CS 252R **DeepCoder** Learning To Write Programs

Eric Zhang – October 15th, 2020

Inductive Program Synthesis

"Inductive" = Inference

Input-output example: <i>Input</i> : 2, [3 5 4 7 5] <i>Output</i> : [7]	Description: A new shop near you is selling n paintings. You have $k < n$ friends and you would like to buy each of your friends a painting from the shop. Return the minimal amount of money you will need to spend.
Input: [6 2 4 7 9], [5 3 6 1 0] Output: 27	Description: In soccer leagues, match winners are awarded 3 points, losers 0 points, and both teams get 1 point in the case of a tie. Com- pute the number of points awarded to the winner of a league given two arrays w, t of the same length, where $w[i]$ (resp. $t[i]$) is the number of times team i won (resp. tied).
Input-output example: <i>Input</i> : [6 2 4 7 9], [5 3 2 1 0] <i>Output</i> : 4	Description: Alice and Bob are comparing their results in a recent exam. Given their marks per ques- tion as two arrays a and b , count on how many questions Alice got more points than Bob.



Neural Networks + Synthesis End-to-end vs. integrated

- **End-to-end**
 - Neuro-Symbolic Program Synthesis (R3NN)
- Integrated
 - Using deep learning to give hints to existing domain-specific search techniques
 - Train neural network on a dataset

Machine translation with some knowledge of the context-free grammar



Learning-Augmented Algorithms Background topic

- **Idea:** take existing algorithms, add machine learning as a black-box, and improve the performance assuming competence of the learned model.
- The best learning-augmented algorithms don't even decrease in worst-case performance when the ML model outputs bad predictions!
- Example: learning-augmented binary search.
 - Binary search takes $O(\log n)$ time, where the size of the array is n.
 - We can start with an *initial guess* from a neural network for the true index, and if this is within *t* from the actual index, then we can get $O(\log t)$ time!

Augmented Alg. **Learned Bloom filter**

- Particularly nice because Bloom filters can get false positives anyway, so we're already assuming some error.
- Even simple neural networks can greatly reduce the storage costs.
- "The Case for Learned Index Structures" (Google, 2018).



LIPS "Learned Inductive Program Synthesis"

Predict Attributes (black-box machine learning) $\Pr(a \mid E) = f(E)$ f is a learned map!

Guide Search

(PL algorithms)

 $\sim \Pr(P \mid a)$

Part I: Machine Learning

Attributes Summarizing P

- DSL has C = 34 functions.
- Program *P* is represented by an attribute vector $a \in \mathbb{R}^{C}$.
- $\mathscr{A}: \mathscr{P} \to \mathbb{R}^C$ attribute function simply gives a 0-1 vector detecting *presence* for each function in the program.



High-Level Attributes ...may be reflected in the input data!

An input-output example: Input: Output: [-12, -20, -32, -36, -68]

- Outputs are sorted => might have the sort, reverse functions.
- Goal: model this human intuition.

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

• All the outputs are even => might have the map, filter, (*n) functions.

Generating Examples I/O examples are represented by the set *E*







Training the Neural Network Deep Learning = Black-Box Function Estimator





```
"examples": [
   "inputs": [
     [-46, -23, -78, 10],
      [125, 105, -69]
    ],
   "output": 82
  },
    "inputs": [
       48, -117, 79, -42, -43, 37, -96],
     [13, -52, 48, 6, -8, -55, 35, 75]
    ],
   "output": null
  },
 // ...
"attribute": [
 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 0, 0, 0
```

"program": "LIST|LIST|COUNT,>0,0|ZIPWITH,+,1,0|ACCESS,2,3",

[90, 103, -57, 13, -45, 28, -30, 68, -113, 60, -71,

0, 0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 1, 0, 1, 0, 0, 0, 0,

	(+1)	(-1)	(*2)	(/3)	(*-1)	(**2)	(£*)	(٤/)	(*4)	(/4)	(0<)	(0>)	(%2==1)	(%2==0)	HEAD	LAST	MAP	FILTER	SORT	REVERSE	TAKE	DROP	ACCESS	ZIPWITH	SCANL1	+		¥	MIN	MAX	COUNT	MINIMUM
1	.0	.2	.0	.1	.4	.0	.0	.2	.0	.1	.0	.2	.1	.0	.1	.0	.3	.4	.2	.1	.5	.2	.2	.6	.5	.2	.4	.0	.9	.1	.0	.1
1	.1	.1	.1	.1	.0	.0	1.0	.0	.1	.0	.2	.1	.1	.1	.0	.3	1.0	.2	.1	.1	.0	.0	.1	1.0	.0	.6	.6	.0	.1	.1	.2	.0
:	.1	.2	.0	.1	.0	.0	.0	.1	.0	.1	.2	.2	.3	.3	.0	.0	.6	.0	.1	.1	.0	.0	.0	1.0	.3	.4	.5	.0	.5	.5	1.0	.0
1	.3	.1	.1	.1	.1	.0	.0	.0	.0	.0	.1	.0	.0	.0	.0	.0	.6	.2	.1	.1	.0	.0	.0	1.0	.3	.3	.3	.1	.2	.7	.0	.0
f	.0	.0	.1	.4	.1	.4	.0	.0	.2	.0	.0	.2	.0	.2	.1	.2	.9	.2	.1	.0	.0	.0	.4	.6	.2	.2	.3	.3	.4	.1	.2	.4
,	.2	.2	.0	.2	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.0	.9	.0	.0	1.0	.0	.0	.0	1.0	.0	.2	.0	.0	1.0	.1	.0	.0
2	.1	.1	.0	.0	.0	.0	.0	.0	.0	.0	.0	.2	.2	.2	.7	.0	.3	.3	.1	.0	.0	.0	.0	1.0	.1	.9	.1	.0	.7	.2	.1	.8
1	.0	.0	.0	.0	.0	.1	.0	.0	.1	.0	.1	.1	.1	.1	.0	.1	.4	.1	.0	.0	.0	.0	.0	1.0	.8	.5	.4	1.0	.1	.0	.2	.0
1	.2	.1	.0	.1	.1	.0	.0	.1	.0	.1	.1	.1	.1	.0	.0	.0	.5	.5	.1	.0	.0	.0	.0	1.0	.4	.4	.5	.0	.3	.6	.0	.0

1: MAP (*3) a | ZIPWITH + b c | MAXIMUM d

2: ZIPWITH - b a | COUNT (>0) c

3: SCANL1 MIN a | ZIPWITH - a b | FILTER (>0) c | SUM d

4: SORT a | SORT b | REVERSE d | ZIPWITH * d e | SUM f

5: REVERSE a | ZIPWITH MIN a b

6: MAP (-1) a | MAP (-1) b | ZIPWITH + c d | MINIMUM e

7: SCANL1 + b | ZIPWITH * a c | SUM d

8: REVERSE a | ZIPWITH - b a | FILTER (>0) c | SUM d



"The key in the LIPS formulation is to ensure that it is feasible to generate a large dataset (ideally millions of programs)."

https://arxiv.org/pdf/1611.01989.pdf, p.3

Menon et al., 2013 A Machine Learning Framework for Programming By Example

(a) Sample of clues used. LIST denotes a list-, E a stringnonterminal in the grammar.

Suggested rule(s) Feature Substring s appears in out- $E \rightarrow \text{``s''}, \text{LIST} \rightarrow \{E\}$ put but not input? Duplicates in input but not $LIST \rightarrow dedup(LIST)$ output? Numbers on each input line $LIST \rightarrow count(LIST)$ but not output line?

Dataset Generation How to pick reasonable input-output pairs

- "We enforce a constraint on the output value bounding integers to some predetermined range, and then propagate these constraints backward through the program to obtain a range of valid values for each input."
- program as a whole?

• Issues: maybe a couple random samples could be unrepresentative of the

Architecture Details of the neural network

- We want to learn a representative map $f: E \mapsto a$.
- Loss function: mean binary cross-entropy.
- Network architecture: fully connected, 3-layer. (Aside: Why not recurrent?)
- Embed each word in the input/output into a vector, combination of one-hot encoding for types, with learned embeddings for integers like "2", "5", "42", up until some maximum representable value.
- Each I/O pair is mapped to an attribute vector, and these are averaged.

Architecture Schematic Time-Distributed Feed-Forward Network + Average Pooling



Learned Embedding for [-256, 255] **Dense embedding in a 20-dimensional space**



First embedding dimension $\phi_1(n)$



Thoughts: Extensions? Slightly less minimal LIPS

- Encoding I/O elements
 - How would you encode other data types? Embeddings.
 - One-hot encoding of arrays versus RNN/Transformer encoding.
 - (Note: authors tried using a GRU but couldn't get it to work well)
- More specific attribute vectors

 - Advantages? Disadvantages?

Context-sensitive: probability of each token with k-token lookbehind

Other Encoder Architectures From the DL community



https://arxiv.org/pdf/1706.05587.pdf



Attention https://arxiv.org/pdf/2004.01800v2.pdf



Attention https://arxiv.org/pdf/2005.10821.pdf



Specificity Where's the right learning interface?

- Divide between "perceptual" and "symbolic".
- **Perceptual:** view E (the set of examples), generate easily-interpretable attribute vector a. More specific attributes are harder to interpret.
 - Extreme case is when a is a complex latent vector with no easy
 - Where have we seen this before?
- the search algorithm wasn't designed necessarily for LIPS.

interpretation, and you just pass this into a neural translation algorithm.

Symbolic: plug in a into the search algorithm, which tends to be easy even if

Discussion **Part I: Machine Learning**

What are the advantages, disadvantages, and alternatives of:

- 1. Choice of attribute vectors.
- 2. Neural network (encoder) architecture.
- 3. Augmented algorithm versus E2E machine learning.

Part II: Search Algorithms

"We use the neural network's predictions to *augment* search techniques from the programming languages community..."

https://arxiv.org/pdf/1611.01989.pdf, p.1 (emphasis mine)

Depth-First Search With Iterative Deepening

- Search through all programs of length $\leq T$.
- Considered a solution if we execute it on M = 5 examples and all are \checkmark .
- C++ implementation runs at 3×10^6 programs / second.



"Sort and Add" **Active function set reduces** memory requirements



Sketch SMT-based synthesis tool, fills in "holes" in code

Synthesis | Published: 02 August 2012 Program sketching

Armando Solar-Lezama 🗠

International Journal on Software Tools for Technology Transfer 15, 475–495(2013) Cite this article

1051 Accesses | 73 Citations | Metrics

Abstract

Sketching is a synthesis methodology that aims to bridge the gap between a programmer's high-level insights about a problem and the computer's ability to manage low-level details. In sketching, the programmer uses a partial program, a sketch, to describe the desired implementation strategy, and leaves the low-level details of the implementation to an automated synthesis procedure. In order to generate an implementation from the programmer provided sketch, the synthesizer uses counterexample-guided inductive synthesis (CEGIS). Inductive synthesis refers to the process of generating candidate implementations from concrete examples of correct or incorrect behavior. CEGIS combines a SAT-based inductive synthesizer with an automated validation procedure, a bounded model-checker, that checks whether the candidate implementation produced by inductive synthesis is indeed correct and to produce new counterexamples. The result is a synthesis procedure that is able to handle complex problems from a variety of domains including ciphers, scientific programs, and even concurrent data-structures.



Enumerative search and deduction with a small library of functions

.lk00K0xc. 'kk;...;kWXc Synthesizin .NN, kMMo 'WMWx Hypotheses kMMk Hypothes ;dkc lWMΧ, .:loc. .OMWx. .okcdWMN, .oX0c. Memoization .xNk' '; .00 kMM0 . ¹. lmmn. Hashcons ta .c0l. .K0 lXWOddddddx0Md Hashcons ta ; MMM, ;kkkkkkkkkkkkk, Hashcons ta oMMM: Hashcons bu .ONWMM1 'XO.0MMo OMMx Signatures: ,Ko • xNc xMM0 dMM0 ;NK, .dNd. lMMX. . . ;XMo : MMM ' ,0. .NMMOlxd. dWNl lK0: ;KMNx. Runtime: 5.4s Found solution: fun b a -> take (reverse (sort (concat a))) b



Alternative Decoders Or "what didn't work for them"

- RNN decoder predicts things token by token + beam search
 - "We combined this... by initializing the RNN using the pooled final layer of the encoder."
 - (This might be a straw man; it's not really "search" at this point.)
 - R3NN paper has a similar, but more reasonable comparison.
- Q: How far did this RNN sequence model get? (length T)



RNN decoder: attributes → program



Discussion **Part II: Search Algorithms**

with machine learning predictions?

1. What other search algorithms would you like to see, augmented

Part III: Experiments

"These programs have been inspired by simple tasks appearing on *real programming competition websites.*"

https://arxiv.org/pdf/1611.01989.pdf, p.12 (emphasis mine)

F. Realistic Gameplay

(Source: Codeforces 1430F)

Recently you've discovered a new shooter. They say it has realistic game mechanics.

- formally, the condition $r_i \leq l_{i+1}$ holds. Take a look at the notes for the examples to understand the process better.

You are confident in yours and your character's skills so you can assume that aiming and shooting are instant and you need exactly one bullet to kill one monster. But reloading takes exactly 1 unit of time.

One of the realistic mechanics is a mechanic of reloading: when you reload you throw away the old magazine with all remaining bullets in it. That's why constant reloads may cost you excessive amounts of spent bullets.

You've taken a liking to this mechanic so now you are wondering: what is the minimum possible number of bullets you need to spend (both used and thrown) to exterminate all waves.

Note that you don't throw the remaining bullets away after eradicating all monsters, and you start with a full magazine.

Input

The first line contains two integers n and k ($1 \le n \le 2000$; $1 \le k \le 10^9$) — the number of waves and magazine size.

- time limit per test: 1 second
- memory limit per test: 256 megabytes
 - input: standard input
 - output: standard output

Your character has a gun with magazine size equal to k and should exterminate n waves of monsters. The i-th wave consists of a_i monsters and happens from the l_i -th moment of time up to the r_i -th moments of time. All a_i monsters spawn at moment l_i and you have to exterminate all of them before the moment r_i ends (you can kill monsters right at moment r_i). For every two consecutive waves, the second wave starts not earlier than the first wave ends (though the second wave can start at the same moment when the first wave ends)





Examples

input	
2 3 2 3 6	
3 4 3	
output	
9	
input	
2 5	
3 7 11	
10 12 15	

output

30

input
5 42
42 42 42
42 43 42
43 44 42
44 45 42
45 45 1
output
-1



Example Program Sum of min *k* elements



Input-output example: *Input*: 2, [3 5 4 7 5] *Output*: [7]



Example Program Given 3-point wins and 1-point ties, find the top-scoring team



Input: [6 2 4 7 9], [5 3 6 1 0] Output: 27

Example Program Given widths and heights, make rectangles of minimum total area

Program 4: $x \leftarrow [int]$ $y \leftarrow [int]$ $c \leftarrow Sort x$ $d \leftarrow SORT y$ $e \leftarrow Reverse d$ $f \leftarrow ZIPWITH (*) de$ $g \leftarrow SUM f$

Input-output example: Input: [7 3 8 2 5], [2 8 9 1 3] Output: 79



Figure 5: Number of test problems solved versus computation time.

DFS: using neural network

- DFS: using prior order
- L2: Sort and add using neural network
- L2: Sort and add in prior order
- Enumeration: Sort and add using neural network
- Enumeration: Sort and add in prior order
- Beam search
- Sketch: Sort and add using neural network
- Sketch: Sort and add in prior order

								10 ³	3				T_t	t	
Timeout needed		DFS		Eı	numera	λ^2	0								
to solve	20%	40%	60%	20%	40%	60%	20%	npa 10'	2		.8.		• 3		
Baseline DeepCoder	163s 24s	$\begin{array}{c} 2887s \\ 514s \end{array}$	6832s 2654s	$8181s \\ 9s$	$>10^4 s$ 264s	$>10^{4}s$ 4640s	463s 48s	ad 10 ³	0	6	· · · · · · · · · · · · · · · · · · ·	• • • • • • • • •		_	
Speedup	6.8 imes	5.6 imes	2.6 imes	907 ×	$>$ 37 \times	$> 2 \times$	9.6 imes	10	1	2	3	4	5	0	
			(a)					L	engt	h of t	est pr (b)	ogran	ns T_{tes}	t	

grams.

Figure 3: Search speedups on programs of length T = 5 and influence of length of training pro-



Discussion Part III: Experiments

- 1. What other experiments might you have wanted to see?
- 2. What was most surprising in the results?
- 3. What could be some practical applications for learned program synthesis, given these results?

https://arxiv.org/abs/1611.01989