



# CHALLENGES AND OPPORTUNITIES IN MACHINE PROGRAMMING (MP)

---

**Justin Gottschlich**

Principal Scientist & Director of Machine Programming Research,  
Intel Labs

**(Incomplete) Active Collaborators:** Maaz Ahmad, Todd Anderson, Saman Amarasinghe, Jim Baca, Regina Barzilay, Michael Carbin, Carlo Carino, Alvin Cheung, Pradeep Dubey, Kayvon Fatahalian, Henry Gabb, Craig Garland, Moh Haghghat, Mary Hall, Adam Herr, Jim Held, Roshni Iyer, Nilesh Jain, Tim Kraska, Brian Kroth, Insup Lee, Geoff Lowney, Shanto Mandal, Ryan Marcus, Tim Mattson, Abdullah Muzahid, Mayur Naik, Paul Petersen, Alex Ratner, Tharindu Rusira, Martin Rinard, Vivek Sarkar, Koushik Sen, Oleg Sokolsky, Armando Solar-Lezama, Julia Sukharina, Yizhou Sun, Joe Tarango, Nesime Tatbul, Josh B. Tenenbaum, Jesmin Tithi, Javier Turek, Rich Uhlig, Anand Venkat, Wei Wang, Jim Weimer, Markus Weimer, Fangke Ye, Shengtian Zhou ... and many others.

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

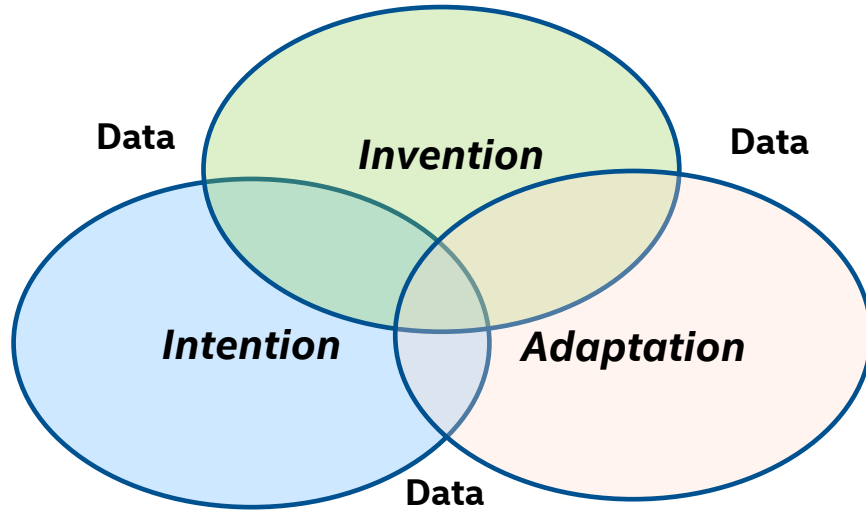
**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# SOME BACKGROUND

# THE THREE PILLARS OF MACHINE PROGRAMMING (MP)



- **MP is the automation of software development**

- **Intention:** Discover the intent of a programmer
- **Invention:** Create new algorithms and data structures
- **Adaption:** Evolve in a changing hardware/software world

Justin Gottschlich, Intel Labs  
Armando Solar-Lezama, MIT  
Nesime Tatbul, Intel Labs  
Michael Carbin, MIT  
Martin Rinard, MIT  
Regina Barzilay, MIT  
Saman Amarasinghe, MIT  
Joshua B Tenenbaum, MIT  
Tim Mattson, Intel Labs

Summarized ~90 works.

Key efforts by Berkeley,  
Google, Microsoft, MIT,  
Stanford, UW and others.

# WHY CALL IT MACHINE PROGRAMMING?

# WHY CALL IT MACHINE PROGRAMMING?

- **Alternatives:**
  - Program Synthesis
  - AI/ML for Code
  - Software 2.0

# WHY CALL IT MACHINE PROGRAMMING?

- **Alternatives:**

- Program Synthesis (**historically w/ formal methods**)
- AI/ML for Code (**it's not just AI/ML**)
- Software 2.0 (**what does this mean?**)
  - And Software 3.0, and 4.0, and 5.0?

- **Aim is to avoid confusion and broaden scope**

# MP USING APPROXIMATE AND PRECISE METHODS

## APPROXIMATE

### Machine Learning

(e.g., Neural Networks, Reinforcement Learning, Genetic Algorithms, Bayesian Networks, etc.)

## PRECISE

### Formal Methods

(e.g., Formal Verifiers, Spatial and Temporal Logics, Formal Program Synthesizers, etc.)

← *Progressively More Approximate* Progressively More Precise →

**Software:** Programming Languages, Algorithms, Data Structures, etc.  
**Hardware:** Compute, Communication, & Memory Architectures, etc.



= Main **Components** Used in MP Systems



= Main **Techniques** Used to Build by MP



# MP USING APPROXIMATE AND PRECISE METHODS

EMERGING SOLUTIONS USING A FUSION OF BOTH

# MP USING APPROXIMATE AND PRECISE METHODS

## EMERGING SOLUTIONS USING A FUSION OF BOTH

---

### Learning to Infer Program Sketches

---

Maxwell Nye<sup>1,2</sup> Luke Hewitt<sup>1,2,3</sup> Joshua Tenenbaum<sup>1,2,4</sup> Armando Solar-Lezama<sup>2</sup>

#### Abstract

Our goal is to build systems which write code automatically from the kinds of specifications humans can most easily provide, such as examples and natural language instruction. The key idea of this work is that a flexible combination of pattern recognition and explicit reasoning can be used to solve these complex programming problems. We propose a method for dynamically integrating these types of information. Our novel intermediate representation and training algorithm allow a program synthesis system to learn, without direct supervision, when to rely on pattern recognition and when to perform symbolic search. Our model matches the memorization and generalization performance of neural synthesis and symbolic search, respectively, and achieves state-of-the-art performance on a dataset of simple English description-to-code programming problems.

way to combine these language constructs to construct an expression with the desired behavior.

A moderately experienced programmer might immediately *recognize*, from learned experience, that because the output list is always a subset of the input list, a `filter` function is appropriate:

```
filter(input, <HOLE>)
```

where `<HOLE>` is a lambda function which filters elements in the list. The programmer would then have to reason about the correct code for `<HOLE>`.

Finally, a programmer very familiar with this domain might immediately recognize both the need for a `filter` function, as well as the correct semantics for the lambda function, allowing the entire program to be written in one shot:

```
filter(input, lambda x: x%2==0)
```

ICLR 2019

### An Abstraction-Based Framework for Neural Network Verification

Yizhak Yisrael Elboher<sup>1</sup>, Justin Gottschlich<sup>2</sup>, and Guy Katz<sup>1(✉)</sup>

<sup>1</sup> The Hebrew University of Jerusalem, Jerusalem, Israel  
{yizhak.elboher,g.katz}@mail.huji.ac.il

<sup>2</sup> Intel Labs, Santa Clara, USA  
justin.gottschlich@intel.com

CAV 2020

**Abstract.** Deep neural networks are increasingly being used as controllers for safety-critical systems. Because neural networks are opaque, certifying their correctness is a significant challenge. To address this issue, several neural network verification approaches have recently been proposed. However, these approaches afford limited scalability, and applying them to large networks can be challenging. In this paper, we propose a framework that can enhance neural network verification techniques by using over-approximation to reduce the size of the network—thus making it more amenable to verification. We perform the approximation such that if the property holds for the smaller (abstract) network, it holds for the original as well. The over-approximation may be too coarse, in which case the underlying verification tool might return a spurious counterexample. Under such conditions, we perform counterexample-guided refinement to adjust the approximation, and then repeat the process. Our approach is orthogonal to, and can be integrated with, many existing verification techniques. For evaluation purposes, we integrate it with the recently proposed Marabou framework, and observe a significant improvement in Marabou's performance. Our experiments demonstrate the great potential of our approach for verifying larger neural networks.

# MACHINE PROGRAMMING + DEEP LEARNING = NEURAL PROGRAMMING?

# MACHINE PROGRAMMING + DEEP LEARNING = NEURAL PROGRAMMING?

---

## Learning to Optimize

---

**Ke Li     Jitendra Malik**

Department of Electrical Engineering and Computer Sciences  
University of California, Berkeley  
Berkeley, CA 94720  
United States

{ke.li,malik}@eecs.berkeley.edu

**ICLR 2017**

### Abstract

Algorithm design is a laborious process and often requires many iterations of ideation and validation. In this paper, we explore automating algorithm design and present a method to *learn* an optimization algorithm, which we believe to be the first method that can automatically discover a better algorithm. We approach this problem from a reinforcement learning perspective and represent any particular optimization algorithm as a policy. We learn an optimization algorithm using guided policy search and demonstrate that the resulting algorithm outperforms existing hand-engineered algorithms in terms of convergence speed and/or the final objective value.

---

## A Zero-Positive Learning Approach for Diagnosing Software Performance Regressions

---

**Mejbah Alam**  
Intel Labs

mejbah.alam@intel.com

**Justin Gottschlich**  
Intel Labs

justin.gottschlich@intel.com

**Nesime Tatbul**  
Intel Labs and MIT  
tatbul@csail.mit.edu

**Javier Turek**  
Intel Labs  
javier.turek@intel.com

**Timothy Mattson**  
Intel Labs  
timothy.g.mattson@intel.com

**Abdullah Muzahid**  
Texas A&M University  
abdullah.muzahid@tamu.edu

**NEURIPS 2019**

### Abstract

The field of *machine programming* (MP), the automation of the development of software, is making notable research advances. This is, in part, due to the emergence of a wide range of novel techniques in machine learning. In this paper,

# MACHINE PROGRAMMING + DEEP LEARNING = NEURAL PROGRAMMING?

## SOME QUESTIONS:

ONLY IMPROVED BY RETRAINING?

UNDERSTANDABLE, INTERPRETABLE, DEBUGGABLE?

(HALIDE'S APPROACH IS DIFFERENT - WHICH YOU'LL HEAR MORE ABOUT)

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS

# EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS



# EXTRACTION OF CODE SEMANTICS

- Why do we care about code semantics?

# EXTRACTION OF CODE SEMANTICS

- **Why do we care about code semantics?**

## HOPPITY: LEARNING GRAPH TRANSFORMATIONS TO DETECT AND FIX BUGS IN PROGRAMS

**Elizabeth Dinella\***  
University of Pennsylvania

**Hanjun Dai\***  
Google Brain

**Ziyang Li**  
University of Pennsylvania

**Mayur Naik**  
University of Pennsylvania

**Le Song**  
Georgia Tech

**Ke Wang**  
Visa Research

**ICLR 2020**

**ABSTRACT**

We present a learning-based approach to detect and fix a broad range of bugs in Javascript programs. We frame the problem in terms of learning a sequence of graph transformations: given a buggy program modeled by a graph structure, our model makes a sequence of predictions including the position of bug nodes and corresponding graph edits to produce a fix. Unlike previous works built upon deep neural networks, our approach targets bugs that are more diverse and complex in nature (i.e. bugs that require adding or deleting statements to fix). We have realized our approach in a tool called HOPPITY. By training on 290,715 Javascript code change commits on Github, HOPPITY correctly detects and fixes bugs in 9,490 out of 36,361 programs in an end-to-end fashion. Given the bug location and type of the fix, HOPPITY also outperforms the baseline approach by a wide margin.

# EXTRACTION OF CODE SEMANTICS

- Why do we care about code semantics?

## HOPPITY: LEARNING GRAPH TRANSFORMATIONS TO DETECT AND FIX BUGS IN PROGRAMS

**Elizabeth Dinella\***  
University of Pennsylvania

**Hanjun Dai\***  
Google Brain

**Ziyang Li**  
University of Pennsylvania

**Mayur Naik**  
University of Pennsylvania

**Le Song**  
Georgia Tech

**Ke Wang**  
Visa Research

ICLR 2020

ABSTRACT

We present a learning-based approach to detect and fix a broad range of bugs in Javascript programs. We frame the problem in terms of learning a sequence of graph transformations: given a buggy program modeled by a graph structure, our model makes a sequence of predictions including the position of bug nodes and corresponding graph edits to produce a fix. Unlike previous works built upon deep neural networks, our approach targets bugs that are more diverse and complex in nature (i.e. bugs that require adding or deleting statements to fix). We have realized our approach in a tool called HOPPITY. By training on 290,715 Javascript code change commits on Github, HOPPITY correctly detects and fixes bugs in 9,490 out of 36,361 programs in an end-to-end fashion. Given the bug location and type or the fix, HOPPITY also outperforms the baseline approach by a wide margin.

# EXTRACTION OF CODE SEMANTICS

- Why do we care about code semantics?
- Hoppity performs well, but infers *bug fix* semantics
- How well would it do if the semantics of *bug fix* are known?

## HOPPITY: LEARNING GRAPH TRANSFORMATIONS TO DETECT AND FIX BUGS IN PROGRAMS

**Elizabeth Dinella\***  
University of Pennsylvania

**Hanjun Dai\***  
Google Brain

**Ziyang Li**  
University of Pennsylvania

**Mayur Naik**  
University of Pennsylvania

**Le Song**  
Georgia Tech

**Ke Wang**  
Visa Research

ICLR 2020

ABSTRACT

We present a learning-based approach to detect and fix a broad range of bugs in Javascript programs. We frame the problem in terms of learning a sequence of graph transformations: given a buggy program modeled by a graph structure, our model makes a sequence of predictions including the position of bug nodes and corresponding graph edits to produce a fix. Unlike previous works built upon deep neural networks, our approach targets bugs that are more diverse and complex in nature (i.e. bugs that require adding or deleting statements to fix). We have realized our approach in a tool called HOPPITY. By training on 290,715 Javascript code change commits on Github, HOPPITY correctly detects and fixes bugs in 9,490 out of 36,361 programs in an end-to-end fashion. Given the bug location and type of the fix, HOPPITY also outperforms the baseline approach by a wide margin.

# WHY EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS?

# WHY EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS?

- Code that is used tends to be maintained
  - “Software that is used is never finished”
- A code snippet may have multiple semantic meanings

# WHY EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS?

- Code that is used tends to be maintained
  - “Software that is used is never finished”
- A code snippet may have multiple semantic meanings

# WHY MULTI-DIMENSIONAL CODE SEMANTICS?

## Software Language Comprehension using a Program-Derived Semantic Graph

Roshni G. Iyer

University of California, Los Angeles, USA  
roshni@cs.ucla.edu

Wei Wang

University of California, Los Angeles, USA  
weiwang@cs.ucla.edu

Yizhou Sun

University of California, Los Angeles, USA  
yzsun@cs.ucla.edu

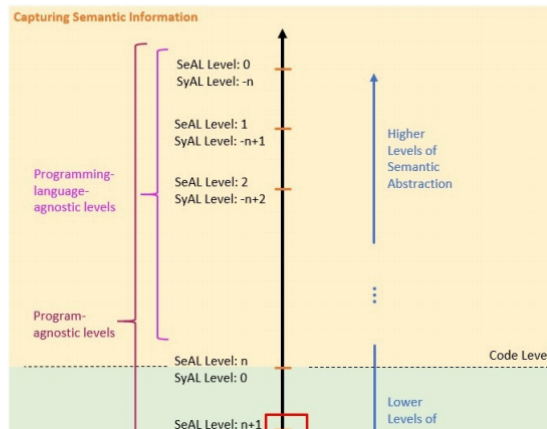
Justin Gottschlich

Intel Labs, USA  
University of Pennsylvania, USA  
justin.gottschlich@intel.com

PREPRINT

### ABSTRACT

Traditional code transformation structures, such as an abstract syntax tree, may have limitations in their ability to extract semantic meaning from code. Others have begun to work on this issue, such as the state-of-the-art Aroma system and its simplified parse tree (SPT). Continuing this research direction, we present a new graphical structure to capture semantics from code using what we refer to as a *program-derived semantic graph* (PSG). The principle behind the PSG is to provide a single structure that can capture program semantics at many levels of granularity. Thus, the PSG is hierarchical in nature. Moreover, because the PSG may have cycles due to dependencies in semantic layers, it is a graph, not a tree. In this paper, we describe the PSG and its fundamental structural differences to the Aroma's SPT. Although our work in the PSG is in its infancy, our early results indicate it is a promising new research direction to explore to automatically extract program semantics.

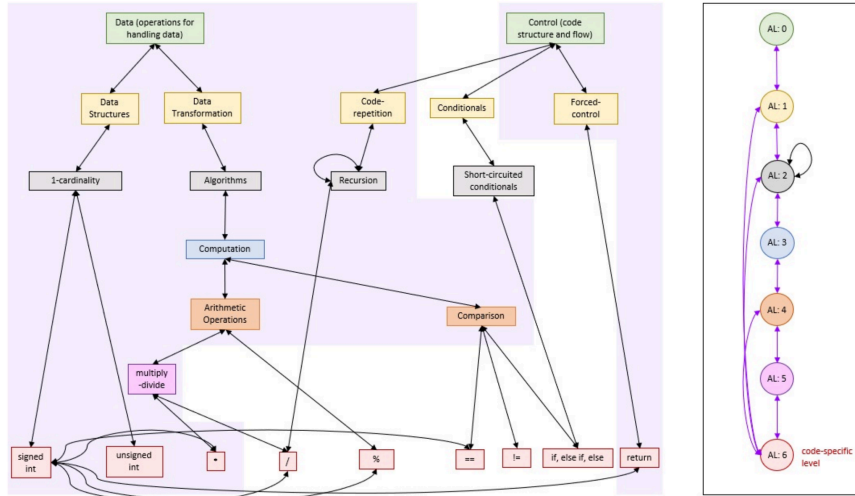




# WHY MULTI-DIMENSIONAL CODE SEMANTICS?

Software Language Comprehension using a Program-Derived Semantic Graph

Preprint, April, 2020



**Figure 5:** PSG of Recursive Power Function. The shaded region denotes overlap in the nodes of the PSG for the iterative power function shown in Figure 6. These total 17 of the 24 total nodes, a 70.83% overlap.

## Implementation 1

```

0 signed int recursive_power(signed int x, unsigned int y)
1 {
2     if (y == 0)
3         return 1;
4     else if (y % 2 == 0)
5         return recursive_power(x, y / 2) *
6             recursive_power(x, y / 2);
7     else
8         return x * recursive_power(x, y / 2) *
9             recursive_power(x, y / 2);

```

# WHY MULTI-DIMENSIONAL CODE SEMANTICS?

Software Language Comprehension using a Program-Derived Semantic Graph

Preprint, April, 2020

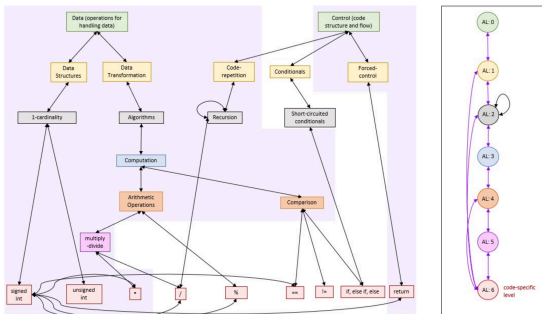


Figure 5: PSG of Recursive Power Function. The shaded region denotes overlap in the nodes of the PSG for the iterative power function shown in Figure 6. These total 17 of the 24 total nodes, a 70.83% overlap.

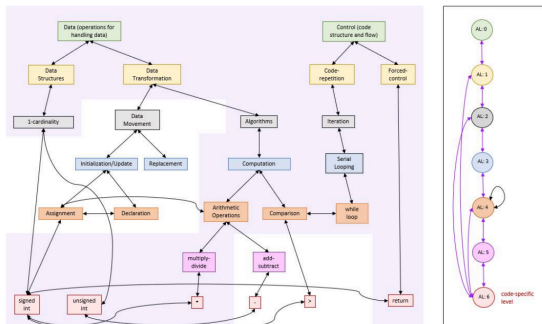


Figure 6: PSG of Iterative Power Function. The shaded region denotes overlap in the nodes of the PSG for the recursive power function shown in Figure 5. These total 19 of the 27 total nodes, a 70.37% overlap.

## Implementation 1

```

0 signed int recursive_power(signed int x, unsigned int y)
1 {
2     if (y == 0)
3         return 1;
4     else if (y % 2 == 0)
5         return recursive_power(x, y / 2) *
6             recursive_power(x, y / 2);
7     else
8         return x * recursive_power(x, y / 2) *
9             recursive_power(x, y / 2);
10 }

```

## Implementation 2

```

0 signed int iterative_power(signed int x, unsigned int y)
1 {
2     signed int val = 1;
3     while (y > 0) {
4         val *= x;
5         y -= 1;
6     }
7     return val;
8 }

```

# WHY MULTI-DIMENSIONAL CODE SEMANTICS?

Both implement exponentiation (only integers)

Both are correct

One is recursive

One is iterative

One has three branch paths

One has one branch path

Each of these properties may be useful  
to know in certain scenarios

Can influence call stacks, speculative  
execution (branch prediction), etc.

## Implementation 1

```
0 signed int recursive_power(signed int x, unsigned int y)
1 {
2     if (y == 0)
3         return 1;
4     else if (y % 2 == 0)
5         return recursive_power(x, y / 2) *
6             recursive_power(x, y / 2);
7     else
8         return x * recursive_power(x, y / 2) *
9             recursive_power(x, y / 2);
10 }
```

## Implementation 2

```
0 signed int iterative_power(signed int x, unsigned int y)
1 {
2     signed int val = 1;
3     while (y > 0) {
4         val *= x;
5         y -= 1;
6     }
7     return val;
8 }
```

# WHY MULTI-DIMENSIONAL CODE SEMANTICS?

Both implement exp  
Both are correct  
One is recursive  
One is iterative  
One has three branch  
One has one branch

## SOME CHALLENGES:

**LOTS OF CODE**  
**FEW SEMANTIC CODE LABELS**

Each of these prop  
to know in certain s  
  
Can influence call s  
execution (branch p

**FIND CLEVER WAYS TO LIFT SEMANTICS?**  
**LIFT SEMANTICS WITHOUT COMPILATION?**  
**FIND SEMANTICS FROM SURROUNDINGS?**

```
igned int x, unsigned int y)
```

```
, y / 2) *  
/ 2);
```

```
er(x, y / 2) *  
/ 2);
```

```
igned int x, unsigned int y)
```

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# NOVEL STRUCTURAL REPRESENTATIONS OF CODE

- Why do we need new code structures?

# NOVEL STRUCTURAL REPRESENTATIONS OF CODE

## ■ Why do we need new code structures?

### Aroma: Code Recommendation via Structural Code Search

SIFEI LUAN, Facebook, USA

DI YANG\*, University of California, Irvine, USA

CELESTE BARNABY, Facebook, USA

KOUSHIK SEN†, University of California, Berkeley, USA

SATISH CHANDRA, Facebook, USA

OOPSLA 2019

Programmers often write code that has similarity to existing code written somewhere. A tool that could help programmers to search such similar code would be immensely useful. Such a tool could help programmers to extend partially written code snippets to completely implement necessary functionality, help to discover extensions to the partial code which are commonly included by other programmers, help to cross-check against similar code written by other programmers, or help to add extra code which would fix common mistakes and errors. We propose Aroma, a tool and technique for code recommendation via structural code search. Aroma indexes a huge code corpus including thousands of open-source projects, takes a partial code snippet as input, searches the corpus for method bodies containing the partial code snippet, and clusters and intersects the results of the search to recommend a small set of succinct code snippets which both contain the query snippet and appear as part of several methods in the corpus. We evaluated Aroma on 2000 randomly selected queries created from the corpus, as well as 64 queries derived from code snippets obtained from Stack Overflow, a popular website for discussing code. We implemented Aroma for 4 different languages, and developed an IDE plugin for Aroma. Furthermore, we conducted a study where we asked 12 programmers to complete programming tasks using Aroma, and collected their feedback. Our results indicate that Aroma is capable of retrieving and recommending relevant code snippets efficiently.

### MISIM: An End-to-End Neural Code Similarity System

Fangke Ye \*

Intel Labs and Georgia Institute of Technology  
yefangke@gatech.edu

Shengtian Zhou \*

Intel Labs  
shengtian.zhou@intel.com

Anand Venkat

Intel Labs  
anand.venkat@intel.com

Ryan Marcus

Intel Labs and MIT  
ryanmarcus@csail.mit.edu

Nesime Tatbul

Intel Labs and MIT  
tatbul@csail.mit.edu

Jesmin Jahan Tithi

Intel Labs  
jesmin.jahan.tithi@intel.com

Paul Petersen

Intel  
paul.petersen@intel.com

Timothy Mattson

Intel Labs  
timothy.g.mattson@intel.com

Tim Kraska

MIT  
kraska@mit.edu

Pradeep Dubey

Intel Labs  
pradeep.dubey@intel.com

Vivek Sarkar

Georgia Institute of Technology  
vsarkar@gatech.edu

Justin Gottschlich

Intel Labs and University of Pennsylvania  
justin.gottschlich@intel.com

### PREPRINT

### Abstract

Code similarity systems are integral to a range of applications from code recommendation to automated construction of software tests and defect mitigation. In this paper, we present *Machine Inferred Code Similarity* (MISIM), a novel end-to-end code similarity system that consists of two core components. First, MISIM uses a novel *context-aware semantic structure*, which is designed to aid in lifting semantic meaning from code syntax. Second, MISIM provides a neural-based code similarity scoring algorithm, which can be implemented with various neural network architectures with learned parameters. We compare MISIM to three state-of-the-art code similarity systems: (i) code2vec, (ii) Neural Code Comprehension, and (iii) Aroma. In our experimental evaluation across 45,780 programs, MISIM consistently outperformed all three systems, often by a large factor (upwards of 40.6×).

# NOVEL STRUCTURAL REPRESENTATIONS OF CODE

- Aroma introduced the **simplified parse tree**
- MISIM introduced the **context-aware semantics structure**
- These structures have led to state-of-the-art accuracy

## Aroma: Code Recommendation via Structural Code Search

SIFEI LUAN, Facebook, USA

DI YANG\*, University of California, Irvine, USA

CELESTE BARNABY, Facebook, USA

KOUSHIK SEN†, University of California, Berkeley, USA

SATISH CHANDRA, Facebook, USA

OOPSLA 2019

Programmers often write code that has similarity to existing code written somewhere. A tool that could help programmers to search such similar code would be immensely useful. Such a tool could help programmers to extend partially written code snippets to completely implement necessary functionality, help to discover extensions to the partial code which are commonly included by other programmers, help to cross-check against similar code written by other programmers, or help to add extra code which would fix common mistakes and errors. We propose Aroma, a tool and technique for code recommendation via structural code search. Aroma indexes a huge code corpus including thousands of open-source projects, takes a partial code snippet as input, searches the corpus for method bodies containing the partial code snippet, and clusters and intersects the results of the search to recommend a small set of succinct code snippets which both contain the query snippet and appear as part of several methods in the corpus. We evaluated Aroma on 2000 randomly selected queries created from the corpus, as well as 64 queries derived from code snippets obtained from Stack Overflow, a popular website for discussing code. We implemented Aroma for 4 different languages, and developed an IDE plugin for Aroma. Furthermore, we conducted a study where we asked 12 programmers to complete programming tasks using Aroma, and collected their feedback. Our results indicate that Aroma is capable of retrieving and recommending relevant code snippets efficiently.

---

## MISIM: An End-to-End Neural Code Similarity System

---

<b>Fangke Ye</b> <sup>*</sup> Intel Labs and Georgia Institute of Technology yefangke@gatech.edu	<b>Shengtian Zhou</b> <sup>*</sup> Intel Labs shengtian.zhou@intel.com	
<b>Anand Venkat</b> Intel Labs anand.venkat@intel.com	<b>Ryan Marcus</b> Intel Labs and MIT ryanmarcus@csail.mit.edu	<b>Nesime Tatbul</b> Intel Labs and MIT tatbul@csail.mit.edu
<b>Jesmin Jahan Tithi</b> Intel Labs jesmin.jahan.tithi@intel.com	<b>Paul Petersen</b> Intel paul.petersen@intel.com	
<b>Timothy Mattson</b> Intel Labs timothy.g.mattson@intel.com	<b>Tim Kraska</b> MIT kraska@mit.edu	<b>Pradeep Dubey</b> Intel Labs pradeep.dubey@intel.com
<b>Vivek Sarkar</b> Georgia Institute of Technology vsarkar@gatech.edu	<b>Justin Gottschlich</b> Intel Labs and University of Pennsylvania justin.gottschlich@intel.com	

### PREPRINT

### Abstract

Code similarity systems are integral to a range of applications from code recommendation to automated construction of software tests and defect mitigation. In this paper, we present *Machine Inferred Code Similarity* (MISIM), a novel end-to-end code similarity system that consists of two core components. First, MISIM uses a novel *context-aware semantic structure*, which is designed to aid in lifting semantic meaning from code syntax. Second, MISIM provides a neural-based code similarity scoring algorithm, which can be implemented with various neural network architectures with learned parameters. We compare MISIM to three state-of-the-art code similarity systems: (i) code2vec, (ii) Neural Code Comprehension, and (iii) Aroma. In our experimental evaluation across 45,780 programs, MISIM consistently outperformed all three systems, often by a large factor (upwards of 40.6×).





# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

- This heterogeneity is generating multiplicative complexity

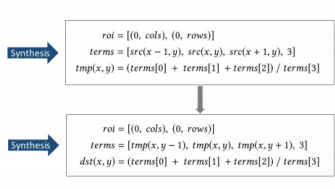
# AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

- This heterogeneity is generating multiplicative complexity

## Automatically Translating Image Processing Libraries to Halide

MAAZ BIN SAFEER AHMAD, University of Washington, Seattle  
JONATHAN RAGAN-KELLEY, University of California, Berkeley  
ALVIN CHEUNG, University of California, Berkeley  
SHOAIB KAMIL, Adobe

```
void blur(uint8_t* dst, uint8_t* src, int rows, int cols,
         int rowbytes) {
    // horizontal blur
    for (int r = 0; r < rows; r++) {
        for (int c = 0; c < cols; c++) {
            tmp[c] = (src[c-1] + src[c] + src[c+1]) / 3;
        }
        tmp += rowbytes;
        src += rowbytes;
    }
    // vertical blur
    for (int r = 0; r < rows; r++) {
        for (int c = 0; c < cols; c++) {
            int sum = (tmp[c-rowbytes] + tmp[c] +
                    tmp[c+rowbytes]) / 3;
            dst[c] = sum / 3;
        }
        dst += rowbytes;
        tmp += rowbytes;
    }
}
```



SIGGRAPH ASIA 2019

```
Func blur(Func dst, Func src, int rows, int cols)
{
    @Body (0, cols, 0, rows);
    Func tmp;
    Var i, j;

    // horizontal blur
    tmp(i, j) = (src(i-1, j) + src(i, j) + src(i+1, j)) / 3;

    // vertical blur
    dst(i, j) = undef<uint8_t>();
    dst(r, x, r, y) = (tmp(r, x, r, y-1) + tmp(r, x, r, y)
                    + tmp(r, x, r, y+1)) / 3;

    return dst;
}
```

Fig. 1. DEXTER parses the input C++ function (shown on the left) into a DAG of smaller stages, then uses our 3-step synthesis algorithm to infer the semantics of each stage, expressed in a high-level IR (middle). Finally, code generation rules compile the IR specifications into executable Halide code (right).

## Automatically Scheduling Halide Image Processing Pipelines

Ravi Teja Mullanpudi\* Andrew Adams† Dillon Sharlet‡ Jonathan Ragan-Kelley† Kayvon Fatahalian\*  
\*Carnegie Mellon University ‡Google †Stanford University

### Abstract

The Halide image processing language has proven to be an effective system for authoring high-performance image processing code. Halide programmers need only provide a high-level strategy for mapping an image processing pipeline to a parallel machine (a *schedule*), and the Halide compiler carries out the mechanical task of generating platform-specific code that implements the schedule. Unfortunately, designing high-performance schedules for complex image processing pipelines requires substantial knowledge of modern hardware architecture and code-optimization techniques. In this paper we provide an algorithm for automatically generating high-performance schedules for Halide programs. Our solution extends the function bounds analysis already present in the Halide compiler to automatically perform locality and parallelism-enhancing global program transformations typical of those employed by expert Halide developers. The algorithm does not require costly (and often impractical) auto-tuning, and, in seconds, generates schedules for a broad set of image processing benchmarks that are performance-competitive with, and often better than, schedules manually authored by expert Halide developers on server and mobile CPUs, as well as GPUs.

**Keywords:** image processing, optimizing compilers, Halide

**Concepts:** •Computing methodologies → Graphics systems and interfaces;

algorithm's execution on a machine (called a *schedule*). The Halide compiler then handles the tedious, mechanical task of generating platform-specific code that implements the schedule (e.g., spawning threads, managing buffers, generating SIMD instructions).

Although Halide provides high-level abstractions for expressing schedules, *designing* schedules that perform well on modern hardware is hard; it requires expertise in modern optimization techniques and hardware architectures. For example, around 70 software engineers at Google currently write image processing algorithms in Halide, but they rely on a much smaller cadre of Halide scheduling experts to produce the most efficient implementations. Further, production image processing pipelines are long and complex, and are difficult to schedule even for the best Halide programmers. Arriving at a good schedule remains a laborious, iterative process of schedule tweaking and performance measurement. Also, in large production pipelines, software engineering considerations (e.g., modularity, code reuse) may preclude experts from having the global program knowledge needed to create optimal schedules.

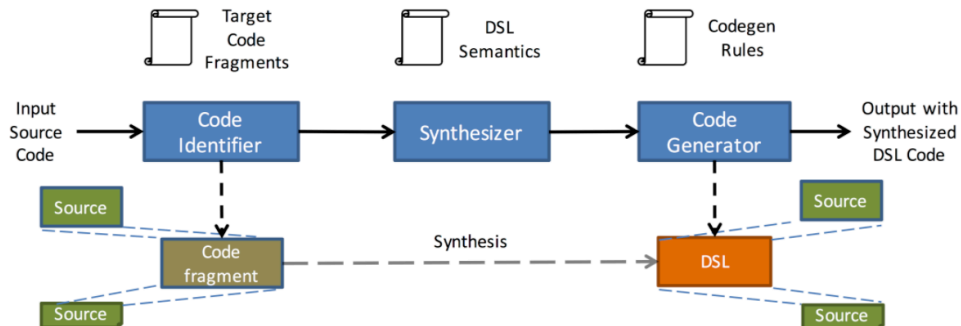
SIGGRAPH 2016

# AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

- This heterogeneity is generating multiplicative complexity

## MetaLift

Leveraging DSLs made easy



### People

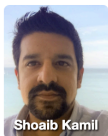
MetaLift is jointly developed by the folks at the University of Washington [Programming Languages and Software Engineering Research Group](#), [Adobe Research](#), and [Intel Labs](#). The following are the main developers of MetaLift:



Maaz Ahmad



Alvin Cheung



Shoalb Kamil



Remy Wang

Verified lifting has been the underlying technology used to build the following compilers:



Dexter is a compiler that translates image processing kernels from C to Halide.



Casper is a compiler that translates sequential Java to Spark and Hadoop.



STNG is a compiler that enables Fortran kernels to leverage GPUs by compiling them into the Halide DSL.

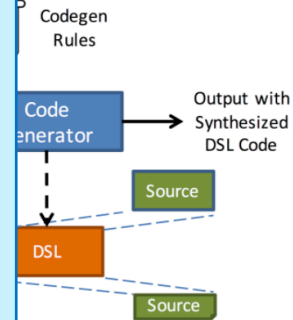
# AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

## SOME QUESTIONS:

WHAT ARE THE QUALITY METRICS FOR TRANSLATION?

CORRECT & PERFORMANCE (ARE PROBABLY OBVIOUS)

WHAT ABOUT SECURITY, MAINTAINABILITY, POWER FOOTPRINT, ETC.?



the following compilers:  
using kernels from C to Halide.  
va to Spark and Hadoop.  
to leverage GPUs by compiling

Peop  
MetaLi  
and Sc  
the ma  
Maa

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# INTENTIONAL PROGRAMMING AND BEHAVIORS

- Focus on what the intention is, not how that intention may manifest



# INTENTIONAL PROGRAMMING AND BEHAVIORS

- Focus on what the intention is, not how that intention may manifest

## LEARNING TO REPRESENT PROGRAMS WITH PROPERTY SIGNATURES

Augustus Odena, Charles Sutton

Google Research

{augustusodena, charlessutton}@google.com

### ABSTRACT

We introduce the notion of property signatures, a representation for programs and program specifications meant for consumption by machine learning algorithms. Given a function with input type  $\tau_{in}$  and output type  $\tau_{out}$ , a property is a function of type:  $(\tau_{in}, \tau_{out}) \rightarrow \text{Bool}$  that (informally) describes some simple property of the function under consideration. For instance, if  $\tau_{in}$  and  $\tau_{out}$  are both lists of the same type, one property might ask ‘is the input list the same length as the output list?’. If we have a list of such properties, we can evaluate them all for our function to get a list of outputs that we will call the property signature. Crucially, we can ‘guess’ the property signature for a function given only a set of input/output pairs meant to specify that function. We discuss several potential applications of property signatures and show experimentally that they can be used to improve over a baseline synthesizer so that it emits twice as many programs in less than one-tenth of the time.

ICLR 2020

## Skip Blocks: Reusing Execution History to Accelerate Web Scripts

SARAH CHASINS, University of California, Berkeley, USA

RASTISLAV BODIK, University of Washington, USA

With more and more web scripting languages on offer, programmers have access to increasing language support for web scraping tasks. However, in our experiences collaborating with data scientists, we learned that two issues still plague long-running scraping scripts: i) When a network or website goes down mid-scrape, recovery sometimes requires restarting from the beginning, which users find frustratingly slow. ii) Websites do not offer atomic snapshots of their databases; they update their content so frequently that output data is cluttered with slight variations of the same information — e.g., a tweet from profile 1 that is retweeted on profile 2 and scraped from both profiles, once with 52 responses then later with 53 responses.

We introduce the *skip block*, a language construct that addresses both of these disparate problems. Programmers write lightweight annotations to indicate when the current object can be considered equivalent to a previously scraped object and direct the program to skip over the scraping actions in the block. The construct is hierarchical, so programs can skip over long or short script segments, allowing adaptive reuse of prior work. After network and server failures, skip blocks accelerate failure recovery by 7.9x on average. Even scripts that do not encounter failures benefit; because sites display redundant objects, skipping over them accelerates scraping by up to 2.1x. For longitudinal scraping tasks that aim to fetch only new objects, the second run exhibits an average speedup of 5.2x. Our small user study reveals that programmers can quickly produce skip block annotations.

CCS Concepts: • Software and its engineering → Control structures;

OOPSLA 2017

# HALIDE: A DOMAIN-SPECIFIC & INTENTIONAL PROGRAMMING LANGUAGE (HETEROGENEOUS HARDWARE)

Learning to Optimize Halide with Tree Search and Random Programs

ANDREW ADAMS, Facebook AI Research  
 KARIMA MA, UC Berkeley  
 LUKE ANDERSON, MIT CSAIL  
 RIYADH BAGHDADI, MIT CSAIL  
 TZU-MAO LI, MIT CSAIL  
 MICHAEL CHARRBI, Adobe  
 BENOIT STEINER, Facebook AI Research  
 STEVEN JOHNSON, Google  
 KAYVON FATRAHALIAN, Stanford University  
 FRÉDÉRIC DURAND, MIT CSAIL  
 JONATHAN RAGAN-KELLEY, UC Berkeley

SIGGRAPH 2019

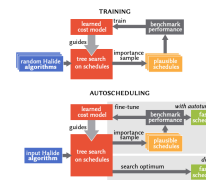
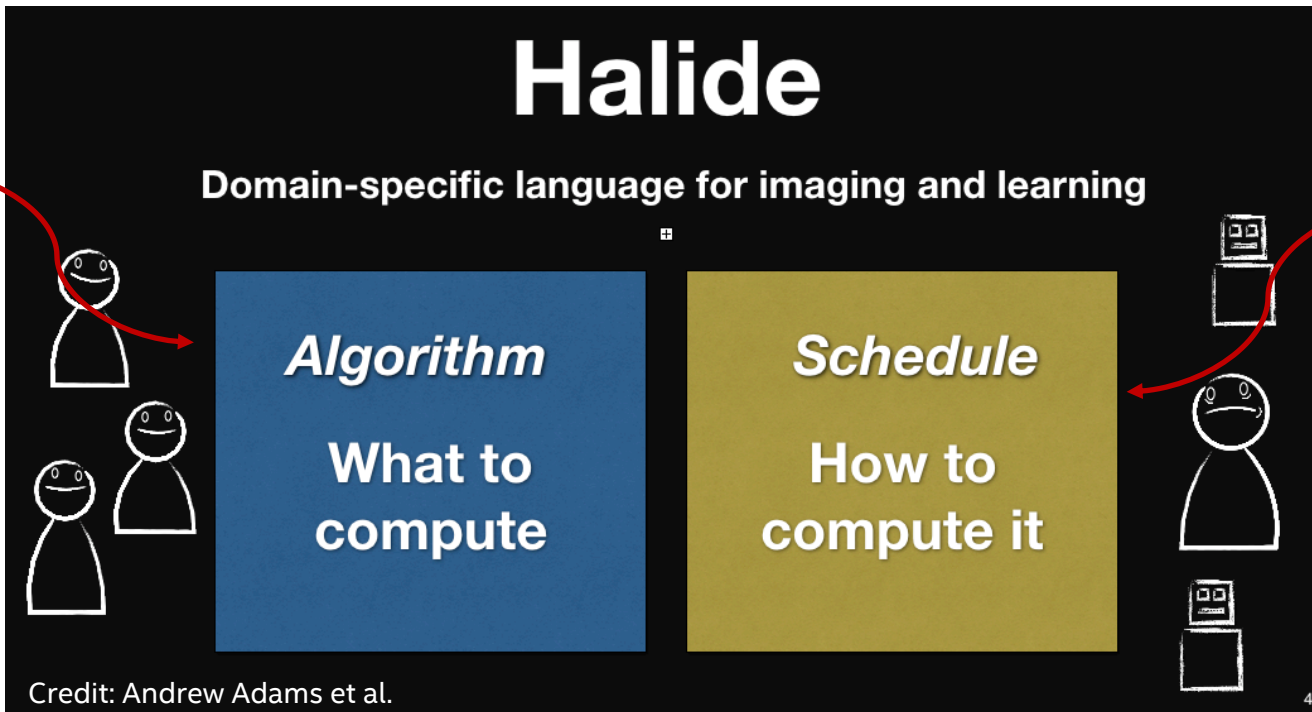


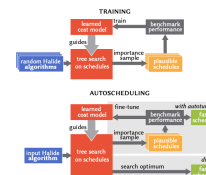
Fig. 1. We generate schedules for Halide programs using tree search over the space of schedules (Sec. 3) guided by a learned cost model and optimal autotuning (Sec. 4). The cost model is trained by benchmarking thousands of randomly generated Halide programs and schedules (Sec. 3). The resulting code significantly outperforms prior work and human experts (Sec. 4).

Intention



Adaptation

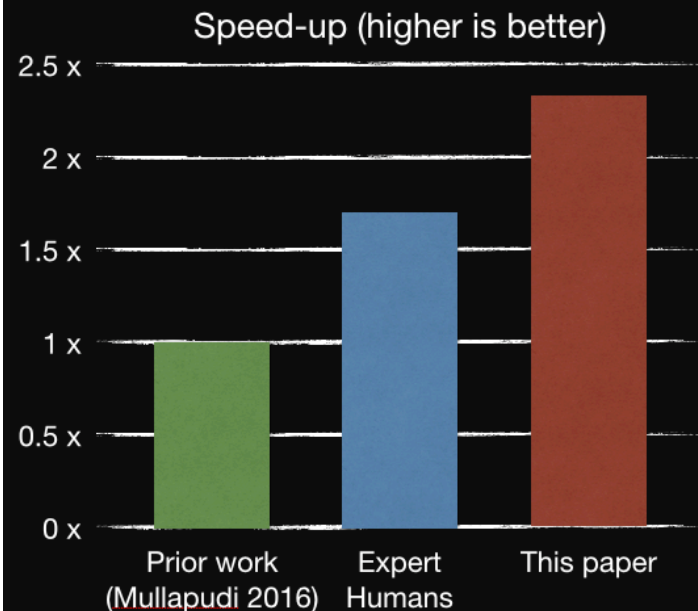
ANDREW ADAMS, Facebook AI Research  
KARIMA MA, UC Berkeley  
LUKE ANDERSON, MIT CSAIL  
RIYADH BAGHDADI, MIT CSAIL  
TZU-MAO LI, MIT CSAIL  
MICHAEL CHARRI, Adobe  
BENOIT STEINER, Facebook AI Research  
STEVEN JOHNSON, Google  
KAYVON FATRAHALIAN, Stanford University  
FRÉDO DURAND, MIT CSAIL  
JONATHAN RAGAN-KELLEY, UC Berkeley



SIGGRAPH 2019

Fig. 1. We generate schedules for Halide programs using tree search over the space of schedules (Sec. 3) guided by a learned cost model and optimal autotuning (Sec. 4). The cost model is trained by benchmarking thousands of randomly generated Halide programs and schedules (Sec. 5). The resulting code significantly outperforms prior work and human experts (Sec. 6).

## A new automatic scheduling algorithm for Halide



### Larger search space

- includes more Halide scheduling features
- extensible

### Hybrid cost model

- Mix of machine learning and hand-designed terms
- Can model complex architectures

Credit: Andrew Adams et al. 12

# HALIDE: SUPER-HUMAN PERFORMANCE

Learning to Optimize Halide with Tree Search and Random Programs

ANDREW ADAMS, Facebook AI Research  
KARIMA MA, UC Berkeley  
LUKE ANDERSON, MIT CSAIL  
RIYADH BAGHDADI, MIT CSAIL  
TZU-MAO LI, MIT CSAIL  
MICHAEL CHARRBI, Adobe  
BENOIT STEINER, Facebook AI Research  
STEVEN JOHNSON, Google

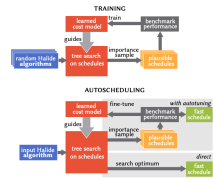
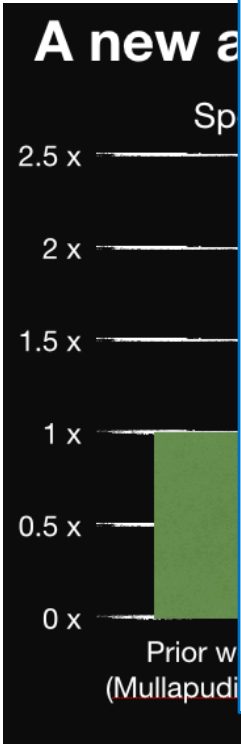


Fig. 1. We generate schedules for Halide programs using tree search over the space of schedules (Sec. 3) guided by a learned cost model and optimal autotuning (Sec. 4). The cost model is trained by benchmarking thousands of randomly generated Halide programs and schedules (Sec. 5). The resulting code significantly outperforms prior work and human experts (Sec. 6).



## SOME QUESTIONS:

**HALIDE IS DOMAIN-SPECIFIC, CAN WE DO THIS GENERALLY?**

**CAN WE PROVIDE INTENTION-BASED INTERFACES TO EXISTING WIDELY USED LANGUAGES (C++, PYTHON, JAVASCRIPT)?**

Credit: Andrew Adams et al. 12

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF MULTI-DIMENSIONAL AND EVOLVING CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP

## Challenges:

- Computational workload via FM and ML may be large
- MP data is large, can be dense, and is mostly unlabeled
- Given this, what does the future MP hardware look like?

# THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP

## Challenges:

- Computational workload via FM and ML may be large
- MP data is large, can be dense, and is mostly unlabeled
- Given this, what does the future MP hardware look like?

I have no idea.

But I do have ideas about things we can think about.

# THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP

## Some open questions:

- **What interfaces do we expect for expression of intention?**
  - What ramifications are associated with those?
- **What are the core techniques used for MP?**
  - What are the data, communication, and compute implications?
- **We have a massive big data problem in front of us**
  - As of summer 2020, there were over 200M+ github repos
  - Code is multi-dimensional by nature
  - The size and density of this data implies new frontiers of hardware



# **(NON-EXHAUSTIVE) TOPICS SKIPPED**

**MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS**

**EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS**

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**ETHICS OF MACHINE PROGRAMMING**

**INTENTIONAL PROGRAMMING AND BEHAVIORS**

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

**THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP**

# (NON-EXHAUSTIVE) TOPICS SKIPPED

MACHINE PROGRAMMING USING APPROXIMATE AND PRECISE METHODS

EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS

NOVEL STRUCTURAL REPRESENTATIONS OF CODE

**HUMAN-INTENDED AND MACHINE-INTENDED PROGRAMMING LANGUAGES**

AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

**ETHICS OF MACHINE PROGRAMMING**

INTENTIONAL PROGRAMMING AND BEHAVIORS

**LEARNING FOR ADAPTIVE SOFTWARE (AND HARDWARE)**

THE FUTURE OF DATA, COMMUNICATION, AND COMPUTATION FOR MP

# THE ERA OF MACHINE PROGRAMMING IS NOW

We are on the verge of a revolutionary shift

# THE ERA OF MACHINE PROGRAMMING IS NOW

**We are on the verge of a revolutionary shift**

**Many institutions are heavily investing in MP**

- Many large tech companies (Amazon, Google, IBM, Intel, Microsoft, etc.)
  - Both research and engineering
- Dozens of startups to solve a single MP problem
- Several leading academic institutions

# THE ERA OF MACHINE PROGRAMMING IS NOW

**We are on the verge of a revolutionary shift**

**Many institutions are heavily investing in MP**

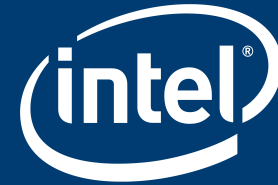
- Many large tech companies (Amazon, Google, IBM, Intel, Microsoft, etc.)
  - Both research and engineering
- Dozens of startups to solve a single MP problem
- Several leading academic institutions

**We can democratize the creation of software with MP**

- Imagine a global population, where everyone can express their creativeness
- Imagine a world where coders only spent time expressing our intentions, not fixing code
- What kind of scientific, artistic, innovative things might we discover?

**Looking forward to working with many (all?) of you!**

# THANK YOU!



Machine Programming Inside

**QUESTIONS / COMMENTS: [JUSTIN.GOTTSCHLICH@INTEL.COM](mailto:JUSTIN.GOTTSCHLICH@INTEL.COM)**