A Data-Centric Optimization Framework for Machine Learning

Oliver Rausch*
Department of Computer Science
ETH Zurich, Switzerland
oliver.rausch@inf.ethz.ch

Tal Ben-Nun*
Department of Computer Science
ETH Zurich, Switzerland
talbn@inf.ethz.ch

Nikoli Dryden
Department of Computer Science
ETH Zurich, Switzerland
nikoli.dryden@inf.ethz.ch

Andrei Ivanov
Department of Computer Science
ETH Zurich, Switzerland
anivanov@inf.ethz.ch

Shigang Li
Department of Computer Science
ETH Zurich, Switzerland
shigang.li@inf.ethz.ch

Torsten Hoefler
Department of Computer Science
ETH Zurich, Switzerland
htor@inf.ethz.ch

arXiv:2110.10802v3 [cs.LG] 29 Aug 2022

ABSTRACT

Rapid progress in deep learning is leading to a diverse set of quickly changing models, with a dramatically growing demand for compute. However, as frameworks specialize performance optimization to patterns in popular networks, they implicitly constrain novel and diverse models that drive progress in research. We empower deep learning researchers by defining a flexible and user-customizable pipeline for optimizing training of arbitrary deep neural networks, based on data movement minimization. The pipeline begins with standard networks in PyTorch or ONNX and transforms computation through progressive lowering. We define four levels of general-purpose transformations, from local intra-operator optimizations to global data movement reduction. These operate on a data-centric graph intermediate representation that expresses computation and data movement at all levels of abstraction, including expanding basic operators such as convolutions to their underlying computations. Central to the design is the interactive and introspectable nature of the pipeline. Every part is extensible through a Python API, and can be tuned interactively using a GUI. We demonstrate competitive performance or speedups on ten different networks, with interactive optimizations discovering new opportunities in EfficientNet.

1 INTRODUCTION

The modern development of deep learning is spearheaded by the conflux of algorithms [7], data [19, 56], hardware [30, 49], and systems [1, 46]. Today, machine learning systems have become critical in light of the increasing complexity of models and the breadth of emerging hardware. Training performance in particular is key to both researcher productivity and reducing the environmental impact of machine learning [47, 55]. To this end, systems and compilers such as XLA [18], ONNX Runtime [39], and TorchScript [48] are widely used to automatically optimize training.

Most such compilers tend to focus on *operator-centric* optimization, that is, defining specific rule-sets based on a set of predefined building blocks. Thus, these frameworks tend to perform very well on popular neural networks, but often have limited performance improvements on new, unseen models — which many researchers would like to explore, but may not be able to train given the high compute costs. Indeed, this effect has led to a situation where the models that perform well on existing hardware and systems are more likely to succeed than others [21].

In this work, we propose to shift the paradigm from operators to their memory access patterns by performing data-centric Deep Neural Network (DNN) optimization. It is well-established that the performance of modern DNNs such as Transformers hinges on data movement minimization [26, 38], and that a certain FLOP reduction does not guarantee a matching speedup [57]. Up until now such optimizations have been performed by specialized engineering teams for specific architectures, remaining out of reach for most research groups [4].

We present **DaCeML**[1], a Data-Centric Machine Learning framework, which provides a simple, flexible, and customizable Python-based pipeline for optimizing training and empowering deep learning research. DaCeML seamlessly integrates with PyTorch [46] and ONNX [44] to enable accelerating and tuning existing models, both in evaluation and backpropagation. The input models are then optimized with a general-purpose transformation pipeline, which reduces data movement at fine and coarse grain, regardless of the internal computation. Lastly, DaCeML provides interfaces to programmatically and interactively guide the optimization process further.

Internally, the framework uses Stateful Dataflow Multigraphs (SDFGs) [5] as a data-centric, hierarchical, graph-based intermediate representation, which enables it to work with operators and data movement at all levels, from registers to distributed memory. The DaCeML optimization pipeline begins with a standard PyTorch or ONNX model and transforms computation through progressive lowering, using four levels of general-purpose transformations: coarse-grained; local data movement reduction; global data layout optimization; and hardware specialization. Data-centric optimizations manifest in different ways, especially when training is considered. One of the several unique controls DaCeML provides, for example, is choosing whether to recompute or retain data for backpropagation. With the included automatic optimizations, DaCeML often *already matches or outperforms modern frameworks*.

The central focus of DaCeML is the ability to then go further — it unpacks the DNN compiler-black box and puts the machine learning practitioner in the driver's seat. This starts with the visual, introspectable intermediate representation that makes finding and understanding performance issues, such as excessive data movement, intuitive. These are subsequently addressed by manipulating data movement or the data layouts of the parameters and intermediate storage. The SDFG IR allows automated or manual searches

*Both authors contributed equally to the paper

[1] https://github.com/spcl/daceml

to be performed on multiple granularities *simultaneously*, rather than in the traditional compiler-based fixed set of passes. These can be performed either using a Python API or interactively using the Visual Studio Code [40] IDE.

We provide an overview of DaCeML's design, its approach to progressive lowering, and its transformations, including novel optimizations not supported by other frameworks. In particular, we demonstrate up to 3.43× speedups over prior best on non mainstream activations; state-of-the-art performance with automatic optimization of a wide range of DNNs from various domains [6, 14, 20, 36, 42, 52, 57, 59, 61, 68], compared with PyTorch, TensorFlow + XLA, and JAX; and two guided optimization case studies on BERT [14] and EfficientNet [57], the former highlighting the importance of data layout and the latter showing that nonlocal data movement minimization can yield 1.33× speedup over cuDNN.

We make the following contributions:

(1) Design and implementation of a framework for data-centric lowering and multi-level optimization of DNNs for training.
(2) A formulation of symbolic automatic differentiation (AD) on a data-centric intermediate representation, which enables novel automatic synthesis of backward-mode operators for arbitrary operators, including lowered kernels.
(3) Integration of differentiation into the optimization pipeline, enabling joint optimization pre- and post-AD for data movement reduction and adaptive recomputation.
(4) Demonstration of a human-guided workflow for DNN compilation, in which performance engineers interactively guide compilation using a graphical interface. This achieves state-of-the-art performance on highly contested deep neural networks [14, 57] and exposes previously untapped optimization opportunities.

2 BACKGROUND AND RELATED WORK

The data-centric view is garnering attention in the DNN optimization community. Ivanov et al. [26] consider a data movement minimization approach to optimize training of transformers [14, 63], observing that the models are largely memory bound. While their work, which focuses exclusively on BERT [14], is entirely manual, it paves the way towards automating such analyses and optimization. Similarly, DeepSpeed [38] also provides manually-optimized primitives for transformers. Numerous other works have focused on specific optimizations [3, 15–17, 22, 29, 33–35, 45, 53, 54, 60, 62, 64, 65, 67, 69] that can be implemented as data-centric transformations.

In the rest of this section, we describe the state of the practice in DNN optimization, and highlight important distinctions between current approaches and the DaCeML multi-level optimization pipeline for training.

Training vs. inference compilation Inference optimization elides several concepts required in training. Most importantly, no automatic differentiation is needed. Secondly, as opposed to training, intermediate activations need not be stored for backpropagation, which drastically changes the optimization search space. Thirdly, several operators (e.g., batch normalization, dropout, convolution) behave differently during training. For example, batch normalization stores running statistics throughout the process, and convolutions that use basis transformations (FFT, Winograd) can pre-transform the weights once, since they will not change during inference. Lastly, apart from backpropagation, gradient updates and stochastic optimizer rules must be applied.

State-of-the-art DNN compilers Researchers have often used custom kernels, linear algebra primitives (e.g., BLAS), and vendor-optimized libraries (e.g., cuDNN [11], oneDNN [23]). Recently, with the modularization and diversity of DNNs, the focus is shifting towards using general compiler infrastructure in DNN optimization, which uses Just-in-Time (JIT) compilation to both optimize individual operators and fuse them.

The methods by which DNN compilers optimize (transform) code can be classified into three categories: *graph rewriting rules*, *expansions*, and *global passes*. Graph rewriting represents DNNs as DAGs and pattern-matches certain subgraphs to replace them with others. Expansions convert a known operation (e.g., batch normalization) directly into an explicit, pre-optimized version. Finally, global passes operate on the entire code (e.g., memory scheduling). Ultimately, the first two categories, and sometimes the third, are implemented on an operator-centric basis. Below, we discuss state-of-the-art DNN compilers and their inner workings.

XLA Advanced compiler infrastructure [18] used by TensorFlow [1] and JAX [8], which contains multiple intermediate representations (e.g., HLO, LLO) and various domain-specific expansion transformations and global passes. Apart from those, general purpose optimizations (such as dead-code elimination) are performed by the LLVM [31] and MLIR [32] (experimental) infrastructure. Prior to XLA, graph rewriting rules in TensorFlow were provided by a component called Grappler.

TorchScript (`torch.jit`) The main static optimization effort in PyTorch [46, 48, 51]. Through either tracing or Python introspection, TorchScript can convert PyTorch modules into a DAG-based IR, and contains a pass manager to manage transformations. As opposed to XLA, optimizations are focused on the forward pass, leaving backpropagation to the `autograd` module.

ONNX Runtime ONNX Runtime [39] provides graph rewrite rules and global passes, mostly focused on cleaning up artifacts resulting from the ONNX [44] format (e.g., constant folding) and algebraic fusion involving chained ONNX operators.

Inference frameworks Frameworks that maintain a graph structure also provide optimizing transformations. TVM [9] provides passes on its two IRs (Relay, TIR) based on statement visitors/mutators, similar to AST manipulation. TASO [28] also provides subgraph pattern rewrite rules, focused on linear algebra. OpenVINO [24] uses nGraph [12] with graph rewrite rules and single-node matching transformations to optimize inference. For interactivity, it allows users to deselect certain optimization passes on given node names through its command-line optimizer. TensorRT [43] and cuDNN [11] provide functionality to fuse elementwise operations to DNN primitives, the former for the purpose of accelerating inference and the latter for training as well.

Limitations Each of the aforementioned DNN compilers is affected by at least one of three limiting factors. First, most of the transformations rely on specific operator types, which change from network to network, and over time. Second, optimization of the backward pass and prior to differentiation are not well-explored in tandem, missing important potential reductions in data movement.

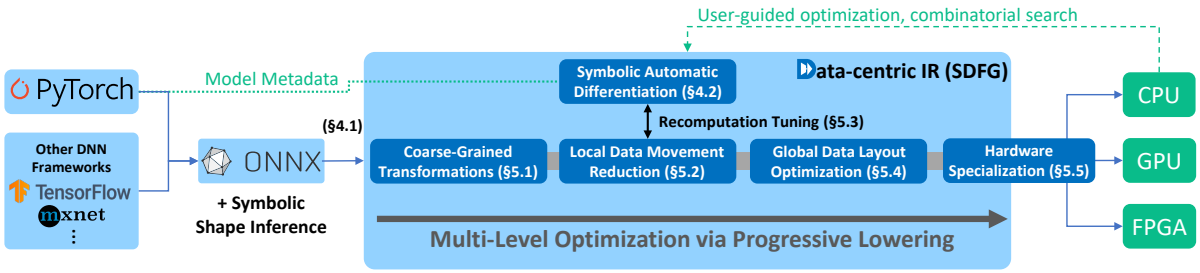

Figure 1: DaCeML system overview.

Lastly, DNN compiler optimizations are usually applied in a fixed set of passes, inhibiting opportunities for further tuning and analysis.

3 SYSTEM OVERVIEW

We propose to tackle the problem of DNN optimization through a holistic and user-centered approach. The Data-Centric Machine Learning framework (**DaCeML**) is a semi-automated, user-extensible pipeline for compiling and optimizing deep learning workloads. It takes PyTorch models and ONNX files, optimizes the program through data-centric means, and generates code with state-of-the-art performance for multiple platforms. The goal of DaCeML is threefold: (1) **Usability**: keeping the simplicity and expressiveness of PyTorch and its powerful model definition; (2) **Generality:** using general-purpose transformations that generalize well to new models, from coarse-grained optimizations to low-level code generation; and (3) **Interactivity**: translating the simplicity of defining models to optimizing them by enabling easy extensibility and interactive reasoning.

To achieve this goal, DaCeML defines a *progressive lowering* pipeline, depicted in Figure 1. DaCeML uses the data-centric Stateful DataFlow multiGraph (SDFG) representation [5] — a parametric graph IR designed for high-performance computing — as its intermediate representation. SDFGs promote separation of concerns between domain scientists and performance engineers by optimizing the data movement of a program separately from the algorithmic part of the input code. The IR and associated transformations are written in Python and have recently powered multiple applications in several domains, including quantum chemistry [70] and numerical weather prediction [13], achieving state-of-the-art performance on CPUs, GPUs, and FPGAs. This data-centric IR is fundamentally different from the existing task graph and multi-level representations. As opposed to LLVM (which uses opaque getelementptr accesses) and MLIR/XLA (which use memrefs that must be propagated globally for analysis), SDFGs enable *constant-time evaluation* of all data movement in a graph (array access indices, sub-tensors, etc.) as symbolic expressions.

Once within that representation, DaCeML can automatically optimize the forward or backward passes of a model to performance that is on par with (or faster than) state-of-the-art DNN compilers. Since most transformations operate on the simple elements of the SDFG IR (memlets, maps, tasklets, and access nodes) rather than coarse-grained operators, optimization generalizes across different operators.

Furthermore, DaCeML's API is geared towards extensibility — from defining new operator implementations, through local and global data movement planning, to low-level work partitioning in code generation — allowing users to assume direct control and guide the optimization process towards faster performance. The process can be done interactively, through APIs and IDE support, or through combinatorial searches over the optimization space.

4 DATA-CENTRIC PROGRESSIVE LOWERING

DaCeML supports loading models from the Open Neural Network eXchange (ONNX) format [44], or by wrapping PyTorch `nn.Modules` with a `@dace_module` decorator. The loaded models can be invoked directly, or used as optimized replacements of sub-networks in existing training and inference codes. Figure 2 provides a summary of the conversion process and SDFG representation.

4.1 From PyTorch/ONNX to Symbolic Dataflow

The ONNX format is a DAG representation of DNNs, which aims to standardize deep learning primitives across frameworks and networks. In the representation (Figure 2, middle), nodes represent operators and edges represent tensors that form data dependencies between them. As of version 1.9, there are 186 operator types in ONNX [44], which cover the vast majority of deep learning models.

Most deep learning frameworks, including PyTorch, have converters to the ONNX representation (typically used for deployment). For this reason, we chose to build DaCeML around the ONNX representation. When a PyTorch module is wrapped, we use PyTorch's ONNX exporter to generate the computational graph. As this is not sufficient for training and parameter management, we augment the ONNX input with metadata such as positional arguments, module hierarchy, and parameter structure. From those elements, we generate a replacement `nn.Module` object that is controlled by DaCeML and seamlessly integrates with the PyTorch `autograd` module for training. Internally, the module keeps either one (inference) or two (forward, backward) SDFGs for execution and optimization.

SDFG We use the example in Figure 2 to briefly explain the SDFG representation. Full details and operational semantics for SDFGs are provided by Ben-Nun et al. [5]. In particular, SDFGs are graphs of graphs, representing state machines of acyclic dataflow multigraphs. Each *state* in the state machine (light blue region) represents data access and computations (the equivalent of the full ONNX graph); the outer, state machine graph describes coarse-grained control flow (e.g., the training loop).

Two central concepts in SDFGs are the explicit separation of data movement from computation, and a multi-level view of data movement. The first concept is aided by the structure of the state graph itself — computations (*tasklets*, octagonal nodes) and *access*

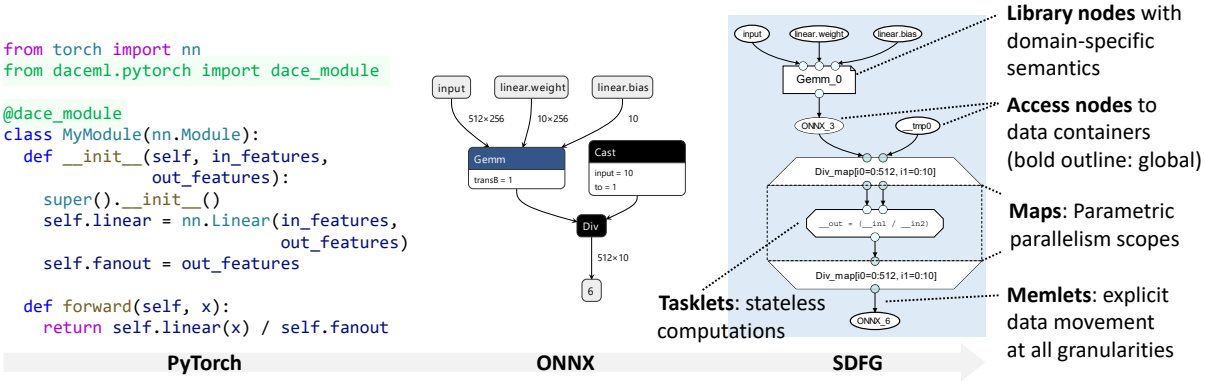

Figure 2: SDFG Intermediate Representation and conversion. ONNX graph rendered in Netron [50].

(oval nodes) to data containers are represented by nodes, and data movement is represented explicitly by the edges, called *memlets*. A tasklet cannot access non-local memory that is not connected to it via a memlet, and memlets are represented as symbolic ranges that connect access nodes and tasklets. Data container names are unique, and indicate disjoint memory regions, but access nodes that refer to them can appear multiple times. Furthermore, containers that are not local to the computation (i.e., cannot be removed in subsequent transformations) are called global memory and visually marked with a bold outline.

The concept of multi-level data movement is facilitated by three other types of nodes: (1) computations with predefined semantics are represented by *library nodes* (e.g., matrix multiplication generated from the `Linear` layer in the figure); (2) parallel regions, such as the `Div` operator in the figure, are grouped together between *map* (trapezoidal) nodes, which contain iteration ranges and a schedule (OpenMP loop, GPU kernel, FPGA region replication, and others); and (3) in case control flow is necessary within a parallel region, SDFGs can be nested in each other with *Nested SDFG* nodes.

ONNX library nodes The predefined semantics of library nodes allows them to be expanded to lower-level implementations, either "native" SDFG elements (tasklets, memlets, access nodes, and maps) or fast library calls (e.g., cuBLAS and cuDNN). As we shall show, native expansions can sometimes be optimized further than such libraries. In DaCeML, each ONNX operator is represented by a library node. Native implementations for most ONNX operators use the DaCe NumPy frontend to generate SDFGs. For example:

```
@python_pure_op_implementation
def Softplus(X, Y):
  Y[:] = numpy.log(1 + numpy.exp(X))
```

Authoring operator implementations using this approach is simple and familiar to ML practitioners, and requires little regard for performance. These unoptimized implementations are later transformed using the recipe described in Section 5.6.

As not all operators are natively implemented, and the ONNX standard is expected to grow, we also support a fallback for any ONNX operator, using ONNX Runtime [39] as a backend. Specifically, the expansions generate code that eagerly calls the operators within the compiled module.

Progressive lowering of library nodes Several library node expansions lower to other ONNX nodes. As an example, the `MatMul` node is lowered first to an Einstein sum (`Einsum`), which enables some tensor-contraction algebraic optimizations. Another is the `im2col` implementation of `Conv`, in which a `Gemm` node is included in the lowering output. This multi-level, progressive lowering reduces implementation redundancies and facilitates operator kernel authoring due to reuse.

4.2 Symbolic Automatic Differentiation

To compile and optimize for training, DaCeML uses source transformation based reverse-mode Automatic Differentiation (AD) to compute Vector-Jacobian Products (VJP) of the operators. Guided by the principle of generality, all lower level optimization on the backward pass is done using the same optimizing transformations as the forward pass. To enable this, the DaCeML AD engine transforms a forward pass SDFG into the corresponding SDFG that computes the required gradients. Critically, the differentiation is performed on the native SDFG, removing the need for manually specified VJPs. Furthermore, since all optimizations and lowering are performed on SDFGs, our AD engine is capable of differentiating scheduled and lowered SDFGs — a capability that, to the best of our knowledge, no existing framework has. By *optimizing pre- and post-AD jointly*, we unlock a novel class of optimizations not expressible in other DNN frameworks.

Differentiating SDFGs is challenging, since nodes not only represent computation, but memory accesses and parametric replication. It is important to differentiate the computation symbolically, even if all sizes are known, since the access pattern (memlets) within each operator is still symbolic.

Given an SDFG, the engine first determines the subgraph to differentiate, based on the output access nodes and the access nodes of inputs that require gradients (part of the metadata obtained from PyTorch). The subgraph is then traversed in reverse topological order and reversed as follows. *Access nodes* are reversed by "inverting" the access of the node, and replacing the data that is written/read with the adjoint of that data. *Maps* can be reversed by simply converting map entry nodes to map exit nodes and vice versa. The reverse of a code node — whether *tasklet, library node, or nested SDFG* — is a node that computes the VJP of the forward node. The inner tasklets

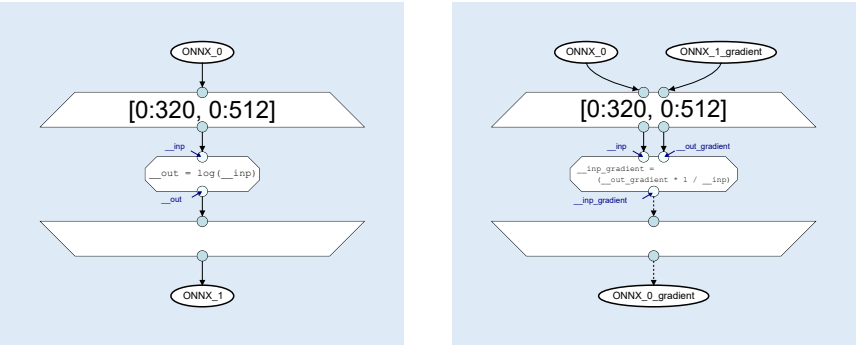

Figure 3: Elementwise logarithm SDFG (left) and symbolic auto-differentiated version (right).

of most machine learning operators work with scalar values. To automatically differentiate these scalar expressions, DaCeML uses the SymPy [37] symbolic differentiation engine (see example in Figure 3). Nested SDFGs are reversed recursively, and in the case of library nodes, they are lowered to their native SDFG implementation for further differentiation, with the possibility to provide manual backward expansions.

The VJPs of code nodes often require values that were inputs to the corresponding forward node. These values are either *forwarded* by storing the value in the forward pass and reading it in the backward pass, or recomputed in the backward pass.

Manual backward expansions DaCeML's AD engine is able to automatically differentiate almost all operator implementations, producing backward passes comparable to the hand-written implementations. However, it is often beneficial to specify manual backward operators (e.g., `Einsum` can be reversed directly for performance or `Softmax`/`LogSoftmax` for numerical stability). DaCeML provides this capability by manual expansions.

5 OPTIMIZATION PIPELINE

DaCeML follows DaCe's white-box optimization approach, in which optimization can be performed automatically as a starting point, and then is interactive and guided by performance engineers. We further simplify the interface and extend its capabilities, in order to allow performance engineers and ML practitioners to productively write transformations of their own.

DaCeML transformations are written in Python, and can use the existing graph rewriting tools from DaCe. A transformation can perform local replacement (complete example in Figure 4), graph pattern-matching, or global modifications. The DaCeML replacement API is graph-based: users provide the subgraph to match (`pattern`), and replacement instructions (`replacement`), which include returning a new node or subgraph and instructions on how to reconnect/augment the pattern graph (via the mapping in the return value). For more complex behavior, any method (exact matching condition, application, reconnection) can be overridden.

The transformations provided in DaCeML's standard set can be categorized into four distinct levels, ordered by their application as the graph is lowered: (1) Coarse-grained transformations on the ONNX-SDFG representation; (2) local data movement reduction, e.g., operator fusion and replication; (3) global (network-wide) data movement optimization; and (4) hardware specialization and workload partitioning. From those categories we sub-select transformations to apply globally on every network in our experiments, forming a "recipe" for DNN optimization, and show how guided application of the rest of them can yield even higher performance.

5.1 Coarse-Grained Transformations

The first step in transforming the SDFG is leveraging the semantics of ONNX operators. We can further categorize this into transformations that clean up the graph from framework- and ONNX-related clutter, and transformations that find subgraph combinations to improve the performance or numerical stability of the model.

Cleanup We provide a general ONNX transformation called `ConstantFolding`, which evaluates nodes whose inputs are not part of the weights or model inputs, and replaces the subsequent paths with the evaluated value. Following expansion to native SDFG, another transformation (`InputToConstant`) complements this by inlining constant inputs into their respective tasklets. This results in a chain of optimizations, for example, a 3.0 exponent in a `Pow` operator is first considered as a constant array; after cleanup, the code generator will subsequently replace the operation with much faster multiplications. Overall, this eliminates many unnecessary copy and computational operations, and enables more complex transformations, such as algebraic fusion.

Algebraic fusion Due to the progressive lowering approach taken in DaCeML, many linear algebra operations (e.g., matrix-vector multiplication, transpose, batched tensor contraction) are lowered into `Einsums`. We provide a set of transformations, such as horizontal and vertical fusions, which can, for example, fuse together a transposition into a single Einsum (e.g., `ij->ji` and `ik,jk->ij` to `ki,jk->ij`). This information potentially feeds into combinatorial searches of optimal data layouts — the shape of a weight can be modified through DaCeML if the Einsum does not generate an optimized BLAS call due to layouts (for which we provide a check). The ability to modify the shape of intermediates and weights is necessary to maximize performance in today's DNN architectures [26] and, to our knowledge, unique to DaCeML.

Lifting and omitting operations Other transformations relate to removing nodes that do not perform computations through symbolic size analysis, such as removing an Adaptive Pooling operator if the input/output sizes are the same, or fusing together padding and convolution operations. Known subgraphs generated

```python
class LiftLayerNorm(ReplacementTransformation):
    @classmethod
    def pattern(cls):
        path = make_onnx_path(nodes.AccessNode, onnx_op.ONNXReduceMean,
                              onnx_op.ONNXSub, onnx_op.ONNXPow,
                              onnx_op.ONNXReduceMean, onnx_op.ONNXAdd,
                              onnx_op.ONNXSqrt, onnx_op.ONNXDiv,
                              onnx_op.ONNXMul, onnx_op.ONNXAdd)
        path.add_edge(path.node(0), path.node(2))  # x as an input of x-mu
        path.add_edge(path.node(2), path.node(7))  # x-mu input of division
        return path

    def replacement(self, path, sdfg, state):
        axis = path[1].axes[0]  # get axis from ReduceMean_0

        ln_node = LayerNorm("LayerNorm", axis)
        return (ln_node,
                dict(X=(path[1], 'data'),     # <-- reconnection
                     weight=(path[8], 'B'),   #     instructions
                     bias=(path[9], 'B'),
                     Y=(path[9], 'C'),
                ))
```

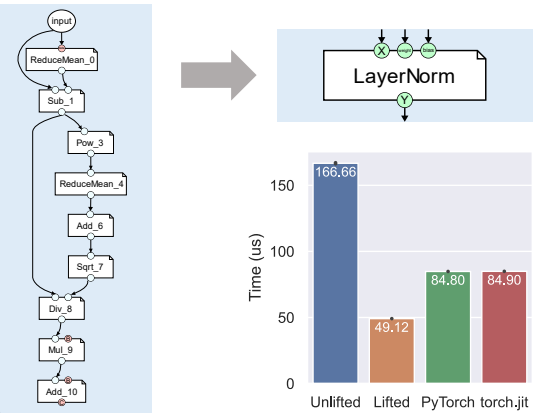

Figure 4: Coarse-grained transformation example (Section 5.1): lifting Layer Normalization from PyTorch.

by frameworks can be also be "lifted" to new nodes (e.g., with custom implementations or backward expansions), as is the case for Layer Normalization [2] in Figure 4.

5.2 Local Data Movement Reduction

Domain-specific transformations on the ONNX representation do not suffice for optimization, as they only capture operator-centric behavior and will not work with unseen operators or unexpected combinations thereof. We thus focus on transformations on the native SDFG (after ONNX library node expansion) that apply directly to ML workloads.

DaCe contains a library of standard transformations that use the structure of the graph to manipulate data movement. Of particular importance are `MapTiling` and `LocalStorage`, where the former partitions the workload of a map into tiles by introducing another nested map, and the latter adds an access node between two such maps in order to create storage only accessible by the current tile (e.g., placing weights in shared memory in Section 6.5). For DaCeML, we develop additional transformations to handle data movement bottlenecks in deep learning workloads.

Transformations Given an arbitrary subgraph, the `SubgraphFusion` transformation tries to combine multiple map scopes into one scope. This reduces write/read roundtrips to global memory and, on GPUs, kernel launch overhead. In elementwise operations where iteration spaces are equal, the operation is trivial. However, in many cases the spaces are permuted (when different data layouts are involved) or do not share the entire iteration domain (e.g., in the Layer Normalization operator, see Section 6.2). For this reason, we extend subgraph fusion to find permutations, offsets (e.g., `start:end` vs. `0:end-start`), parallel regions in reduction, and the greatest common subset of map iteration domains to fuse over. The transformation extracts them out first and then fuses the requested outer maps, greatly extending the realm of fusion possibilities.

Fusion space exploration To automate the process of optimizing data movement in graphs, we develop automatic optimization heuristics that enumerate the space of fusible subgraphs. Currently, DaCeML supports greedy enumeration with pruning (based on path constraints for fusion), and ranking according to scoring functions. For evaluation, we rank by largest neighboring regions to minimize data movement and GPU kernels.

5.3 Backpropagation and Data Movement

When optimizing the forward and backward passes at the same time, data-centric optimizations can control aspects of auto-differentiation. As AD requires intermediate value forwarding, the memory footprint can become infeasible and movement too demanding [10, 27]. However, fusing maps on the forward pass means omitting intermediate values, and replicating them means recomputation on the backward pass. This process can create a tunable "knob" to optimize backpropagation.

For this purpose, we develop two transformations: `TasskletFusion`, which fuses two computations into one symbolic tasklet; and `MapReplication`, a transformation that detects access nodes being read more than once, and replicates the immediate map leading to it. Example uses of the two on the Mish activation [41] and statistical normalization are shown in Section 6.

5.4 Global Data Layout Optimizations

After tuning local data movement, one can look beyond the traditional "peephole optimizations" and schedule data assignment and movement on the entire graph with data-centric transformations. Briefly, such transformations do not have a pattern to match, but view the entire graph and generate multiple options. One such optimization is `TransientReuse`, which detects arrays that have the same volume but are used in non-overlapping segments of time (using DAG level analysis). References to these arrays are replaced by a reference to a single memory region, and unused arrays are removed. Another potential example is finding an optimal set of data layouts for weights and intermediate values. Since DaCeML generates native SDFGs, it can generate optimized code for each layout, constrain the search space by ensuring layouts match, and prune it based on domain-specific constraints, e.g., by only considering BLAS-optimized layouts for Einsums.

Allocation and deallocation of memory can create significant overheads when happening within critical code. To avoid them, the SDFG provides fine-grained control over their lifetimes. For

example, when storage is annotated as Persistent, it will outlive multiple SDFG invocations.

5.5 Hardware Specialization

Lastly, we need to consider the underlying system we are compiling for. This may mean the accelerator architecture(s), or whether we are running on one or more nodes.

GPU specialization Data-centric transformations can offload code to different platforms by introducing copies and kernel code [5]. In DaCeML, we extend the transformation capabilities to partition workload efficiently on GPUs. `WarpTiling` is a variant of `MapTiling` that takes a GPU kernel map and divides its work across a warp. It detects reductions inside the map and inserts efficient warp-collective reductions as necessary. `Vectorization` allows using target-specific vector instructions. We extend the system to support reduced-precision vector types, provide fast implementations of vector-to-scalar reductions, and "fill in gaps" in math library functions by calling sequences of scalar operations.

Distributed computing DaCeML supports data-parallelism for distributed training. `DistDataParallel` adds a distributed `all-reduce` library node after each weight gradient access node in the backward-pass SDFG. Since the code generator traverses SDFGs in topological order, communication is performed as soon as the data are ready, promoting pipelining.

5.6 Automatic Transformations

Now that we have established a set of transformations, we can create an automated sequence of their application for arbitrary DNNs. Our transformation recipe used in the results is as follows: (1) Clean up the ONNX-based SDFG with constant folding and algebraic fusion; (2) lower ONNX nodes to implementations, choosing from ONNXRuntime, cuDNN, or PyTorch implementations depending on the node type; (3) in the resulting SDFG, inline nested SDFGs, remove all redundant copies, and greedily fuse subgraphs; (4) specialize the program to the hardware: in fused maps, hoist out initialization to kernel start, eliminate trivial (0, 1 element) maps, specialize memory types of transients and call `WarpTiling`; (5) flatten map dimensions to coalesce memory accesses, and vectorize maps to 128-bit accesses if applicable. When manually applying guided transformations, this recipe can be used as a starting point.

6 EVALUATION

We first demonstrate DaCeML's ability to optimize single operators using the recipe described in Section 5.6, and then examine how these perform when used as part of larger models.

Experimental setup We run each experiment at least 100 times and report the median value with a 95% nonparametric confidence interval. For measurement, we use a server with an Intel Xeon Gold 6140 CPU (2.30GHz), 768 GiB RAM, and an NVIDIA Tesla V100 GPU (16 GiB RAM). We use Python 3.8.8, DaCe 0.10.8, CUDA 11.4, cuDNN 8.2.0.54. PyTorch 1.8.1, ONNXRuntime 1.7.0, TensorFlow 2.5.0, TensorRT 8.0.1.6, TVM 0.8.0dev0 (commit dbfbebe), Triton 1.1.1, and JAX 0.2.13. The torch-based frameworks were all run from the same model source code, unmodified except for the addition of a decorator for `torch.jit` and DaCeML.

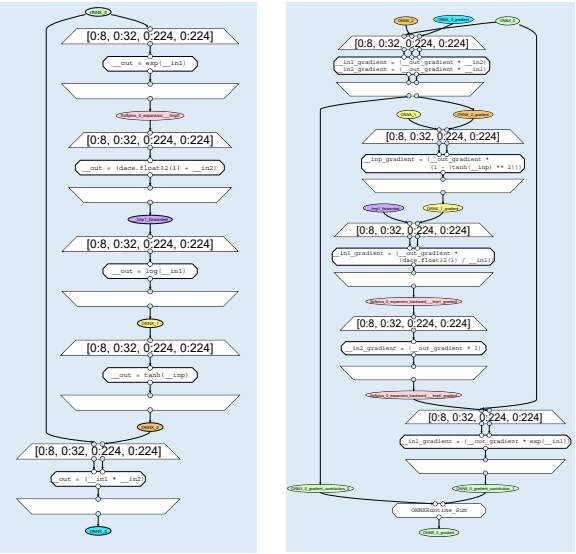

(a) Before Optimization (intermediate memory: 134.75MiB)

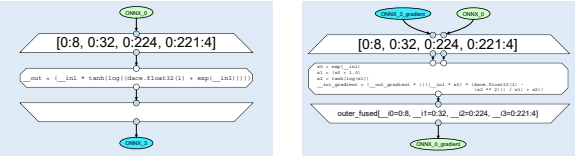

(b) After Optimization (intermediate memory: 0MiB)

Figure 5: Mish operator SDFG forward pass (left) and the resulting backpropagated graph (right). Colored access nodes correspond to matching data.

6.1 Mish Activation

YOLOv4 [6] is a widely-used network for object detection. Among other features, it relies on the Mish [41] activation: $\text{mish}(x) = x \tanh(\text{softplus}(x))$. As with many operators, Mish is not natively supported by ONNX or ONNX Runtime.

Running Mish for training in PyTorch yields the graph shown in Figure 5a. The operator consists of sequential elementwise application of the individual primitives used in the function (i.e., exp, tanh, log, +, $*$), with intermediate arrays generated to be retained for the backward pass, as can be seen on the right-hand side of the figure. These arrays are global and cannot be optimized once the AD process is complete, which is when they would typically be optimized on most DNN compilers.

The DaCeML generic transformation recipe lives both before and after AD, and yields an optimized SDFG prior to AD as shown in Figure 5b (left). As the kernels are entirely fused, the decision is made to recompute intermediate values for differentiation on a scalar level, rather than stash them for the backward pass. Differentiating the fused and scheduled kernel, a capability that current deep learning frameworks lack, produces a backpropagated graph that minimizes data movement, as shown in the figure (right). DaCeML is currently, to the best of our knowledge, *the only deep learning framework with this capability.*

```
@op_implementation(op="Softmax", name="native")
class NativeSoftmax(ONNXForward):
  @staticmethod
  def forward(node, state, sdfg):
    # Get property into scope
    axis = node.axis

    # Operator implementation in numpy
    def prog(input, output):
      max = input.max(axis=axis, keepdims=True)
      exp = np.exp(input - max)
      sum = exp.sum(axis=axis, keepdims=True)
      output[:] = exp / sum

    return program_for_node(prog, sdfg,
                            state, node)
```

(a) NumPy Implementation

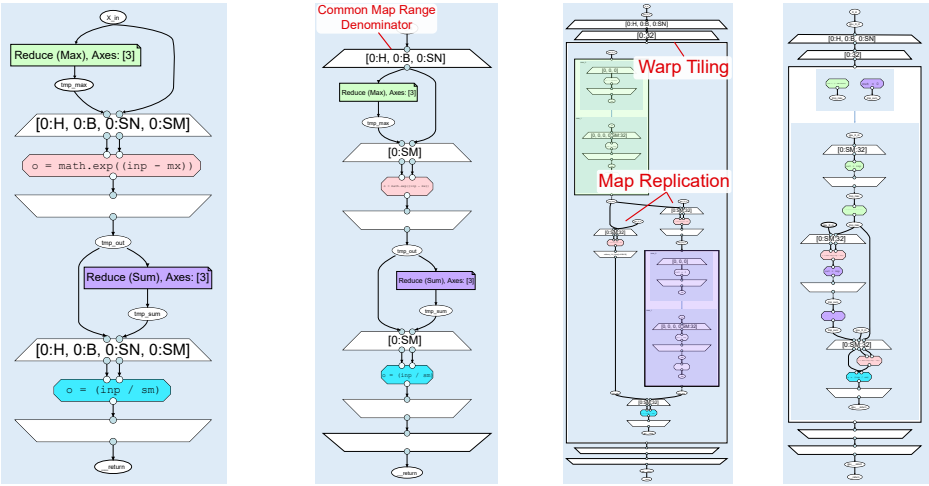

(b) Initial SDFG (c) Subgraph Fusion (d) Specialization (e) Final

Figure 6: Softmax implementation in DaCeML (colored sections correspond to same operators).

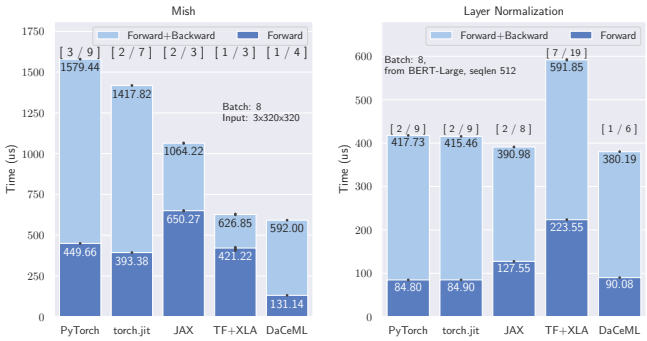

Figure 7: Runtime results for DNN operators. Brackets [·] indicate number of kernels launched in forward / fwd+bwd.

We evaluate Mish in isolation in Figure 7 (left). By fully fusing the graph *before* differentiation, DaCeML improves performance over all other tested frameworks. For forward evaluation, PyTorch produces three kernels: tanh, softplus, and multiplication. `torch.jit` and JAX are able to fuse two out of the three operations. TensorFlow+XLA manages to produce a single fused kernel but fails to eliminate global loads/stores that stash intermediate values for the backward pass. DaCeML achieves 3.43× improvements over PyTorch and 3× over `torch.jit`. In backpropagation, we similarly achieve 2.67× improvements over PyTorch and 1.06× over TF+XLA.

6.2 Statistical Normalization

Statistical normalization operations (e.g., Batch Normalization [25], Layer Normalization [2], Group Normalization [66]) are commonplace in modern DNNs. Despite the fact that all normalization operations are similar, their performance varies. All the aforementioned operations perform the general computation $\gamma \cdot \frac{x - \mathbb{E}[x]}{\sqrt{\mathrm{Var}(x) + \epsilon}} + \beta$, but on different dimensions or subsets thereof. We examine layer normalization [2], a primitive widely used in transformer models.

We find that PyTorch has two separate implementations for layer and batch normalization. The implementation of batch normalization can use cuDNN, while layer normalization uses a manually-optimized function. Thus, despite batch normalization being more expensive than layer normalization (due to saving running mean/-variance statistics to memory), the former is faster than the latter.

We evaluate the implementation of layer normalization following the DaCeML recipe, listing its performance in Figure 7 (right). The generated code outperforms all frameworks on the forward-and-backward pass, where JAX closely matches in backpropagation but misses several fusion opportunities. On the forward part, PyTorch is slightly faster due to DaCeML storing more information for backpropagation, which ends up being faster overall. Without this storage, DaCeML takes 49 μs (Figure 4). DaCeML's native SDFG is also fusible with neighboring operators, which delivers improved overall performance when used in larger models.

6.2.1 Softmax. This operator is present in many networks, both for classification and as a key component of attention mechanisms [63]. Even though it is computationally different from normalization, it exhibits a similar data movement pattern. Thus, the same data-centric transformations performed on layer normalization apply to softmax directly. Below, we dive into the transformations applied, showing how the representation and new DaCeML transformations can reduce data movement in nontrivial ways, regardless of the underlying computation.

The input NumPy implementation of the operator (Figure 6a) lowers the ONNX library node to four separate operations: two reductions (computing the maximum value to subtract; and summarizing the contributions for the denominator), and two elementwise operations (exponentiation, division). In the initial graph (Figure 6b), three extraneous tensors sized as the input would be generated. The second step of our recipe locally reduces this data movement.

Fusion of these operators is not trivial, as the dimensions of the maps and reductions differ. The greedy fusion strategy applied in DaCeML automatically expands reductions and finds common dimensions to extract out of the maps in the subgraph, which in

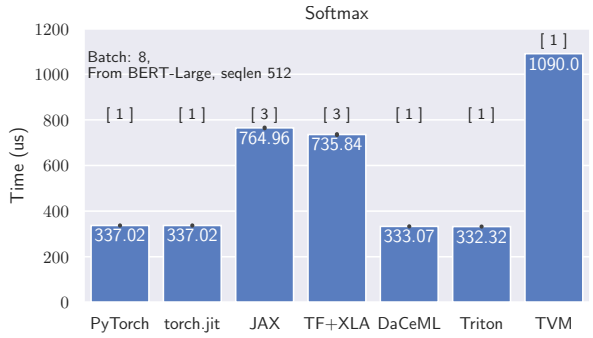

Figure 8: Runtime results for the softmax operator. Brackets [·] indicate number of kernels launched.

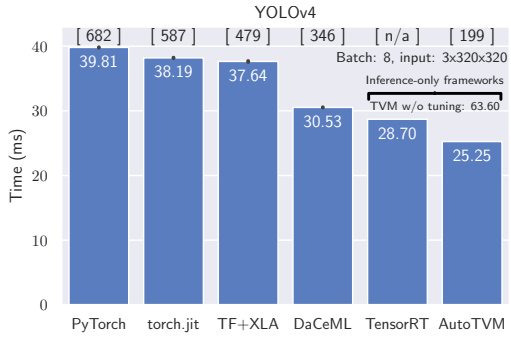

Figure 9: Runtime results for inference with YOLOv4. Brackets [·] indicate number of kernels launched. Due to profiler issues, kernel counts for TensorRT could not be obtained.

turn allows to fuse them all to one map with the range `0:H, 0:B, 0:SN`. This stores all memory locally (within the map) and will be allocated as GPU registers upon code generation (Figure 6c). As the global data movement optimization step does not apply in this case, the recipe proceeds to specialize the code to the GPU.

The optimization pipeline gradually lowers the representation, and applies the DaCeML `WarpTiling` transformation to partition the work in the internal maps across warps (Figure 6d), and automatically generates warp-level reductions for max/sum. To avoid extra data allocation, the recipe *replicates* all computations that have multiple results dependent on reductions by default, minimizing data movement at the cost of redundant computations. Lastly, the SDFG is cleaned up via another pass of transformations — front-loading reduction initialization, map fusion, and vectorization (Figure 6e).

Due to the popularity of the operator, many frameworks include optimized implementations of softmax. In Figure 8, we evaluate our generated implementation against several competing frameworks, including ones that do not provide learning capabilities. Of the tested frameworks, Triton [58] (which uses a hand-written kernel) and DaCeML achieve the best performance. The hand-written kernel included in PyTorch hardcodes a similar warp-tiling scheme as the DaCeML pipeline optimizes to, yet runs slightly slower. The XLA-based frameworks fail to fully fuse the operator, resulting in 3 kernels. While TVM contains a Relay operator for softmax, it is not tunable with AutoTVM, and does not contain warp-level primitives, both of which contribute to its high runtime relative to the warp-optimized frameworks. As detailed in Section 6.6, the DaCeML implementation is fully fusible with neighboring operations, and applying the recipe in the BERT transformer DNN results exploits this for an additional 196 μs gain per call. Handwritten kernel implementations of other frameworks (e.g., Triton, PyTorch) are generally not fusible.

6.3 Inference with Optimized Operators

Figure 9 lists inference results for the YOLOv4 DNN with the automated recipe, which uses the optimized Mish activation (Section 6.1) in the context of the full network. The figure shows a reduction in kernel counts, as well as 1.22–1.3× speedup over training frameworks. This nearly reaches the optimization levels of inference-only frameworks, such as TensorRT and TVM. Upon deeper inspection, the two frameworks make use of inference-specific optimizations, such as custom implementations for convolutions and pre-transforming weights. As DaCeML is designed to optimize training workloads, such techniques are outside the scope of this work.

6.4 Automatic Optimization for Training

While DaCeML's major strength lies in its manual tuning capabilities, the automated transformation recipe from Section 5.6 already yields performance improvements over state-of-the-art frameworks. We demonstrate networks from different domains, with varying data movement patterns, as well as utilizing different hardware units (e.g., tensor cores). We run the following models with the recipe alone, and list the results in Table 1 (bottom): ResNet-50 [20], Wide ResNet-50-2 [68], MobileNet V2 [52], EfficientNet-B0's MBConv block [57], MLP Mixer [59], Fully Convolutional Network [36], WaveNet [61], $\text{BERT}_{\text{LARGE}}$ [14] encoder block in single (32-bit) and mixed (16-bit) precision, and the DLRM [42] recommendation system. An interested ML practitioner could then use DaCeML to further optimize performance as necessary.

The table shows that no single framework operates best across all DNNs. Additionally, with the automatic optimizations alone, we see that DaCeML roughly matches and in multiple cases outperforms the other compiler frameworks. This is especially interesting in variants of popular networks, such as Wide ResNets. DaCeML is roughly 2× slower on Wide ResNet-50-2 than on ResNet-50, as expected for performing twice the computations; yet other frameworks are up to 2.89× slower. This potentially indicates specialization for certain operators and sizes, which does not occur with the data-centric transformations in DaCeML.

We now proceed with two case studies that highlight the possibilities enabled by DaCeML's user-guided optimization workflow, where the results are summarized in Table 1 (top).

6.5 Guided Optimization Case Study: EfficientNet

In this case study, we consider the EfficientNet-B0 [57] convolutional neural network. Since EfficientNets use repeated blocks, we can optimize one such block and reuse the techniques for the rest

		PyTorch		torch.jit		JAX		TF+XLA		DaCeML	
		→	⇄	→	⇄	→	⇄	→	⇄	→	⇄
Guided	EfficientNet [57]	2.05	6.90	2.04	6.94	2.39	7.40	1.54	6.37	**1.40**	**5.97**
	$\text{BERT}_{\text{LARGE}}$ (mixed) [14]	2.94	8.18	2.92	8.20	3.19	8.11	3.80	10.76	**2.74**	**7.62**
Automatic	ResNet-50 [20]	14.55	32.04	**9.98**	**31.94**	14.17	33.93	12.33	35.57	10.03	32.45
	Wide ResNet-50-2 [68]	22.50	70.94	22.45	70.83	40.49	98.13	32.79	99.06	**20.62**	**67.99**
	MobileNet V2 [52]	9.98	18.45	6.22	15.53	—	—	7.42	20.29	**4.74**	**14.77**
	EfficientNet [57]	2.05	6.90	2.04	6.94	2.39	7.40	**1.54**	**6.37**	1.57	15.00
	MLP Mixer [59]	1.63	**3.65**	**1.36**	3.66	1.77	4.01	—	—	1.48	4.25
	FCN8s [36]	46.85	158.42	46.82	**158.40**	—	—	—	—	**45.97**	166.30
	WaveNet [61]	23.21	46.39	**18.67**	41.49	—	—	—	—	26.16	**41.07**
	$\text{BERT}_{\text{LARGE}}$ (single) [14]	11.05	31.76	11.05	31.82	**10.93**	**29.94**	11.14	38.73	11.44	32.98
	$\text{BERT}_{\text{LARGE}}$ (mixed) [14]	2.94	8.18	**2.92**	8.20	3.19	**8.11**	3.80	10.76	3.34	9.25
	DLRM [42]	118.07	126.55	**117.38**	126.83	—	—	—	—	117.69	**126.42**

Table 1: Median runtime in milliseconds of the forward (→) and forward + backward (⇄) passes for convolutional vision (👁), non-convolutional vision (🙈), audio (🔊), image segmentation (🧩), transformer (🤖) and recommendation system (📈) models. $\text{BERT}_{\text{LARGE}}$ corresponds to a single encoder layer; EfficientNet to the first MBConv layer. A — indicates an implementation was not found for the model.

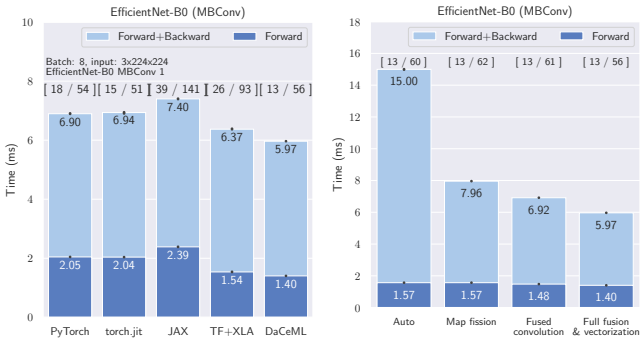

Figure 10: Runtime results (left) and transformation breakdown (right) of EfficientNet-B0 guided optimization.

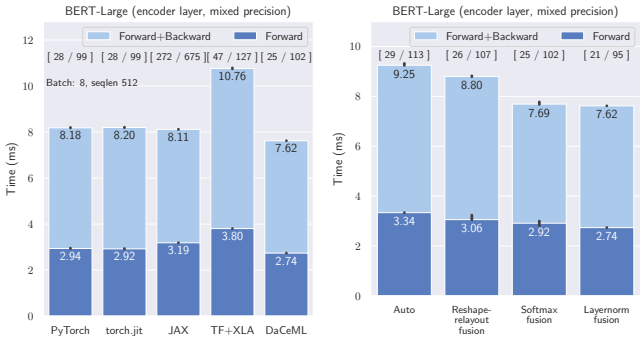

Figure 11: Runtime results (left) and transformation breakdown (right) of $\text{BERT}_{\text{LARGE}}$ guided optimization.

of the network. We focus on the first MBConv block, where performance results are reported in Figure 10, as well as the progressive improvements gained by various optimizations, using the automatic recipe as a base that was further improved. For this case study, we perform the guided optimizations using the DaCe Visual Studio Code plugin (see Figure 12), and report the number of clicks performed in the UI.

Map fission (6 clicks) We use two fission transformations to split a map in the backward pass. This allows us to avoid atomic operations by swapping iteration order and accumulating into thread-local memory.

Fused convolution (10 clicks) We discover a new fusion opportunity by expanding the convolutional operators to native SDFG, rather than cuDNN. Finding this automatically is non-trivial, requiring an exchange of the loop iteration variables in two operators to enable fusion. The data-centric GUI aids in this process — viewing data movement volumes in the graph highlights a potential bottleneck even without running the program. This fused convolution nets a ~1.33× speedup over cuDNN and PyTorch's hand-optimized versions in the forward pass.

Full fusion and vectorization (2 clicks) The map fission performed enables further fusion in the backward pass. We also tune the backpropagation to recompute intermediate values rather than loading and storing them. This is done by applying `TaskletFusion` before the AD engine runs (in 27 lines of code). We fully fuse, flatten and vectorize the maps where possible.

6.6 Guided Optimization Case Study: BERT

Our second case study is the widely-used BERT [14] transformer. We optimize the $\text{BERT}_{\text{LARGE}}$ architecture with a batch size of 8 and sequence length 512. We run in mixed precision: here the advantages of the data-centric optimization are more pronounced as data movement becomes more important. For example, we already see in Figure 11 that TF+XLA exhibits slower performance due to suboptimal data layouts that yield slower tensor contractions.

All transformations are performed using the Python API starting from the automatic recipe; we report lines of code.

Reshape-relayout Fusion (25 lines) After applying the algebraic fusion heuristic, constrained to generate BLAS operations only, we observed extraneous data relayout (transposition) calls

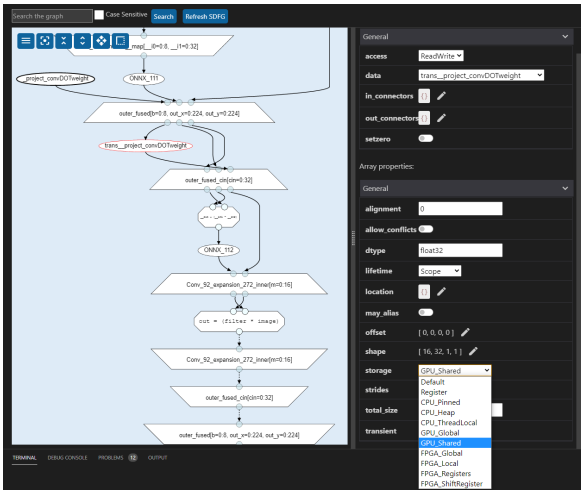

Figure 12: EfficientNet optimization in Visual Studio Code.

with reshapes. We then wrote a simple transformation that detects this pattern and fuses the output write into the prior (or subsequent) computation, if elementwise. This applies six times (thrice forward, thrice backward) in the graph.

Softmax Fusion (39 lines) Instead of calling a pre-optimized library for softmax, we decide to "break out of the library jail" and expand it to the native SDFG, followed by applying the our generic recipe (detailed in Section 6.2). This results in fusing scaling into softmax, at a 27 μs overhead instead of 223 μs in the forward pass, and 1.1 ms gain in total.

Layer-normalization Fusion (47 lines) Here we use the lifted layer normalization scheme to nest the preceding linear layer bias into normalization. We extend and adapt the lifting transformation to include the prior/subsequent operations. For the backpropagation, since bias is reduced into 1,024 elements, we use a warp-based reduction schedule and combine it with the layer normalization weight gradient kernel, which has the same iteration space.

In both case studies, the resulting code outperforms *all* compiler infrastructures. This demonstrates the strength of guided data-centric optimization — inspecting complex DNN models from a bird's eye view for data movement bottlenecks, and mitigating them via transformations.

7 CONCLUSION

We explore the concept of data-centric optimization for deep learning with DaCeML. The framework enables general-purpose data layout and movement transformations to be applied on arbitrary networks written in PyTorch, matching and outperforming state-of-the-art compilers. The two key insights of the data-centric view are focusing on data movement minimization based on memory access patterns rather than operator types, and allowing performance engineers to further tune global and local movement. The former can perform a superset of the optimizations applied by DNN compilers; and the latter can turn massive engineering effort into a click of a button in the analysis and transformation UI. Either automatic or human-in-the-loop, DaCeML helps practitioners speed up training without sacrificing productivity.

ACKNOWLEDGEMENTS

This project received funding from the European Research Council under the European Union's Horizon 2020 programme (Project PSAP, No. 101002047); and receives EuroHPC-JU funding under grants EUPILOT, No. 101034126; MAELSTROM, No. 955513; and DEEP-SEA, No. 955606, with support from the Horizon 2020 programme. T.B.N. is supported by the Swiss National Science Foundation (Ambizione Project #185778). N.D. is supported by the ETH Postdoctoral Fellowship. The authors wish to acknowledge the support from the PASC program (Platform for Advanced Scientific Computing), as well as the Swiss National Supercomputing Center (CSCS) for providing computing infrastructure.

REFERENCES

[1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. https://www.tensorflow.org/ Software available from tensorflow.org.

[2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).

[3] Riyadh Baghdadi, Abdelkader Nadir Debbagh, Kamel Abdous, Benhamida Fatima Zohra, Alex Renda, Jonathan Elliott Frankle, Michael Carbin, and Saman Amarasinghe. 2019. TIRAMISU: A Polyhedral Compiler for Dense and Sparse Deep Learning. In *Workshop on Systems for ML at NeurIPS*.

[4] Paul Barham and Michael Isard. 2019. Machine learning systems are stuck in a rut. In *Proceedings of the Workshop on Hot Topics in Operating Systems*.

[5] Tal Ben-Nun, Johannes de Fine Licht, Alexandros N Ziogas, Timo Schneider, and Torsten Hoefler. 2019. Stateful dataflow multigraphs: A data-centric model for performance portability on heterogeneous architectures. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*.

[6] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection.

[7] Léon Bottou, Frank E Curtis, and Jorge Nocedal. 2018. Optimization methods for large-scale machine learning. *Siam Review* 60, 2 (2018).

[8] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. JAX: composable transformations of Python+NumPy programs. http://github.com/google/jax

[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 end-to-end optimization stack for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*.

[10] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174* (2016).

[11] Sharan Chetlur, Cliff Woolley, Philippe Vandermersch, Jonathan Cohen, John Tran, Bryan Catanzaro, and Evan Shelhamer. 2014. cuDNN: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759* (2014).

[12] Scott Cyphers, Arjun K Bansal, Anahita Bhiwandiwalla, Jayaram Bobba, Matthew Brookhart, Avijit Chakraborty, Will Constable, Christian Convey, Leona Cook, Omar Kanawi, et al. 2018. Intel nGraph: An intermediate representation, compiler, and executor for deep learning. *arXiv preprint arXiv:1801.08058* (2018).

[13] Johannes de Fine Licht, Andreas Kuster, Tiziano De Matteis, Tal Ben-Nun, Dominic Hofer, and Torsten Hoefler. 2020. StencilFlow: Mapping large stencil programs to distributed spatial computing systems. In *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*.

[14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.

[15] Xiao Dong, Lei Liu, Peng Zhao, Guangli Li, Jiansong Li, Xueying Wang, and Xiaobing Feng. 2019. Acorns: A framework for accelerating deep neural networks with input sparsity. In *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*.

[16] Venmugil Elango, Norm Rubin, Mahesh Ravishankar, Hariharan Sandanagobalane, and Vinod Grover. 2018. Diesel: DSL for linear algebra and neural net computations on GPUs. In *Proceedings of the 2nd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*.

[17] Roy Frostig, Matthew James Johnson, and Chris Leary. 2018. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning* (2018).

[18] Google. 2022. XLA: Optimizing Compiler for Machine Learning. https://www.tensorflow.org/xla

[19] Alon Halevy, Peter Norvig, and Fernando Pereira. 2009. The unreasonable effectiveness of data. *IEEE Intelligent Systems* 24, 2 (2009).

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep Residual Learning for Image Recognition. arXiv:1512.03385 [cs.CV]

[21] Sara Hooker. 2020. The hardware lottery. *arXiv preprint arXiv:2009.06489* (2020).

[22] Yuwei Hu, Zihao Ye, Minjie Wang, Jiali Yu, Da Zheng, Mu Li, Zheng Zhang, Zhiru Zhang, and Yida Wang. 2020. FeatGraph: A flexible and efficient backend for graph neural network systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*.

[23] Intel. 2021. oneAPI Deep Neural Network Library. https://01.org/oneDNN

[24] Intel. 2022. Open Visual Inference and Neural Network Optimization Toolkit. https://01.org/openvinotoolkit

[25] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*. PMLR, 448–456.

[26] Andrei Ivanov, Nikoli Dryden, Tal Ben-Nun, Shigang Li, and Torsten Hoefler. 2021. Data movement is all you need: A case study on optimizing transformers. In *Proceedings of the Fourth Conference on Machine Learning and Systems (MLSys)*.

[27] Paras Jain, Ajay Jain, Aniruddha Nrusimha, Amir Gholami, Pieter Abbeel, Kurt Keutzer, Ion Stoica, and Joseph E Gonzalez. 2020. Checkmate: Breaking the Memory Wall with Optimal Tensor Rematerialization. In *Proceedings of the Third Conference on Machine Learning and Systems (MLSys)*.

[28] Zhihao Jia, Oded Padon, James Thomas, Todd Warszawski, Matei Zaharia, and Alex Aiken. 2019. TASO: optimizing deep learning computation with automatic generation of graph substitutions. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP)*.

[29] Zhihao Jia, James Thomas, Todd Warszawski, Mingyu Gao, Matei Zaharia, and Alex Aiken. 2019. Optimizing DNN computation with relaxed graph substitutions. In *Proceedings of the 2nd Conference on Systems and Machine Learning (SysML)*.

[30] Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*.

[31] Chris Lattner and Vikram Adve. 2004. LLVM: A compilation framework for lifelong program analysis & transformation. In *International Symposium on Code Generation and Optimization, 2004. CGO 2004*.

[32] 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 preprint arXiv:2002.11054* (2020).

[33] Richard Lethin. 2017. Polyhedral Optimization of TensorFlow Computation Graphs. In *Sixth Workshop on Extreme-scale Programming Tools (ESPT)*.

[34] Chao Li, Yi Yang, Min Feng, Srimat Chakradhar, and Huiyang Zhou. 2016. Optimizing memory efficiency for deep convolutional neural networks on GPUs. In *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 633–644.

[35] Shigang Li, Tal Ben-Nun, Salvatore Di Girolamo, Dan Alistarh, and Torsten Hoefler. 2020. Taming Unbalanced Training Workloads in Deep Learning with Partial Collective Operations. In *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming* (San Diego, California) *(PPoPP '20)*. Association for Computing Machinery, New York, NY, USA, 45–61. https://doi.org/10.1145/3332466.3374528

[36] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[37] Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig, Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. 2017. SymPy: symbolic computing in Python. *PeerJ Computer Science* 3 (2017), e103.

[38] Microsoft. 2022. DeepSpeed. https://www.deepspeed.ai

[39] Microsoft. 2022. ONNX Runtime. https://www.onnxruntime.ai

[40] Microsoft. 2022. Visual Studio Code - Code editing. Refined. https://code.visualstudio.com/.

[41] Diganta Misra. 2020. Mish: A Self Regularized Non-Monotonic Activation Function. In *Proceedings of the 31st British Machine Vision Conference (BMVC)*.

[42] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Mallevich, Ilia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. 2019. Deep Learning Recommendation Model for Personalization and Recommendation Systems. arXiv:1906.00091 [cs.IR]

[43] NVIDIA. 2022. NVIDIA TensorRT. https://developer.nvidia.com/tensorrt

[44] ONNX. 2021. ONNX: Open Neural Network Exchange. https://onnx.ai/

[45] Yosuke Oyama, Tal Ben-Nun, Torsten Hoefler, and Satoshi Matsuoka. 2018. Accelerating deep learning frameworks with micro-batches. In *2018 IEEE International Conference on Cluster Computing (CLUSTER)*.

[46] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*.

[47] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. 2021. Carbon Emissions and Large Neural Network Training. *arXiv preprint arXiv:2104.10350* (2021).

[48] PyTorch Team. 2022. TorchScript. https://pytorch.org/docs/stable/jit.html

[49] Rajat Raina, Anand Madhavan, and Andrew Y Ng. 2009. Large-scale deep unsupervised learning using graphics processors. In *Proceedings of the 26th annual international conference on machine learning (ICLR)*.

[50] Lutz Roeder. 2022. Netron. https://netron.app/

[51] Nadav Rotem, Jordan Fix, Saleem Abdulrasool, Garret Catron, Summer Deng, Roman Dzhabarov, Nick Gibson, James Hegeman, Meghan Lele, Roman Levenstein, et al. 2018. Glow: Graph lowering compiler techniques for neural networks. *arXiv preprint arXiv:1805.00907* (2018).

[52] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2019. MobileNetV2: Inverted Residuals and Linear Bottlenecks. arXiv:1801.04381 [cs.CV]

[53] Muthian Sivathanu, Tapan Chugh, Sanjay S Singapuram, and Lidong Zhou. 2019. Astra: Exploiting predictability to optimize deep learning. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*.

[54] Benoit Steiner, Chris Cummins, Horace He, and Hugh Leather. 2020. Value Function Based Performance Optimization of Deep Learning Workloads. arXiv:2011.14486 [cs.LG]

[55] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*.

[56] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. 2017. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision (ICCV)*.

[57] Mingxing Tan and Quoc Le. 2019. EfficientNet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*. PMLR, 6105–6114.

[58] Philippe Tillet, Hsiang-Tsung Kung, and David Cox. 2019. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages*.

[59] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. 2021. MLP-Mixer: An all-MLP Architecture for Vision. arXiv:2105.01601 [cs.CV]

[60] Leonard Truong, Rajkishore Barik, Ehsan Totoni, Hai Liu, Chick Markley, Armando Fox, and Tatiana Shpeisman. 2016. Latte: A language, compiler, and runtime for elegant and efficient deep neural networks. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*.

[61] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. 2016. WaveNet: A Generative Model for Raw Audio. arXiv:1609.03499 [cs.SD]

[62] 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).

[63] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*.

[64] Anand Venkat, Tharindu Rusira, Raj Barik, Mary Hall, and Leonard Truong. 2019. SWIRL: High-performance many-core CPU code generation for deep neural networks. *The International Journal of High Performance Computing Applications* 33, 6 (2019).

[65] Richard Wei, Lane Schwartz, and Vikram Adve. 2018. DLVM: A modern compiler infrastructure for deep learning systems. In *Proceedings of the Sixth International Conference on Learning Representations - Workshop Track (ICLR)*.

[66] Yuxin Wu and Kaiming He. 2018. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*.

[67] Yichen Yang, Phitchaya Mangpo Phothilimtha, Yisu Remy Wang, Max Willsey, Sudip Roy, and Jacques Pienaar. 2021. Equality Saturation for Tensor Graph Superoptimization. arXiv:2101.01332 [cs.AI]

[68] Sergey Zagoruyko and Nikos Komodakis. 2017. Wide Residual Networks. arXiv:1605.07146 [cs.CV]

[69] Lianmin Zheng, Chengfan Jia, Minmin Sun, Zhao Wu, Cody Hao Yu, Ameer Haj-Ali, Yida Wang, Jun Yang, Danyang Zhuo, Koushik Sen, et al. 2020. Ansor: Generating high-performance tensor programs for deep learning. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*.

[70] Alexandros Nikolaos Ziogas, Tal Ben-Nun, Guillermo Indalecio Fernández, Timo Schneider, Mathieu Luisier, and Torsten Hoefler. 2019. Optimizing the data movement in quantum transport simulations via data-centric parallel programming. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*.