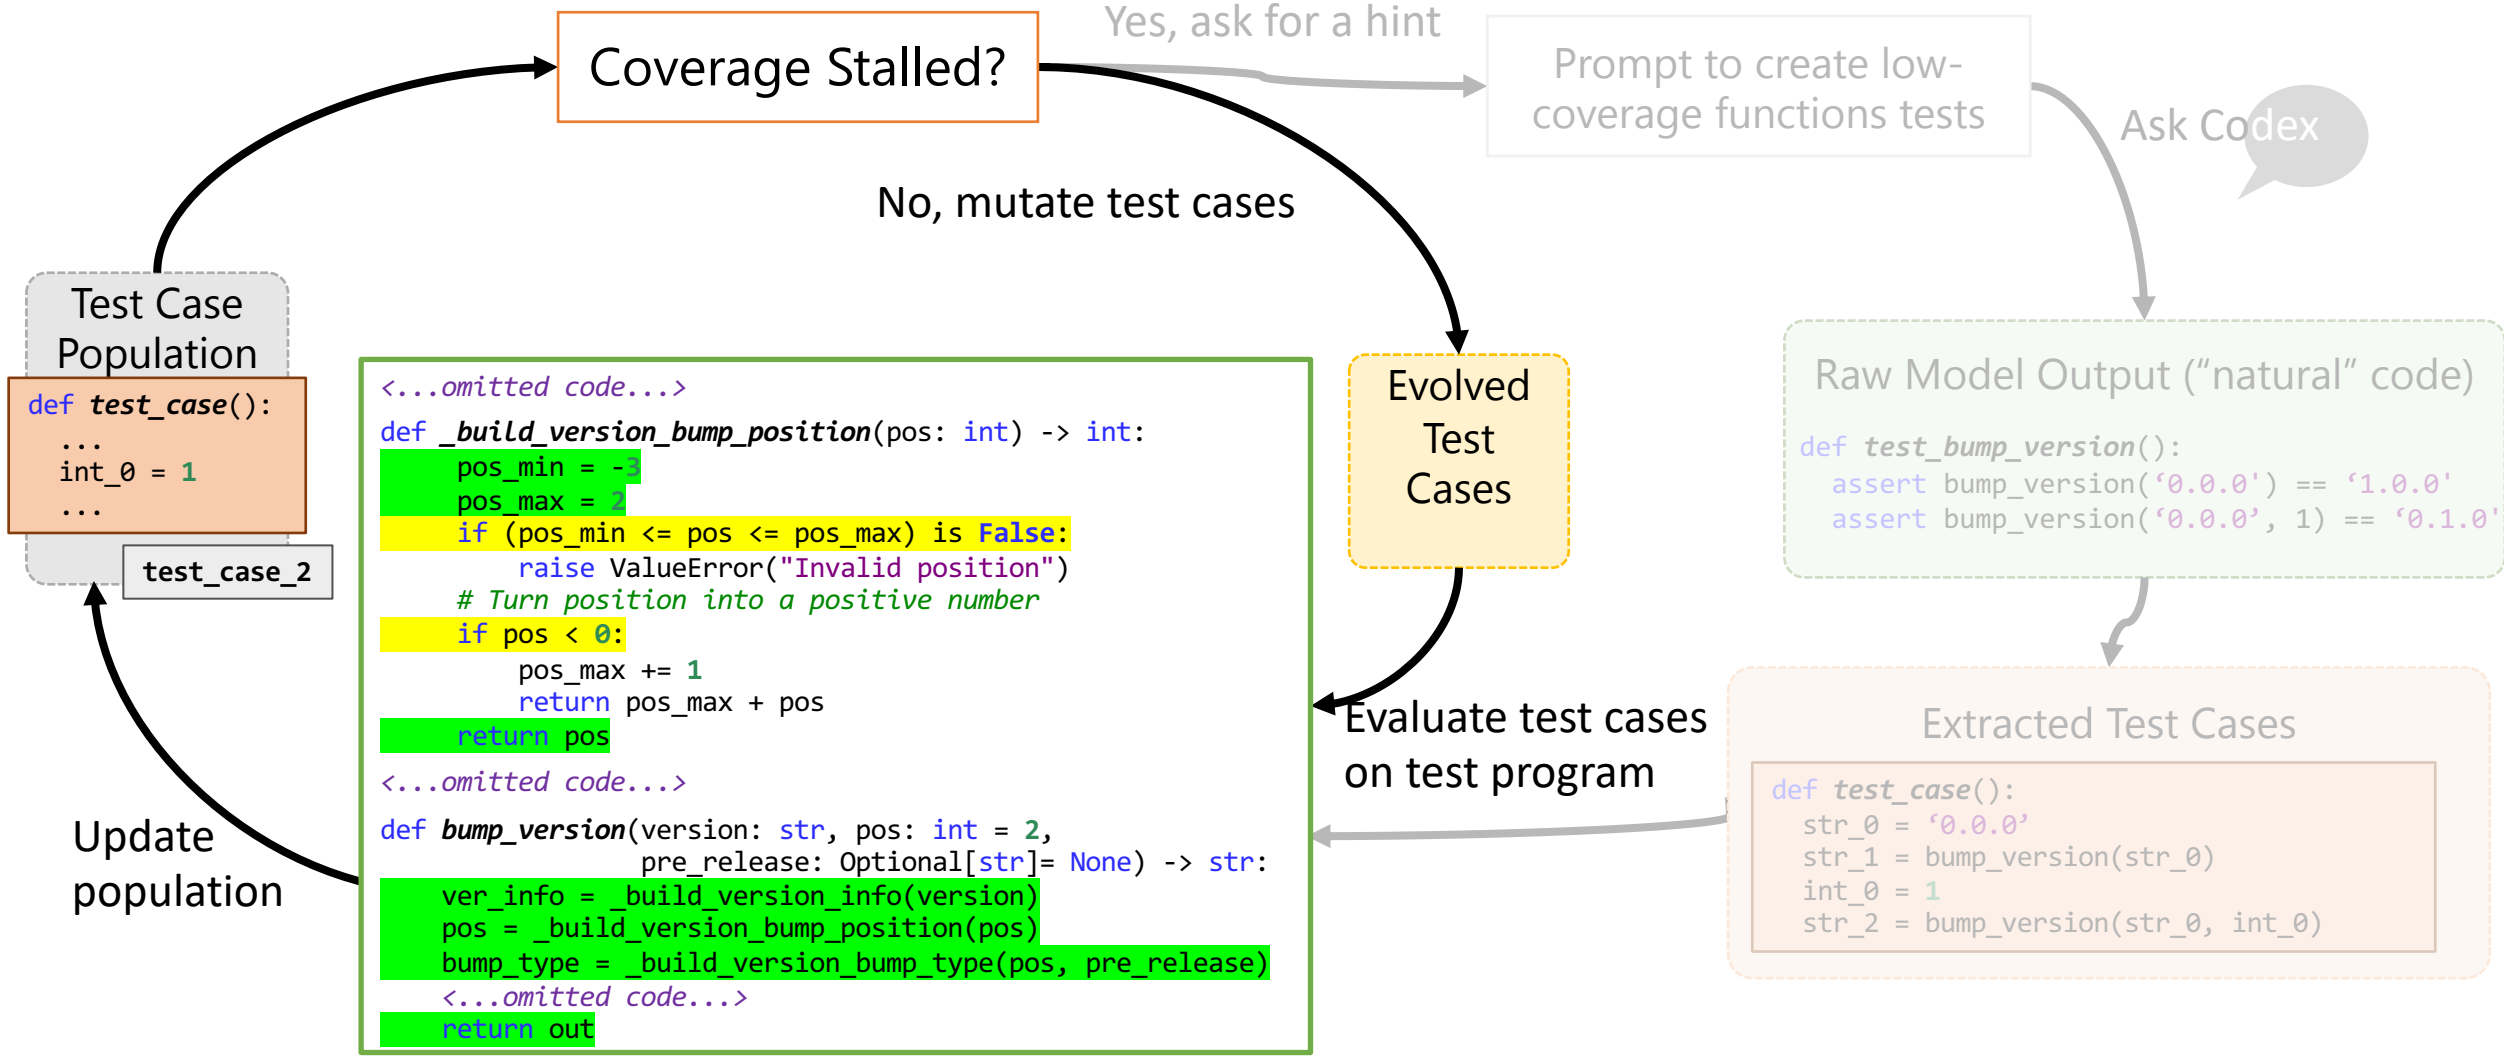

Search No Longer Stalled

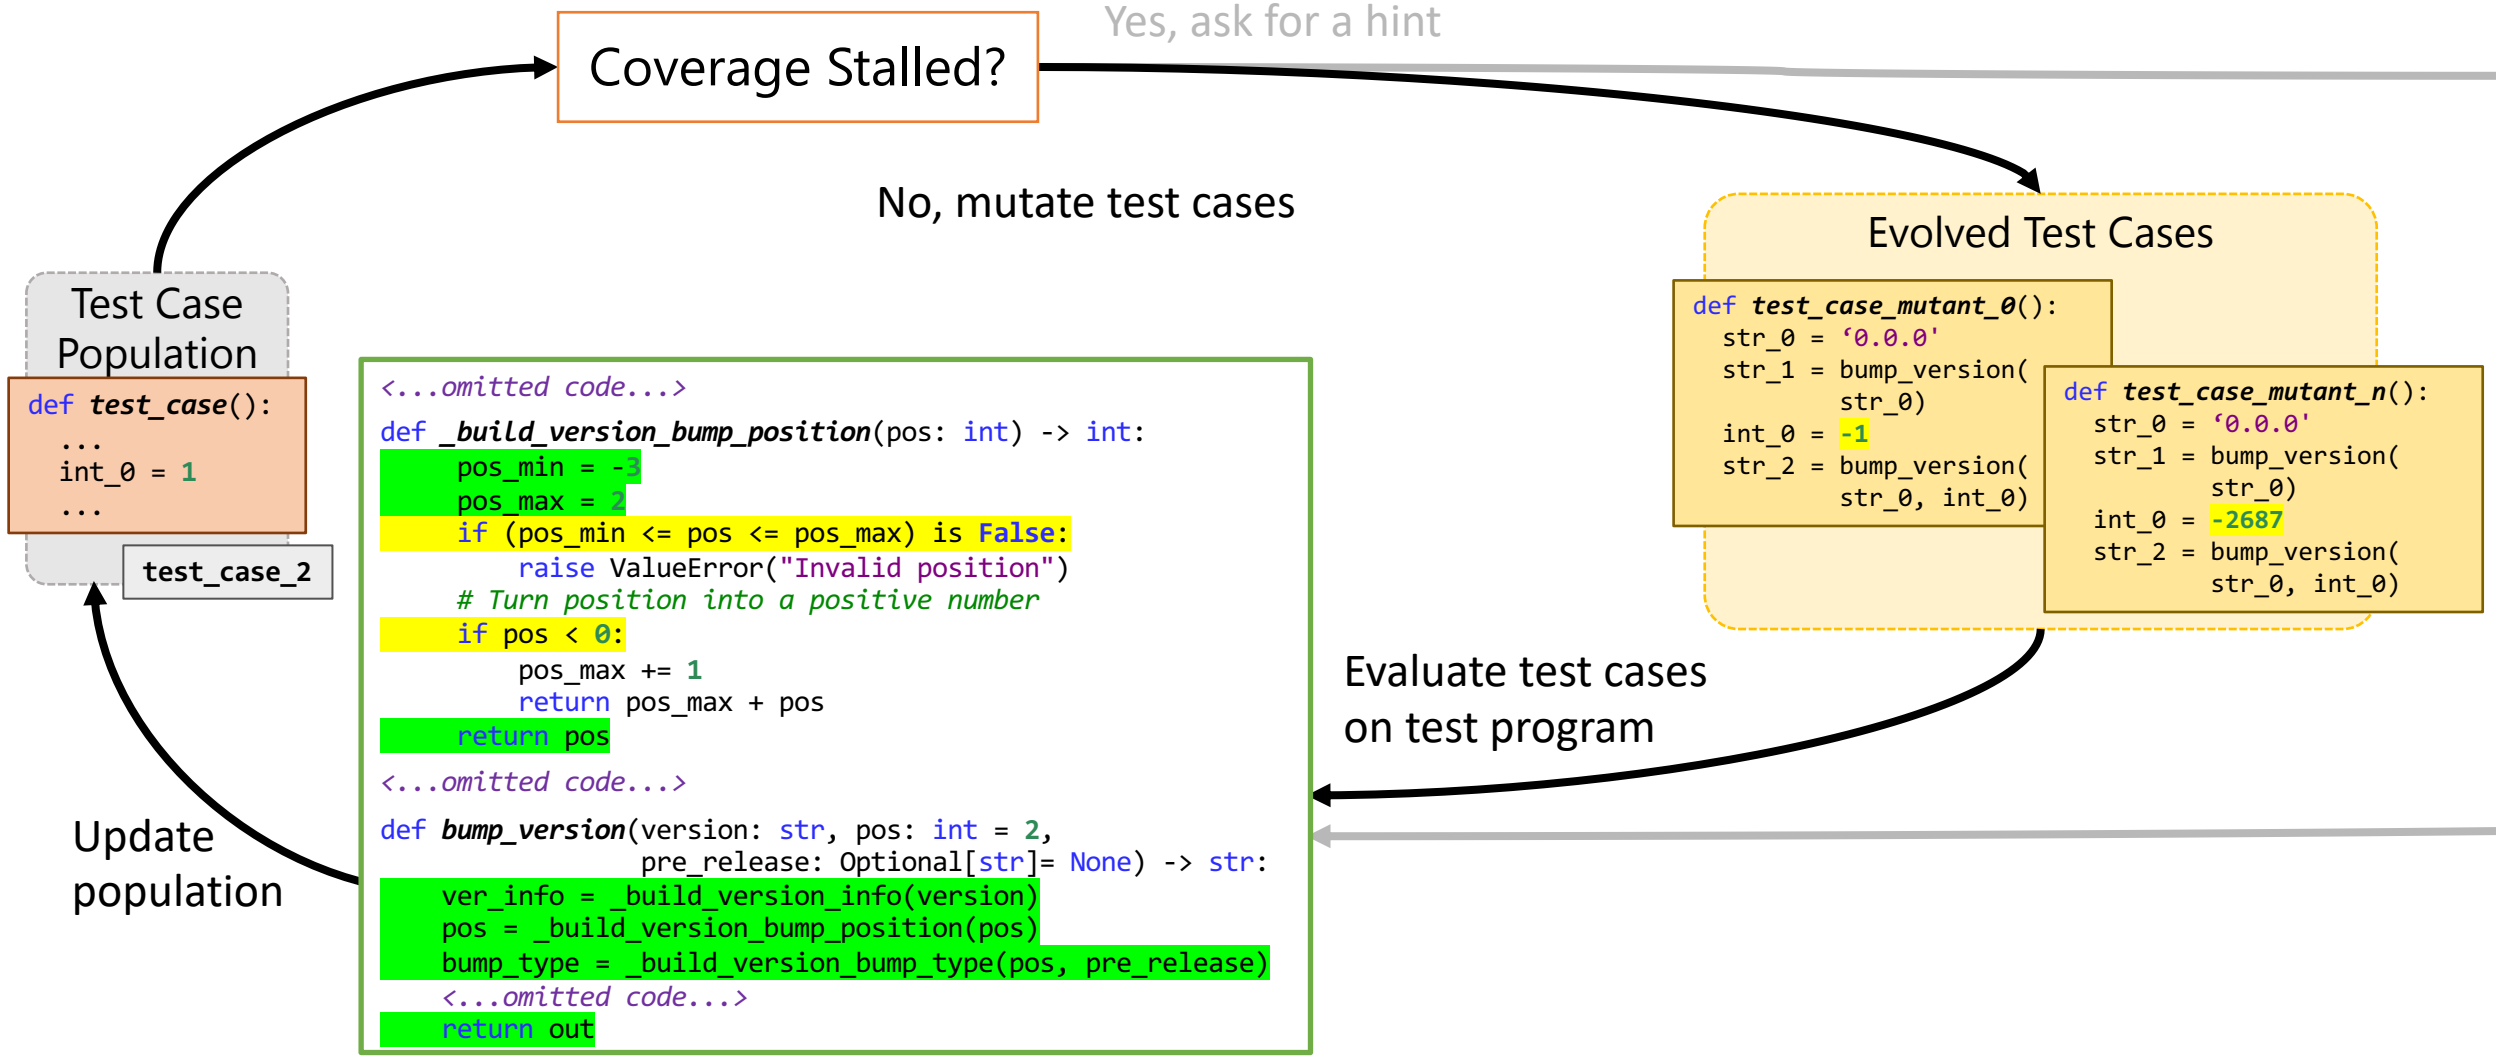

Search No Longer Stalled

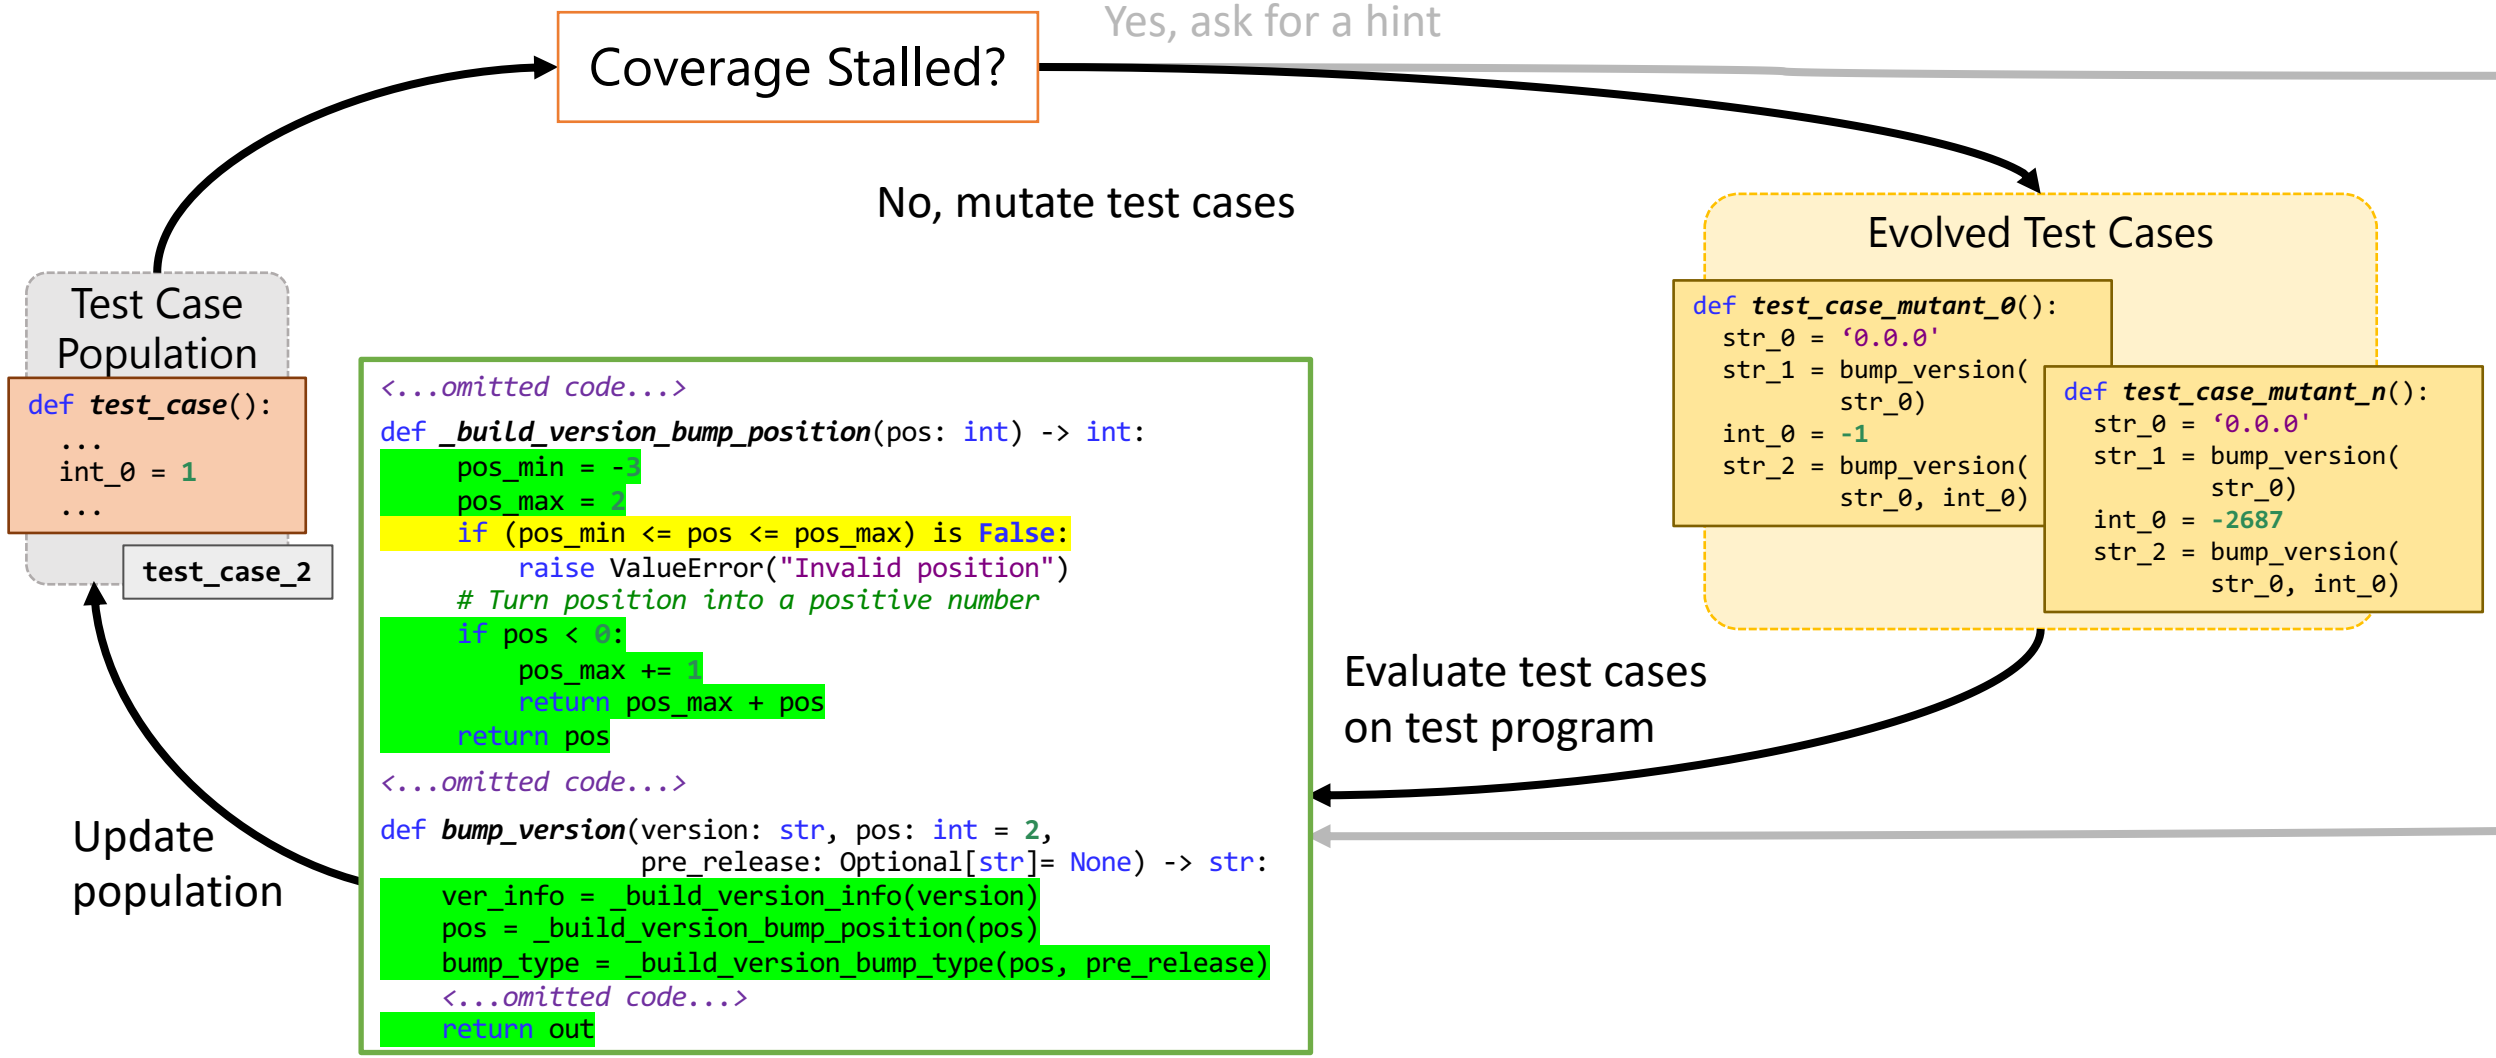

Search No Longer Stalled

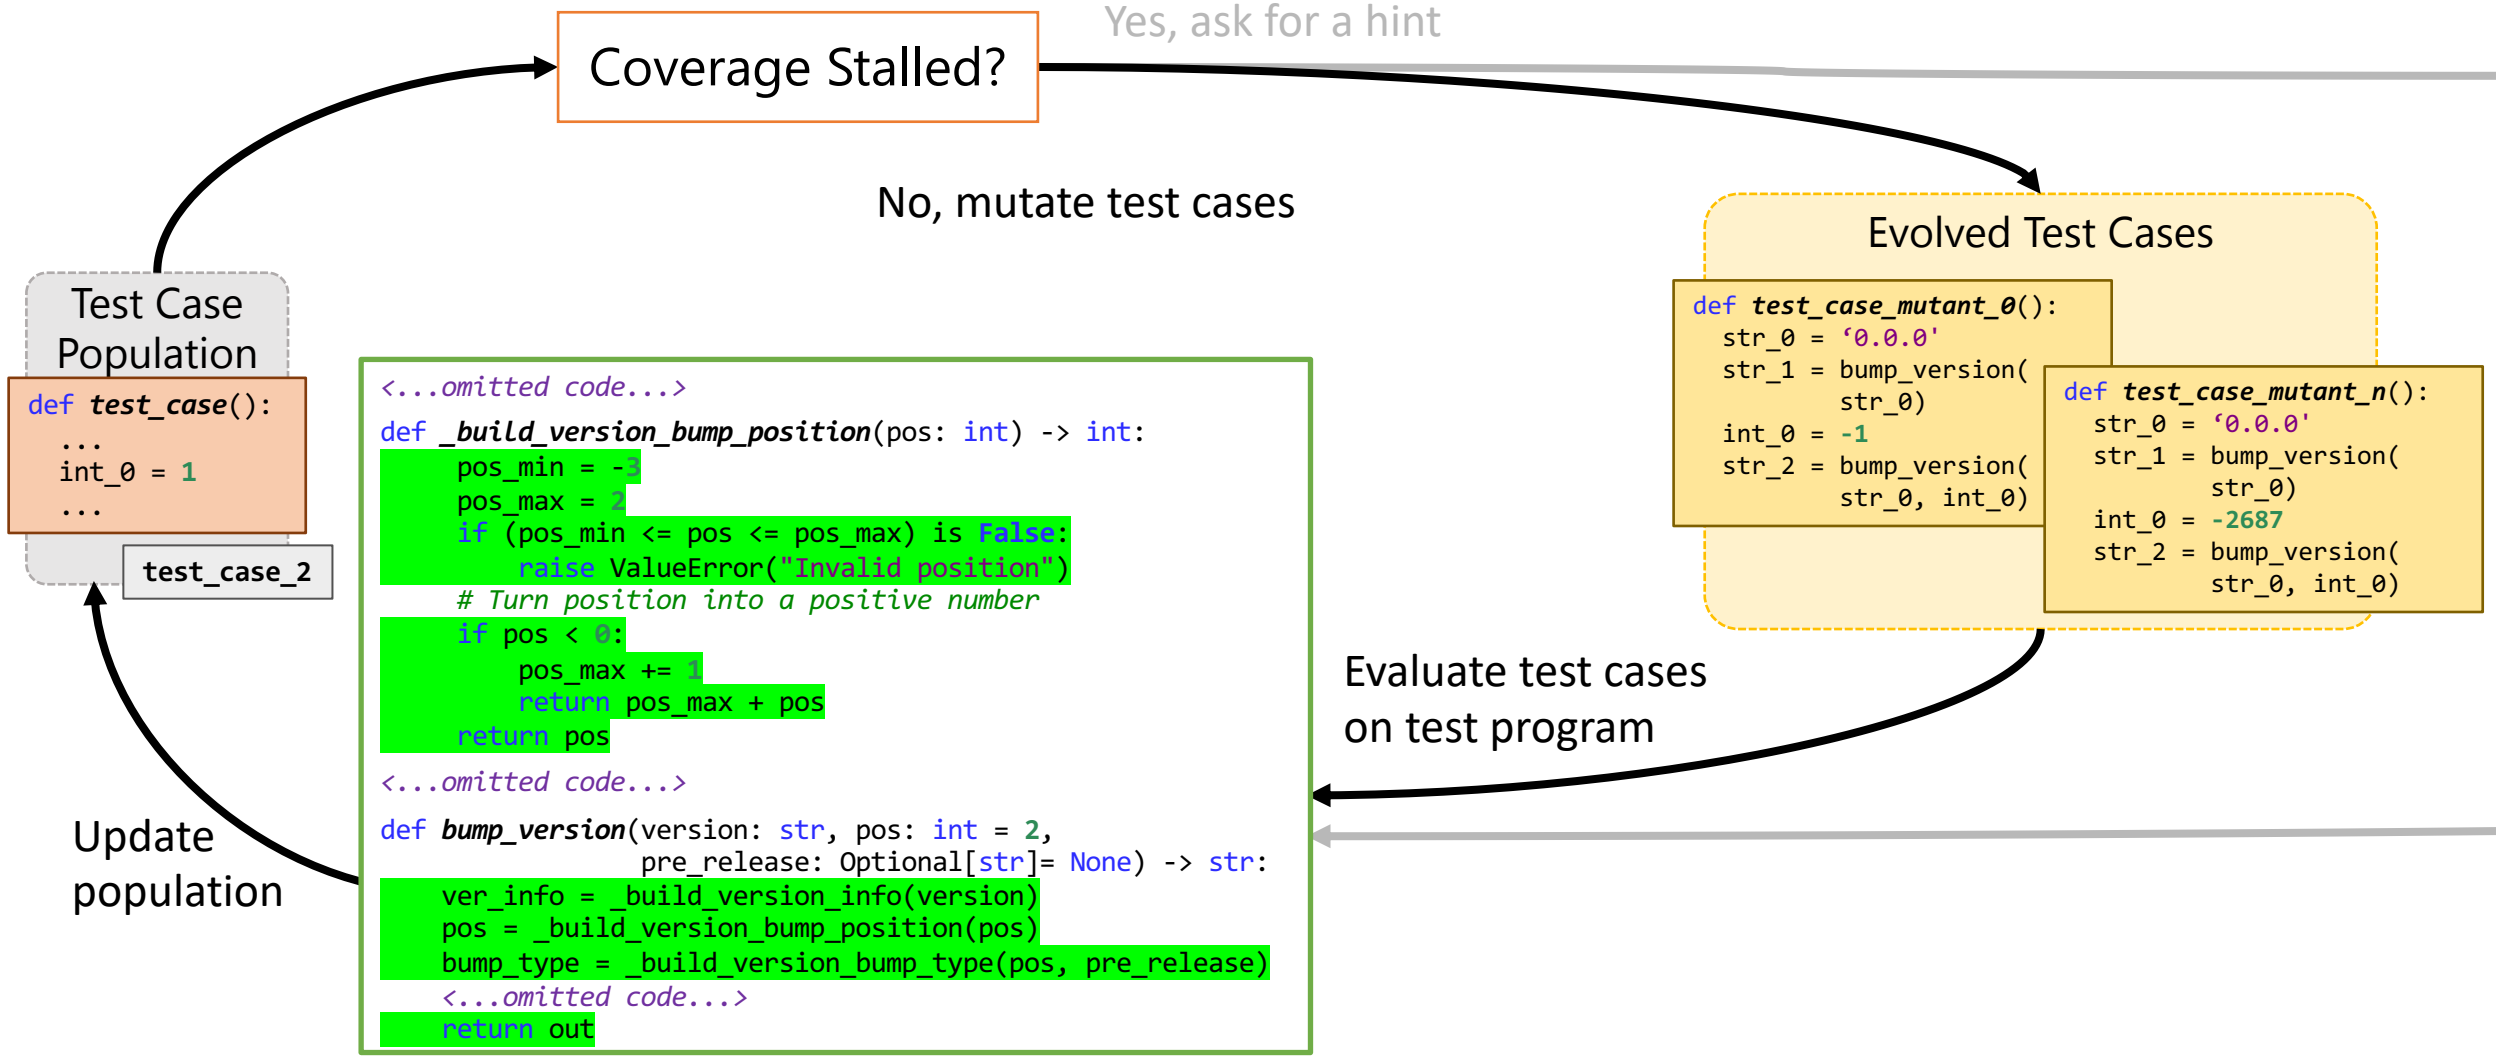

Search No Longer Stalled

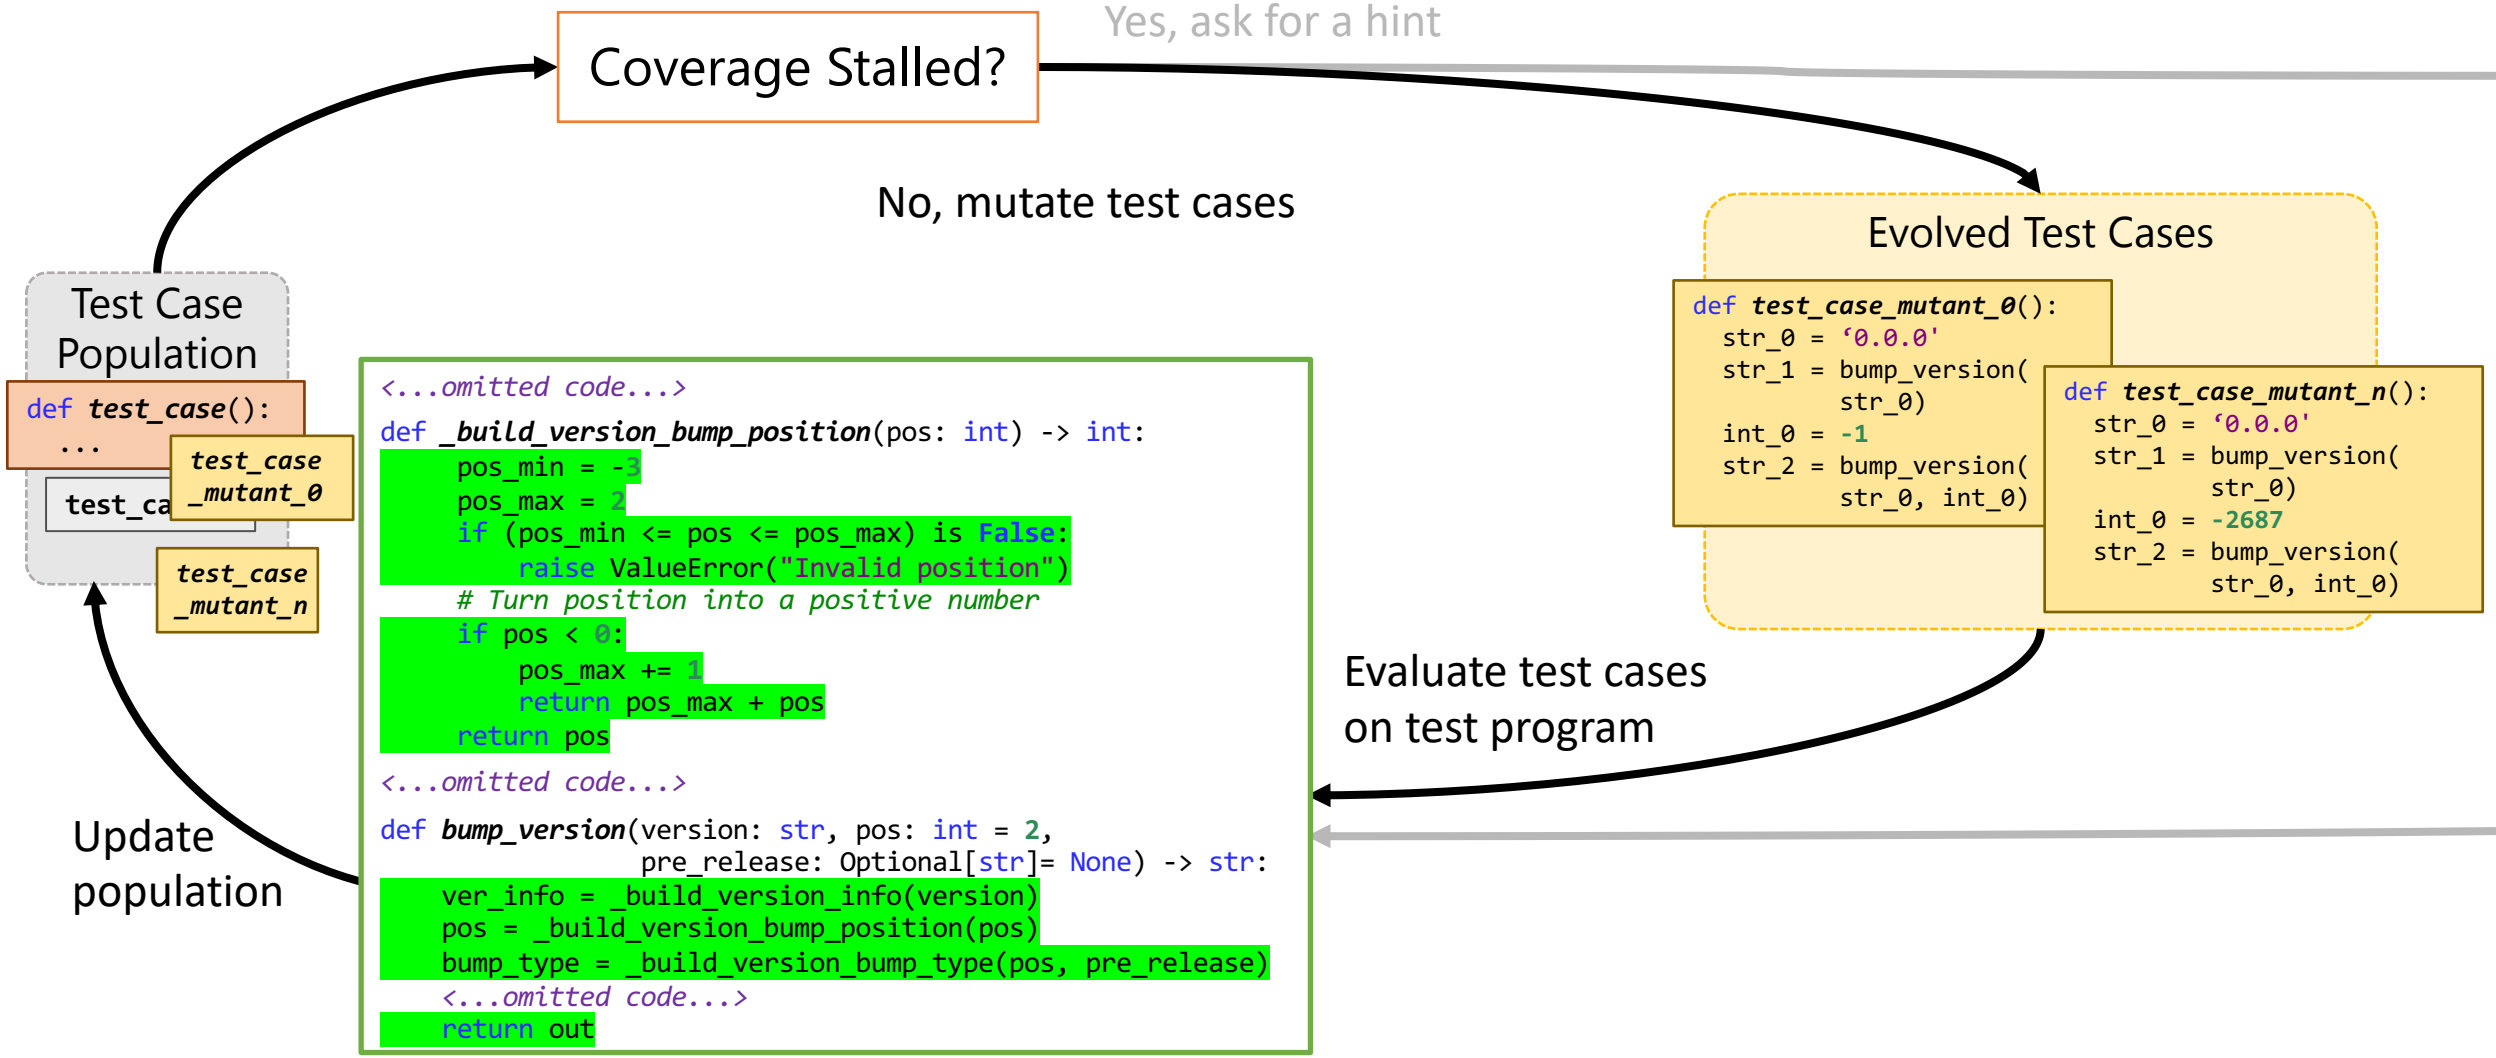

Spoiler: Results on this Benchmark

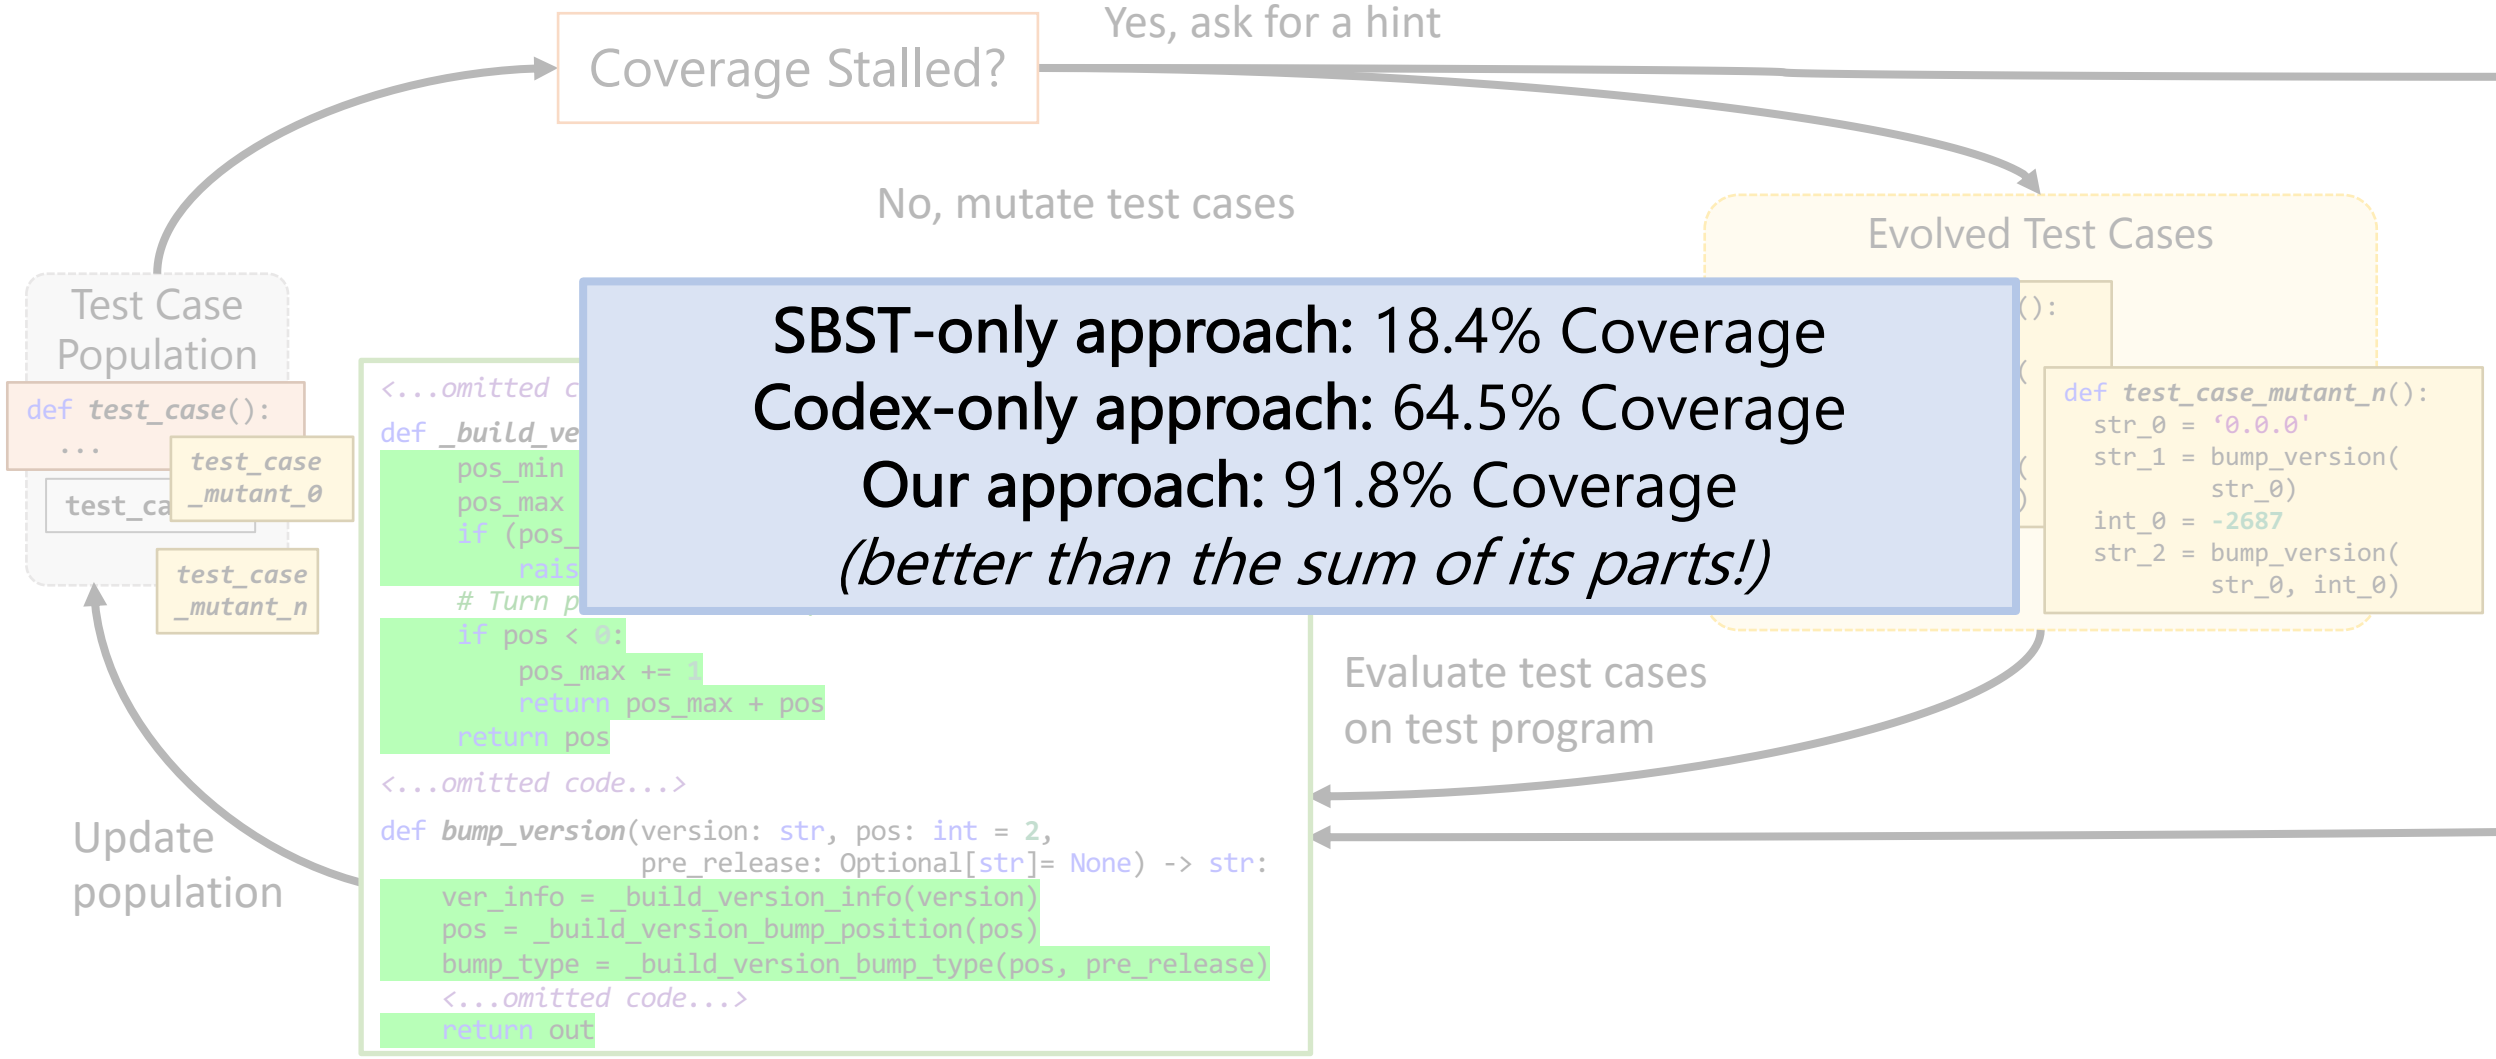

CodaMOSA Approach

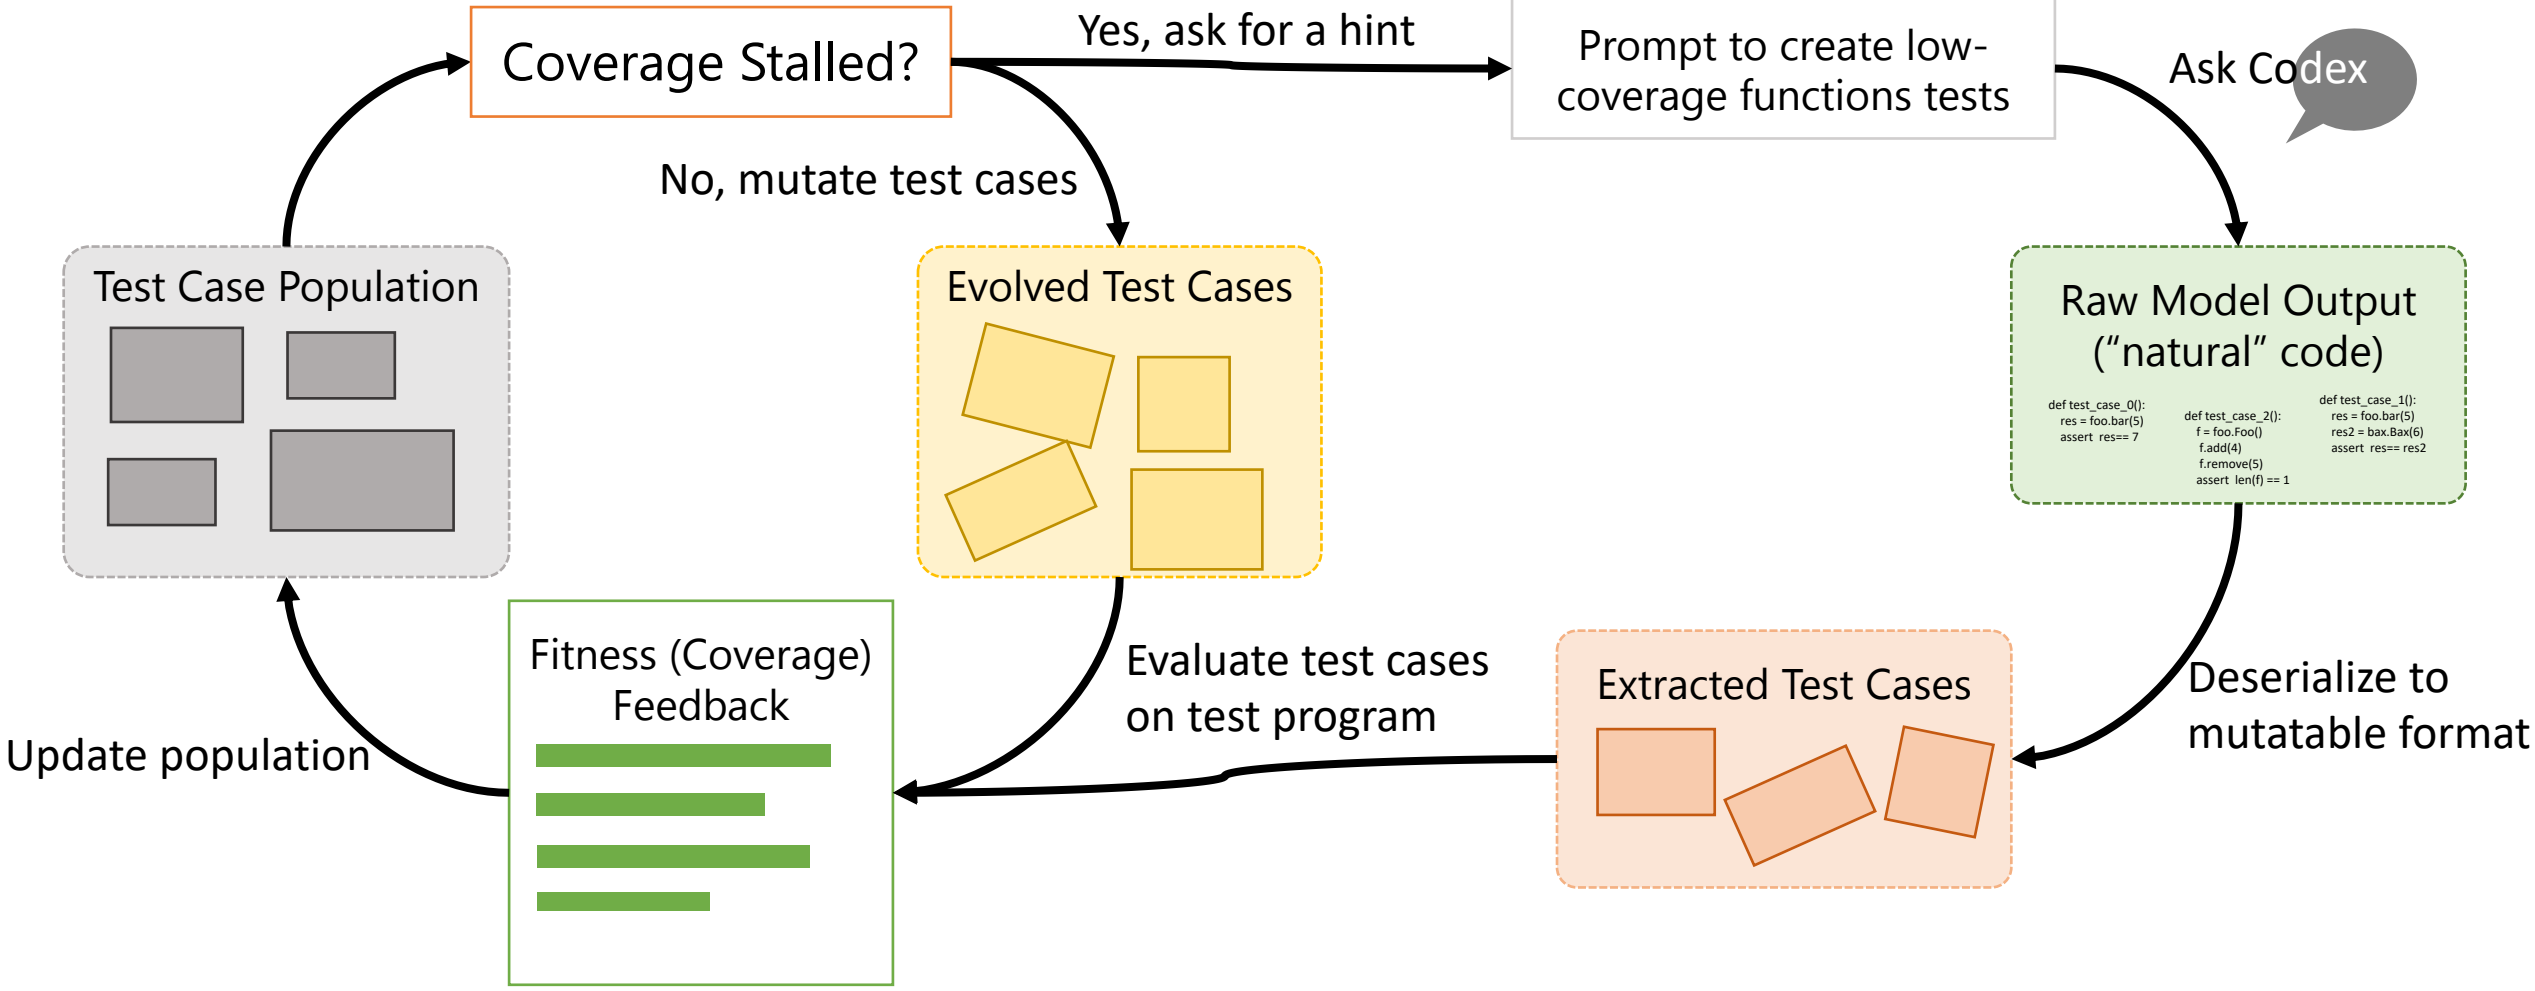

Challenge: When to ask for a hint?

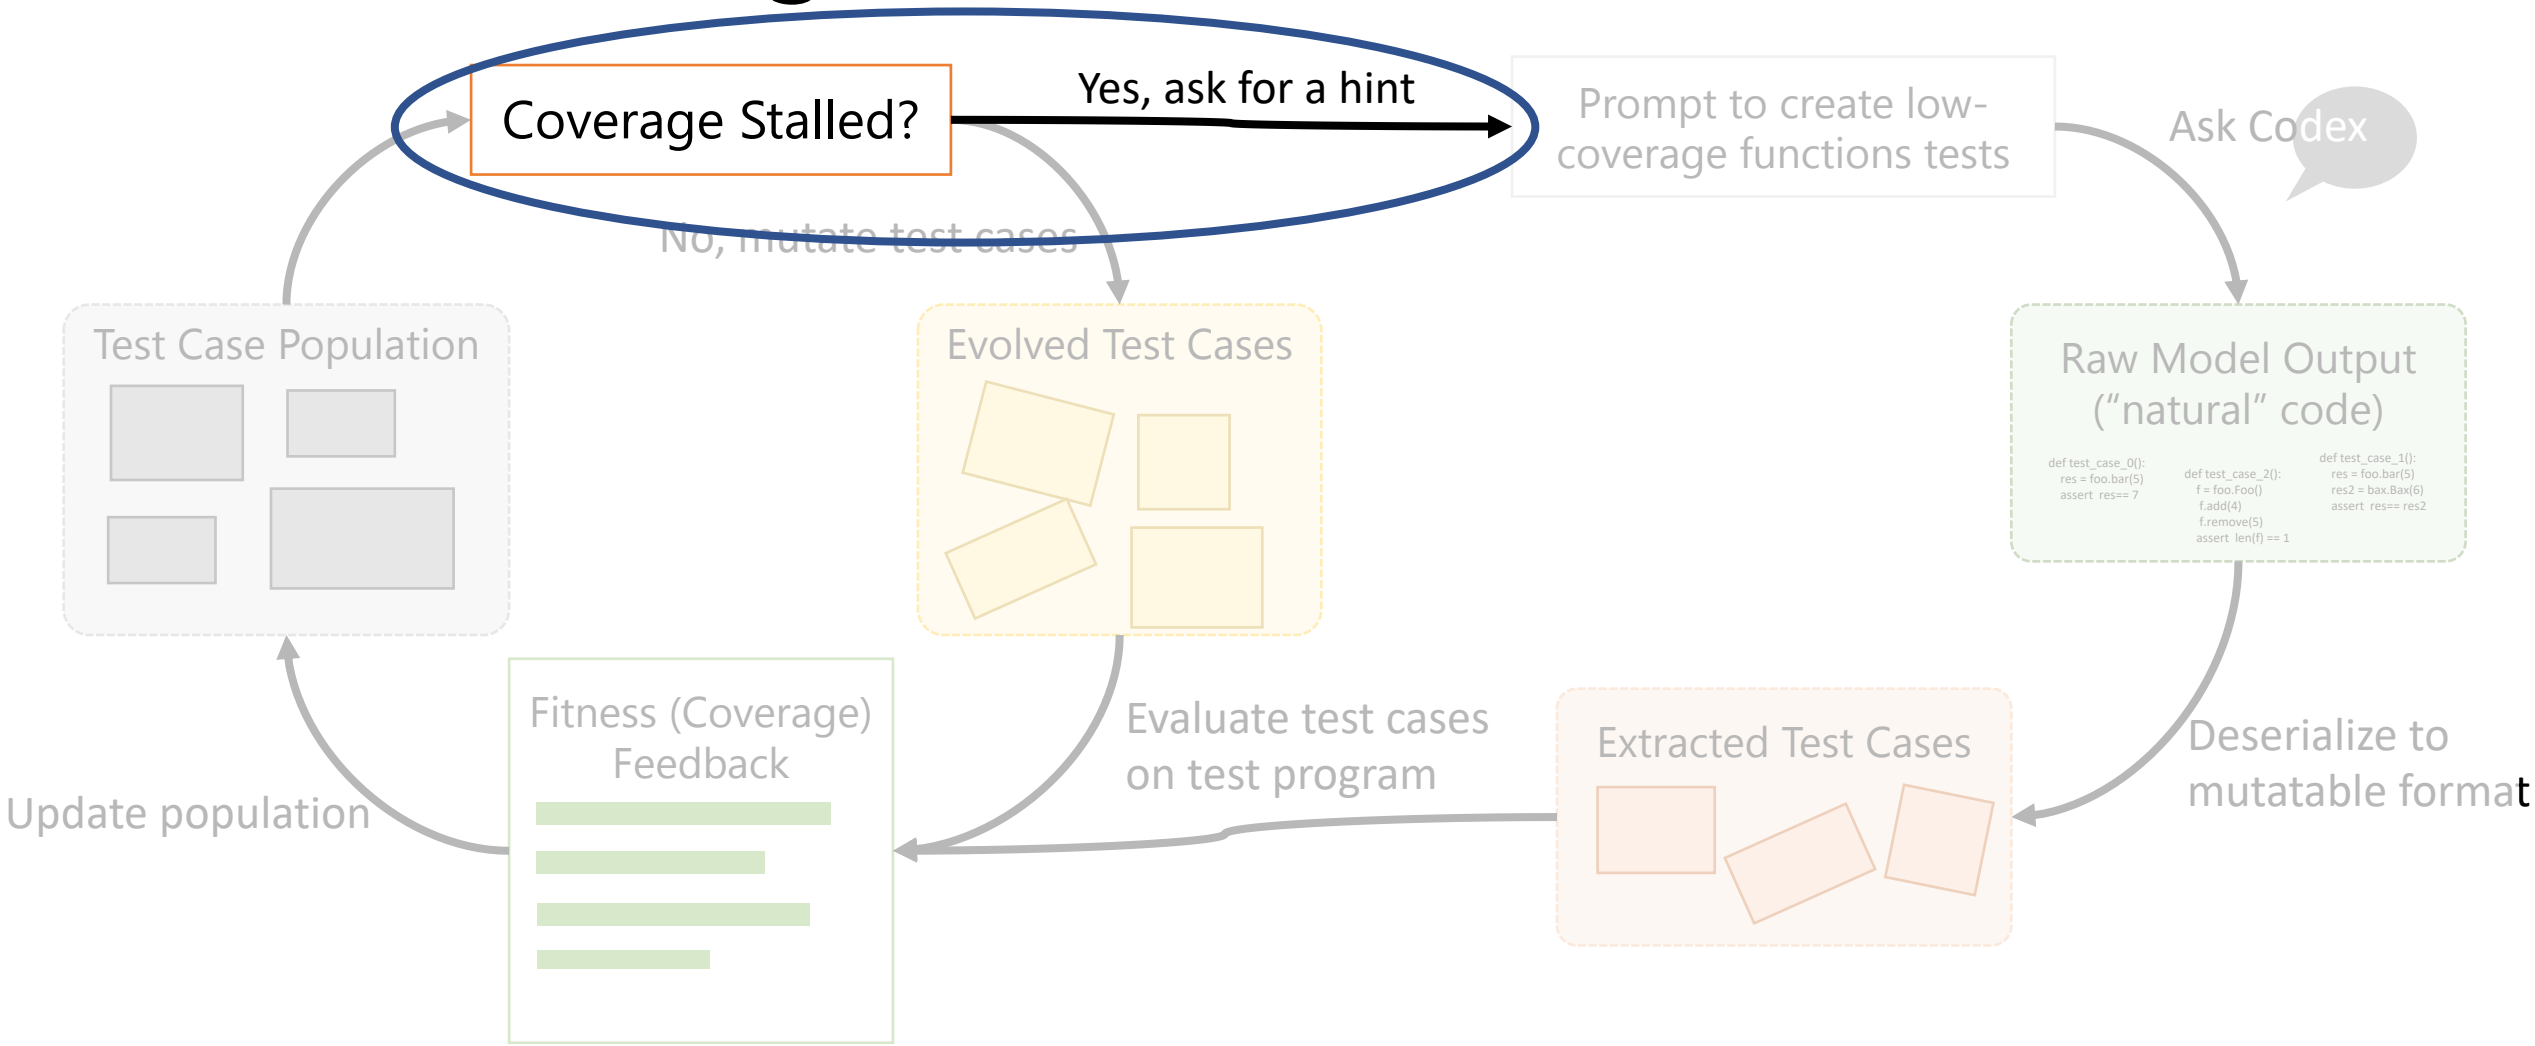

Challenge: How to ask for a hint?

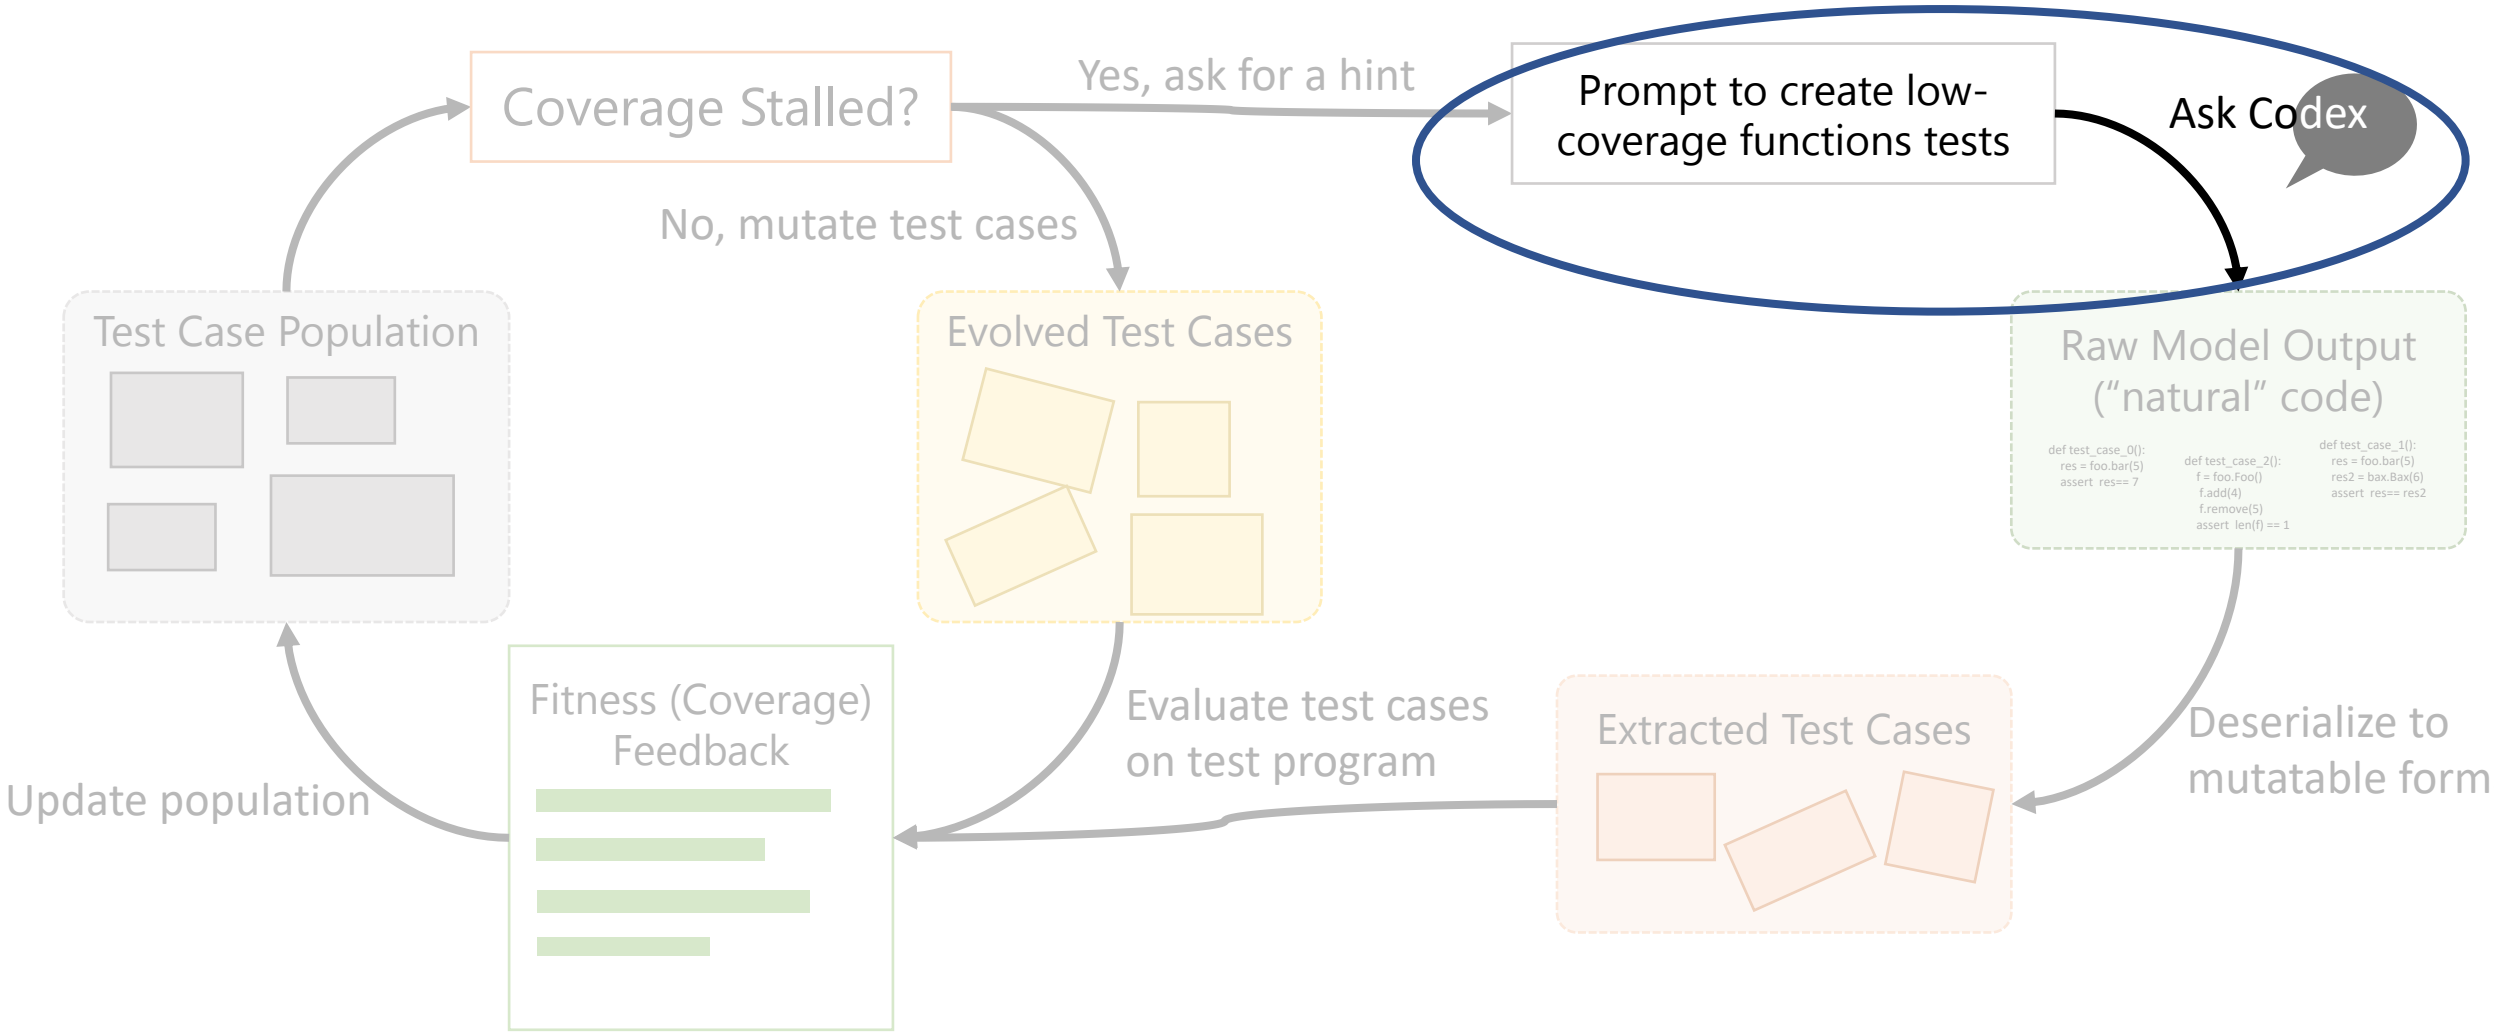

Challenge: How to handle (potentially) arbitrary output from Codex?

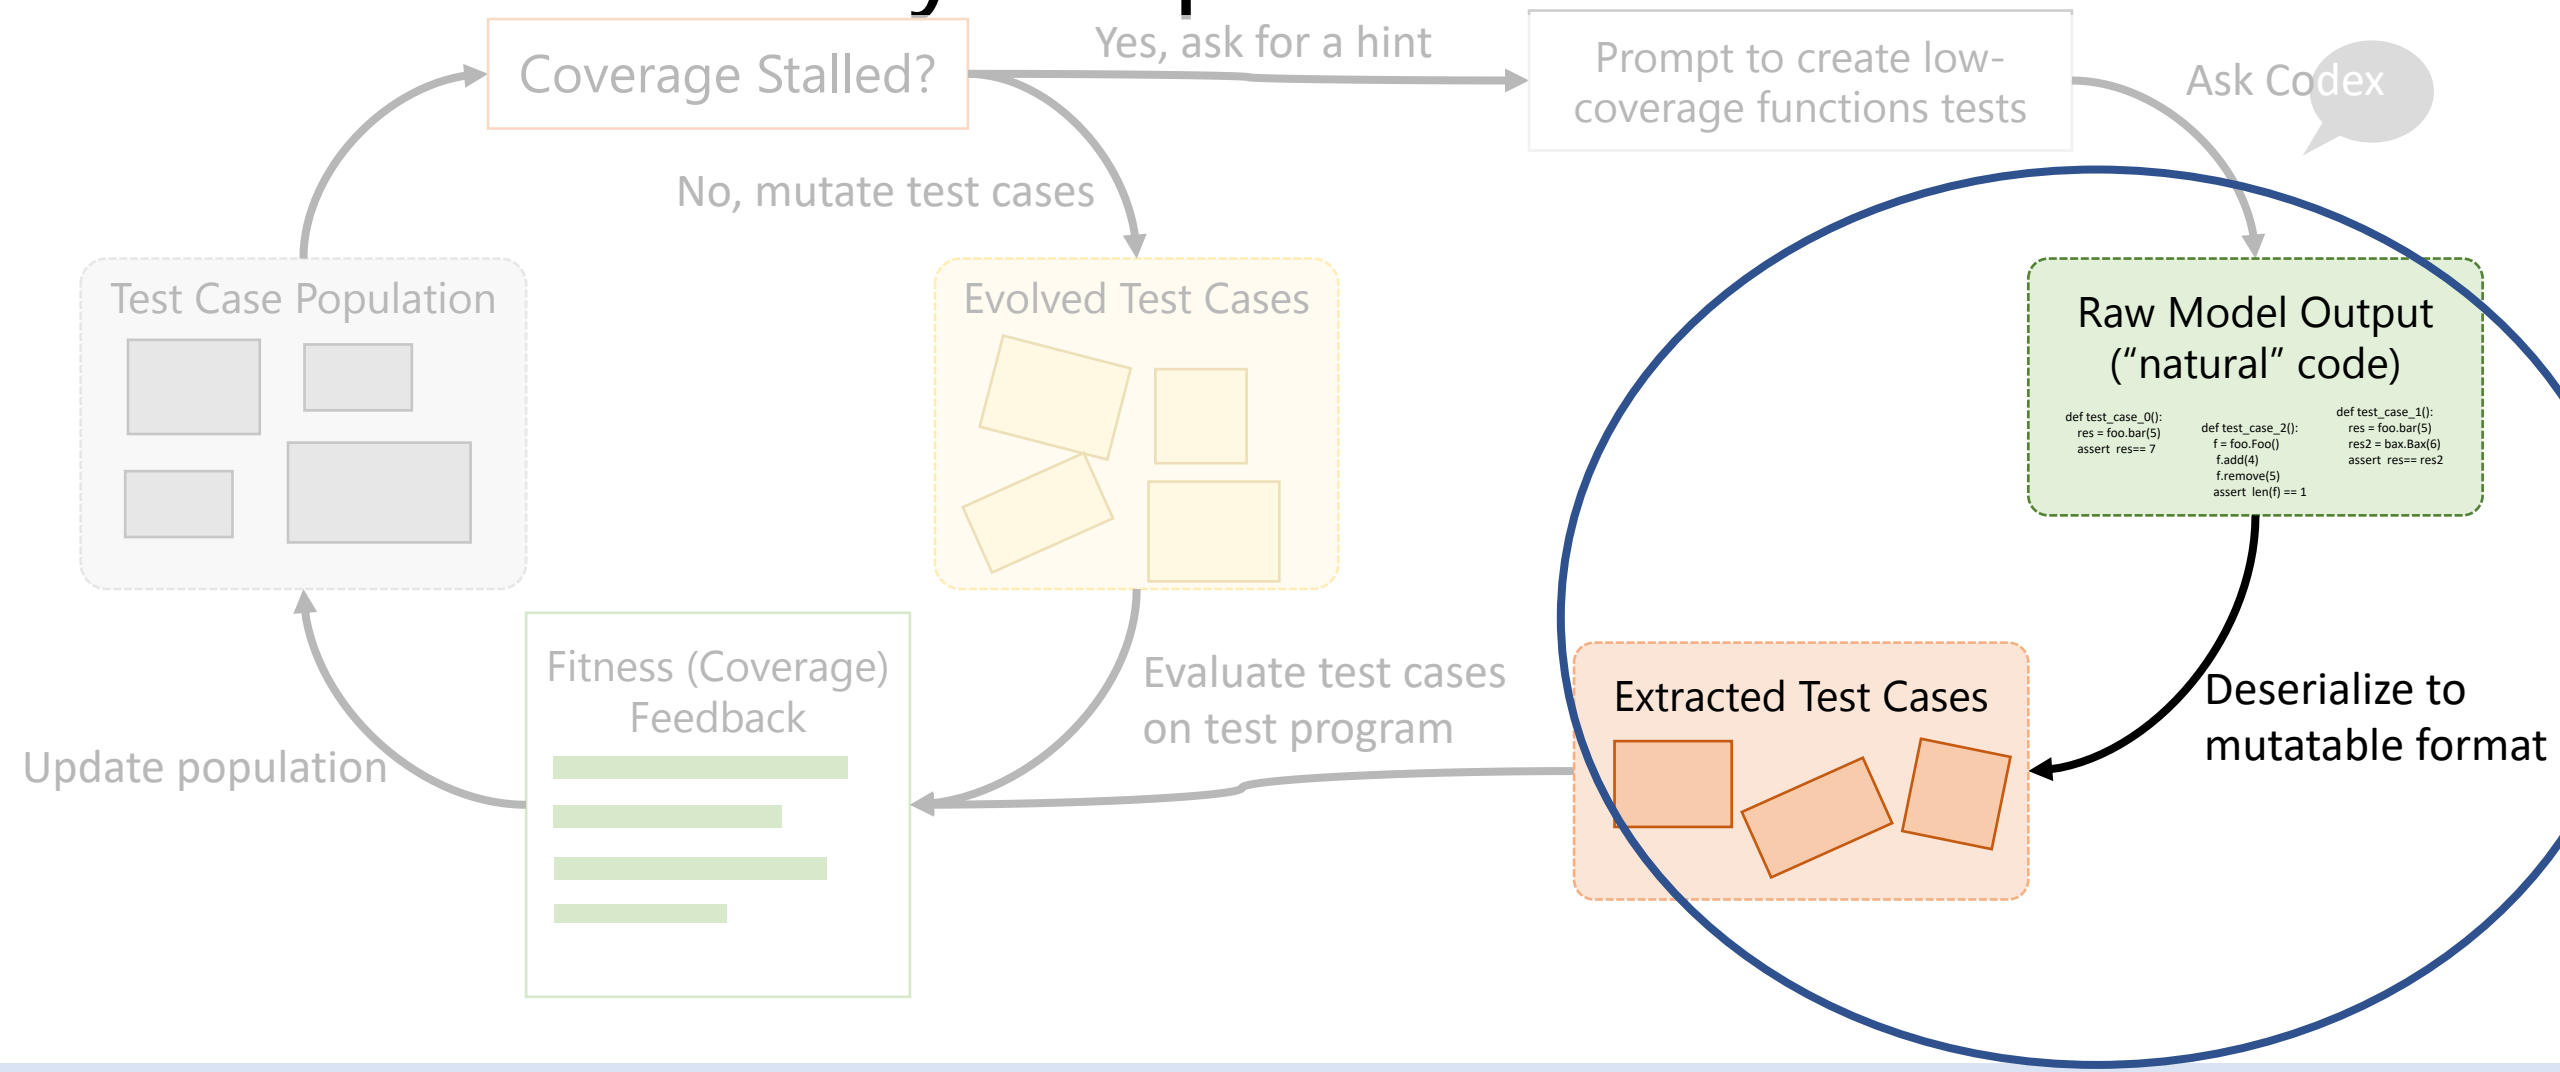

Solutions Discussed Further in Paper

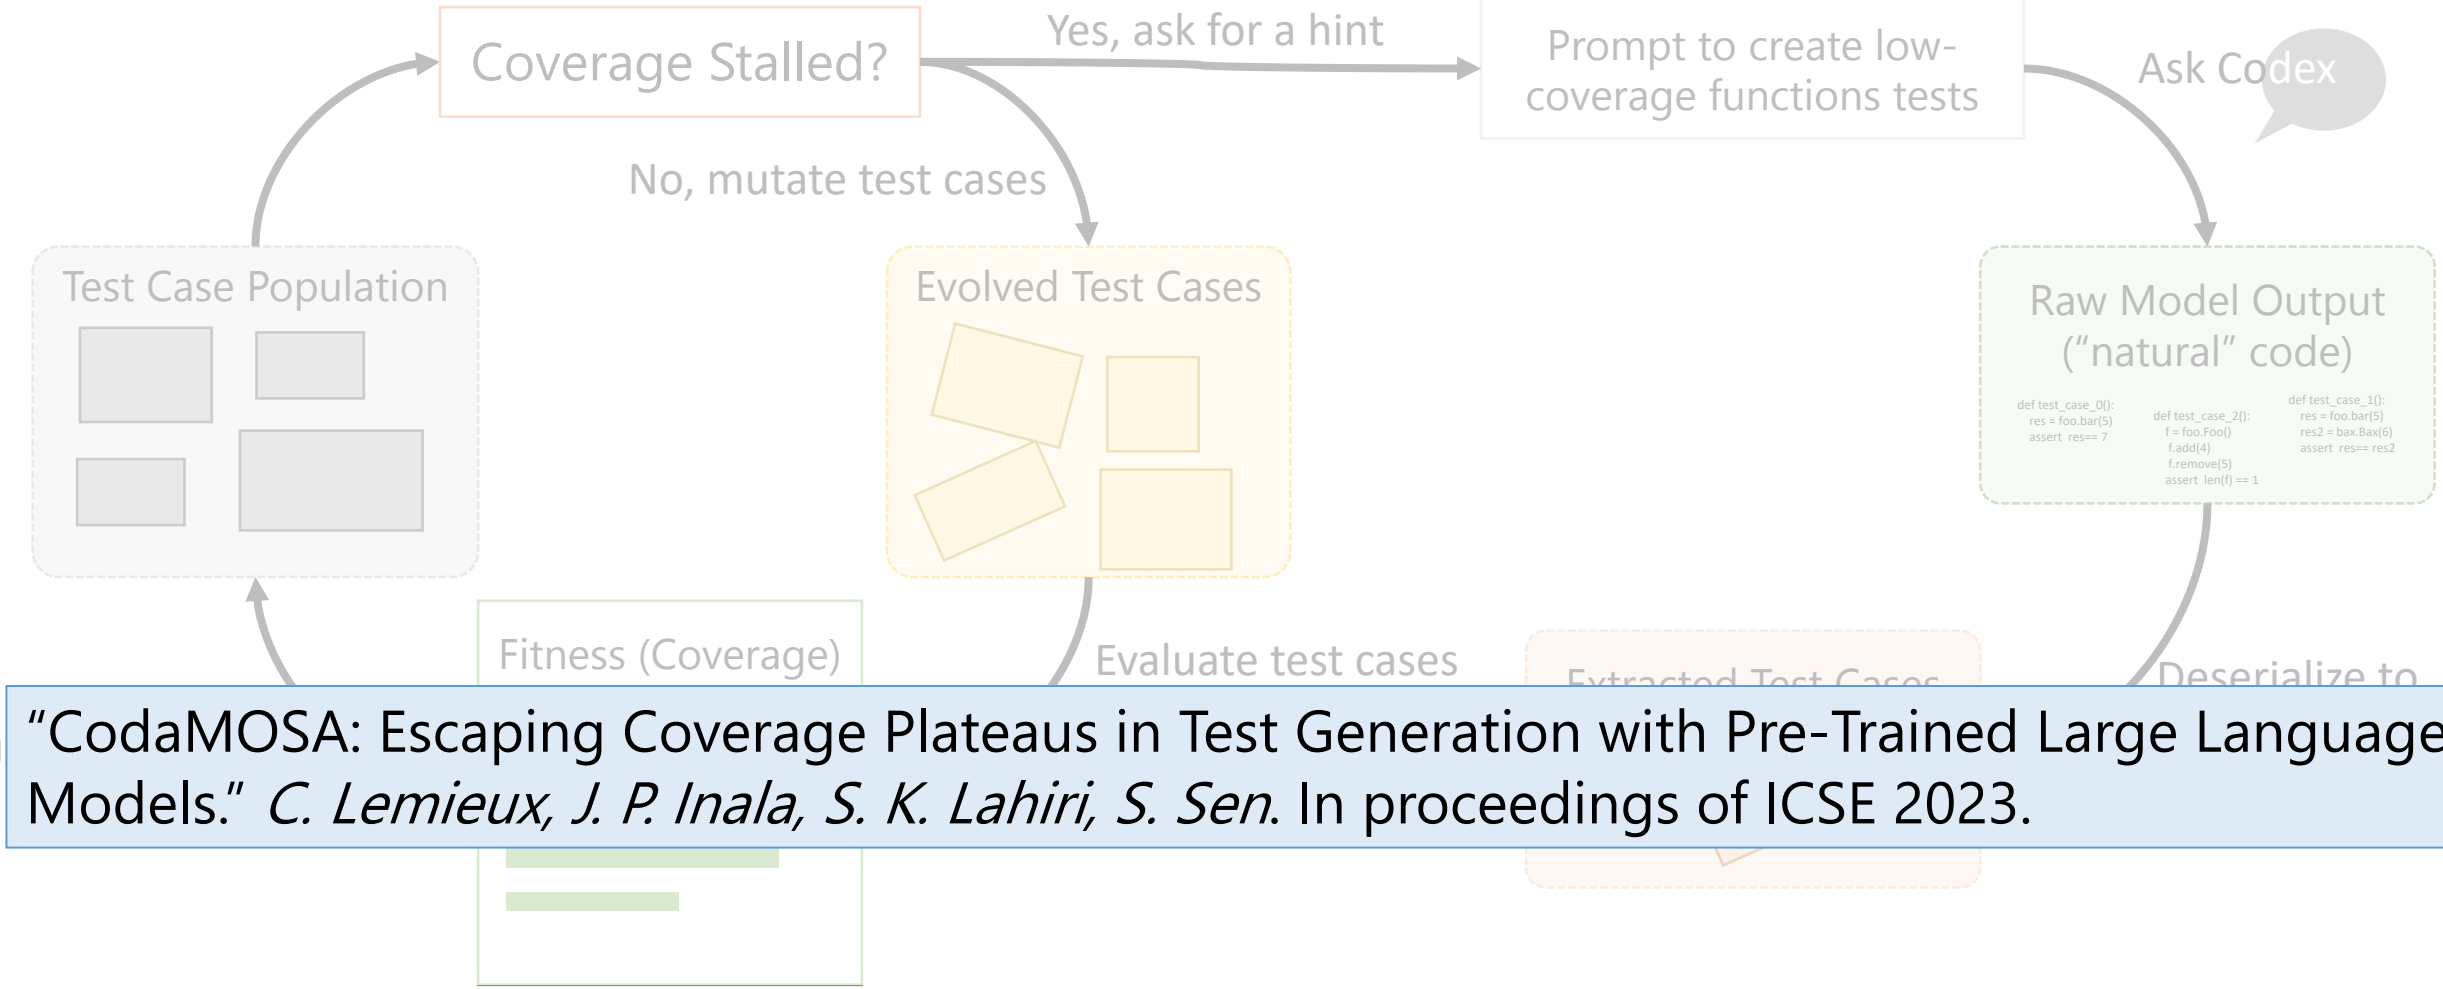

"CodaMOSA: Escaping Coverage Plateaus in Test Generation with Pre-Trained Large Language Models." *C. Lemieux, J. P. Inala, S. K. Lahiri, S. Sen*. In proceedings of ICSE 2023.

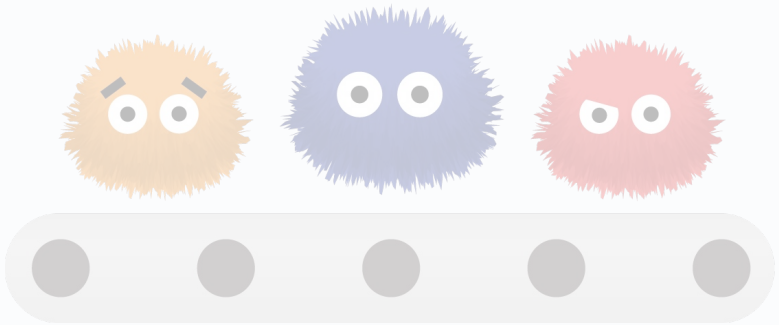

Using *generalized feedback maps* to expand *bugs findable by fuzz testing*

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

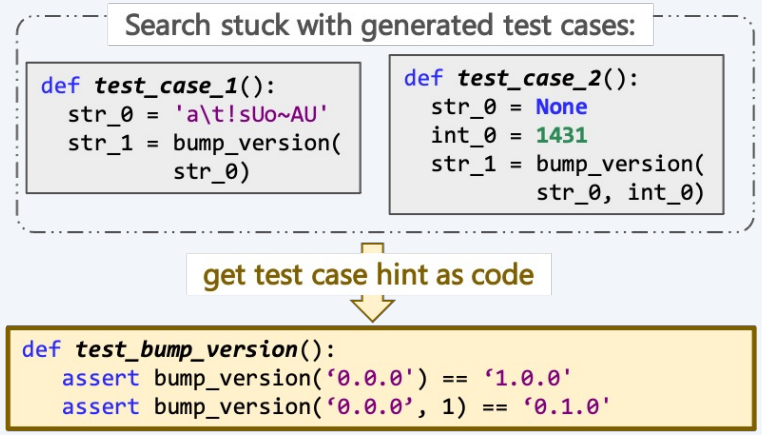

Using *large language models* to improve *automated test suite generation*

CodaMOSA (ICSE'23)

Using *generalized feedback maps* to

PerfFuzz (ISSTA'18)

FuzzFactory (OOPSLA'19)

→ Powerful synergy between LLMs ("what is most expected") and mutative search ("something close, but unexpected")

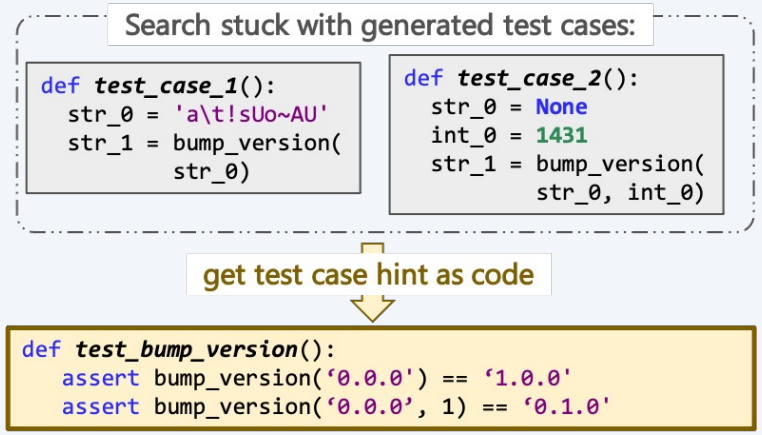

Using ***large language models*** to improve ***automated test suite generation***

CodaMOSA (ICSE'23)

So Far: Innovations in Fuzzing Algorithms

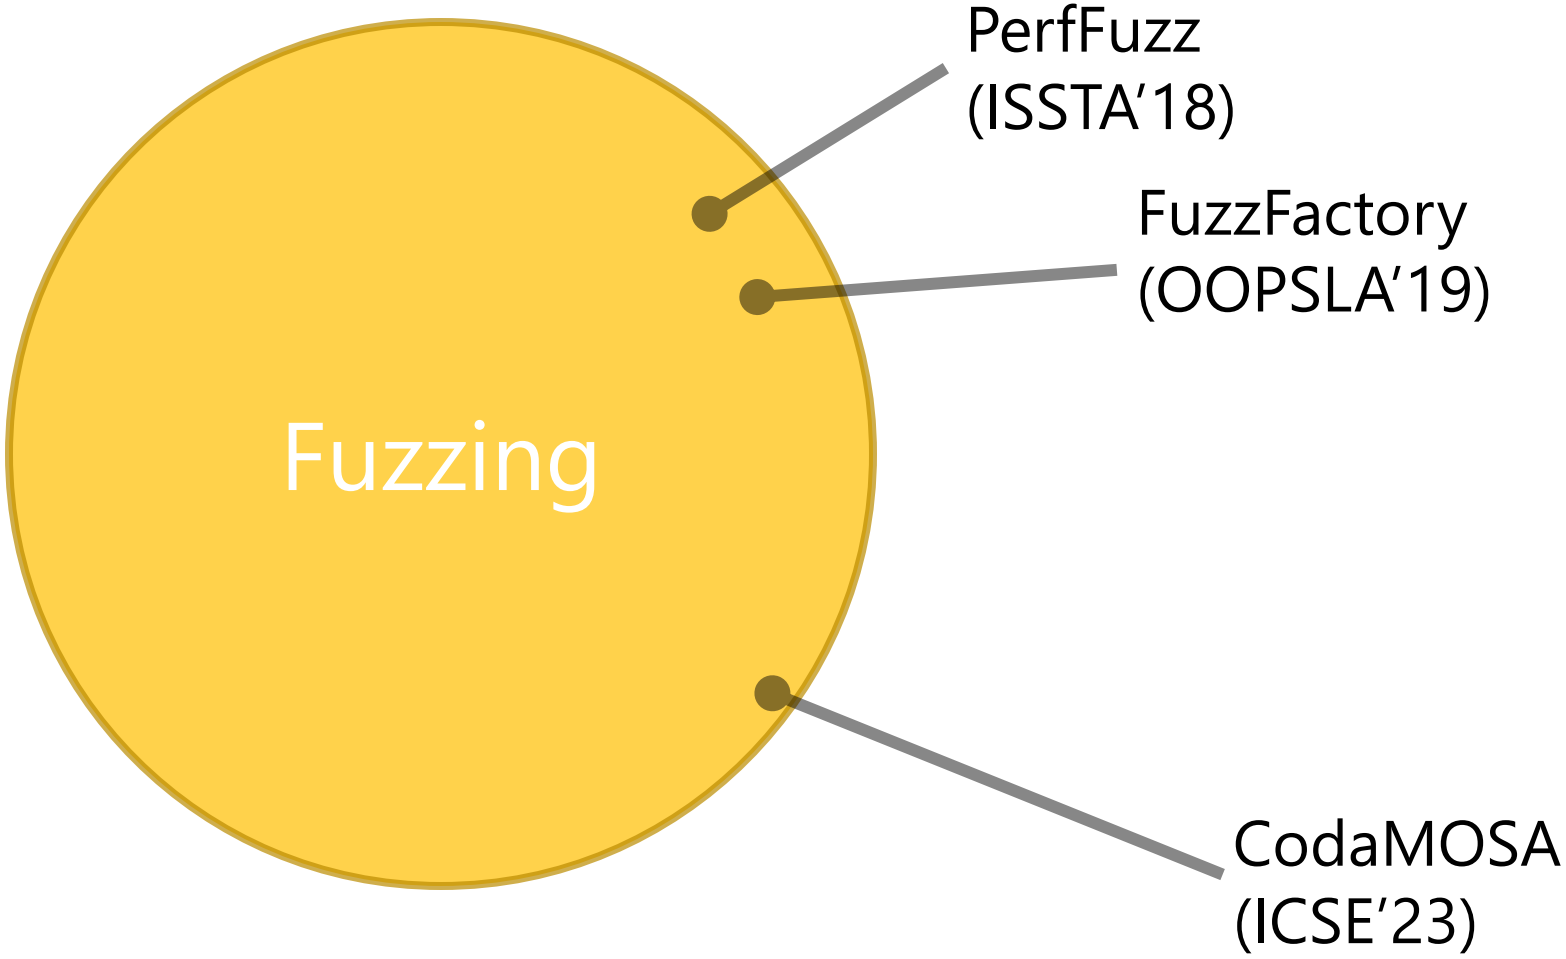

Next: Solving Problems **Around** Fuzzing

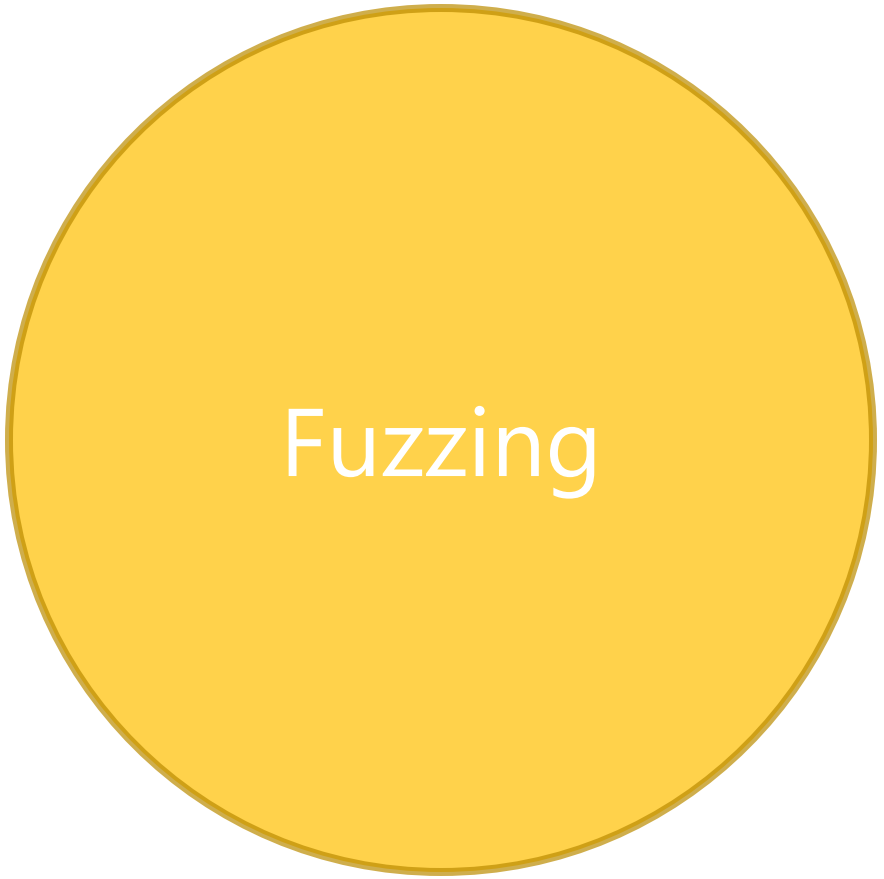

Next: Solving Problems **Around** Fuzzing

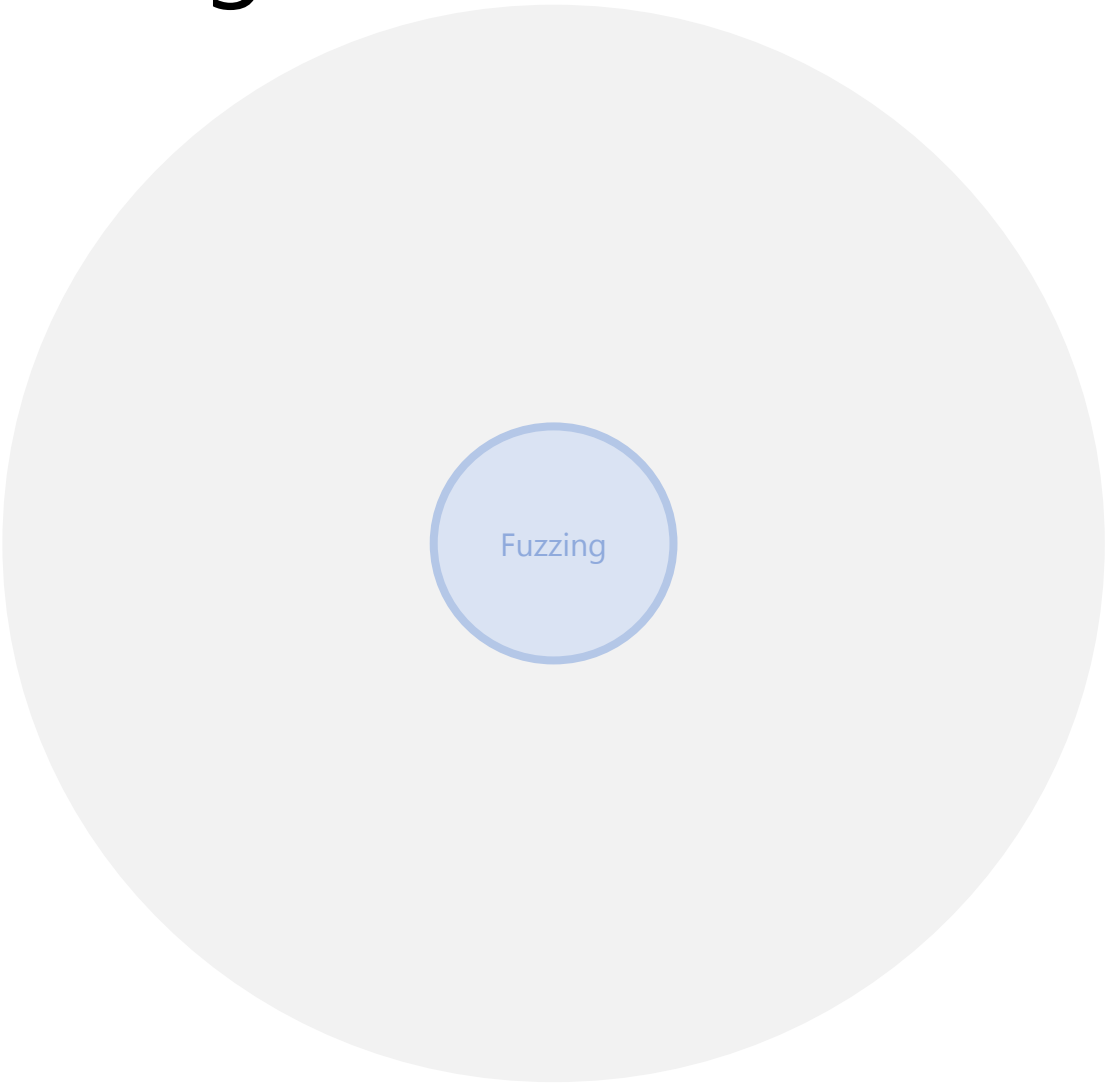

Enabling Fuzz-Driven Development

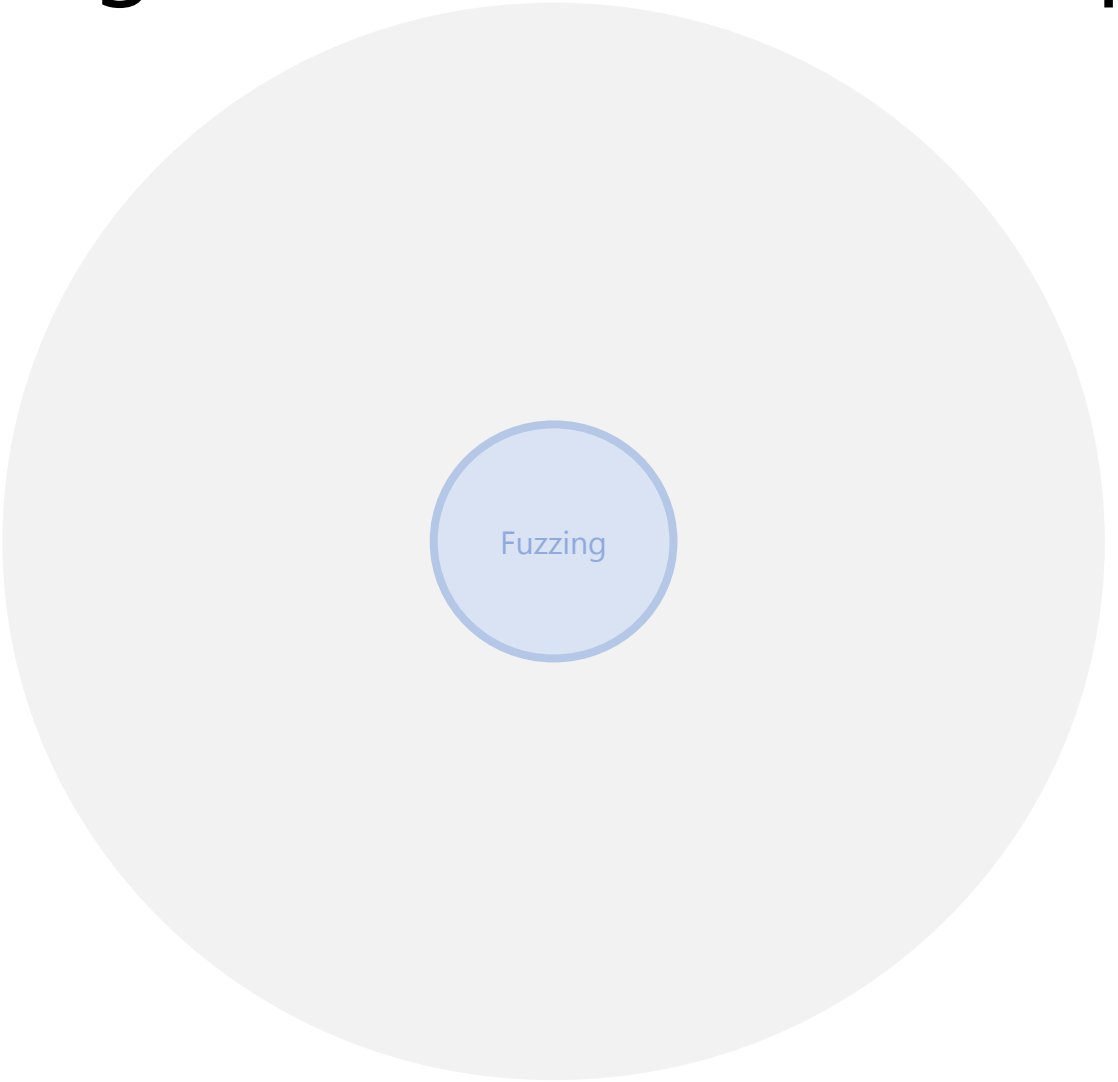

Fuzz-Driven Development: Three Pillars

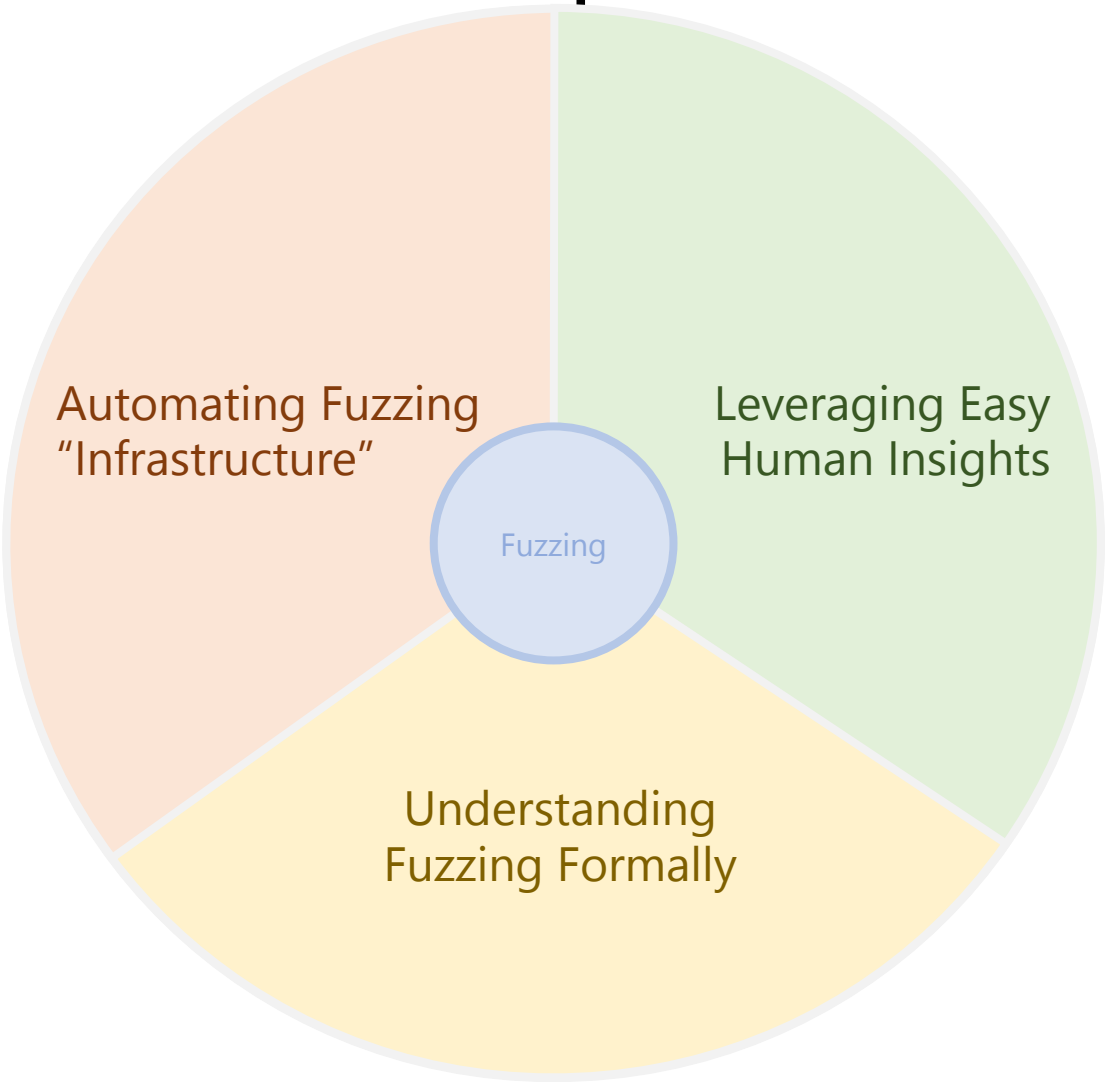

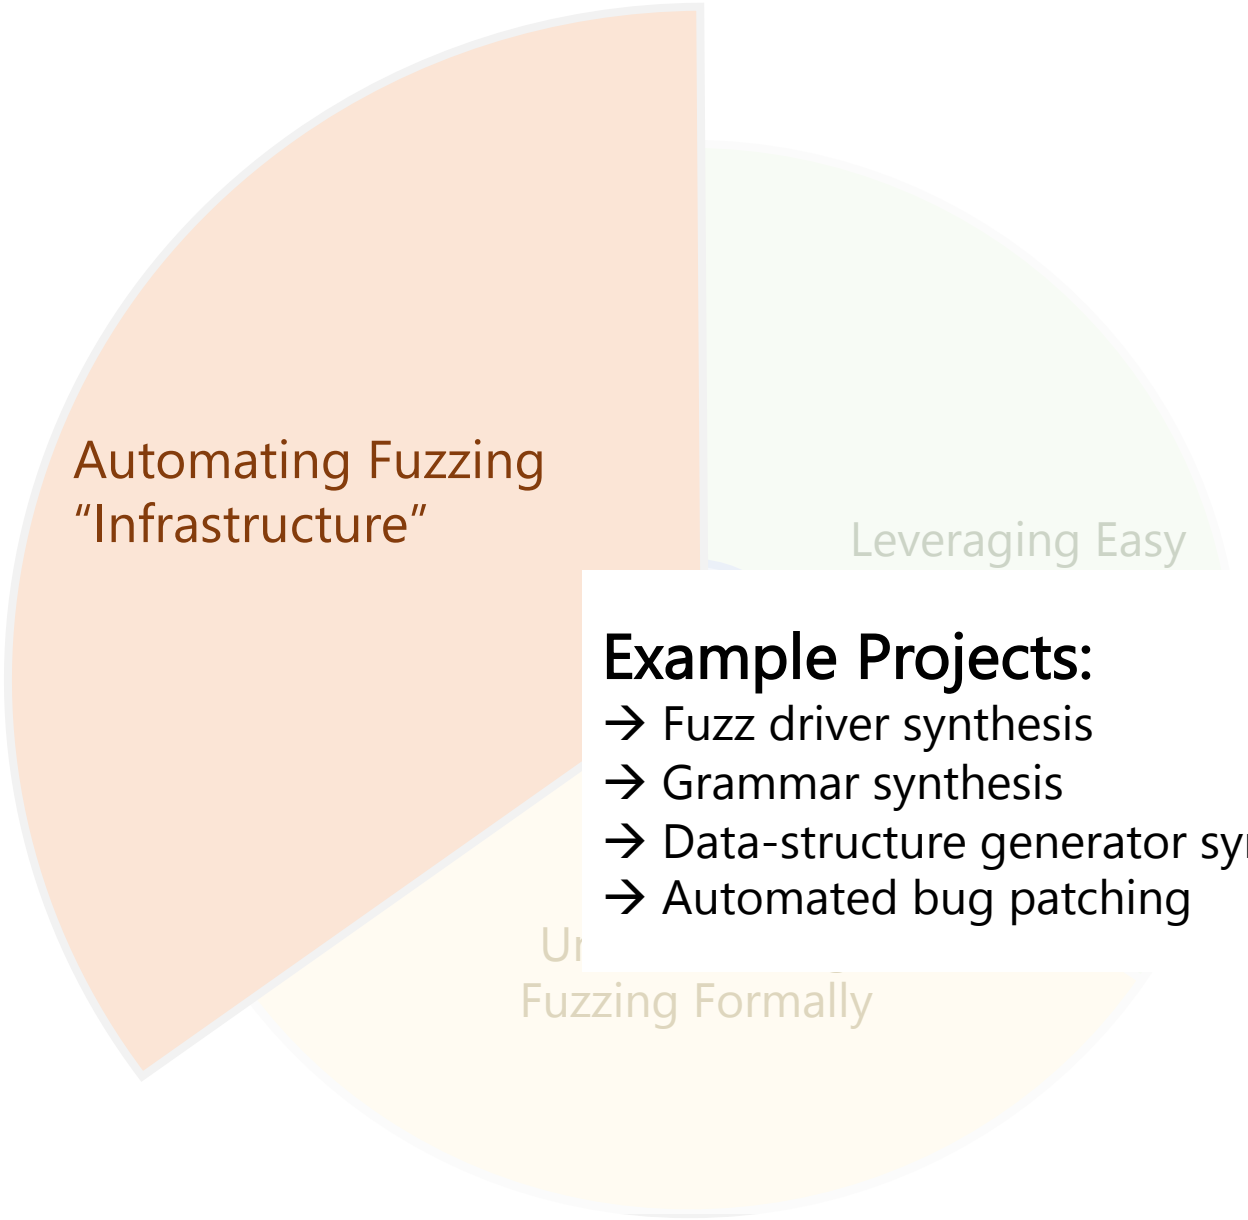

Example Projects:

→ Fuzz driver synthesis
→ Grammar synthesis
→ Data-structure generator synthesis
→ Automated bug patching

Automating Fuzzing
"Infrastructure"

Leveraging Easy

FUDGE

Babic, Bucur, Chen, Ivancic, King, Kusano, **Lemieux**, Szekeres, Wang
ESEC/FSE'19 **(Industry Track)**

Best Paper Award
(Industry Track)

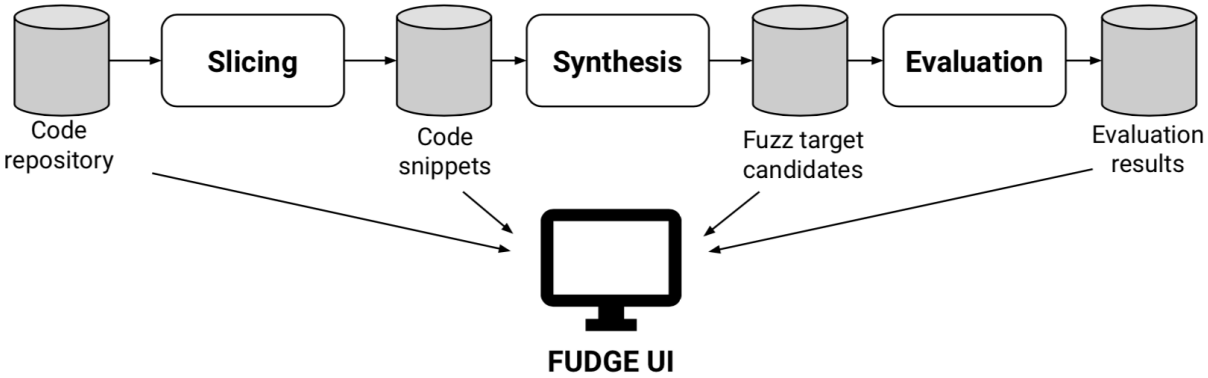

200 drivers integrated into open-source projects
→**150 security-improving fixes**

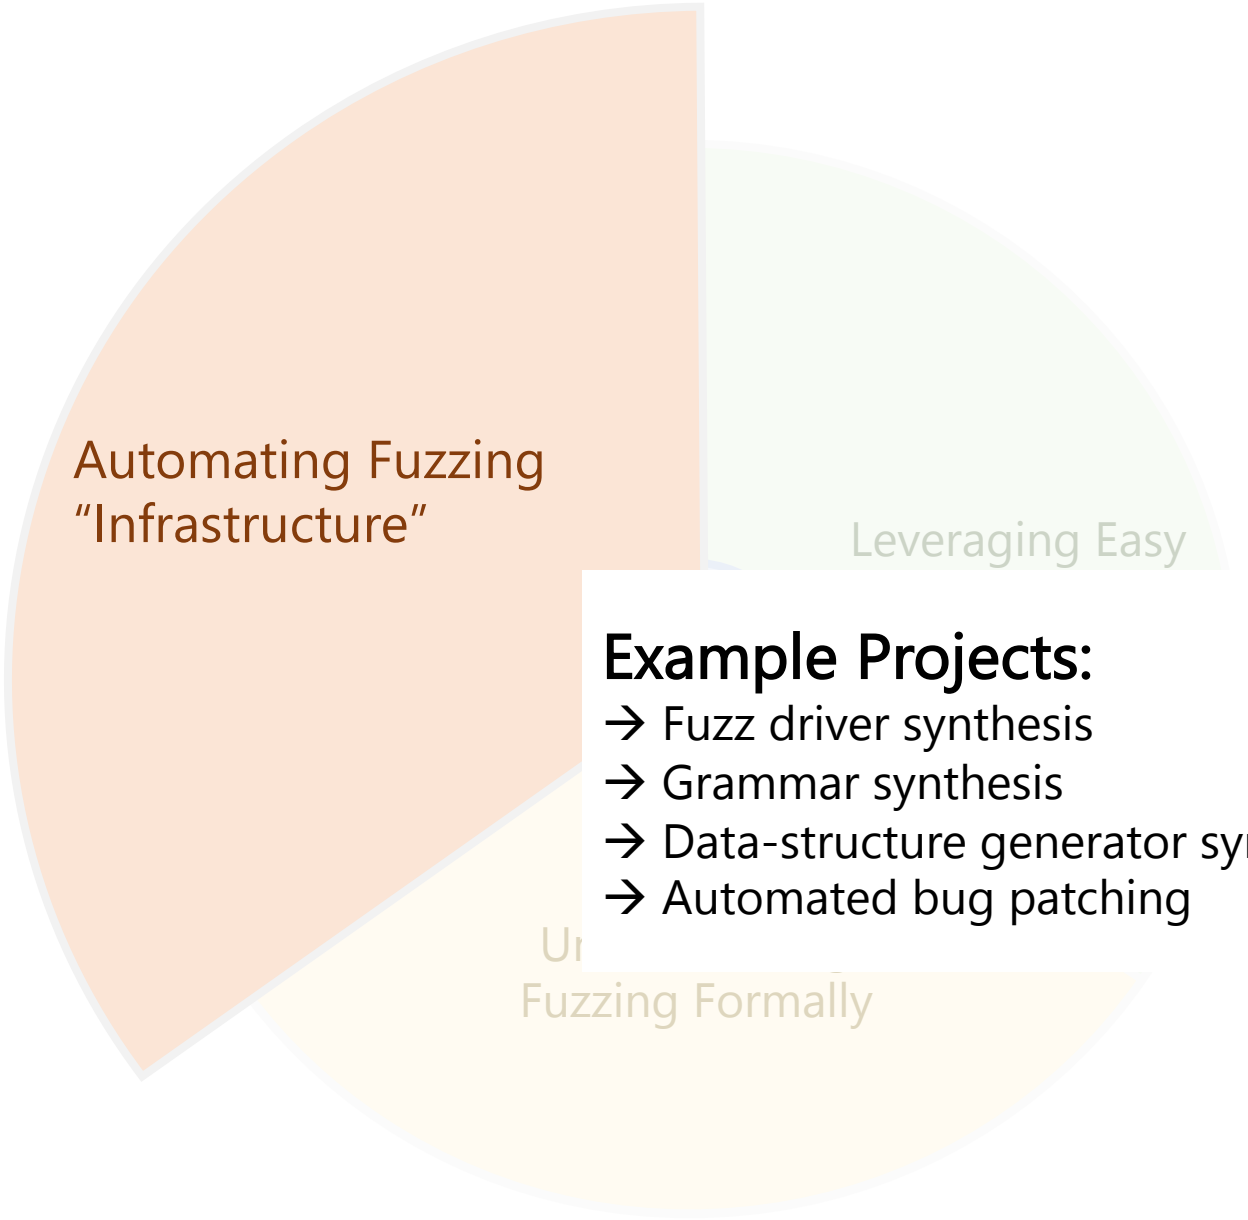

Example Projects:

→ Fuzz driver synthesis
→ Grammar synthesis
→ Data-structure generator synthesis
→ Automated bug patching

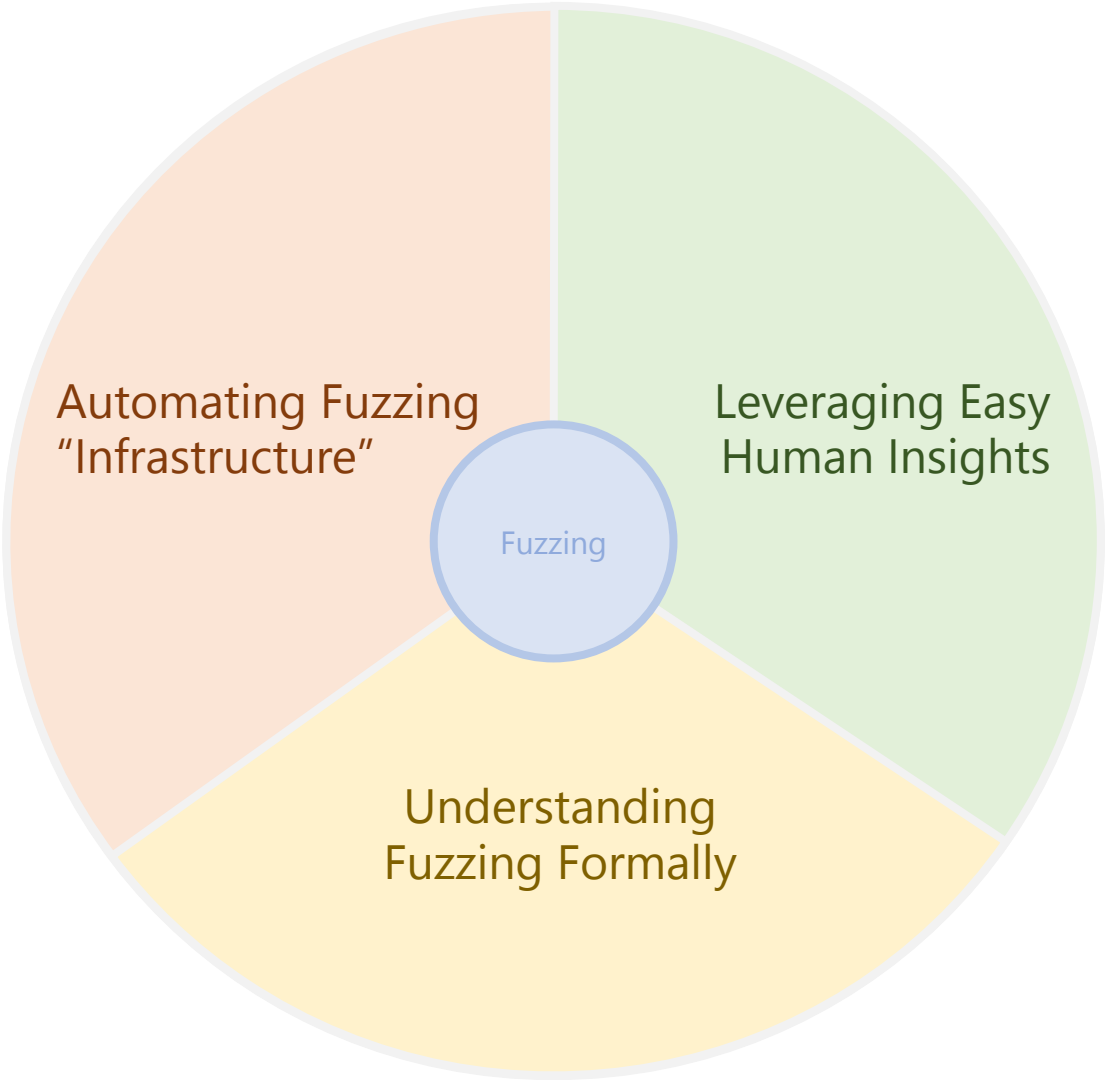
Automating Fuzzing
“Infrastructure”
Leveraging Easy
Human Insights
Fuzzing
Understanding
Fuzzing Formally

PerfFuzz (ISSTA'18) + FuzzFactory (OOPSLA'19)

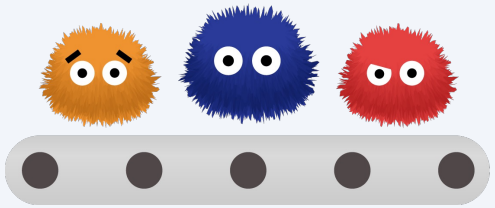

Using ***generalized feedback maps*** to expand ***bugs findable by fuzz testing***

CodaMOSA (ICSE'23)

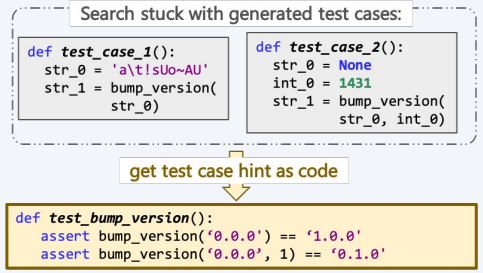

Using ***large language models*** to improve ***automated test suite generation***

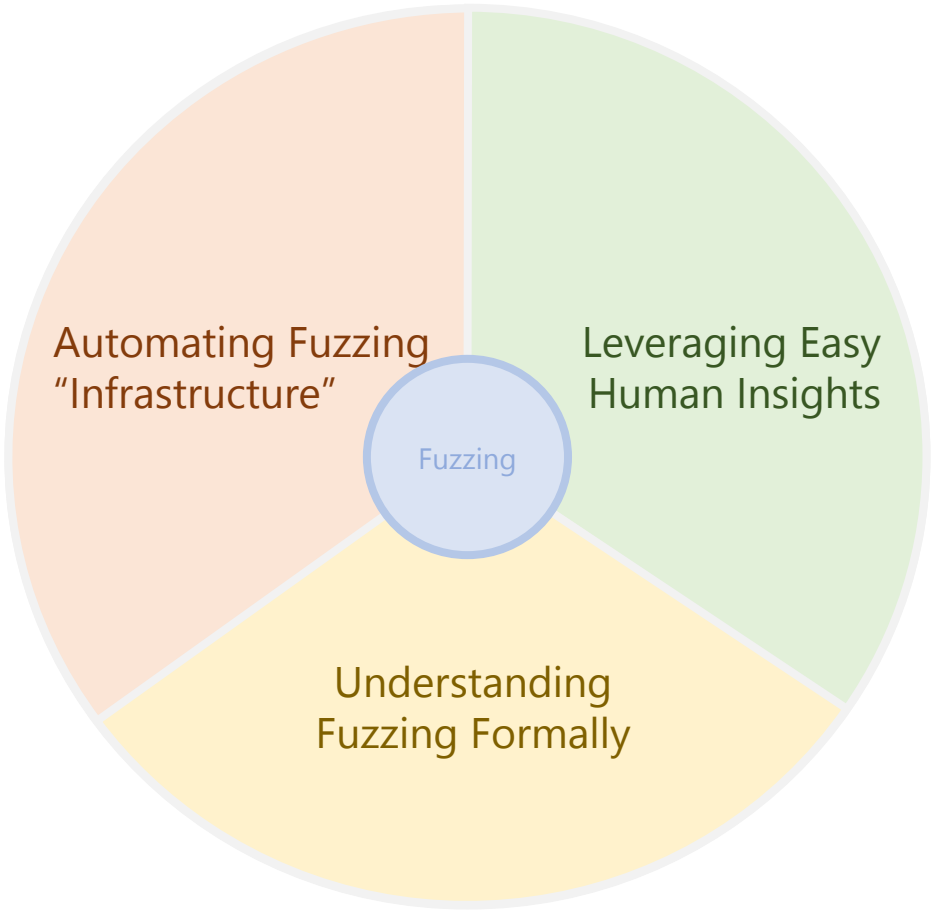

clemieux@cs.ubc.ca carolemieux.com @cestlemieux