



# Type-Constrained Code Generation with Language Models

Niels Mündler\*, Jingxuan He\*, Hao Wang, Koushik Sen, Dawn Song, Martin Vechev

ACM PLDI 2025



# AI is Redefining Software Development

## Completion



## Chat



## Agent



**Microsoft has over a million paying Github Copilot users:**  
**CEO Nadella**



Written by [Tierman Ray](#)  
Oct. 25, 2023 at 5:52 a.m. PT

**AI means everyone can now be a programmer, Nvidia chief says**



By Reuters  
May 29, 2023 3:40 PM PDT

**Over 25% of Google's code is now written by AI—and CEO Sundar Pichai says it's just the start**



BY [GREG MCKENNA](#)  
October 30, 2024 at 11:14 AM EDT

**ANTHROPIC**

37.2% of queries sent to Claude were in the “computer and mathematical” category, which in large part covers software engineering roles.

Feb 10, 2025

# *A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse*

A new wave of “reasoning” systems from companies like OpenAI is producing incorrect information more often. Even the companies don’t know

The image consists of two news snippets. The top snippet is from Forbes, featuring a black header with the word 'Forbes' in white. Below it is a dark gray bar with the word 'INNOVATION' in white. The main title 'How AI-Generated Code Is Unleashing A Tsunami Of Security Risks' is in large, bold, white font. The bottom snippet is from WIRED, with its logo in the top left corner. Below the logo, there's a small navigation bar with 'SECURITY', 'POLITICS', and 'THE BIG'. The author 'DAN GOODIN, Ars Technica' and the date 'APR 30, 2025 3:08 PM' are at the top. The main title 'AI Code Hallucinations Increase the Risk of ‘Package Confusion’ Attacks' is in large, bold, black font. A descriptive subtitle below it reads: 'A new study found that code generated by AI is more likely to contain made-up information that can be used to trick software into interacting with malicious code.'

**SOTA LLMs generate semantic errors and unsafe code**

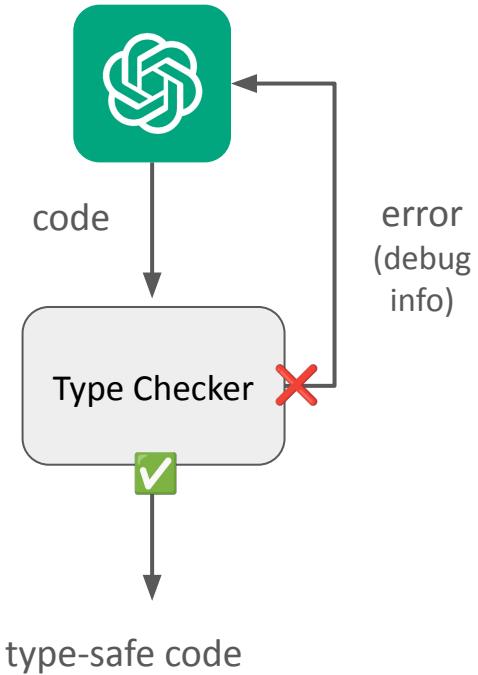
## Our Work

**Leveraging Type Systems for safe LLM code generation**

# Direct approach: Feedback loop with type checker

Treat type checker as a black box:

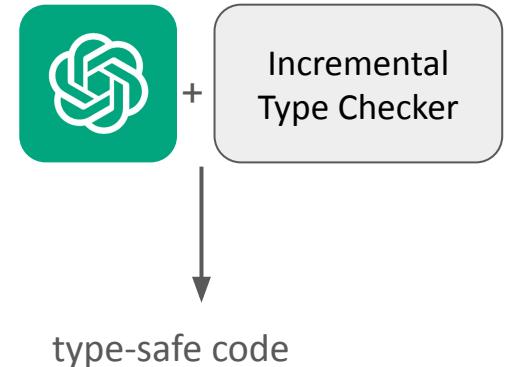
- No control during LLM decoding
- LLM needs to figure out a well-typed alternative from the error messages
- LLM passes: unpredictable, can require many iterations!



# Our work: Constraining LLMs' next-token decoding with types

## Potential Benefits:

- Full control during decoding
- Generated code is guaranteed to be type safe
- LLM passes: predictable, one pass



Wanted: new **incremental** type checkers for LLM decoding

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

function

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is
```

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int
```

Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”



# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt
```

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num
```

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num,
```

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num,
```



num is of type number  
parseInt expects type string

Generated by  
CodeLlama 32B

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num, 10)  
}
```



num is of type number  
parseInt expects type string

Generated by  
CodeLlama 32B

# Background: LLM generates code left-to-right

“Write a function to check if a string represents an integer”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num.toString()) === num  
}
```

There is a valid completion!

# Incremental type checker

“Can current output be completed to be well-typed?”

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num)
```

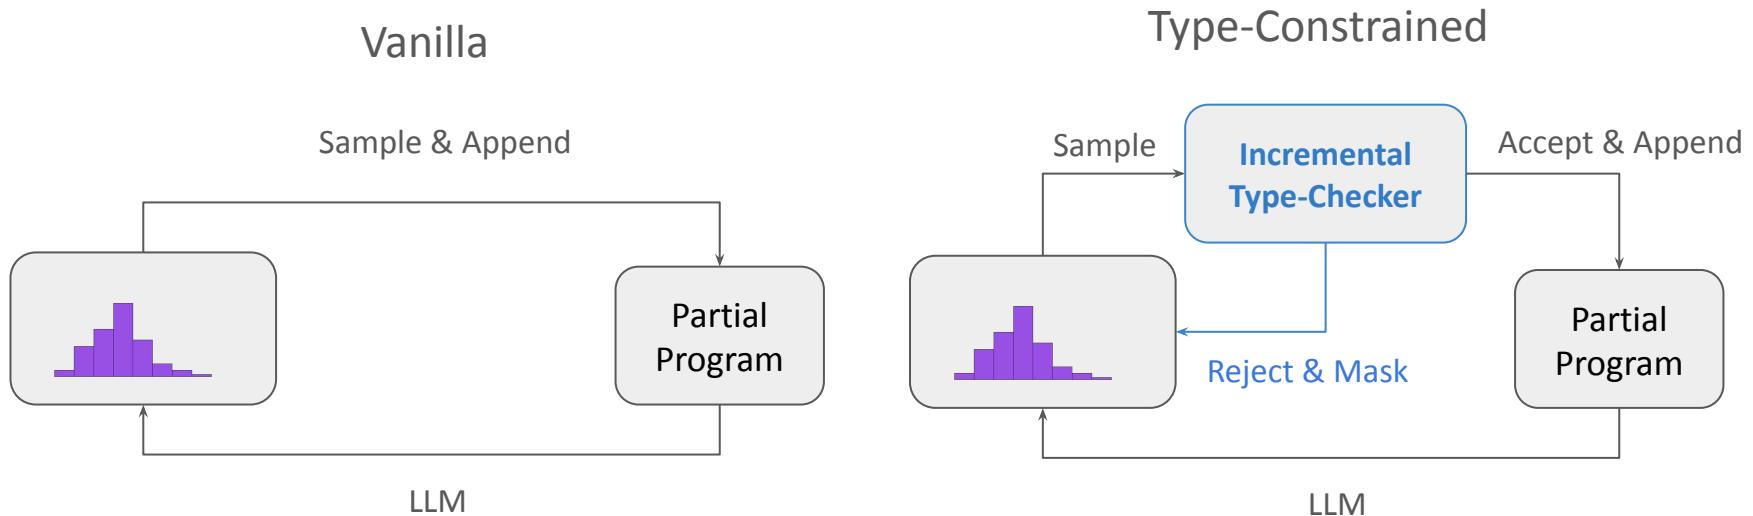
ok!

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num,
```

not ok!



# Vanilla vs. Type-Constrained LLM code generation



# Building an Incremental Type checker

# Key ingredient: Incremental type-checking automaton

Non-deterministic automaton that only accepts well-typed programs

- Abstract syntax tree
- Augmented with typing environment

Ensure *prefix property* on automaton

- Every state leads to an accepting state
- If we are in state → current parsed output is prefix
- If in no state → current parsed output is not prefix

# Applying our method to a subset of TypeScript

## Generic simply typed language

$l ::=$	Literal
$\backslash d+$	Numeric Literal
$"\w*$	String Literal
$true   false$	Boolean Literal
$x ::= \w+$	Identifier
$e ::= e_0   e_1$	Expression
$e_0 ::=$	Base Expression
$l$	Literal
$x$	Identifier
$(\bar{p}) \Rightarrow e$	Function Expression
$(e)$	Grouped Expression
$e_1 ::=$	Extension Expression
$e \odot e$	Binary Operator
$e(\bar{e})$	Function Call
$e.n$	Member Access

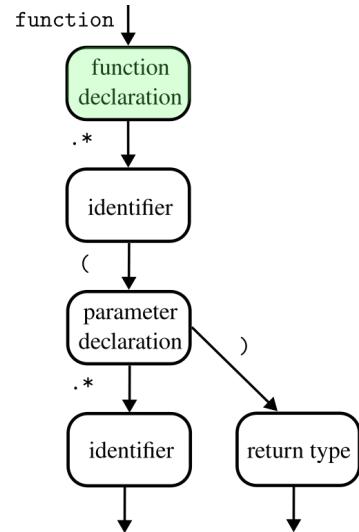
$p ::= x : T$	Typed Identifier
$T ::=$	Type
$\text{number}$	Numeric Type
$\text{string}$	String Type
$\text{boolean}$	Boolean Type
$(\bar{p}) \Rightarrow T$	Function Type
$s ::=$	Statement
$\text{let } x : T;$	Variable Declaration
$e;$	Expression Statement
$\text{return } e;$	Return Statement
$\{\bar{s}\}$	Statement Block
$\text{function } x(\bar{p}) : T \{ \bar{s} \}$	Function Definition
$\text{if } (e) s \text{ else } s$	If-Then-Else Statement
$M ::= \bar{s}$	Program

## TypeScript extensions

- Mutable vs Immutable Variables
- Arrays
- Loops
- More Operators and Types
- Operator Precedence
- Global Identifiers
- Imports
- Polymorphic Built-Ins
- ...

# Incremental type checking

Automaton                    Typing Environment

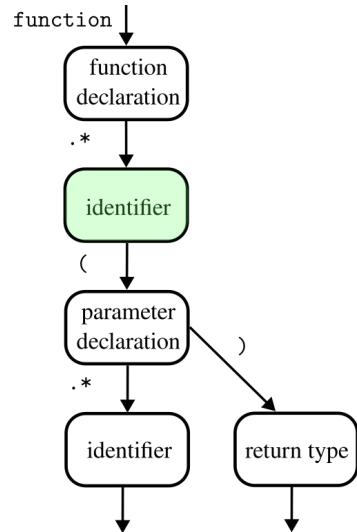


Code

Identifier	Type	Code
parselnt	(string) => number	function

# Incremental type checking

Automaton



Typing Environment

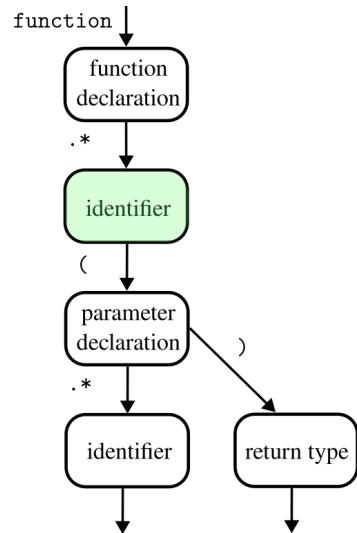
Identifier	Type
parseInt	(string) => number

Code

function is

# Incremental type checking

Automaton



Typing Environment

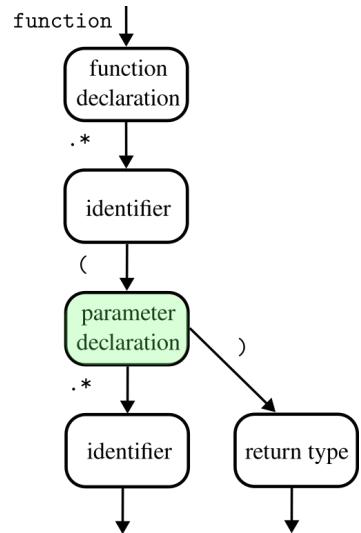
Identifier	Type
parselnt	(string) => number

Code

function is\_int

# Incremental type checking

Automaton



Typing Environment

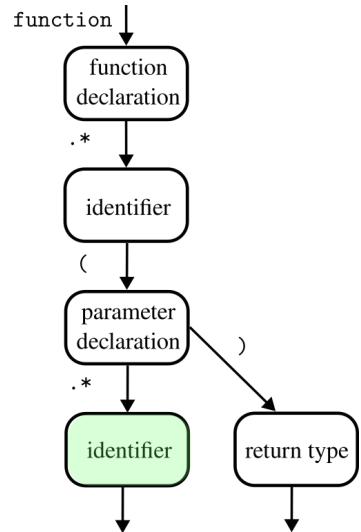
Identifier	Type
parselnt	(string) => number

Code

```
function is_int(text
```

# Incremental type checking

Automaton



Typing Environment

Identifier	Type
parselnt	(string) => number

Code

```
function is_int(text
```

# Incremental type checking

Automaton

Typing Environment

Code

Identifier	Type	Code
parseInt	(string) => number	
text	string	
is_int	(string) => boolean	<code>function is_int(text: string): boolean {</code>

# Incremental type checking

Automaton

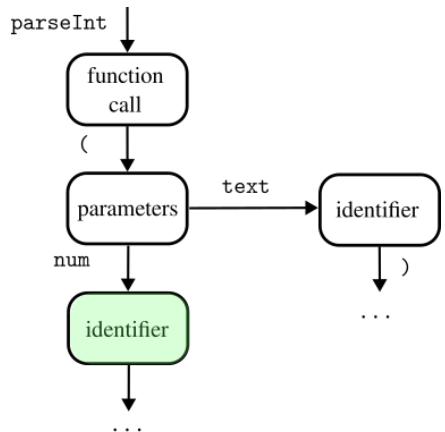
Typing Environment

Code

Identifier	Type	
parseInt	(string) => number	<code>function is_int(text: string): boolean {</code>
text	string	<code>const num = Number(text);</code>
is_int	(string) => boolean	
num	number	

# Incremental type checking

Automaton



Typing Environment

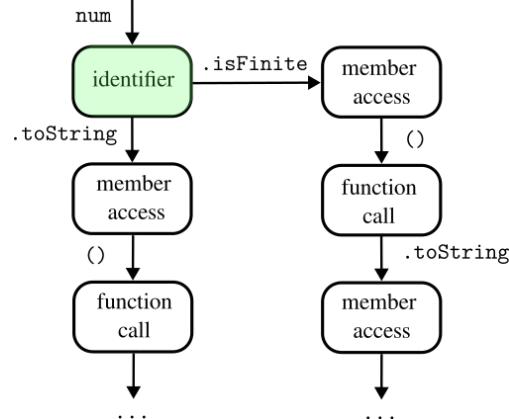
Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

Code

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num)
```

# Incremental type checking

Automaton



Typing Environment

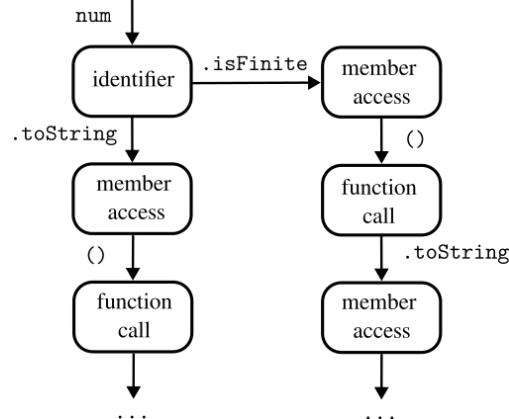
Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

Code

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num)
```

# Incremental type checking

Automaton



Typing Environment

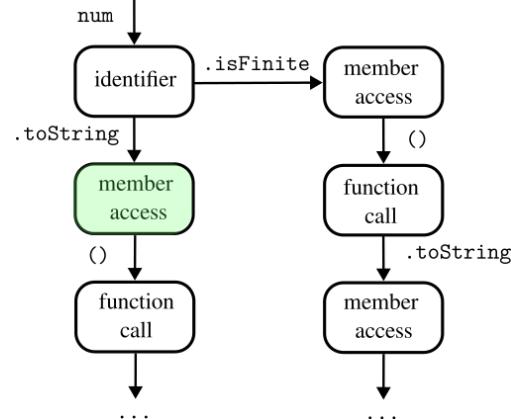
Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

Code

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num,
```

# Incremental type checking

Automaton



Typing Environment

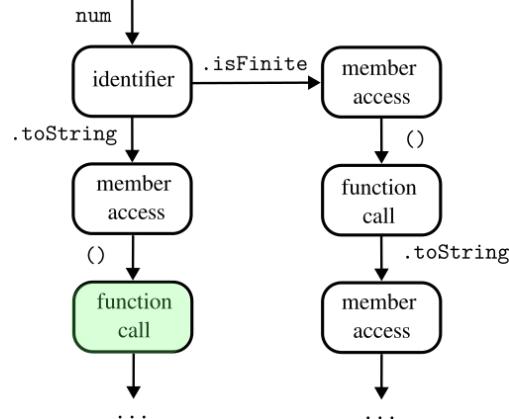
Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

Code

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num.toString())
```

# Incremental type checking

Automaton



Typing Environment

Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

Code

```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num.toString())
```

# Incremental type checking

Automaton

Typing Environment

Code

Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

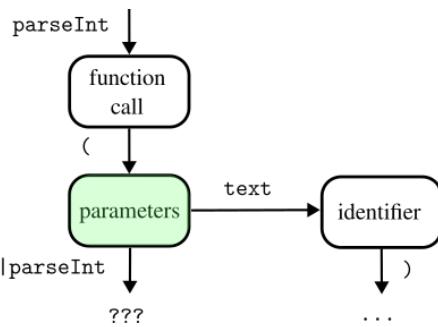
```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num.toString()) === num  
}
```

# Rewind: Addressing type inhabitation

Automaton

Typing Environment

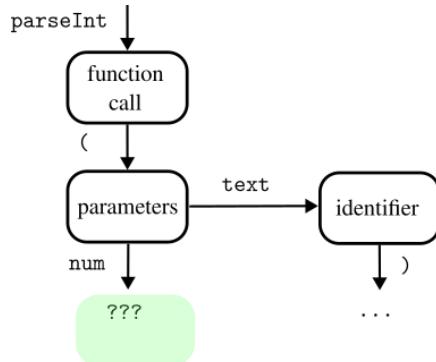
Code



Identifier	Type
parseInt	(string) => number
text	string
is_int	(string) => boolean
num	number

```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Rewind: Addressing type inhabitation

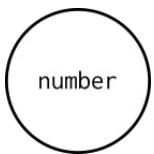
Automaton	Typing Environment	Code										
	<table><thead><tr><th>Identifier</th><th>Type</th></tr></thead><tbody><tr><td>parseInt</td><td>(string) =&gt; number</td></tr><tr><td>text</td><td>string</td></tr><tr><td>is_int</td><td>(string) =&gt; boolean</td></tr><tr><td>num</td><td>number</td></tr></tbody></table>	Identifier	Type	parseInt	(string) => number	text	string	is_int	(string) => boolean	num	number	<pre>function is_int(text: string): boolean {     const num = Number(text);     return !isNaN(num) &amp;&amp;         parseInt(num)</pre>
Identifier	Type											
parseInt	(string) => number											
text	string											
is_int	(string) => boolean											
num	number											

# Rewind: Addressing type inhabitation

Automaton	Typing Environment	Code										
<pre>graph TD; parseInt --&gt; funcCall[function call]; funcCall --&gt; params[parameters]; params -- text --&gt; identifier[identifier]; identifier -- num --&gt; ???[???]; params --&gt; ???;</pre>	<table><thead><tr><th>Identifier</th><th>Type</th></tr></thead><tbody><tr><td>parseInt</td><td>(string) =&gt; number</td></tr><tr><td>text</td><td>string</td></tr><tr><td>is_int</td><td>(string) =&gt; boolean</td></tr><tr><td>num</td><td>number</td></tr></tbody></table>	Identifier	Type	parseInt	(string) => number	text	string	is_int	(string) => boolean	num	number	<pre>function is_int(text: string): boolean {     const num = Number(text);     return !isNaN(num) &amp;&amp;         parseInt(num)</pre> <p>Check if we <i>can inhabit</i> the goal type with a completion of the current partial expression</p>
Identifier	Type											
parseInt	(string) => number											
text	string											
is_int	(string) => boolean											
num	number											

# Type reachability search

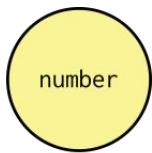
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num
```

# Type reachability search

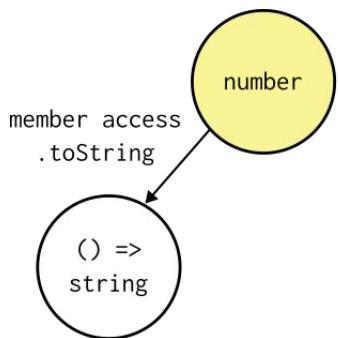
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num
```

# Type reachability search

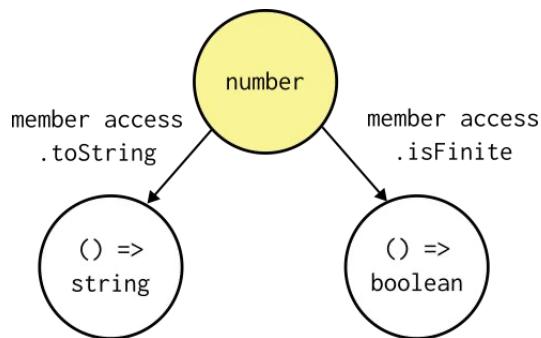
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num)
```

# Type reachability search

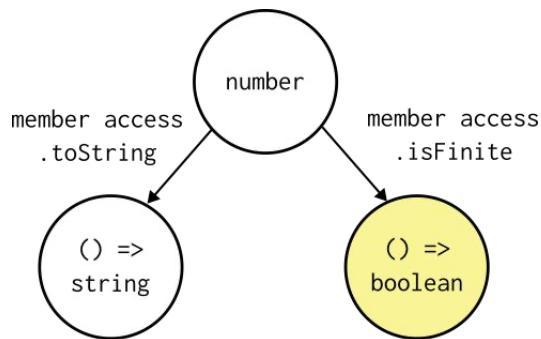
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type reachability search

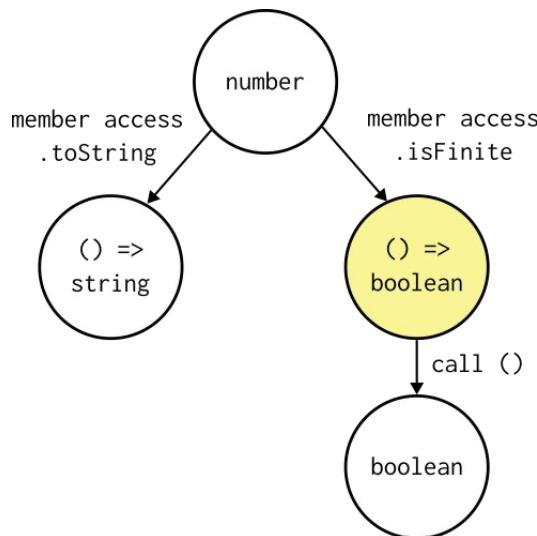
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type reachability search

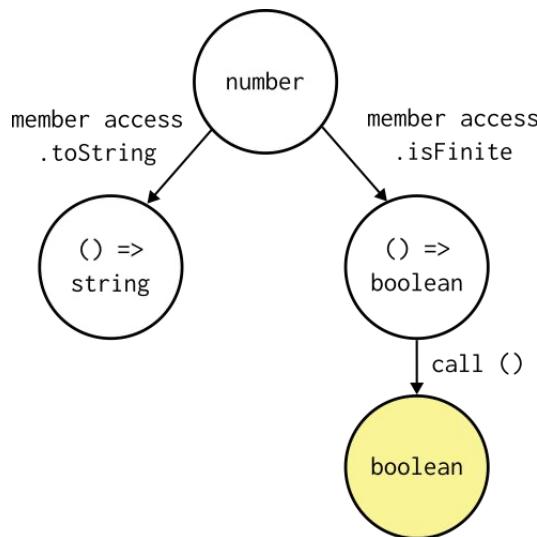
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type reachability search

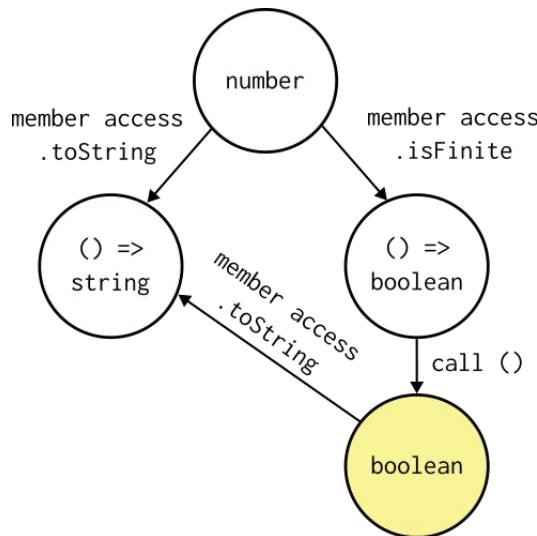
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type reachability search

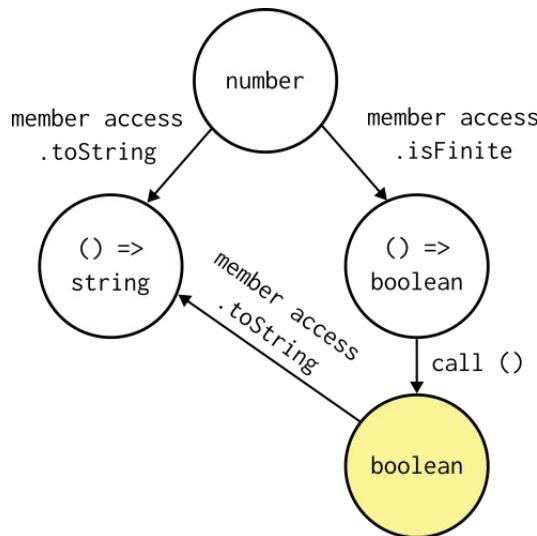
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type reachability search

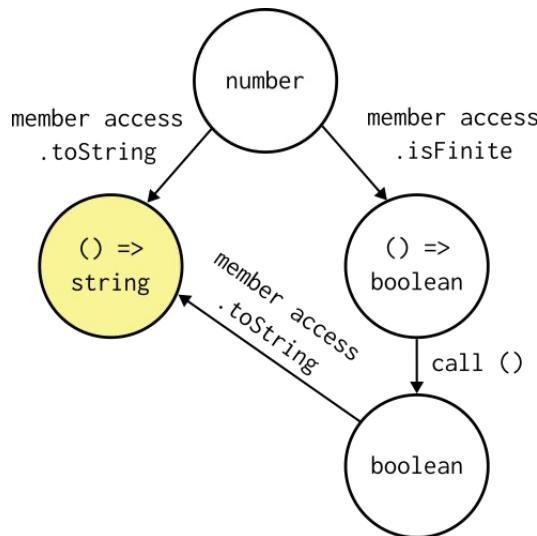
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type reachability search

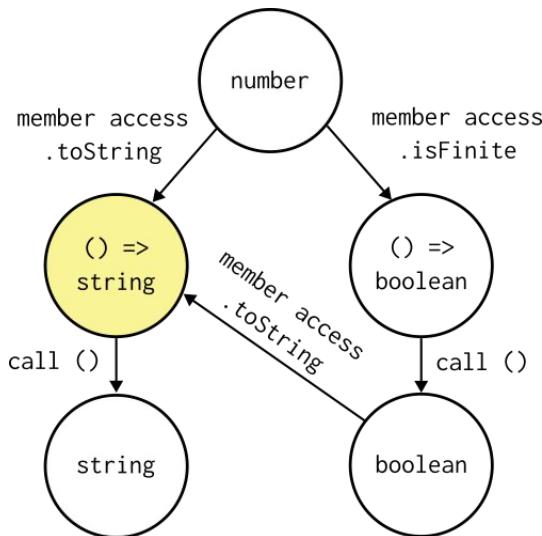
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
    const num = Number(text);  
    return !isNaN(num) &&  
        parseInt(num
```

# Type reachability search

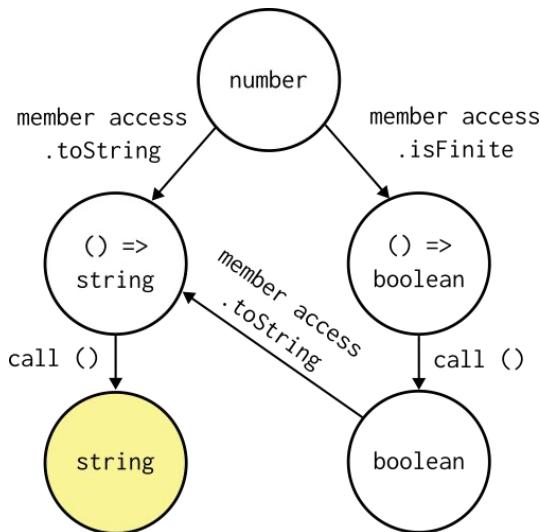
Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

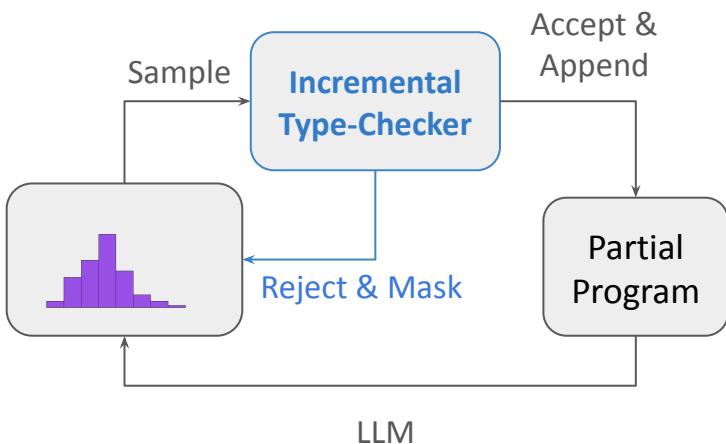
# Type reachability search

Can we append operators to turn type number into string?



```
function is_int(text: string): boolean {  
  const num = Number(text);  
  return !isNaN(num) &&  
    parseInt(num)
```

# Type-Constrained LLM code generation: benefits



One forward LLM inference pass, no backtracking

Guaranteed type safety on termination

Fully leverages type system during decoding

Requires manually coming up with the automata  
from the language and its type system definitions

# Experimental Evaluation

# Drastic reduction of compiler errors

Model	Synthesis		Translation		Repair		
	Vanilla	Types	Vanilla	Types	Vanilla	Types	
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Drastic reduction of compiler errors

Model	Synthesis		Translation		Repair		
	Vanilla	Types	Vanilla	Types	Vanilla	Types	
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Drastic reduction of compiler errors

	Model	Synthesis		Translation		Repair	
		Vanilla	Types	Vanilla	Types	Vanilla	Types
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Drastic reduction of compiler errors

Model	Synthesis		Translation		Repair		
	Vanilla	Types	Vanilla	Types	Vanilla	Types	
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Drastic reduction of compiler errors

Model	Synthesis		Translation		Repair		
	Vanilla	Types	Vanilla	Types	Vanilla	Types	
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Drastic reduction of compiler errors

	Model	Synthesis		Translation		Repair	
		Vanilla	Types	Vanilla	Types	Vanilla	Types
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Drastic reduction of compiler errors

Model	Synthesis		Translation		Repair		
	Vanilla	Types	Vanilla	Types	Vanilla	Types	
HumanEval	Gemma 2 2B	103	<b>44</b> <sub>↓57.3%</sub>	177	<b>80</b> <sub>↓54.8%</sub>	194	<b>103</b> <sub>↓46.9%</sub>
	Gemma 2 9B	45	<b>13</b> <sub>↓71.1%</sub>	75	<b>16</b> <sub>↓78.7%</sub>	113	<b>52</b> <sub>↓54.0%</sub>
	Gemma 2 27B	15	<b>2</b> <sub>↓86.7%</sub>	20	<b>3</b> <sub>↓85.0%</sub>	45	<b>22</b> <sub>↓51.1%</sub>
	DS Coder 33B	26	<b>5</b> <sub>↓80.8%</sub>	18	<b>7</b> <sub>↓61.1%</sub>	36	<b>15</b> <sub>↓58.3%</sub>
	CodeLlama 34B	86	<b>28</b> <sub>↓67.4%</sub>	158	<b>59</b> <sub>↓62.7%</sub>	153	<b>48</b> <sub>↓68.6%</sub>
	Qwen2.5 32B	17	<b>2</b> <sub>↓88.2%</sub>	24	<b>5</b> <sub>↓79.2%</sub>	36	<b>13</b> <sub>↓63.9%</sub>
MBPP	Gemma 2 2B	67	<b>27</b> <sub>↓59.7%</sub>	126	<b>79</b> <sub>↓37.3%</sub>	194	<b>108</b> <sub>↓44.3%</sub>
	Gemma 2 9B	30	<b>10</b> <sub>↓66.7%</sub>	67	<b>33</b> <sub>↓50.7%</sub>	129	<b>63</b> <sub>↓51.2%</sub>
	Gemma 2 27B	20	<b>7</b> <sub>↓65.0%</sub>	37	<b>22</b> <sub>↓40.5%</sub>	71	<b>32</b> <sub>↓54.9%</sub>
	DS Coder 33B	32	<b>19</b> <sub>↓40.6%</sub>	29	<b>13</b> <sub>↓55.2%</sub>	90	<b>43</b> <sub>↓52.2%</sub>
	CodeLlama 34B	80	<b>41</b> <sub>↓48.8%</sub>	126	<b>54</b> <sub>↓57.1%</sub>	157	<b>76</b> <sub>↓51.6%</sub>
	Qwen2.5 32B	19	<b>13</b> <sub>↓31.6%</sub>	22	<b>16</b> <sub>↓27.3%</sub>	55	<b>29</b> <sub>↓47.3%</sub>

# Improvement of functional correctness

	Model	Synthesis		Translation		Repair	
		Vanilla	Types	Vanilla	Types	Vanilla	Types
HumanEval	Gemma 2 2B	29.1	<b>30.2</b> ↑3.8%	50.2	<b>53.9</b> ↑7.5%	11.6	<b>20.9</b> ↑79.4%
	Gemma 2 9B	56.6	<b>58.3</b> ↑3.1%	73.7	<b>78.3</b> ↑6.2%	24.0	<b>34.9</b> ↑45.7%
	Gemma 2 27B	69.5	<b>71.2</b> ↑2.5%	86.6	<b>87.7</b> ↑1.3%	38.4	<b>41.1</b> ↑7.1%
	DS Coder 33B	68.9	<b>71.1</b> ↑3.2%	88.7	<b>90.1</b> ↑1.6%	47.6	<b>50.7</b> ↑6.5%
	CodeLlama 34B	41.0	<b>43.4</b> ↑5.7%	58.6	<b>63.5</b> ↑8.3%	17.5	<b>27.4</b> ↑56.9%
	Qwen2.5 32B	79.6	<b>81.8</b> ↑2.8%	92.1	<b>93.9</b> ↑1.9%	65.4	<b>71.2</b> ↑8.9%
MBPP	Gemma 2 2B	40.4	<b>42.4</b> ↑5.2%	52.3	<b>56.0</b> ↑7.0%	12.1	<b>22.6</b> ↑86.7%
	Gemma 2 9B	65.4	<b>67.4</b> ↑3.2%	71.4	<b>75.8</b> ↑6.2%	24.2	<b>31.9</b> ↑31.7%
	Gemma 2 27B	70.6	<b>72.1</b> ↑2.2%	83.1	<b>84.4</b> ↑1.6%	39.1	<b>45.2</b> ↑15.5%
	DS Coder 33B	65.4	<b>67.2</b> ↑2.8%	85.9	<b>89.1</b> ↑3.6%	35.1	<b>43.1</b> ↑23.0%
	CodeLlama 34B	42.2	<b>45.6</b> ↑8.0%	55.7	<b>63.3</b> ↑13.6%	15.7	<b>26.6</b> ↑69.2%
	Qwen2.5 32B	76.3	<b>76.6</b> ↑0.3%	89.6	<b>90.4</b> ↑0.9%	48.0	<b>54.0</b> ↑12.6%

# Future work

Extend to more TypeScript features

Construct incremental type-checker for other languages

Characterize the *kind* of type systems that *can* be incrementally verified

# Conclusion

## A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse

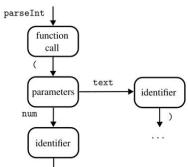
The screenshot shows a news article from Forbes. The headline is "How AI-Generated Code Is Unleashing A Tsunami Of Security Risks". Below the headline, there's a sub-headline "AI Code Hallucinations Increase the Risk of 'Package Confusion' Attacks". The text discusses a study showing that AI-generated code is more likely to contain made-up information that can be used to trick software into interacting with malicious code.

SOTA LLMs generate semantic errors and unsafe code

## Type-constrained code generation in LLMs: key ingredient

### An automata tracked during LLM decoding that only accepts type safe programs

Technically, the automata is an AST augmented with type relevant context



```
function is_int(text: string): boolean {
    const num = Number(text);
    return !isNaN(num) &&
        parseInt(num)
```

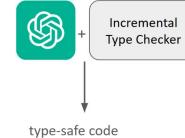
Type-Constrained Code Generation with Language Models

24

## Our work: Type-augmented LLM decoders

### Potential Benefits:

- Full control during decoding
- Generated code is guaranteed to be type safe
- Inference-time costs: predictable, one pass



Wanted: new **incremental** type checkers for LLM decoding

## Drastic reduction of compiler errors

Model	Synthesis		Translation		Repair		
	Vanilla	Types	Vanilla	Types	Vanilla	Types	
HumanEval	Gemma 2 2B	103	44 <sub>57.3%</sub>	177	80 <sub>54.8%</sub>	194	103 <sub>46.9%</sub>
	Gemma 2 9B	45	13 <sub>71.1%</sub>	75	16 <sub>78.7%</sub>	113	52 <sub>54.0%</sub>
	Gemma 2 27B	15	2 <sub>86.7%</sub>	20	3 <sub>85.0%</sub>	45	22 <sub>51.1%</sub>
	DS Coder 33B	26	5 <sub>80.8%</sub>	18	7 <sub>61.1%</sub>	36	15 <sub>58.3%</sub>
	CodeLlama 34B	86	28 <sub>67.4%</sub>	158	59 <sub>62.7%</sub>	153	48 <sub>68.6%</sub>
	Qwen2.5 32B	17	2 <sub>88.2%</sub>	24	5 <sub>79.2%</sub>	36	13 <sub>63.9%</sub>
MBPP	Gemma 2 2B	67	27 <sub>59.7%</sub>	126	79 <sub>37.3%</sub>	194	108 <sub>44.3%</sub>
	Gemma 2 9B	30	10 <sub>66.7%</sub>	67	33 <sub>50.7%</sub>	129	63 <sub>51.2%</sub>
	Gemma 2 27B	20	7 <sub>65.0%</sub>	37	22 <sub>40.5%</sub>	71	32 <sub>54.9%</sub>
	DS Coder 33B	32	19 <sub>40.6%</sub>	29	13 <sub>55.2%</sub>	90	43 <sub>52.2%</sub>
	CodeLlama 34B	80	41 <sub>48.8%</sub>	126	54 <sub>57.1%</sub>	157	76 <sub>51.6%</sub>
	Qwen2.5 32B	19	13 <sub>31.6%</sub>	22	16 <sub>27.3%</sub>	55	29 <sub>47.3%</sub>

<https://github.com/eth-sri/type-constrained-code-generation>

