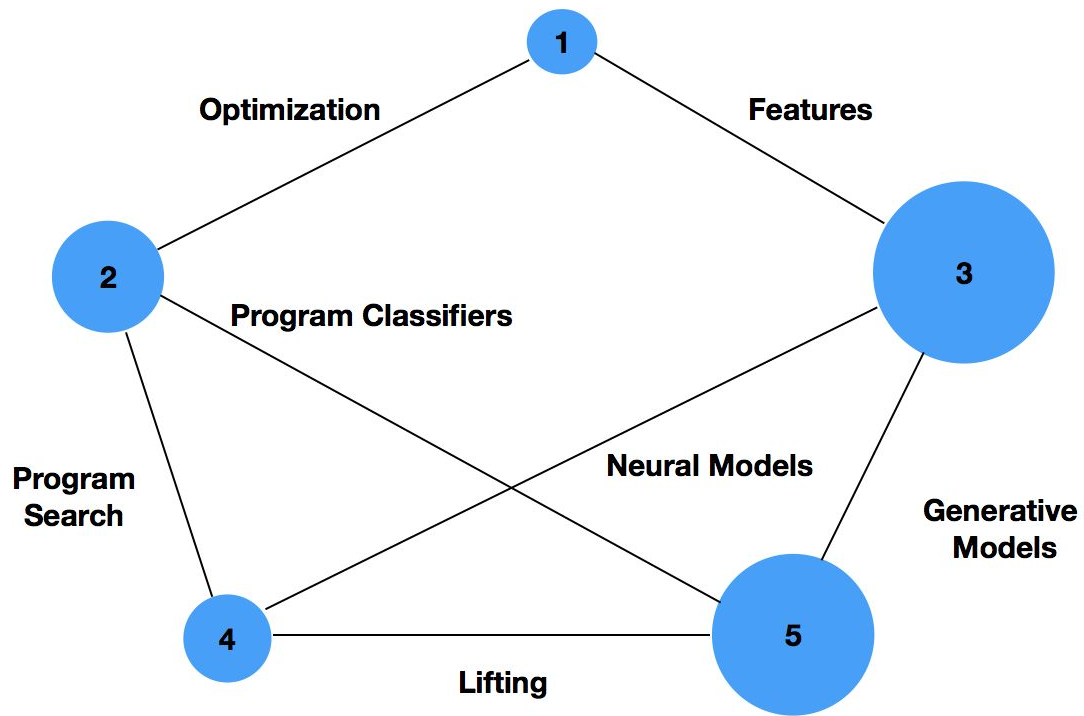


Rethinking Compilation: L4



Agenda

1. What is program synthesis?
2. Synthesis in compilers context
3. Fundamental techniques
4. Pruning strategies
5. The future
6. Bibliography

What is program synthesis?

What is program synthesis?

- Synthesis is **NOT**:
 - Compilation
 - Machine Learning
- According to Solar-Lezama [1]


“Program Synthesis correspond to a class of techniques that are able to generate a program from a collection of artifacts that establish semantic and syntactic requirements for the generated code. ”

What is program synthesis?


- WHERE do we want to automatically generate software?



Data wrangling



Coding assistance




Program Lifting

Program \longrightarrow x86 \longrightarrow Hardware

Program \longrightarrow OpenCL \longrightarrow Hardware

Program \longrightarrow cBLAS \longrightarrow Hardware
 \longrightarrow Halide \longrightarrow Hardware
 \longrightarrow

LIFT code to API or DSL




Vendor responsibility to map API/DSL to hardware - already the case

Our job - automatically lift it to API/DSL enabling hardware utilisation

What is program synthesis?

- WHY IS IT HARD to automatically generate software?
- Writing software:



Demands expertise



Search space is enormous



Hard to ensure
correctness

What is program synthesis?

*“Program Synthesis correspond to a class of techniques that are able to generate a program from a collection of artifacts that establish **semantic** and **syntactic** requirements for the generated code. ”*

Input/Output exs

```
in: [1,2,3,4], [2,5,7,8]  
out: [5,11,16,20]
```

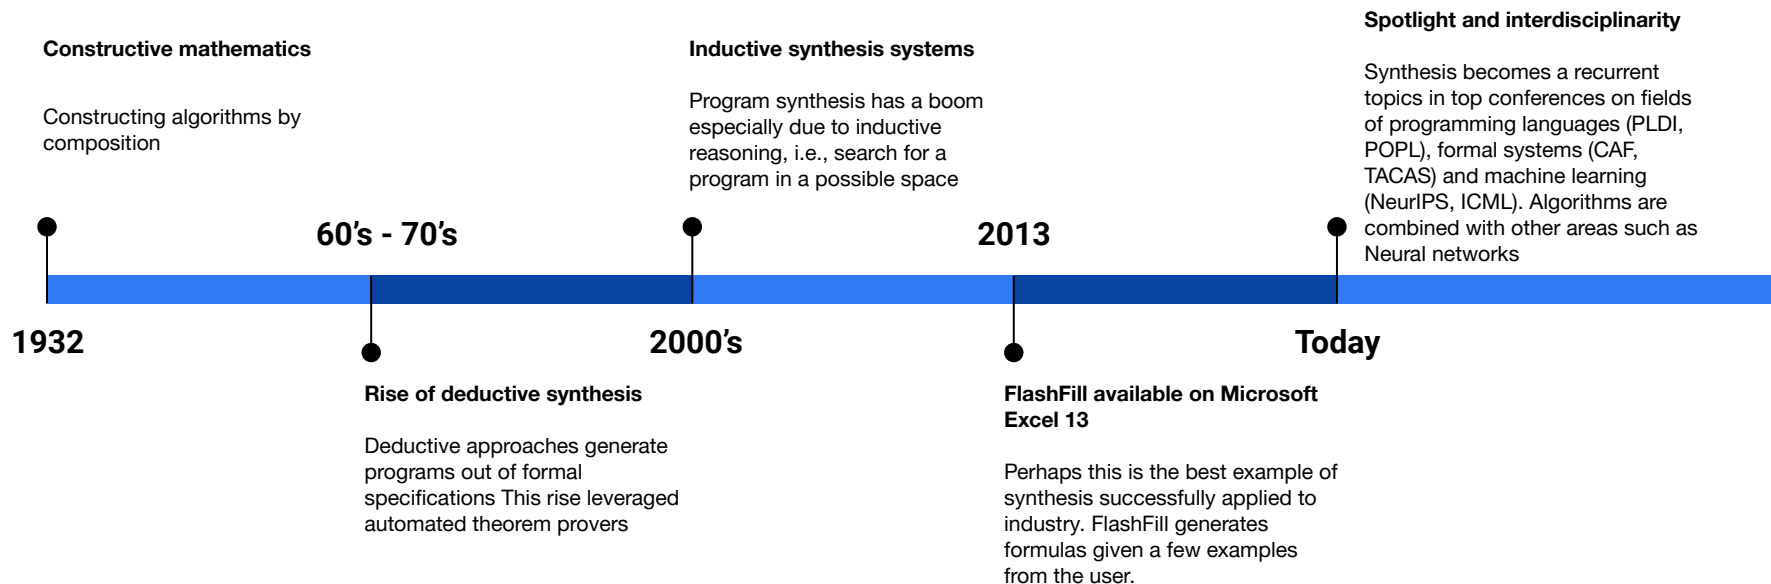
Logic specification

```
post(a, b)  $\equiv \forall$  imin+1  $\leq$  i  $\leq$  imax,  
jmin  $\leq$  j  $\leq$  jmax.  
a(i,j) = b(i-1,j) + b(i,j)
```

Grammar

```
<prog> ::= <var> = <expr>  
<exp> ::= <exp> + <exp>  
<exp> ::= <exp> * <exp>  
<exp> ::= ( <exp> )  
<var> ::= a | b | c
```

What is program synthesis?



Synthesis in compilers context

Synthesis in compiler context

- High-level code translation

```
def max(a,b):  
    return a if a > b else b
```




```
int max(int a, int b)  
{  
    return a > b ? a : b;  
}
```

Synthesis in compiler context

- High-level code translation for:

Security:



Performance:

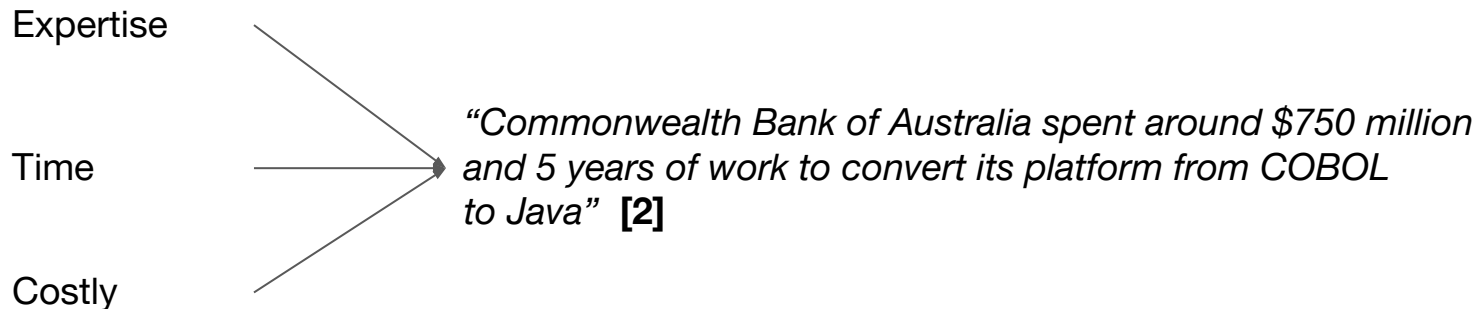


Maintenance:



Synthesis in compiler context


- High-level code translation involves



Synthesis can automate tasks

Synthesis in compiler context


- Optimizing Database-Backed Applications with Query Synthesis (PLDI' 2013)
[3]



- Synthesis from verification conditions
- Performance obtained

Synthesis in compiler context


- Decompile





Synthesis in compiler context

- Decompile:

Malware detection:




Port code to
new devices:



Synthesis in compiler context

- Helium: Lifting High-Performance Stencil Kernels from Stripped x86 Binaries to Halide DSL Code (PLDI' 2015) [4]




```
RDom r_0(input);  
output(input(r_0.x,r_0.y)) =  
cast<uint64_t>(output(input(r_0.x,r_0.y)) + 1);
```

- Synthesis leveraging binary traces
- 75% performance on Adobe PhotoShop

Synthesis in compiler context

- Api migration




TensorFlow



- Maintenance
- Usability

Synthesis in compiler context

- SOAR: A Synthesis Approach for Data Science API Refactoring (ICSE' 2021) [5]




- Synthesis with NLP

Fundamental techniques

Fundamental techniques - Bottom-up

- Bottom-up search
 - Starts from small constructs
 - Combine them into larger programs

```
if(a > cond) {  
    a = b + 2  
}else{  
    b = b + a  
}
```



Bottom-up

Fundamental techniques - Bottom-up

Algorithm Bottom-up Enumerative Synthesis

```
1: function BOTTOM-UP(grammar, specification)
2:   c_list  $\leftarrow$  TERMINALS(grammar)
3:   while true do
4:     for  $r \in$  RULES(grammar) do
5:       c_list  $\leftarrow$  EXPAND(c_list, r)
6:     end for
7:     c_list  $\leftarrow$  ELIM_EQUIVALENTS(c_list)
8:     for  $c \in$  c_list do
9:       if IS_CORRECT(c, specification) then
10:        return c
11:      end if
12:    end for
13:  end while
14: end function
```

Fundamental techniques - Bottom-up

- Find the maximum of two numbers

```
<program> ::= <expr>
<expr> ::= <expr> + <expr>
          | <expr> > <expr>
          | if <expr> then <expr> else
<expr>
          | <var>
          | 1
<var> ::= x | y
```

iteration

1

x

y

1


Fundamental techniques - Bottom-up

- Find the maximum of two numbers


```
<program> ::= <expr>
<expr> ::= <expr> + <expr>
        | <expr> > <expr>
        | if <expr> then <expr> else
<expr>
        | <var>
        | 1
<var> ::= x | y
```

iteration

1



2




Fundamental techniques - Bottom-up

- Find the maximum of two numbers


```
<program> ::= <expr>
<expr> ::= <expr> + <expr>
          | <expr> > <expr>
          | if <expr> then <expr> else
<expr>
          | <var>
          | 1
<var> ::= x | y
```

iteration

1



2



3




Fundamental techniques - Bottom-up

- Pros
 - Simple
 - Explores small programs first
 - Black-box language building
- Cons
 - Scalability
 - Poor performance for synthesizing constants

Fundamental techniques - Top-down

- Top-down search
 - Starts from high-level constructs
 - Fill the holes in temporary candidates

```
if(a > cond) {  
    a = b + 2  
}else{  
    b = b + a  
}
```



Fundamental techniques - Top-down

Algorithm Top-down Enumerative Synthesis

```
1: function TOP-DOWN(grammar, specification)
2:   c_list  $\leftarrow \emptyset$ 
3:   r_list  $\leftarrow \text{RULES}(\textit{grammar})$ 
4:   while true do
5:     for r  $\in$  r_list do
6:       for t  $\in$   $\text{TERMINALS}(\textit{grammar})$  do
7:         if r can expand to t then
8:           c_list  $\leftarrow \text{EXPAND}(\textit{c\_list}, t)$ 
9:         end if
10:      end for
11:      for r'  $\in$   $\text{RULES}(\textit{grammar})$  do
12:        if r can expand to r' then
13:          r_list  $\leftarrow \text{EXPAND}(\textit{r\_list}, r')$ 
14:        end if
15:      end for
16:    end for
17:    c_list  $\leftarrow \text{ELIM\_EQUIVALENTS}(\textit{c\_list})$ 
18:    for c  $\in$  c_list do
19:      if IS_CORRECT(c, specification) then
20:        return c
21:      end if
22:    end for
23:  end while
24: end function
```

Fundamental techniques - Top-down

- Find the maximum of two numbers

```
<program> ::= <expr>
<expr> ::= <expr> + <expr>
          | <expr> > <expr>
          | if <expr> then <expr> else
<expr>
          | <var>
          | 1
<var> ::= x | y
```

iteration

1

if <expr> then <expr> else
 <expr>

<expr> + <expr>

...

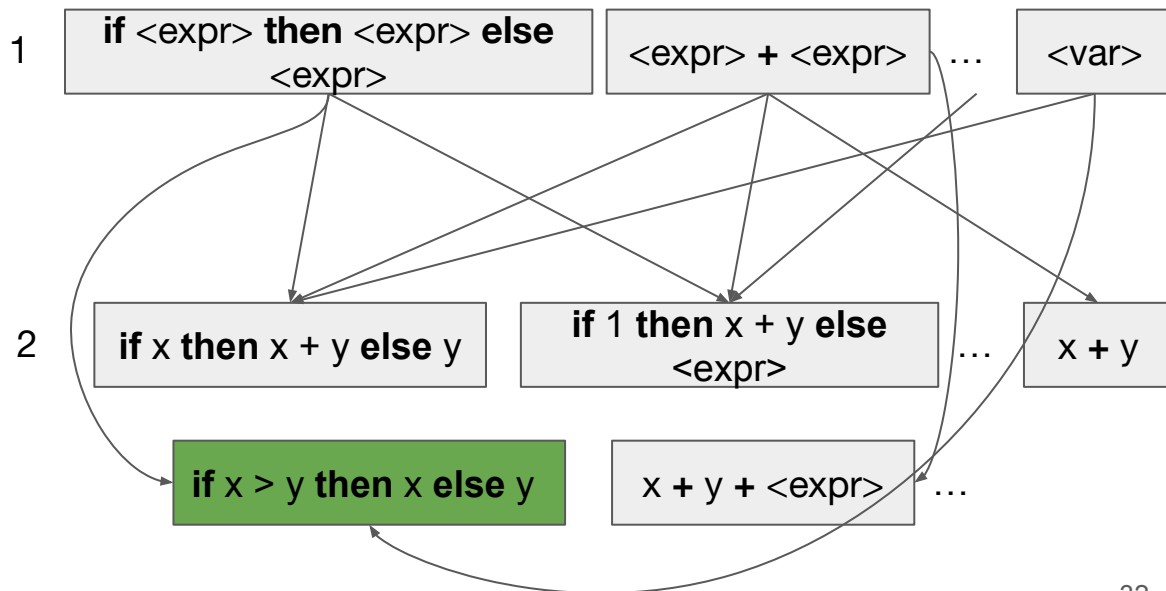
<var>

Fundamental techniques - Top-down

- Find the maximum of two numbers

```
<program> ::= <expr>
<expr> ::= <expr> + <expr>
          | <expr> > <expr>
          | if <expr> then <expr> else
<expr>
          | <var>
          | 1
<var> ::= x | y
```

iteration



Fundamental techniques - Top-down

- Pros


- Efficient with a good initial guess
- Effective when dealing with recursion

- Cons

- It is hard to execute temporary programs
- Scalability

Fundamental techniques - CEGIS

- Counterexample Guided Inductive Synthesis [6]
- Synthesizer + Verifier
- Feedback to synthesizer



Fundamental techniques - Verification conditions

- Convert a program into a logical statement:
 - $\forall x.Q(x) \Rightarrow$ program is correct
- Verification conditions (vc's) can be used to synthesize programs
- Given a verification condition, the goal is to synthesize a predicate that makes such a verification valid
- Usually, VC's contain a *loop invariant*
- A loop invariant is a formula that must hold in every iteration of a loop

Fundamental techniques - Verification conditions

- Verified Lifting of Stencil Computations (PLDI' 2015) [7]
- The goal is to synthesize a verification condition that summarizes the computation above.

```
procedure sten(imin,imax,jmin,jmax,a,b)
  real (kind=8), dimension(imin:imax,jmin:jmax) :: a
  real (kind=8), dimension(imin:imax,jmin:jmax) :: b
  do j=jmin,jmax
    t = b(imin, j)
    do i=imin+1,imax
      q = b(i,j)
      a(i,j) = q + t
      t = q
    enddo
  enddo
end procedure
```

(a)

Fundamental techniques - Verification conditions

- Verification condition is formed by:
 - Pre-condition
 - Loop invariant
 - Pos-condition
- Such that:
 - $\forall s. pre(s) \rightarrow invariant(s)$
 - $\forall s. invariant(s) \wedge cond(s) \rightarrow invariant(body(s))$
 - $\forall s. invariant(s) \wedge \neg cond(s) \rightarrow post(s)$

Fundamental techniques - Verification conditions

- A pre-condition can be obtained by syntax-guided algorithms
 - Pre-condition (outermost loop)

$$I_j(a, b, j_{\min})$$

- Loop invariant and post-condition must be synthesized
- Verified lifting uses CEGIS with a SMT solver as verifier to find the last two
- Invariant:
 - Post-condition (outermost loop):

$$\begin{aligned} invariant(a, b, j) \equiv & j \leq j_{\max} + 1 \wedge \\ & \forall i_{\min} + 1 \leq i \leq i_{\max}, j_{\min} \leq j' < j. \\ & a(i, j') = b(i-1, j') + b(i, j') \end{aligned}$$

(c)

$$\begin{aligned} post(a, b) \equiv & \forall i_{\min} + 1 \leq i \leq i_{\max}, j_{\min} \leq j \leq j_{\max}. \\ & a(i, j) = b(i-1, j) + b(i, j) \end{aligned}$$

(b)

Fundamental techniques - Verification conditions

- Once loop invariant and post-condition are synthesized, SMT proves that the VC holds
- The VC is then translated to the target language

```
int main() {  
    ImageParam b(type_of<double>(),2);  
    Func func; Var i, j;  
    func(i,j) = b(i-1,j) + b(i,j);  
    func.compile_to_file("ex1", b);  
    return 0; }
```

(d)

Fundamental techniques - Verification conditions

- Synthesized code is faster:

Benchmark	Kernel	Halide Speedup	icc Before Speedup	icc After Speedup	Halide GPU Speedup	Halide GPU Speedup (no transfer)
	akl81	7.44	1.73	4.53	4.09	10.88
	akl83	4.57	0.95	0.95	3.29	8.27
	akl84	4.51	0.71	0.94	2.94	7.88
	akl85	4.05	0.95	0.77	2.68	6.85
	akl86	4.04	0.97	0.79	2.27	6.00
	ackl95	4.19	1.00	1.00	1.94	7.47
	amkl100	3.84	0.93	0.89	1.57	6.72
	amkl101	3.64	1.01	1.00	1.52	5.44
	amkl103	3.42	0.93	0.93	1.99	6.77
	amkl105	3.37	0.98	0.98	1.50	5.17
	amkl107	3.95	0.94	0.96	2.31	8.71
	amkl97	4.19	1.00	1.01	1.96	6.98


Fundamental techniques - Verification conditions

- Limitations of verified lifting:
 - Conditional statements
 - Loop invariables that increase non-monotonically
 - Demands a lot of user specification

Pruning Strategies


Pruning strategies

- Problem: large search space
- Brute force is unfeasible



Pruning strategies

- Solution: pruning search space
- Several techniques
- Focus on 3:
 - Type-directed
 - Sketch
 - Neural guidance



Pruning strategies

- Type-directed synthesis
- Use type checking to discard candidates
- Synthesis looks for compatible programs

Pruning strategies

- TF-Coder: Program Synthesis for Tensor Manipulations (ACM Transactions on Programming Languages 2021) [8]

I/O Example  **TF-Coder**  TensorFlow Code

Input tensor(s):

```
in1 = [[0, 1, 0, 0],  
       [0, 1, 1, 0],  
       [1, 1, 1, 1]]
```

Output tensor:

```
output = [[0.0, 1.0, 0.0, 0.0],  
          [0.0, 0.5, 0.5, 0.0],  
          [0.25, 0.25, 0.25, 0.25]]
```

Natural language description (*optional*):

“Normalize the rows of a tensor”

Scalar constants (*optional*):


(none needed for this problem)

Solution (found in *1.3 seconds*):

```
output = tf.cast(  
    tf.divide(in1, tf.expand_dims(  
        tf.reduce_sum(in1, axis=1),  
        axis=1)),  
    tf.float32)
```

Pruning strategies

- Determine the arguments of TensorFlow API calls
- Type mismatch and compatibility among arguments themselves



Pruning strategies

- Sketch
- Provide a *draft* of solution
 - Given by user
 - Learnt some preprocessing step
- Synthesizer needs to fill the draft

Pruning strategies

- Program synthesis by sketching. [9]
- Sketch basics:
 - Unknown constants
 - Test harness
 - Generator functions

Pruning strategies

- Program synthesis by sketching. [9]
- Consider the problem of transposing a 5x5 matrix
- Sketch:

```
int[25] transpose5x5(int[25] mat){  
    int[25] out;  
    for(int i=0; i<5; ++i) for(int j=0; j<5; ++j){  
        out[ ??*i + ??*j + ??] = mat[??*i + ??*j + ??];  
    }  
    return out;  
}
```

Pruning strategies

- Expression $??*i + ??*j + ??$ is repeated twice
- Generator encompasses the set of possible linear algebra expressions involving i and j

```
generator int legen(int i, int j){  
    return ??*i + ??*j + ??;  
}  
  
int[25] transpose5x5(int[25] mat){  
    int[25] out;  
    for(int i=0; i<5; ++i) for(int j=0; j<5; ++j){  
        out[ legen(i,j) ] = mat[ legen(i,j) ];  
    }  
    return mat;  
}
```


Pruning strategies

- Each call to the generator will resolve to a different expression, resulting in a correct implementation

```
int[25] transpose5x5(int[25] mat){  
    int[25] out;  
    for(int i=0; i<5; ++i) for(int j=0; j<5; ++j){  
        out[ 5*i + j ] = mat[ i + 5*j ];  
    }  
    return mat;  
}
```

Pruning strategies


- Neural-guided synthesis



- Use of neural networks to help synthesis

Pruning strategies


- Neural-guided synthesis



- Neural network outputs a probability distribution
- Indicates likelihood of tokens in target language

Pruning strategies

- Neural-guided synthesis



- Distribution leads the search

Pruning strategies

- Deepcoder: Learn to Write Programs (ICLR 2017) [10]
- Maps IO examples to program properties

```
a ← [int]
b ← FILTER (<0) a
c ← MAP (*4) b
d ← SORT c
e ← REVERSE d
```

An input-output example:

Input:

`[-17, -3, 4, 11, 0, -5, -9, 13, 6, 6, -8, 11]`

Output:

`[-12, -20, -32, -36, -68]`

Pruning strategies

```
a ← [int]
b ← FILTER (<0) a
c ← MAP (*4) b
d ← SORT c
e ← REVERSE d
```

An input-output example:

Input:

`[-17, -3, 4, 11, 0, -5, -9, 13, 6, 6, -8, 11]`

Output:

`[-12, -20, -32, -36, -68]`

NN

(+1)	(-1)	(*2)	(/2)	(*1)	(**2)	(*3)	(/3)	(*4)	(/4)	(>0)	(>0)	(%2==1)	(%2==0)	HEAD	LAST	MAP	FILTER	SORT	REVERSE	TAKE	DROP	ACCESS	ZIPWITH	SCANL1	+	.	*	MIN	MAX	COUNT	MINIMUM	MAXIMUM	SUM
.0	.0	.1	.0	.0	.0	.0	.0	1.0	.0	.0	1.0	.0	.2	.0	.0	1.0	1.0	1.0	.7	.0	.1	.0	.4	.0	.0	.1	.0	.2	.1	.0	.0	.0	.0

Pruning Strategies

- Deepcoder uses Depth-First Search
- Extends programs based on the probabilities
- Speedup on search

Table 1: Search speedups on programs of length $T = 3$ due to using neural network predictions.

Timeout needed to solve	DFS			Enumeration			λ^2			Sketch		Beam
	20%	40%	60%	20%	40%	60%	20%	40%	60%	20%	40%	20%
Baseline	41ms	126ms	314ms	80ms	335ms	861ms	18.9s	49.6s	84.2s	$>10^3s$	$>10^3s$	$>10^3s$
DeepCoder	2.7ms	33ms	110ms	1.3ms	6.1ms	27ms	0.23s	0.52s	13.5s	2.13s	455s	292s
Speedup	15.2×	3.9×	2.9×	62.2×	54.6×	31.5×	80.4×	94.6×	6.2×	$>467\times$	$>2.2\times$	$>3.4\times$

Pruning Strategies

- Limitations of neural-guided synthesis
 - Performance of synthesizer depends on the model
 - Therefore, it depends on the quality of training data
 - Languages targeted are still domain-specific/small
 - Need to find the most relevant features

The future

The future

- Sketching = how to find the best initial sketch?
- Machine Learning + Synthesis
- Loop invariants
- How to strength verification?
- How to generate good specifications (IO)?

Summary

- Program synthesis emerge as a technique to generate programs automatically with more control to the user
- As computers evolve, the inductive approach has risen as the main synthesis method
- Formal verification and machine learning incorporated into synthesis
- Challenges for the future include finding the best starting point and handle more complex languages and domains

Overview

- L1: Motivation and brief survey of auto-tuning/machine learning for compilers
- L2: Program rewriting schemes - e-graphs and equality saturation
- L3: Program embeddings and Graph Neural Networks
- This lecture: Program synthesis and neural synthesis
- Next L5: Neural Machine Translation, Transformers and Large language models

Bibliography

Survey:

- GULWANI, Sumit et al: **Program synthesis**. Foundations and Trends® in Programming Languages, v. 4, n. 1-2, p. 1-119, 2017.

Papers

1. Armando Solar-Lezama: **Introduction to Program Synthesis**. Lecture Notes. 2018
2. Baptiste Rozière, Marie-Anne Lachaux, Lowik Chanussot, Guillaume Lample: **Unsupervised Translation of Programming Languages**. NeurIPS 2020
3. Alvin Cheung, Armando Solar-Lezama, Samuel Madden: **Optimizing Database-Backed Applications with Query Synthesis**. PLDI 2013
4. Charith Mendis, Jeffrey Bosboom, Kevin Wu, Shoaib Kamil, Jonathan Ragan-Kelley, Sylvain Paris, Qin Zhao, Saman P. Amarasinghe: **Helium: lifting high-performance stencil kernels from stripped x86 binaries to halide DSL code**. PLDI 2015
5. Ansong Ni, Daniel Ramos, Aidan Z. H. Yang, Inês Lynce, Vasco M. Manquinho, Ruben Martins, Claire Le Goues: **SOAR: A Synthesis Approach for Data Science API Refactoring**. ICSE 2021
6. Armando Solar-Lezama, Liviu Tancau, Rastislav Bodík, Sanjit A. Seshia, Vijay A. Saraswat: **Combinatorial sketching for finite programs**. ASPLOS 2006

Bibliography

7. Shoaib Kamil, Alvin Cheung, Shachar Itzhaky, Armando Solar-Lezama: **Verified lifting of stencil computations**. PLDI 2016
8. Kensen Shi, David Bieber, Rishabh Singh: **TF-Coder: Program Synthesis for Tensor Manipulations**. ACM Trans. Program. Lang. Syst. 44(2): 10:1-10:36 (2022)
9. Armando Solar-Lezama: **Program Synthesis By Sketching**. Ph.D. Thesis. University of California, Berkeley 2008
10. Matej Balog, Alexander L. Gaunt, Marc Brockschmidt, Sebastian Nowozin, Daniel Tarlow: **DeepCoder: Learning to Write Programs**. ICLR 2017