

Pure Tensor Program Rewriting via Access Patterns (Representation Pearl)

Gus Henry Smith
gussmith@cs.washington.edu
University of Washington
Seattle, WA, USA

Scott Davidson
stdavids@cs.washington.edu
University of Washington
Seattle, WA, USA

Luis Ceze
luisceze@cs.washington.edu
University of Washington
Seattle, WA, USA

Andrew Liu
andy99@cs.washington.edu
University of Washington
Seattle, WA, USA

Joseph McMahan
jmcman@cs.washington.edu
University of Washington
Seattle, WA, USA

Zachary Tatlock
ztatlock@cs.washington.edu
University of Washington
Seattle, WA, USA

Steven Lyubomirsky
sslyu@cs.washington.edu
University of Washington
Seattle, WA, USA

Michael Taylor
profmbt@cs.washington.edu
University of Washington
Seattle, WA, USA

Abstract

Tensor kernels in machine learning (ML) often correspond to pure mathematical expressions, making term rewriting an attractive strategy for optimization and mapping to specialized hardware accelerators. However, existing ML intermediate representations (IRs) tend to either be *pure but high-level*, making low-level rewrites to hardware targets inexpressible, or *low-level but impure*, hampering the use of term rewriting altogether.

This paper introduces Glenside, a pure IR whose core abstraction—the *access pattern*—enables low-level, layout-aware, hardware-centric program rewrites. We demonstrate how term rewriting in Glenside can be used to map program fragments to hardware accelerator invocations and automatically discover classic data layout transformations like `im2col`. Glenside establishes a new foundation for exploring further term rewriting techniques in optimizing low-level tensor programs.

CCS Concepts: • Software and its engineering → Domain specific languages; • Computing methodologies → Machine learning; • Theory of computation → Equational logic and rewriting.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MAPS '21, June 21, 2021, Virtual, Canada

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8467-4/21/06...\$15.00
<https://doi.org/10.1145/3460945.3464953>

Keywords: machine learning compilers, term rewriting

ACM Reference Format:

Gus Henry Smith, Andrew Liu, Steven Lyubomirsky, Scott Davidson, Joseph McMahan, Michael Taylor, Luis Ceze, and Zachary Tatlock. 2021. Pure Tensor Program Rewriting via Access Patterns (Representation Pearl). In *Proceedings of the 5th ACM SIGPLAN International Symposium on Machine Programming (MAPS '21)*, June 21, 2021, Virtual, Canada. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3460945.3464953>

1 Introduction

Machine learning (ML) and other high-performance computing (HPC) applications increasingly rely on specialized hardware accelerators to provide speed and energy efficiency [17, 19, 32]. This trend has highlighted the need for flexible accelerator support in domain-specific compilers like Halide [30], TVM [9], TensorFlow/MLIR [1, 20], and PyTorch [28].

Adding accelerator support to an existing compiler typically uses custom pattern matching to map expensive tensor operations from applications down to accelerator invocations [11, 43]. Pattern matching often additionally relies on various other transformations to canonicalize intermediate representations (IRs) and massage data layouts into formats matching accelerator requirements [26]. Even with these changes, users may need to manually modify their application to help the compiler discover opportunities for dispatching operations to accelerators, such as by changing data types or unrolling loops.

In principle, term rewriting techniques [3] should be able to facilitate many of these transformation and mapping tasks within a compiler. Halide and TVM already rely on extensive rewrite systems for optimizing scalar computations and simplifying loop bounds in order to support further downstream optimizations [15, 24].




Figure 1. Four access patterns, representing different ways a tensor program (or *kernel*) might access the same 3D tensor. For example, (c) represents accessing a 3D tensor as a vector of 2D matrices.

Unfortunately, existing IRs in compilers for array/tensor programming DSLs tend to present abstraction and granularity mismatches that hamper term rewriting approaches. Term rewriting is most easily applied in *pure* (side effect-free) IRs that support equational reasoning. At the same time, mapping to accelerators requires considering low-level hardware details like data layout. Existing pure IRs for ML frameworks are used primarily for high-level transformations (e.g., type elaboration and inlining) and do not expose low-level data layout details [33]. On the other hand, IRs used for crucial lower-level optimizations like operator fusion must support precise reasoning about memory use, and therefore are typically impure, hampering term rewriting.

To help mitigate such impedance mismatches, we present *Glenside*,¹ a pure tensor program IR that enables hardware-level term rewriting. *Glenside* is based on a simple *access pattern* abstraction that supports expressing and reasoning about data layout transformations via syntactic rewrite rules. When combined with standard arithmetic rewrites for per-tensor-element computations, access patterns enable implementing complex transformations for accelerator support as compositions of simple rewrites.

Tensors are traditionally characterized by their *shape*, an n -tuple of positive integers indicating the size of each of a tensor’s dimensions. Access patterns instead characterize each tensor with two shapes, e.g., $((x), (y, z))$, separating the dimensions which are *iterated over* from the dimensions which are *computed on*. Figure 1(c) depicts an example where a 3D tensor’s first dimension is iterated over and some computation applied to each corresponding 2D matrix.

We demonstrate how *Glenside* enables implementing representative hardware-level transformation via term rewriting, including mapping computations to systolic arrays [17] (a common hardware module in ML accelerators) and automatically discovering the *im2col* data layout transformation [7], which enables mapping 2D convolutions to matrix multiplication hardware. In particular, by employing *equality saturation* [42], these transformations “fall out for free” (i.e., without any carefully crafted rewrite orderings [41]), from a

¹Publicly available at <https://github.com/gussmith23/glenside>.

handful of general rewrites concerning tensor transposition, Cartesian product, dot product, etc., expressed in terms of access patterns.

To summarize, our contributions include:

- *Access patterns*, a tensor representation that employs a simple, extended tensor shape type to distinguish iteration and computation dimensions
- The *Glenside* IR, a pure compiler IR that facilitates term rewriting to enable support for specialized accelerators
- A library of generic rewrites over *Glenside* programs
- Case studies demonstrating how *Glenside* enables automatically discovering key transformations for mapping applications to custom accelerators via equality saturation with the *egg* [42] library.

The rest of the paper is organized as follows: Section 2 provides background and briefly surveys closely related work. Section 3 motivates *Glenside* via a running example exploring pure matrix multiplication. Section 4 details the design and implementation of *Glenside*. Section 5 details case studies showing the potential benefits of *Glenside*’s term rewriting approach to low-level tensor program transformations.

2 Background and Related Work

Glenside is designed to help target tensor hardware accelerators and builds on past work in tensor IRs and term rewriting.

2.1 Machine Learning Accelerators

A variety of accelerators [12, 17, 21, 22, 25] have been developed to provide efficient implementations of tensor operators for ML applications. These devices accelerate tensor operators through hardware parallelism, simultaneously applying related operations across many tensors in the accelerator’s memory (which are often laid out according to custom rules that facilitate hardware optimization). Tensor program compilers must translate expensive application code fragments down to accelerator invocations that adhere to these layout rules, which often involves both (a) higher-level transformations like tensor reshaping to match accelerator size bounds and loop unrolling to expose optimization opportunities, and (b) lower-level transformations like operator fusion and *im2col* to match accelerator calling conventions and even implement different operations using the same accelerator, e.g., on systolic arrays [7, 16].

2.2 Tensor IRs and Compilers

Tensor compilers for ML and HPC applications strive to balance clear, high-level operator semantics and support for the low-level optimizations necessary to target specialized accelerators. *Halide* [31] achieves this balance by separating operator *specifications* (what is computed) from *schedules* (how, when, and where each output element is generated). This style of separation has proven highly effective across both application domains and hardware targets; numerous

compilers including TVM [8], FireIron [14], LIFT [35], and Accelerate [5] follow variations of this strategy.

The specification/schedule separation approach allows the same high-level program (specification) to be flexibly optimized for and mapped to different hardware targets by applying different schedules. From this perspective, schedules represent different rewriting strategies to explore various loop ordering and memory layouts; in LIFT and Accelerate these take the form of functional combinators closely related to Glenside’s approach. As in classic term rewriting, experts must often carefully craft schedules for each target to achieve the best performance and mitigate phase ordering challenges [41], though recent projects have produced promising results in automatic scheduling [2, 10, 45].

Other tensor IRs like TACO [18], Keops [6], and COMET [37] rely on *index notation*² to concisely express tensor operators and simplify optimization by uniformly representing per-output-element computations. These approaches also rely on rewriting passes to generate kernel implementations specialized to tensor sparsity / density, operator combinations arising in the source application, and details of the target hardware. In Section 3 we discuss some of the tradeoffs of these approaches with respect to other rewriting strategies.

Finally, polyhedral compilers [34] like Tensor Comprehensions [39] and Tiramisu [4] optimize loop-filled programs by modeling loop nests as polyhedra and applying geometric transformations. The polyhedral approach exploits regular loop structure, but is also restricted to geometrically affine transformations. In contrast, term rewriting is neither guided nor restricted by geometric constraints, making these approaches broadly complementary.

2.3 Term Rewriting and Equality Saturation

Term rewriting is a classic program optimization technique [3] that relies on iteratively applying rewrite rules of the form $\ell \rightarrow r$: when part of a program matches the pattern ℓ under substitution σ , it is rewritten into $\sigma(r)$. This approach is ubiquitous, frequently used in both mainstream and DSL compilers to implement features including preprocessing, strength reductions, and peephole optimizations [13].

Classic term rewriting systems where rewrites are applied destructively suffer phase ordering problems [41]: the order in which rewrites are applied can enhance or severely diminish performance. Recent work has shown how program synthesis can help address this challenge in peephole optimizers like Halide’s scalar expression rewriter [24].

Advances in alternate rewriting techniques like equality saturation [36] also mitigate phase ordering by exploiting the e-graph data structure from SMT solvers to repeatedly

apply all rewrites simultaneously, thus obviating rule ordering considerations. In particular, the egg library [42] has been used to develop new synthesis and optimization tools across diverse domains [23, 27, 40], including DSP compiler vectorization [38] and tensor computation graphs [44].

Glenside provides the first tensor IR amenable to equality saturation by introducing access patterns to provide pure, higher order tensor kernel combinators that support rank-polymorphism without the need for binding structures like anonymous functions or index notation.

3 From Pure matMul to IR Design Goals

Applying functional techniques and term rewriting to tensor IRs requires careful design. For example, we must ensure that operators be compositional with respect to tensor shapes and that the representation support generic rules within the target rewrite engine. To highlight such constraints and motivate access patterns in Glenside, this section illustrates potential pitfalls with a simple matrix multiplication example.

3.1 Pure Matrix Multiplication

We write `f64` for the type of 64-bit floats and `[A]` for vectors over type `A`. Using this notation, we can specify operators like dot product and 2D matrix transpose as:

```
dotProd : [f64] * [f64] -> f64
trans2  : [[f64]] -> [[f64]]
```

Implementing 2D matrix multiplication on inputs P and Q requires computing an output matrix R where $R_{ij} = \sum_k P_{ik} Q_{kj} = P_i \cdot Q_j^T$. The need to compute `dotProd` for every pair of a row from P and a column from Q suggests `map` and Cartesian product operators which we might specify with:

```
map : (A -> B) * [A] -> [B]
cartProd : [A] * [B] -> [A * B]
```

Naively, we can almost implement matrix multiplication as:

```
matMul(P, Q) :=
  map(dotProd, cartProd(P, trans2(Q)))
```

However, the result type will have been flattened to just `[f64]`, making it impossible to compose with other matrix operators that expect `[[f64]]` inputs.

Our first problem is that the `cartProd` specification above “forgets” the shape of its arguments. We could change this specification to arrange the output as a matrix:

```
cartProd2D : [A] * [B] -> [[A * B]]
```

But this result type prevents directly mapping `dotProd`.³ Now the problem is that `map` only applies a computation by

²Index notation is closely related to “Einstein notation” where reduction indices are implicit.

³This simple type does not specify how `cartProd2D` orders its output relative to its input vectors. We assume the order expected for matrix multiplication.

iterating over the first (outermost) dimension of a tensor. If we specialize `map` to iterate over the second dimension:

```
mapAt2 : (A -> B) * [[A]] -> [[B]]
```

then we can implement a compositional `matMul` operator that correctly produces results of type `[[f64]]` as:

```
matMul(P, Q) :=
  mapAt2(dotProd, cartProd2D(P, trans2(Q)))
```

3.2 Glenside Design Constraints and Goals

This style of pure, higher-order functional program representation enables term rewriting and equational reasoning via rules like:

```
dotProd(P, Q) ↔ dotProd(Q, P)
trans2(trans2(P)) ↔ P
map(f, map(g, P)) ↔ map(f ∘ g, P)
mapAt2(f, trans2(P)) ↔ trans2(mapAt2(f, P))
```

However, some of these rules depend on the shapes of dimension-specific operators aligning. What happens when we need to support higher-dimensional tensors? Without a mechanism to abstract which dimensions of a tensor are being iterated as opposed to computed over, we would have to generate versions of each rule for every combination of dimensions. Worse, these problems do not only affect rewrite rules; they also lead to code blowup just to specify all the variants of tensor kernels that arise in practice.

One strategy to address these challenges is adding support for anonymous functions (“lambdas”), currying, and closures to the tensor program representation. These features can provide sufficient flexibility to handle shape alignment issues that otherwise may require dimension-specific operators like `cartProd2D` and `mapAt2` above. For example, given curried versions of `dotProd` and `map`, we could have used such features to implement a curried `matMul` as:

```
matMul' P Q :=
  map' (λ r => map' (dotProd' r) (trans2 Q)) P
```

Alternatively, some IRs rely on index notation for even pithier implementations like:

```
matMul(P,Q)[i,j] := dotProd(P[i], trans2(Q)[j])
```

Unfortunately, these approaches all rely on some form of *name binding* which can significantly complicate term rewriting. Rewriting under binders, whether explicitly in the form of lambdas or implicitly with index notation, requires additionally analyzing the potential *contexts* (what names are bound to) of every subexpression. While it is still technically possible to apply state-of-the-art rewrite engines like `egg` [42] via explicit variable substitution rules and free variable analyses, we have found the additional complexity and rewrite search space blow up substantially eliminate the potential advantages of term rewriting in such IR designs.

All the above constraints inform Glenside’s key design goal: providing an IR that flexibly supports specifying and composing higher-order tensor operators⁴ over arbitrary dimensions while still enabling high-performance term rewriting techniques like equality saturation. In the rest of this paper, we show how *access patterns* enable achieving these goals with a focus on applications to mapping application fragments down to specialized hardware accelerators.

4 Glenside

This section details Glenside’s implementation, focusing on its core abstraction, *access patterns*. We use Section 3’s matrix multiplication as a running example throughout.

4.1 Access Patterns

Access patterns encode common tensor IR patterns where some tensor dimensions are *iterated over* (accessed) while others are *computed on*.⁵ Section 3’s `matMul` example *iterates over* dimension 0 of input P , while *computing on* dimension 1, effectively viewing P as a 1D vector of 1D vectors.

Access patterns are specified by their *shape* — a pair of tuples of positive integers (S_A, S_C) . An access pattern of shape (S_A, S_C) is, in turn, a tensor T whose shape is given by the concatenation of the access pattern shape tuples $S_A ++ S_C$; we refer to S_A and S_C as the *access* and *compute* dimensions of T , respectively.

Access patterns represent the view of an $(|S_A| + |S_C|)$ -dimensional tensor as a tensor of shape S_A , each of whose elements has shape S_C . For an access pattern T of shape (S_A, S_C) where $|S_A| = n_A$, we use the syntax `(access T n_A)` to represent T in Glenside. For example, if a 2D matrix T has shape (m, n) , then the Glenside expression `(access T 1)` yields an access pattern of shape $((m), (n))$.

The matrix multiplication example from Section 3 directly accesses the rows of P , but uses `trans2` to iterate over the columns of Q . Instead of requiring an explicit transpose operator, Glenside provides access pattern *transformers*.

4.2 Access Pattern Transformers

Access pattern transformers manipulate one or more access patterns to produce a new access pattern, allowing Glenside to support more complex patterns like slicing, transposing, and interleaving. Table 1 lists Glenside’s transformers.

To produce an access pattern representing the columns of Q for matrix multiplication, we employ the transpose transformer. It takes an access pattern and a list of dimension indices, and rearranges the dimensions of the access pattern in the order specified by the indices. If Q has shape (N, O) , `(transpose (access Q 1) (list 1 0))` produces an access pattern of shape $((O), (N))$.

⁴As `map` and `mapAt2` in Section 3.1 illustrate, an IR can support higher-order operators without necessarily providing lambdas, currying, or closures.

⁵This is similar to NumPy’s concept of *universal functions*.

Table 1. Glenside’s access pattern transformers.

Transformer	Input(s)	Output Shape
access	$((a_0, \dots), (\dots, a_n))$ and non-negative integer i	$((a_0, \dots, a_{i-1}), (a_i, \dots, a_n))$
transpose	$((a_0, \dots), (\dots, a_n))$, ℓ (a permutation of $(0, \dots, n-1)$)	$((a_{\ell_0}, \dots), (\dots, a_{\ell_n}))$
cartProd	$((a_0, \dots, a_m), (c_0, \dots, c_p)), ((b_0, \dots, b_m), (c_0, \dots, c_p))$	$((a_0, \dots, a_m, b_0, \dots, b_m), (2, c_0, \dots, c_p))$
windows	$((a_0, \dots, a_m), (b_0, \dots, b_n))$, window shape (w_0, \dots, w_n) , strides (s_0, \dots, s_n)	$((a_0, \dots, a_m, b'_0, \dots, b'_n), (w_0, \dots, w_n))$, where $b'_i = \lceil (b_i - (k_i - 1)) / s_i \rceil$
slice	$((a_0, \dots), (\dots, a_n))$, dimension index d , bounds $[l, h)$	$((a'_0, \dots), (\dots, a'_n))$ with $a'_i = a_i$ except $a'_d = h - l$
squeeze	$((a_0, \dots), (\dots, a_n))$, index d where $a_d = 1$	$((a_0, \dots), (\dots, a_n))$ with a_d removed
flatten	$((a_0, \dots, a_m), (b_0, \dots, b_n))$	$((a_0 \cdots a_m), (b_0 \cdots b_n))$
reshape	$((a_0, \dots, a_m), (b_0, \dots, b_n))$, access pattern shape literal $((c_0, \dots, c_p), (d_0, \dots, d_q))$	$((c_0, \dots, c_p), (d_0, \dots, d_q))$, if $a_0 \cdots a_m = c_0 \cdots c_p$ and $b_0 \cdots b_n = d_0 \cdots d_q$
pair	two access patterns of shape $((a_0, \dots), (\dots, a_n))$	$((a_0, \dots), (2, \dots, a_n))$

Table 2. Glenside’s access pattern operators.

Operator	Type	Description
reduceSum	$(\dots) \rightarrow ()$	sum values
reduceMax	$(\dots) \rightarrow ()$	max of all values
dotProd	$(t, s_0, \dots, s_n) \rightarrow ()$	eltwise mul; sum

The `cartProd` transformer takes access patterns of shapes $((a_0, \dots, a_n), (c_0, \dots, c_p))$ and $((b_0, \dots, b_m), (c_0, \dots, c_p))$ respectively, and produces an access pattern of the shape $((a_0, \dots, a_n, b_0, \dots, b_m), (2, c_0, \dots, c_p))$, where $(2, c_0, \dots, c_p)$ represents a 2-tuple of the input access patterns’ compute dimensions. The access dimensions of the input access patterns are simply concatenated. In the matrix multiplication example, the Cartesian product of the rows of P with the columns of Q is an access pattern of shape $((M, O), (2, N))$, where the second shape represents a 2-tuple of a row from P with a column from Q .

We have nearly re-implemented matrix multiplication example in Glenside. The final step is to implement the dot product, for which Glenside uses access pattern *operators*.

4.3 Access Pattern Operators

Operators are the only Glenside constructs which perform computation. They are invoked only in compute expressions, which map the operator over the compute dimensions of an access pattern. For an input access pattern A of shape $((s_0, \dots, s_{m-1}), (s_m, \dots, s_n))$, and an operator f with type $(s_m, \dots, s_n) \rightarrow (s'_m, \dots, s'_n)$, the result of $(\text{compute } f \ A)$ will have shape $((s_0, \dots, s_{m-1}), (s'_m, \dots, s'_n))$; that is, a compute expression cannot change the access dimensions of the input access pattern. Table 2 lists the operators in Glenside.

Recall where we are in converting our matrix multiplication example: we have accessed the rows of P and the columns of Q and taken their Cartesian product, resulting in an access pattern of shape $((M, O), (2, N))$, and we need

now to compute the dot product of these row-column pairs. In Glenside, the `dotProd` operator (see Table 2) does just that. To compute the dot product over our row-column pairs, we need only to apply compute `dotProd` to our access pattern, to produce an access pattern with final shape $((M, N), ())$. The entire Glenside specification of matrix multiplication is shown in Figure 2b.

5 Case Studies

To demonstrate Glenside’s utility, we first show how it enables concise specifications of several critical ML kernels (Section 5.1). We then show how Glenside’s pure, binder-free representation enables mapping kernels to an example accelerator via direct application of generic rewrite rules (Section 5.2). Finally, we highlight how Glenside enables the flexible mapping of larger, more diverse kernels to our accelerator, utilizing the power of equality saturation to automatically discover a variety of program transformations. Specifically, we show how Glenside can automatically map convolutions to matrix multiplications (Section 5.3) and automatically map large matrix multiplications into a sequence of smaller matrix multiplications (Section 5.4).

5.1 Representation of Common ML Kernels

Figure 2 lists the Glenside specifications of three common ML kernels: 2D convolution, matrix multiplication, and max pooling. Below, we discuss the specifications of 2D convolution and max pooling; see Section 4 for a description of matrix multiplication.

2D Convolution. 2D convolution (`conv2d`) is a core kernel in deep learning, defined element-by-element over tensors storing activations A , strides S , and weights W as:

$$\text{out}[n, o, x, y] = \sum_{dx, dy, c} (A[n, c, S[0] \cdot x + dx, S[1] \cdot y + dy] \cdot W[o, c, dx, dy])$$

<pre> (transpose ; ((N, O, H', W'), ()) (squeeze ; ((N, H', W', O), ()) (compute dotProd ; ((N, 1, H', W', O), ())) (cartProd ; ((N, 1, H', W', O), (2, C, K_h, K_w)) (windows ; ((N, 1, H', W'), (C, K_h, K_w)) (access activations 1) ; ((N), (C, H, W)) (shape C Kh Kw) (shape 1 Sh Sw)) (access weights 1))) ; ((O), (C, K_h, K_w)) 1) (list 0 3 1 2)) </pre> <p style="text-align: center;">(a) 2D convolution.</p>	<pre> (compute dotProd ; ((M, O), ())) (cartProd ; ((M, O), (2, N)) (access activations 1) ; ((M), (N)) (transpose ; ((O), (N)) (access weights 1) ; ((N), (O)) (list 1 0))) </pre> <p style="text-align: center;">(b) Matrix multiplication.</p> <pre> (compute reduceMax ; ((N, C, H', W'), ())) (windows ; ((N, C, H', W'), (K_h, K_w)) (access activations 2) ; ((N, C), (H, W)) (shape Kh Kw) (shape Sh Sw))) </pre> <p style="text-align: center;">(c) Max pooling.</p>
--	--

Figure 2. Common tensor kernels from machine learning expressed in Glenside. Lines containing access patterns are annotated with their access pattern shape. N is batch size; H/W are spatial dimension sizes; C/O are input/output channel count; K_h/K_w are filter height/width; S_h/S_w are strides.

where n indexes the output batch, o indexes output channels, x/y index spatial dimensions, dx/dy index the convolutional window spatial dimensions, and c indexes input channels. 2D convolution slides each of the o filters of shape (c, dx, dy) through each possible (c, dx, dy) -shaped window of the input images. At each of these locations, an elementwise multiplication and reduction sum is computed.

The Glenside specification of `conv2d` is shown in Figure 2a. We access the weights as a vector of O filters and the activations as a vector of N images. We leave the filters as they are, but form windows of shape (C, K_h, K_w) over the activations using the `windows` access pattern transformer (Table 1). This produces an access pattern of shape $((N, 1, H', W'), (C, K_h, K_w))$, i.e., a batch of “images” of new spatial shape (H', W') , where every location is a window of the original input. Finally, we take the Cartesian product of the filters and the windows, compute their dot product, and squeeze and transpose the output into the correct layout.

Max Pooling. Max pooling, commonly used in ML to condense intermediate activations (“act” below), is defined as:

$$\text{out}[n, c, x, y] = \max_{dx, dy} (\text{act}[n, c, \text{strides}[0] \cdot x + dx, \text{strides}[1] \cdot y + dy])$$

Max pooling slides a window of shape (dx, dy) over all possible locations within the spatial (i.e., x and y) dimensions. At each window location, it reduces the window to a scalar with the max operator. The Glenside specification merely applies `reduceMax` over each two-dimensional window.

Discussion. Glenside separates the *computation* from the *data access patterns* in these kernels while exposing the simplicity of their computation—and the relative complexity of their data access. In all three kernels, the computation can

```

(compute dotProd (cartProd ?a0 ?a1)) ==>
  (systolicArray ?rows ?cols
   ?a0 (access (transpose ?a1 (list 1 0)) 0))
where ?a0 is of shape ((?batch), (?rows))
and ?a1 is of shape ((?cols), (?rows))

```

Figure 3. Our rewrite rewriting matrix multiplication to a systolic array invocation.

be described with a single operator; most of the specification entails setting up the data access pattern.

Furthermore, Glenside exposes similar structure between kernels; for example, both `conv2d` and matrix multiplication feature the expression `(compute dotProd (cartProd . . .))`. At their core, these kernels are performing the same computation, but with different patterns of data access. In Section 5.3, we exploit this similarity in structure when mapping kernels to hardware.

These kernels highlight the expressive power of access patterns. Consider the use of `windows` in `conv2d` and max pooling. Both kernels form windows differently: `conv2d` forms three-dimensional windows over the channels, height, and width dimensions, while max pooling forms two-dimensional windows over the height and width. Rather than passing configuration parameters to `windows`, Glenside attaches this information to the tensors themselves.

5.2 Mapping matMul to Accelerators

Glenside can be used to uncover opportunities to invoke accelerator components. Consider a weight-stationary systolic array, a common matrix multiplication architecture. A weight-stationary systolic array with r rows and c columns takes two lists of length- r vectors (the activations and

```

?a ⇒ (reshape (flatten ?a) ?shape)
(cartProd (reshape ?a0 ?shape0) (reshape ?a1 ?shape1)) ⇒ (reshape (cartProd ?a0 ?a1) ?newShape)
(compute dotProd (reshape ?a ?shape)) ⇒ (reshape (compute dotProd ?a) ?newShape)

```

Figure 4. Rewrites enabling the discovery of the `im2col` transformation.

weights, respectively), pairing each vector from one list with each vector from the other, and computes a dot product over each pair. The second list contains c vectors, while the first can be of any length.

Glenside’s purity allows us to implement this hardware mapping task using a term rewriting system, in which we rewrite a matching program pattern to an invocation of our systolic array. Our rewrite is shown in Figure 3, mimicking egg’s rewrite syntax. Tokens starting with a question mark (such as `?a0` in Figure 3) are variables in the pattern, bound by the left-hand side (LHS), and then used on the right-hand side (RHS). egg also allows for conditions on rewrites, which we print below our rewrites.

To design our rewrite, we first must design the LHS to match program patterns that resemble the data access pattern and compute pattern of our systolic array. Glenside is eminently suitable for this task, as it can express exactly the data access and computation pattern we described for the systolic array. Pairing all vectors from one list with all vectors from another and computing the dot product of the pairs is represented as `(compute dotProd (cartProd ?a0 ?a1))`, binding `?a0` and `?a1` to the input access patterns. We encode the systolic array’s constraints on the input shapes as a condition on the rewrite. Patterns which match the LHS are mapped to the RHS; in this case, we introduce a new `systolicArray` construct to represent the functioning of our systolic array. The shape of the systolic array is given by the `?rows` and `?cols` parameters, and the inputs are given as access patterns. Note how we also transform the second access pattern to more accurately convey how the actual systolic array hardware accesses the weight tensor: it reads it all at once (hence, `(access . . . 0)`), and expects it to be laid out in transposed form in memory. This added information—enabled by Glenside’s access patterns—provides richer data layout information, potentially helping future rewrites or code generation steps.

5.3 Flexible Mapping: Discovering `im2col`

The `im2col` transformation is a data layout optimization which enables computing `conv2d` on matrix multiplication hardware. The transformation involves instantiating the convolutional windows over the input activations directly in memory [7]. This leads to data duplication, but the resulting speedup more than offsets that overhead. In this case study, we show how a few general rewrites within Glenside lead to the *automatic rederivation* of the `im2col` transformation.

```

(transpose
 (squeeze
 (reshape ; ((N, 1, H', W'), O), ())
 (compute dotProd ; ((N · 1 · H' · W'), O), ()))
 (cartProd
 (flatten ; ((N · 1 · H' · W'), (C · Kh · Kw))
 (windows (access activations 1)
 (shape C Kh Kw) (shape 1 Sh Sw)))
 (flatten ; ((O), (C · Kh · Kw))
 (access weights 1))))
 ?shape) 1) (list 0 3 1 2))

```

Figure 5. An `im2col`-transformed `conv2d`, after the application of the rewrites in Figure 4 and just before the application of the systolic array rewrite.

Glenside’s representation underscores the structural similarity between `conv2d` and matrix multiplication, reflected also by the shared `(compute dotProd (cartProd . . .))` between `conv2d` and the LHS of the systolic array rewrite in Figure 3. Using this rewrite on `conv2d` would permit mapping it to the systolic array; however, the restrictions on the shape of `?a0` and `?a1` prevent its application. The systolic array has an activation access pattern of shape $((a), (b))$ and a weight access pattern of shape $((c), (d))$, while `conv2d` operates over access patterns of shape $((N, 1, H', W'), (C, K_h, K_w))$ and $((O), (C, K_h, K_w))$, respectively. Transforming the access pattern into a lower-dimensional form would enable the systolic array rewrite.

Figure 4 shows the rewrites which enable this transformation. We call the first rewrite an *exploratory* rewrite as it optimistically matches any access pattern expression. It flattens an access pattern and immediately reshapes it back to its original shape, thus preserving equality (see Table 1 for formal definitions). This exploratory rewrite introduces the flattening necessary to convert the higher-dimensional access patterns of `conv2d` into the access patterns matched by the systolic array rewrite. However, the reshape operator will still need to be moved before we can fire Figure 3’s systolic array rewrite. The second and third rewrites in Figure 4 take care of this; they implement *composition commutativity* of reshape with `cartProd` and `compute dotProd`, which “bubble” reshape operators up and out of expressions. These rewrites express general properties of these operators and are not specific to this task.

```

?a ⇒ (concat (slice ?a ?dim ?b0 ?b1)
        (slice ?a ?dim ?b1 ?b2) ?dim)

(cartProd ?a (concat ?b0 ?b1 ?dim)) ⇒
  (concat (cartProd ?a ?b0) (cartProd ?a ?b1) ?newDim)
if ?dim is an access dimension

(cartProd (concat ?a0 ?a1 ?dim0)
  (concat ?a2 ?a3 ?dim1)) ⇒
  (concat (cartProd ?a0 ?a2)
    (cartProd ?a1 ?a3) ?newDim)
if ?dim0 and ?dim1 are the same shape dimension

(compute dotProd (concat ?a0 ?a1 ?dim)) ⇒
  (concat (compute dotProd ?a0)
    (compute dotProd ?a1) ?dim)
if ?dim is an access dimension

(compute dotProd (concat ?a0 ?a1 ?dim)) ⇒
  (compute reduceSum (pair (compute dotProd ?a0)
    (compute dotProd ?a1)))
if ?dim is a shape dimension

```

Figure 6. Rewrites for blocking matMul.

These three rewrites work in concert to map conv2d to a systolic array. First,⁶ the exploratory rewrite flattens and reshapes all access pattern expressions. This includes the inputs to conv2d’s cartProd subexpression, which are flattened to shapes $((N \cdot 1 \cdot H' \cdot W')$, $(C \cdot K_h \cdot K_w)$) and $((O)$, $(C \cdot K_h \cdot K_w)$) and reshaped back to their original shapes. Next, the composition commutativity rewrites for cartProd and compute dotProd fire in sequence, bubbling the reshape up through the cartProd and dotProd expressions (shown in Figure 5). Finally, the systolic array rewrite completes the im2col transform. Glenside’s equality saturation based rewrite engine discovers these rewrites as the exploratory rewrite fires on every term and no rewrites are missed due to the absence of phase ordering.

This example highlights how, *with straightforward, generally applicable rewrites defined over Glenside*, equality saturation can emergently discover useful transformations that previously required expert insights to apply.

5.4 Flexible Mapping: matMul Blocking

Equality saturation can also be used with Glenside to emergently discover a matrix multiplication blocking scheme. Matrix multiplication blocking is the common strategy of breaking up a single, large matrix multiplication into smaller multiplications, by multiplying subsets of the input matrices and assembling the results to form the output matrix. This

⁶Since equality saturation explores rewrites non-destructively, the rewriting order here is purely for explanatory purposes.

```

(concat ; ((32,32), ()))
(concat ; ((16,32), ()))
  (compute reduceSum ; ((16,16), ()))
    (pair ; ((16,16), (2)))
      (compute dotProd ; ((16,16), ()))
        (cartProd ; ((16,16), (2,16)))
          (slice ; ((16), (16)))
            (slice (access activations 1) 0 0 16) 1 0 16)
          (transpose ; ((16), (16)))
            (slice
              (slice (access weights 1) 0 0 16) 1 0 16)
              (list 1 0))))
      (compute dotProd ; ((16,16), ()))
        (cartProd ; ((16,16), (2,16)))
          (slice ; ((16), (16)))
            (slice (access activations 1) 0 16 32) 1 0 16)
          (transpose ; ((16), (16)))
            (slice
              (slice (access weights 1) 0 16 32) 1 0 16)
              (list 1 0)))))) ...

```

Figure 7. A 32×32 matMul blocked into 16×16 matMuls. Only two of the eight total multiplications are shown.

is essential in practice, as systolic arrays are small (often between 16×16 and 256×256) while matrices in ML and HPC applications can be much larger.

As in Section 5.3, this transformation follows from an exploratory rewrite and associated “cleanup” rewrites. The exploratory rewrite used for blocking is shown at the top of Figure 6. Given an access pattern, this rewrite slices the access pattern into two pieces along a dimension and then concatenates them back together. The dimension as well as the division strategy are configurable. For this example, we assume for simplicity that we run this rewrite on every available dimension, that we divide each dimension perfectly in half, and that all dimensions are powers of 2 in size. Figure 6 gives rewrites for bubbling the introduced concat operators up through the expression, namely the compositional commutativity of concat with cartProd and compute dotProd. Starting from the matrix multiplication in Figure 2b, assuming input shapes of $(32, 32)$, the exploratory rewrite first slices and concatenates the access patterns at the input of cartProd. Then, using the commutativity rewrites, the resulting concats are bubbled up to produce the final expression in Figure 7. The effect of these rewrites is that the single 32×32 matMul becomes eight separate 16×16 matMuls, which are summed and concatenated to form the full output matrix. This case study demonstrates yet again that Glenside’s expressiveness allows a small set of rewrites to produce interesting and useful emergent transformations.

6 Conclusion

In this paper, we proposed *access patterns* as an abstraction to enable equality saturation style term rewriting for low-level tensor program mapping to hardware accelerators. Crucially, access patterns support specifying and composing higher-order, arbitrary dimension tensor operators without the need for binding structures like anonymous functions or index notation. We demonstrated the potential utility of access patterns in the Glenside IR through case studies showing how rewrites in Glenside can automatically uncover common layout transformations like `im2col` used for accelerator mapping. We are excited for the community to join in further exploring the potential applications of access patterns and to build additional optimizations on Glenside's foundations.

Acknowledgments

This work was sponsored by the Real Time Machine Learning (RTML) DARPA project, and the Applications Driving Architectures (ADA) and Center for Research in Intelligent Storage and Processing in Memory (CRISP) JUMP Research Centers, co-sponsored by SRC and DARPA. We thank Chandrakana Nandi for her extensive and enthusiastic editing. We also thank Sorawee Porncharoenwase, Luis Vega, and Max Willsey for their helpful comments.

A Broader Impact Statement

The ability to develop effective compiler support for specialized hardware accelerators in ML, and HPC more broadly, has generally been restricted to a handful of elite, well-resourced teams. This restriction slows hardware development and creates barriers to entry for teams in less privileged environments to contribute to and help guide the development of the field.

We believe that the access pattern abstraction and Glenside's approach to term rewriting for improving compiler support for custom accelerators will help advance both near-term practical and longer-term principled approaches to building flexible compiler infrastructure. In turn, we hope that this infrastructure will help contribute to a broader, more diverse, and more inclusive community of folks working together to build efficient technologies for social good.

Of course, all technology is political and it can be difficult to anticipate how future researchers and practitioners may apply Glenside. While the most obvious consequence of more efficient hardware utilization is better performance for users and lower environmental impact via decreased power consumption, it is also possible that access patterns and Glenside would enable the rapid obsolescing of current hardware platforms and therefore contribute to harmful electronic waste. This work could also stimulate demand for hardware customization by removing compiler development-related

overheads and ultimately lead to higher negative environmental impact similar to the situation with respect to custom ASICs for bitcoin mining [29].

Also, any improvement to ML efficiency or applicability may contribute to economic and privacy concerns arising from increased technology company monopolization as discussed in Zuboff's *The Age of Surveillance Capitalism* [46].

References

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 265–283. <https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf>
- [2] Luke Anderson, Andrew Adams, Karima Ma, Tzu-Mao Li, and Jonathan Ragan-Kelley. 2020. Learning to Schedule Halide Pipelines for the GPU. arXiv:2012.07145 [cs.PL]
- [3] Franz Baader and Tobias Nipkow. 1998. *Term Rewriting and All That*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139172752>
- [4] Riyadh Baghdadi, Jessica Ray, Malek Ben Romdhane, Emanuele Del Sozzo, Abdurrahman Akkas, Yunming Zhang, Patricia Suriana, Shoaib Kamil, and Saman Amarasinghe. 2019. Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code. *2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)* (Feb 2019). <https://doi.org/10.1109/cgo.2019.8661197>
- [5] Manuel M T Chakravarty, Gabriele Keller, Sean Lee, Trevor L. McDonnell, and Vinod Grover. 2011. Accelerating Haskell array codes with multicore GPUs. In *DAMP '11: The 6th workshop on Declarative Aspects of Multicore Programming*. ACM.
- [6] Benjamin Charlier, Jean Feydy, Joan Alexis Glaunès, François-David Collin, and Ghislain Durif. 2021. Kernel Operations on the GPU, with Autodiff, without Memory Overflows. *Journal of Machine Learning Research* 22, 74 (2021), 1–6. <http://jmlr.org/papers/v22/20-275.html>
- [7] Kumar Chellapilla, Sidd Puri, and Patrice Simard. 2006. High Performance Convolutional Neural Networks for Document Processing. In *Tenth International Workshop on Frontiers in Handwriting Recognition*, Guy Lorette (Ed.). Université de Rennes 1, Suvisoft, La Baule (France). <https://hal.inria.fr/inria-00112631> <http://www.suvisoft.com>.
- [8] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. 2018. {TVM}: An automated end-to-end optimizing compiler for deep learning. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*. 578–594.
- [9] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. 2018. TVM: An Automated End-to-End Optimizing Compiler for Deep Learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. USENIX Association, Carlsbad, CA, 578–594. <https://www.usenix.org/conference/osdi18/presentation/chen>
- [10] Tianqi Chen, Lianmin Zheng, Eddie Yan, Ziheng Jiang, Thierry Moreau, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. 2018. Learning to Optimize Tensor Programs. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (Carlsbad, CA, USA) (NIPS'18)*. USENIX Association, USA, 3393–3404.
- [11] Zhi Chen and Cody Yu. 2020. How to Bring Your Own Codegen to TVM. <https://tvm.apache.org/2020/07/15/how-to-bring-your-own->

- codegen-to-tvm.
- [12] Chen, Yu-Hsin and Krishna, Tushar and Emer, Joel and Sze, Vivienne. 2016. Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks. In *IEEE International Solid-State Circuits Conference, ISSCC 2016, Digest of Technical Papers*. 262–263.
- [13] Hubert Garavel, Mohammad-Ali Tabikh, and Imad-Seddik Arrada. 2018. Benchmarking Implementations of Term Rewriting and Pattern Matching in Algebraic, Functional, and Object-Oriented Languages - The 4th Rewrite Engines Competition. In *Proceedings of the 12th International Workshop on Rewriting Logic and its Applications (WRLA'18)*. Thessaloniki, Greece. <https://hal.inria.fr/hal-01883212>
- [14] Bastian Hagedorn, Archibald Samuel Elliott, Henrik Barthels, Rastislav Bodik, and Vinod Grover. 2020. Fireiron: A Scheduling Language for High-Performance Linear Algebra on GPUs. arXiv:2003.06324 [cs.PL]
- [15] Bastian Hagedorn, Johannes Lenfers, Thomas Kundenedhler, Xueying Qin, Sergei Gorchatch, and Michel Steuwer. 2020. Achieving High-Performance the Functional Way: A Functional Pearl on Expressing High-Performance Optimizations as Rewrite Strategies. *Proc. ACM Program. Lang.* 4, ICFP, Article 92 (Aug. 2020), 29 pages. <https://doi.org/10.1145/3408974>
- [16] Yangqing Jia. 2014. *Learning Semantic Image Representations at a Large Scale*. Ph.D. Dissertation. EECS Department, University of California, Berkeley. <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-93.html>
- [17] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, Rick Boyle, Pierre-luc Cantin, Clifford Chao, Chris Clark, Jeremy Coriell, Mike Daley, Matt Dau, Jeffrey Dean, Ben Gelb, Tara Vazir Ghaemmaghami, Rajendra Gottipati, William Gulland, Robert Hagmann, C. Richard Ho, Doug Hogberg, John Hu, Robert Hundt, Dan Hurt, Julian Ibarz, Aaron Jaffey, Alek Jaworski, Alexander Kaplan, Harshit Khaitan, Daniel Killebrew, Andy Koch, Naveen Kumar, Steve Lacy, James Laudon, James Law, Diemthu Le, Chris Leary, Zhuyuan Liu, Kyle Lucke, Alan Lundin, Gordon MacKean, Adriana Maggiore, Maire Mahony, Kieran Miller, Rahul Nagarajan, Ravi Narayanaswami, Ray Ni, Kathy Nix, Thomas Norrie, Mark Omernick, Narayana Penukonda, Andy Phelps, Jonathan Ross, Matt Ross, Amir Salek, Emad Samadiani, Chris Severn, Gregory Sizikov, Matthew Snelham, Jed Souter, Dan Steinberg, Andy Swing, Mercedes Tan, Gregory Thorson, Bo Tian, Horia Toma, Erick Tuttle, Vijay Vasudevan, Richard Walter, Walter Wang, Eric Wilcox, and Doe Hyun Yoon. 2017. In-Datacenter Performance Analysis of a Tensor Processing Unit. *SIGARCH Comput. Archit. News* 45, 2 (June 2017), 1–12. <https://doi.org/10.1145/3140659.3080246>
- [18] Fredrik Kjolstad, Shoaib Kamil, Stephen Chou, David Lugato, and Saman Amarasinghe. 2017. The Tensor Algebra Compiler. *Proc. ACM Program. Lang.* 1, OOPSLA, Article 77 (Oct. 2017), 29 pages. <https://doi.org/10.1145/3133901>
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*.
- [20] Chris Lattner, Mehdi Amini, Uday Bondhugula, Albert Cohen, Andy Davis, Jacques Pienaar, River Riddle, Tatiana Shpeisman, Nicolas Vasilache, and Oleksandr Zinenko. 2020. MLIR: A Compiler Infrastructure for the End of Moore's Law. arXiv:2002.11054 [cs.PL]
- [21] Stefano Markidis, Steven Wei Der Chien, Erwin Laure, Ivy Bo Peng, and Jeffrey S. Vetter. 2018. NVIDIA Tensor Core Programmability, Performance & Precision. 2018 *IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (May 2018). <https://doi.org/10.1109/ipdpsw.2018.00091>
- [22] T. Moreau, T. Chen, L. Vega, J. Roesch, L. Zheng, E. Yan, J. Fromm, Z. Jiang, L. Ceze, C. Guestrin, and A. Krishnamurthy. 2019. A Hardware-Software Blueprint for Flexible Deep Learning Specialization. *IEEE Micro* (2019), 1–1. <https://doi.org/10.1109/MM.2019.2928962>
- [23] Chandrakana Nandi, Max Willsey, Adam Anderson, James R. Wilcox, Eva Darulova, Dan Grossman, and Zachary Tatlock. 2020. Synthesizing Structured CAD Models with Equality Saturation and Inverse Transformations. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (London, UK) (PLDI 2020)*. Association for Computing Machinery, New York, NY, USA, 31–44. <https://doi.org/10.1145/3385412.3386012>
- [24] Julie L. Newcomb, Andrew Adams, Steven Johnson, Rastislav Bodik, and Shoaib Kamil. 2020. Verifying and Improving Halide's Term Rewriting System with Program Synthesis. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 166 (Nov. 2020), 28 pages. <https://doi.org/10.1145/3428234>
- [25] Nvidia. 2018. The NVIDIA Deep Learning Accelerator (NVDLA). <http://nvidia.org/>.
- [26] NVIDIA. 2020. Convolutional Layers User Guide. <https://docs.nvidia.com/deeplearning/performance/dl-performance-convolutional/index.html>.
- [27] Pavel Panchekha, Alex Sanchez-Stern, James R. Wilcox, and Zachary Tatlock. 2015. Automatically Improving Accuracy for Floating Point Expressions. *SIGPLAN Not.* 50, 6 (June 2015), 1–11. <https://doi.org/10.1145/2813885.2737959>
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. arXiv:1912.01703 [cs.LG] <https://arxiv.org/abs/1912.01703>
- [29] Shize Qin, Lena Klaußen, Ulrich Gellersdörfer, Christian Stoll, and Da Zhang. 2020. Bitcoin's future carbon footprint. arXiv:2011.02612 [econ.GN]
- [30] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. In *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation (Seattle, Washington, USA) (PLDI '13)*. ACM, New York, NY, USA, 519–530. <https://doi.org/10.1145/2491956.2462176>
- [31] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. *Acm Sigplan Notices* 48, 6 (2013), 519–530.
- [32] Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Sidharth Samsi, and Jeremy Kepner. 2019. Survey and Benchmarking of Machine Learning Accelerators. 2019 *IEEE High Performance Extreme Computing Conference (HPEC)* (Sep 2019). <https://doi.org/10.1109/hpec.2019.8916327>
- [33] Jared Roesch, Steven Lyubomirsky, Marisa Kirisame, Josh Pollock, Logan Weber, Ziheng Jiang, Tianqi Chen, Thierry Moreau, and Zachary Tatlock. 2019. Relay: A High-Level IR for Deep Learning. *CoRR abs/1904.08368* (2019). arXiv:1904.08368 <http://arxiv.org/abs/1904.08368>
- [34] A. Simbürger, S. Apel, A. Größlinger, and C. Lengauer. 2013. The potential of polyhedral optimization: An empirical study. In *2013 28th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. 508–518. <https://doi.org/10.1109/ASE.2013.6693108>
- [35] Michel Steuwer, Toomas Rimmelg, and Christophe Dubach. 2017. Lift: A Functional Data-Parallel IR for High-Performance GPU Code Generation. In *Proceedings of the 2017 International Symposium on Code Generation and Optimization (Austin, USA) (CGO '17)*. IEEE Press, 74–85.
- [36] Ross Tate, Michael Stepp, Zachary Tatlock, and Sorin Lerner. 2009. Equality Saturation: A New Approach to Optimization. In *Proceedings of the 36th Annual ACM Symposium on Principles of Programming Languages (POPL '09)*. 264–276. <https://doi.org/10.1145/1480881.1480915>

- [37] Ruiqin Tian, Luanzheng Guo, Jiajia Li, Bin Ren, and Gokcen Kestor. 2021. A High-Performance Sparse Tensor Algebra Compiler in Multi-Level IR. *arXiv:2102.05187* [cs.DC]
- [38] Alexa VanHattum, Rachit Nigam, Vincent T Lee, James Bornholt, and Adrian Sampson. 2021. Vectorization for Digital Signal Processors via Equality Saturation. (2021).
- [39] Nicolas Vasilache, Oleksandr Zinenko, Theodoros Theodoridis, Priya Goyal, Zachary DeVito, William S Moses, Sven Verdoolaege, Andrew Adams, and Albert Cohen. 2018. Tensor comprehensions: Framework-agnostic high-performance machine learning abstractions. *arXiv preprint arXiv:1802.04730* (2018).
- [40] Yisu Remy Wang, Shana Hutchison, Dan Suciu, Bill Howe, and Jonathan Leang. 2020. SPORES: Sum-Product Optimization via Relational Equality Saturation for Large Scale Linear Algebra. *Proc. VLDB Endow.* 13, 11 (2020), 1919–1932. <http://www.vldb.org/pvldb/vol13/p1919-wang.pdf>
- [41] Deborah L. Whitfield and Mary Lou Soffa. 1997. An Approach for Exploring Code Improving Transformations. *ACM Trans. Program. Lang. Syst.* 19, 6 (Nov. 1997), 1053–1084. <https://doi.org/10.1145/267959.267960>
- [42] Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha. 2021. egg: fast and extensible equality saturation. *Proceedings of the ACM on Programming Languages* 5, POPL (2021), 1–29.
- [43] Xuan Yang, Mingyu Gao, Qiaoyi Liu, Jeff Setter, Jing Pu, Ankita Nayak, Steven Bell, Kaidi Cao, Heonjae Ha, Priyanka Raina, Christos Kozyrakis, and Mark Horowitz. 2020. Interstellar: Using Halide’s Scheduling Language to Analyze DNN Accelerators. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (Lausanne, Switzerland) (*ASPLOS ’20*). Association for Computing Machinery, New York, NY, USA, 369–383. <https://doi.org/10.1145/3373376.3378514>
- [44] Yichen Yang, Phitchaya Mangpo Phothilimtha, Yisu Remy Wang, Max Willsey, Sudip Roy, and Jacques Pienaar. [n.d.]. Equality Saturation for Tensor Graph Superoptimization. *arXiv preprint arXiv:2101.01332* ([n. d.]).
- [45] Lianmin Zheng, Chengfan Jia, Minmin Sun, Zhao Wu, Cody Hao Yu, Ameer Haj-Ali, Yida Wang, Jun Yang, Danyang Zhuo, Koushik Sen, Joseph E. Gonzalez, and Ion Stoica. 2020. Anso: Generating High-Performance Tensor Programs for Deep Learning. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. USENIX Association, Banff, Canada, 863–879. <https://www.usenix.org/conference/osdi20/presentation/zheng>
- [46] Shoshana Zuboff. 2018. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (1st ed.).