



# CHALLENGES AND OPPORTUNITIES IN MACHINE PROGRAMMING (MP)

---

**Justin Gottschlich**

Principal Scientist & Director/Founder 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 Haghighat, Mary Hall, Niranjana Hasabnis, Adam Herr, Jim Held, Roshni Iyer, Nilesch 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) MP TOPICS**

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**


**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**


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

**BUT FIRST – SOME BACKGROUND**

# SOME BACKGROUND



***INTRODUCING:***



# MACHINE PROGRAMMING RESEARCH (MPR)

**A NEW PIONEERING RESEARCH INITIATIVE @ INTEL**



# INTEL LABS' MPR GOALS

---

**Automation of software (and hardware) to improve:**

**1. *productivity: minimal human effort\****


*\*Measured as 1000x+ improvement over human work performed today*

**2. *quality: better software than the best human programmers\****

*\*Measured as superhuman correctness, performance, security, etc.*

**We speculate this end-point to be at least 2+ decades away.**

# 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
- **Adaptation:** 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?

Names matter. Should infer meaning from name.

# WHY CALL IT MACHINE PROGRAMMING?

Names matter. Should infer meaning from name.

Why isn't this seminar series called?

- AI for Computer Science
- Neural Networks for Optimization
- Machine Learning for Software

None of these name precisely match the intention of this seminar series (as I understand it 😊).

# WHY CALL IT MACHINE PROGRAMMING?

- Likewise, our alternatives were:
  - Program Synthesis
  - AI/ML for Code
  - Software 2.0

# WHY CALL IT MACHINE PROGRAMMING?

- Likewise, our alternatives were:
  - Program Synthesis (historical w/ formal methods; not always synthesizing)
  - AI/ML for Code (it's not just AI/ML – this is important)
  - Software 2.0 (what does this mean?)
    - And Software 3.0, and 4.0, and 5.0?
- The machine programming name was coined to avoid confusion and broaden scope.

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**

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



# MP USES STOCHASTIC AND DETERMINISTIC METHODS

## STOCHASTIC

### Machine Learning

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

## DETERMINISTIC

### Formal Methods

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

←----- More stochastic, but may have a larger solution space      More deterministic, but may have smaller solution space -----→

**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 USES STOCHASTIC AND DETERMINISTIC METHODS

EMERGING SOLUTIONS USING A FUSION OF BOTH

# MP USES STOCHASTIC AND DETERMINISTIC 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.

# MP USES STOCHASTIC AND DETERMINISTIC METHODS

## FORMAL METHODS FOR INCREASED DETERMINISM OF NEURAL NETS

### 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

**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.

CAV 2020



## NEURON COALESCENCE VIA MATHEMATICAL TRANSITIVITY USING COUNTEREXAMPLE-GUIDED ABSTRACTION REFINEMENT

# MACHINE PROGRAMMING + DEEP LEARNING = NEURAL PROGRAMMING?

**Neural Programming:** use of neural networks as a replacement of code.

# MACHINE PROGRAMMING + DEEP LEARNING = NEURAL PROGRAMMING?

Neural Programming: use of neural networks as a replacement of code.

---

## 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,

# AUTOPERF: PERFORMANCE REGRESSION TESTING

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


AutoPerf invents and adapts these tests

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,

NEURIPS 2019

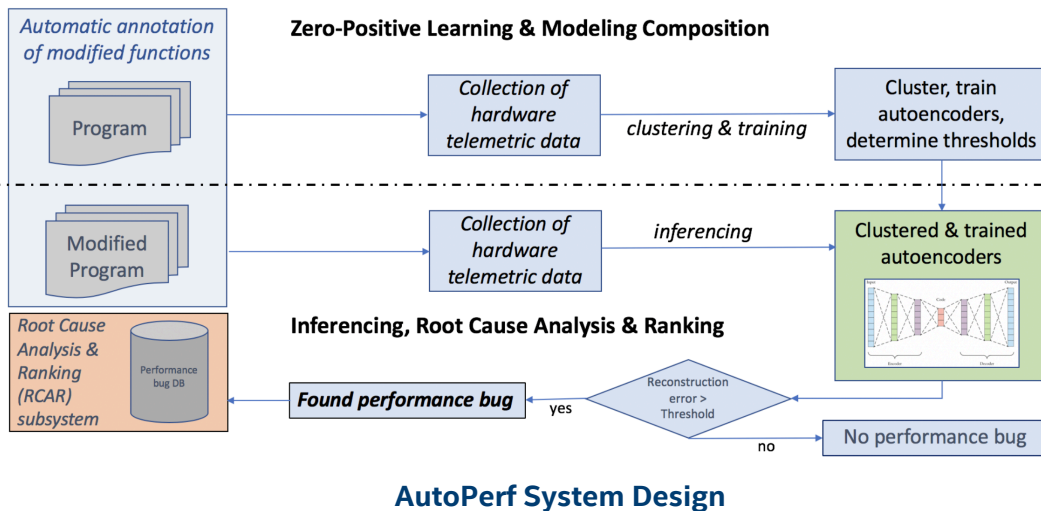
## Performance Regression Testing



Test for detecting **performance anomaly** introduced by a change in software

# AUTOPERF: PERFORMANCE REGRESSION TESTING

- Uses zero-positive learning (ZPL), autoencoders, hardware telemetry
- Emits no false negatives (no missed performance bugs)
- Negligible (4%) performance overhead using hardware performance counters (HWPCs)



## How is this neural programming?

ML *invents* the regression tests and *adapts* them to the specialized hardware to analyze performance.

NN *is* the code/test.



# MACHINE PROGRAMMING + DEEP LEARNING = NEURAL PROGRAMMING?

## SOME CONCERNS W/ NEURAL PROGRAMMING

ONLY IMPROVED BY RETRAINING?

UNDERSTANDABLE, INTERPRETABLE, DEBUGGABLE?

THERE ARE OTHER MP APPROACHES THAT GENERATE ACTUAL CODE  
(WE'LL SEE SOME EXAMPLES TODAY)

# (NON-EXHAUSTIVE) TOPICS

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**

**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 care about code semantics?

# EXTRACTION OF CODE SEMANTICS

- Why 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.

# HOPPITY: CODE REPAIR AS GRAPH TRANSFORMATIONS

- Example of Hoppity's bug repair graph transformation

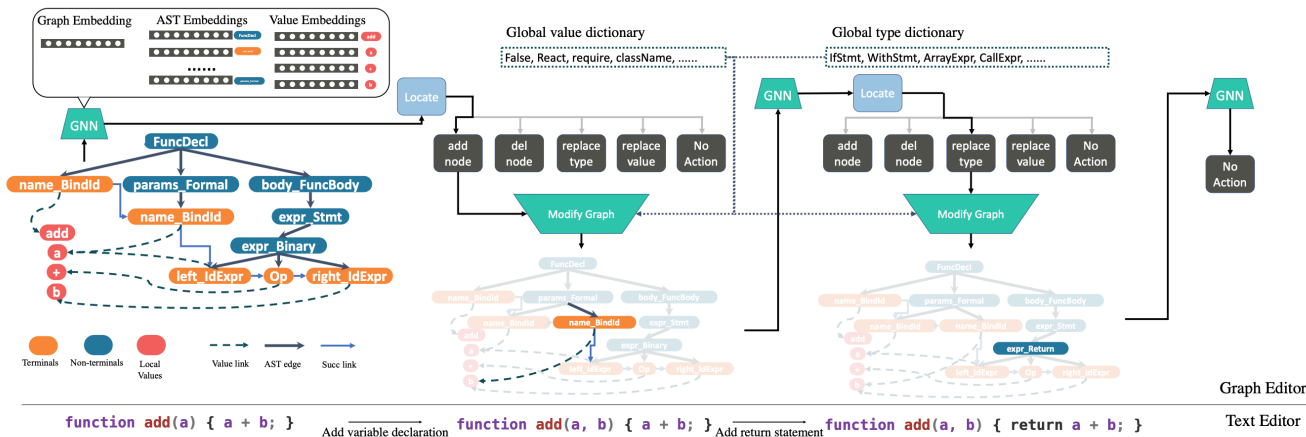


Figure 2: Code repair as graph transformation. Each step the source code graph is edited via one of the operator module until STOP is triggered by controller.

- How does Hoppity find bug fixes to learn from?

# HOW DOES HOPPITY FIND BUG FIXES TO LEARN FROM?

Given a commit, we download the Javascript file before and after the change:  $(src_{buggy}, src_{fixed})$ . Commits can contain many types of changes such as feature additions, refactorings, bug fixes, etc. In an attempt to filter our dataset to only include bug fixes, we use a heuristic based on the number of changes to the AST. Our insight is that a commit with a smaller number of AST differences is more likely to be a bug fix than a commit containing large changes. Thus for the experiments, we use three different datasets: OneDiff with precisely one edit; ZeroOneDiff with zero and one edit and ZeroOneTwoDiff with zero, one or two edits. We additionally filter out data points with ASTs larger than 500 nodes as a parameter in our system. A detailed overview of our corpus crawler is available in Appendix B.

- Looks at repo changesets – if small enough, deem a potential bug fix
  - Infers *bug fix* semantics on repository delta size
- How would Hoppity perform if the semantics of *bug fix* are known?
  - What about other environmental factors that could be inferred?



# WHY EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS?

# WHY EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS?

- **Evolving:**

- Code that is used, tends to be maintained
  - “Software that is used is never finished”
- Evolving code == evolving semantics?

- **Multi-dimensional:**

- A code snippet may have multiple semantic meanings
- A bit more challenging to understanding ...

# WHY MULTI-DIMENSIONAL CODE SEMANTICS?

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

Roshni G. Iyer

University of California, Los Angeles, USA  
roshnigiyer@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

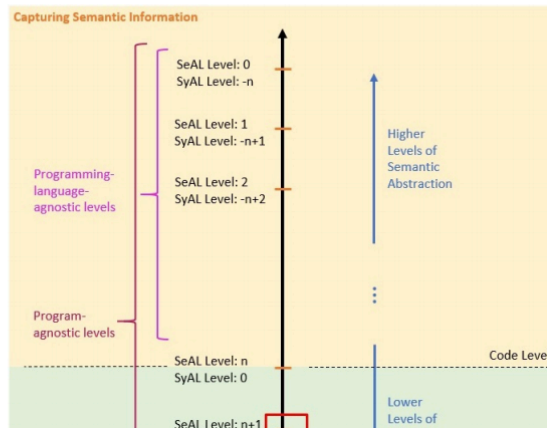
Justin Gottschlich

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

PREPRINT

### ABSTRACT

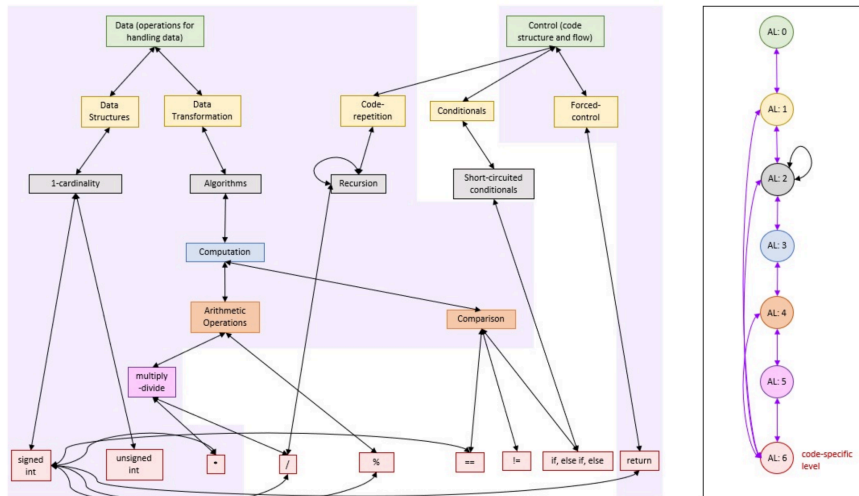
Traditional code transformation structures, such as an abstract syntax tree, may have limitations in their ability to extract semantic meaning from code. Others have begun to work on this issue, such as the state-of-the-art Aroma system and its simplified parse tree (SPT). Continuing this research direction, 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);
10 }
```

# 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.

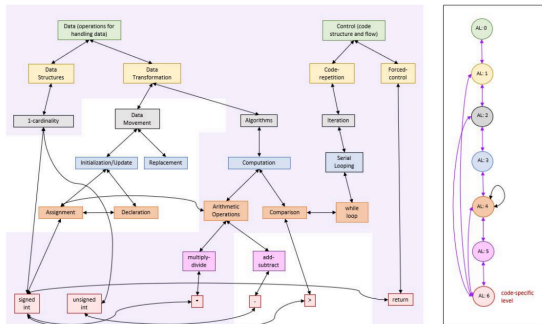


Figure 6: PSG of Iterative Power Function. The shaded region denotes overlap in the nodes of the PSG for the recursive power function shown in Figure 5. These total 19 of the 27 total nodes, a 70.37% overlap.

## 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?

## Some semantics:

Both implement exponentiation (only integers)

Both are correct

One is recursive

One is iterative

One has multiple branches

One has one branch path

## Each semantic may be useful.

Can influence code comprehension,  
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 }
```

# EXTRACTION OF EVOLVING AND MULTI-DIMENSIONAL CODE SEMANTICS

## SOME THOUGHTS:

**BIG AND DENSE CODE, FEW SEMANTIC CODE LABELS**

**DISCOVER NOVEL WAYS TO LIFT SEMANTICS**

**LIFT SEMANTICS WITHOUT COMPILATION (WORKS WITH BROKEN CODE)?**

**FIND SEMANTICS FROM SURROUNDINGS?**

# (NON-EXHAUSTIVE) TOPICS

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**

**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

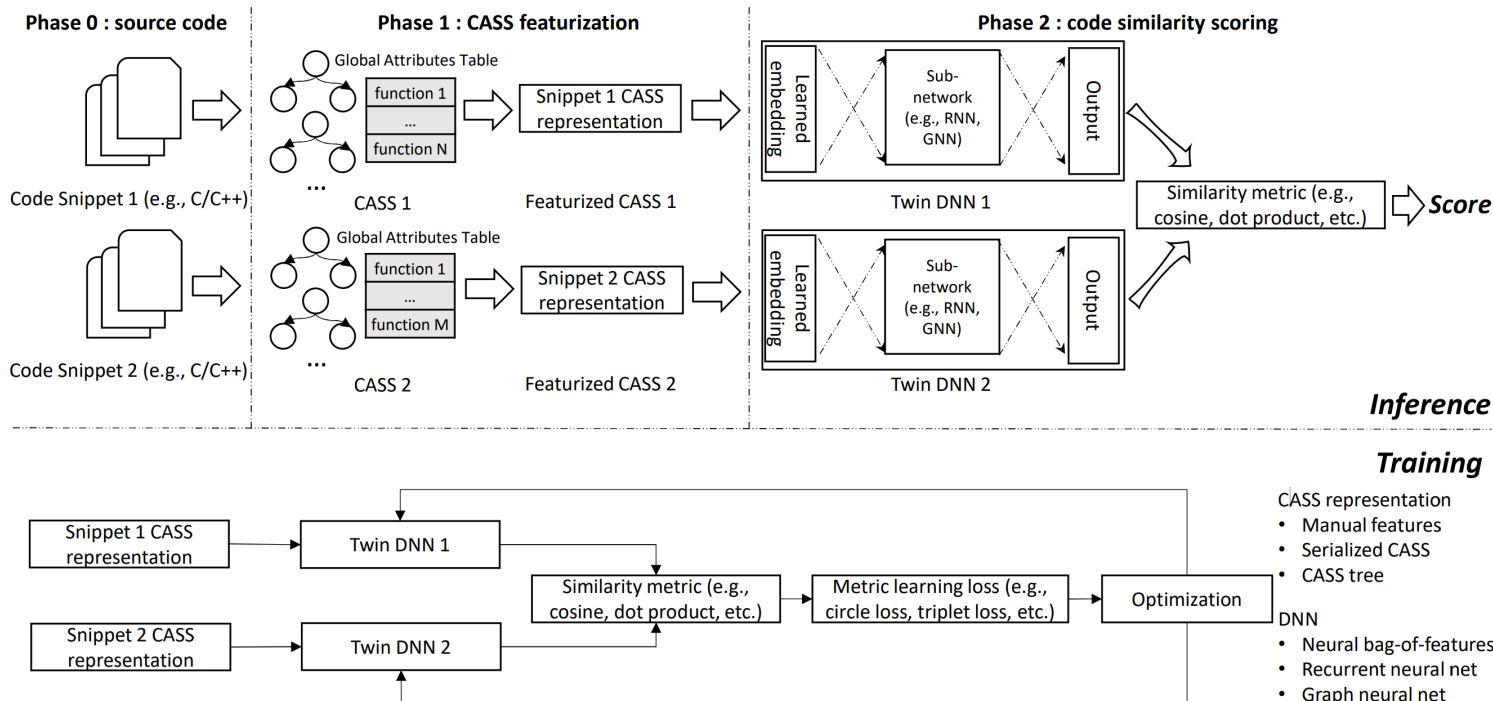
### PREPRINT

### 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×).

# MACHINE INFERRED CODE SIMILARITY (MISIM)

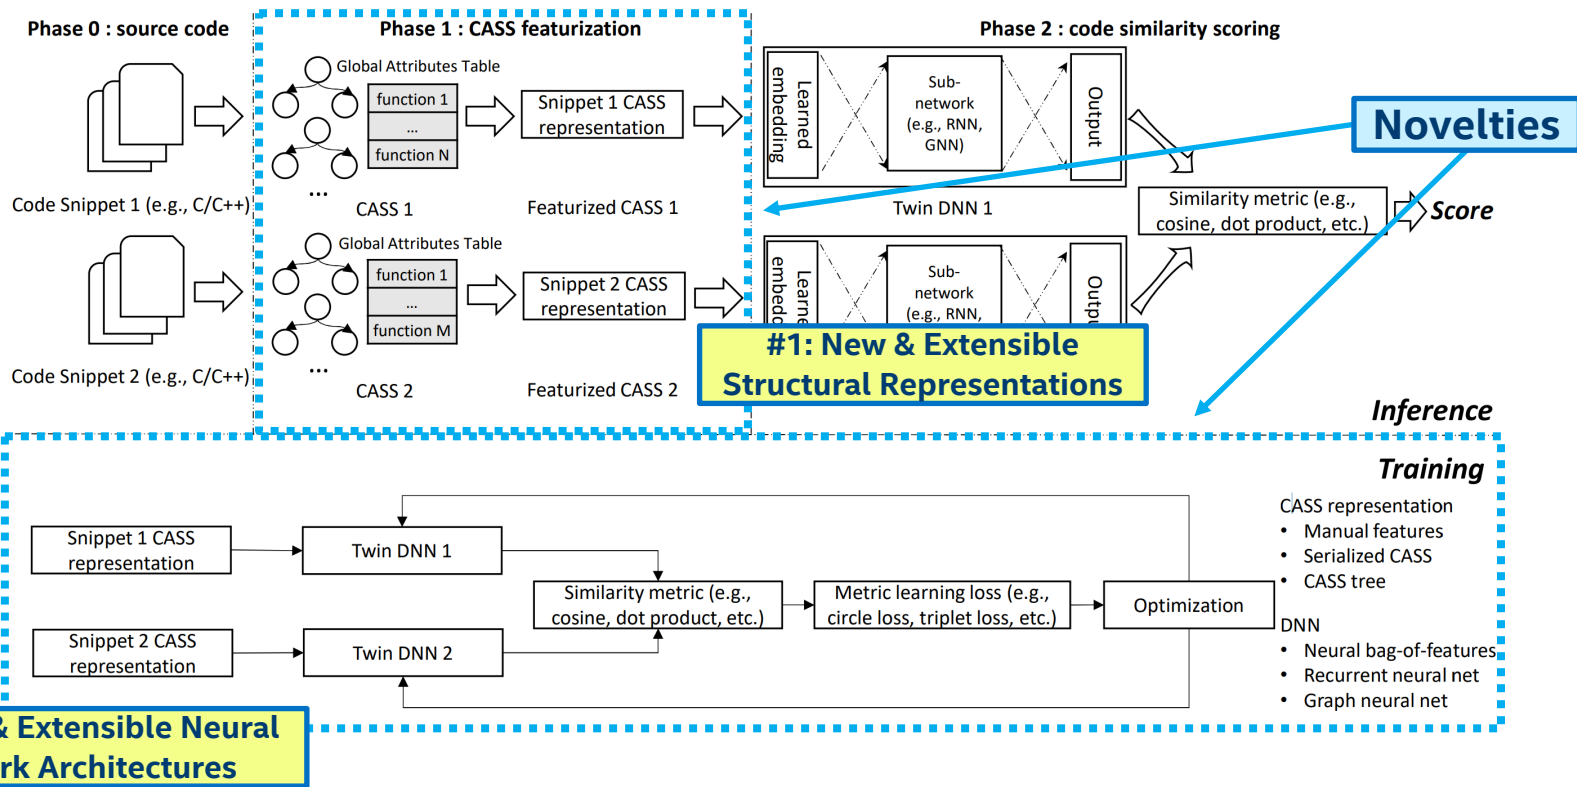
## MISIM System Design



MISIM created by Intel, Georgia Tech, and MIT

# MACHINE INFERRED CODE SIMILARITY (MISIM)

## MISIM System Design



# NOVEL STRUCTURAL REPRESENTATIONS OF CODE

- Aroma introduced the *simplified parse tree (SPT)*
- MISIM introduced the *context-aware semantics structure (CASS)*
- Both intentionally moved away from classical structures (like AST)

*These structures have led to state-of-the-art accuracy*

## Take-away:

Historical code representations may restrict our thinking for pioneering research in MP. Let's not do that. 😊

QUESTIONS / COMMENTS: JUSTIN.GOTTSCHLICH@INTEL.COM

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

<b>Fangke Ye *</b> Intel Labs and Georgia Institute of Technology yefangke@gatech.edu		<b>Shengtian Zhou *</b> 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 Tathai</b> Intel Labs and MIT tathai@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	

### 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 semantics 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%).

PREPRINT

## Aroma: Code Recommendation via Structural Code Search

SIFEI LUAN, Facebook, USA  
DI YANG\*, University of California, Irvine, USA  
CELESTE BARNABY, Facebook, USA  
KOLUSHIK SEN†, University of California, Berkeley, USA  
SATISH CHANDRA, Facebook, USA

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.

OOPSLA 2019

## STRUCTURES THAT CAN ONLY BE LEARNED

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**

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

# AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

- SW / HW heterogeneity is creating multiplicative complexity



# AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY

## Verified Lifting: Target multiple programming languages

### 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) {
    int tmp = new 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;
        }
        tmp += rowbytes;
        src += rowbytes;
    }
}
```




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).

SIGGRAPH ASIA 2019

```
Func blur(Func dst, Func src, int rows, int cols) {
    // 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, y) = (tmp(r, x, y-1) + tmp(r, x, y) +
                  tmp(r, x, y+1)) / 3;
    return dst;
}
```

## Halide: Target multiple hardware compute devices

### Automatically Scheduling Halide Image Processing Pipelines

Ravi Teja Mullapudi\* 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.

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.

In this paper we address this problem by providing an algorithm

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

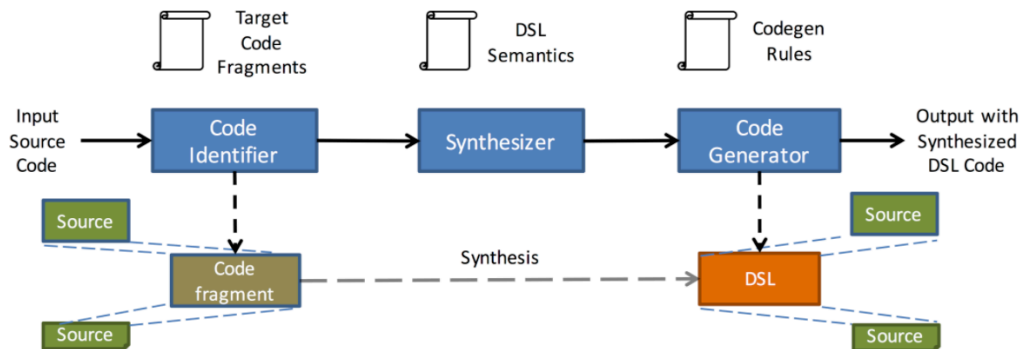
**Concepts:** \*Computing methodologies → Graphics systems and interfaces;

SIGGRAPH 2016

# VERIFIED LIFTING EVOLUTION: METALIFT


# 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



Shoaib 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

**AN OPEN QUESTION:**

**WHAT ARE THE QUALITY METRICS FOR  
HETEROGENEOUS TRANSLATION?**

**CORRECT & PERFORMANCE (OF COURSE)**

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

# **(NON-EXHAUSTIVE) TOPICS**

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**

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

# INTENTIONAL PROGRAMMING

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

# INTENTIONAL PROGRAMMING

- 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

# INTENTIONAL PROGRAMMING

- 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

If you haven't read this  
paper, please read it!

### 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

- Identify semantics of code
- Provide function semantics signatures
- Powerful & elegant

# INTENTIONAL PROGRAMMING WITH HALIDE

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

- Halide is a domain-specific language (DSL)
- Separation of concerns
  - Splits programming intention from programming adaptation

## 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

MICHAËL GHARBI, Adobe

BENOIT STEINER, Facebook AI Research

STEVEN JOHNSON, Google

KAYVON FATAHALIAN, Stanford University

FRÉDO DURAND, MIT CSAIL

JONATHAN RAGAN-KELLEY, UC Berkeley

We present a new algorithm to automatically schedule Halide programs for high-performance image processing and deep learning. We significantly improve upon the performance of previous methods, which considered a limited subset of schedules. We define a parameterization of possible schedules much larger than prior methods and use a variant of beam search to search over it. The search optimizes runtime predicted by a cost model based on a combination of new derived features and machine learning. We train the cost model by generating and featurizing hundreds of thousands of random programs and schedules. We show that this approach operates effectively with or without autotuning. It produces schedules which are on average almost twice as fast as the existing Halide autoscheduler without autotuning, or more than twice as fast with, and is the first automatic scheduling algorithm to significantly outperform human experts on average.

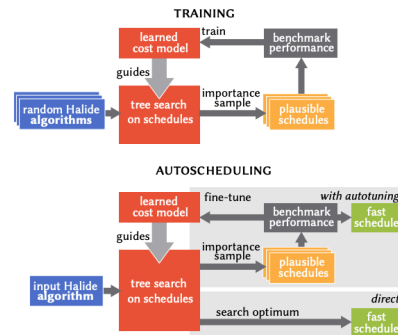


Fig. 1. We generate schedules for Halide programs using tree search over the space of schedules (Sec. 3) guided by a learned cost model and optional autotuning (Sec. 4). The cost model is trained by benchmarking thousands of randomly-generated Halide programs and schedules (Sec. 5). The resulting code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

SIGGRAPH 2019



# INTENTIONAL PROGRAMMING WITH HALIDE

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  
KAYVAN FATAHLIAN, Stanford University  
FRÉDO DURAND, MIT CSAIL  
JONATHAN RAGAN-KELLEY, UC Berkeley

SIGGRAPH 2019


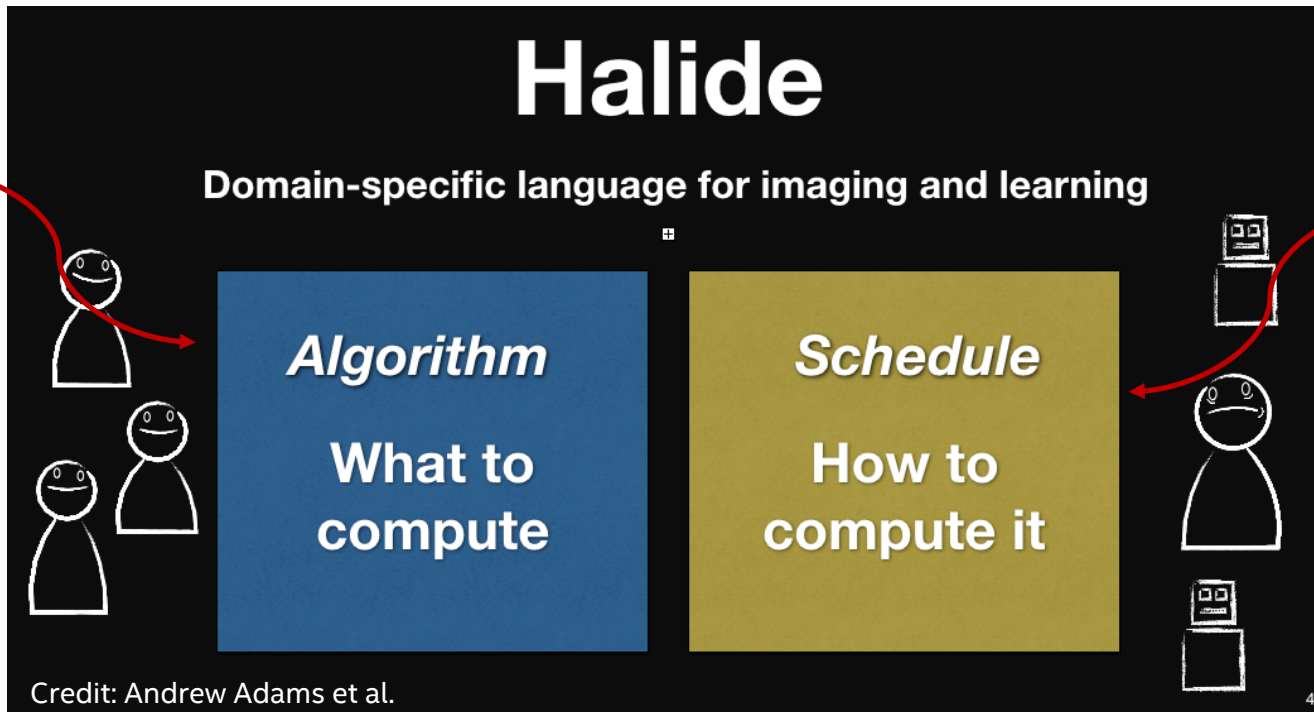


Fig. 1. We generate schedules for Halide programs using tree search over the space of schedules (Sec. 3) guided by a learned cost model and optional autotuning (Sec. 4). The cost model is trained by benchmarking thousands of randomly generated Halide programs and schedules (Sec. 3). The resulting code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


Intention



Adaptation

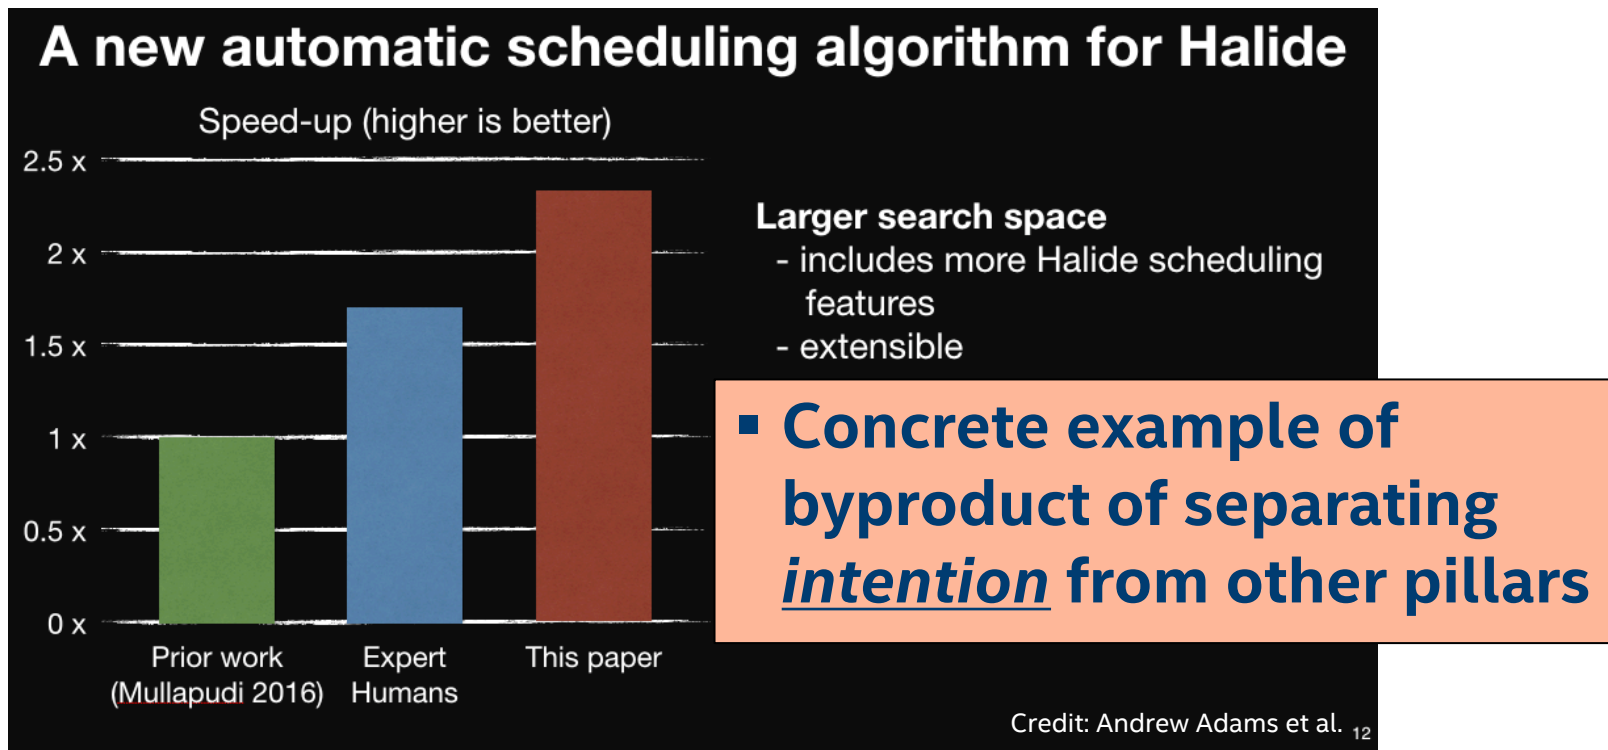
# INTENTIONAL PROGRAMMING CAN LEAD TO SUPER-HUMAN PERFORMANCE

## A new automatic scheduling algorithm for Halide



Credit: Andrew Adams et al. <sup>12</sup>

# INTENTIONAL PROGRAMMING CAN LEAD TO SUPER-HUMAN PERFORMANCE



# INTENTIONAL PROGRAMMING CAN LEAD TO SUPER-HUMAN PERFORMANCE

## MY BURNING QUESTIONS:

CAN WE BUILD GENERAL-PURPOSE INTENTIONAL PROGRAMMING SYSTEMS?

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

# (NON-EXHAUSTIVE) TOPICS

**MACHINE PROGRAMMING USES STOCHASTIC AND DETERMINISTIC METHODS**

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

**NOVEL STRUCTURAL REPRESENTATIONS OF CODE**

**AUTOMATION FOR SOFTWARE AND HARDWARE HETEROGENEITY**

**INTENTIONAL PROGRAMMING**

**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?



# 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?

# 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 and dense data problem in front of us
  - As of summer 2020, there were over 200M+ github repos
  - Code is multi-dimensional by nature
  - This data implies new frontiers of compute, communication, and data hardware

# THE ERA OF MACHINE PROGRAMMING IS NOW

We are on the verge of a revolutionary shift

Machine Programming is a Pioneering Research Initiative at Intel

# THE ERA OF MACHINE PROGRAMMING IS NOW

We are on the verge of a revolutionary shift

Machine Programming is a Pioneering Research Initiative at Intel

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 (like UWisc)

MP has the potential to change the rules for (almost) everything

# LET'S BECOME THE 100%

# LET'S BECOME THE 100%

## 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?

**A great way forward is to build a community (like Remzi et al. are)!**

**We need more of this; please help us spread the word**

**Please reach out to me if you are interested in collaborating!**

# THANK YOU!



Machine Programming Inside

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