
SWIFT FOR TENSORFLOW: A PORTABLE, FLEXIBLE PLATFORM FOR DEEP LEARNING

Brennan Saeta¹ Denys Shabalín¹ Marc Rasi¹ Brad Larson¹ Xihui Wu¹ Parker Schuh¹ Michelle Casbon¹
Daniel Zheng¹ Saleem Abdulrasool¹ Aleksandr Efremov¹ Dave Abrahams¹ Chris Lattner² Richard Wei²

ABSTRACT

Swift for TensorFlow is a deep learning platform that scales from mobile devices to clusters of hardware accelerators in data centers. It combines a language-integrated automatic differentiation system and multiple `Tensor` implementations within a modern ahead-of-time compiled language oriented around mutable value semantics. The resulting platform has been validated through use in over 30 deep learning models and has been employed across data center and mobile applications.

1 INTRODUCTION

Deep learning has demonstrated remarkable gains in performance across diverse tasks including game playing (Silver et al., 2016; Vinyals et al., 2019), image understanding (He et al., 2016; Szegedy et al., 2016), natural language understanding (Devlin et al., 2018), and beyond. Training modern deep learning models from scratch—a requirement during the development of new neural network architectures—requires enormous amounts of compute (Amodei et al., 2018). In practice, neural networks are trained on clusters containing up to thousands of hardware accelerators spread across a supporting data center.

Deploying machine-learned models on edge devices can deliver functionality and low latency without internet connectivity. The popularity of modern edge devices has motivated research on models for limited hardware (compute & memory capacity), and energy constraints (Howard et al., 2017; Tan & Le, 2019). Additionally, many models can be fine-tuned directly on a user’s device without copying personal data over a network. While datacenter-scale training focuses on peak throughput, mobile apps optimize for startup and execution time to deliver fluid user experiences.

Day to day, machine learning practitioners are likely to develop their models on a single machine with a dedicated GPU for hardware acceleration. Thanks to advancements in transfer learning, recent models (e.g. BERT (Devlin et al., 2018)) have been explicitly designed with pre-training in mind. By starting from a pre-trained checkpoint, effective

models can be trained on one desktop GPU.

Today, tools for training deep neural networks at datacenter scale are often embedded in dynamically typed languages (Abadi et al., 2016; Paszke et al., 2019; Bradbury et al., 2020; Innes et al., 2019) and rely on Just-In-Time (JIT) compilation to obtain optimal performance. On the other end of the spectrum, libraries for neural network inference on mobile devices are compiled before installation. Deploying neural networks onto devices often involves a translation step from the datacenter system to the on-device execution engine.

In this paper, we present Swift for TensorFlow, a platform for machine learning. The combination of Swift’s Ahead-of-Time (AOT) compilation, support for mutable value semantics, and convenient syntax extended with language-integrated automatic differentiation yields a surprisingly effective platform that scales from mobile phones to distributed accelerator clusters. This work demonstrates the value of exploring beyond dynamically typed languages for deep learning. More concretely, our contributions are:

- **Language-integrated automatic differentiation (AD)** (Section 2). We extend Swift with a source-to-source compile-time transformation to automatically generate derivatives for arbitrary Swift functions. The AD system is not coupled with the `Tensor` implementation; it can be used with any type that conforms to the `Differentiable` protocol and any `@differentiable` annotated function.
- **Mutable value semantics** (Section 4). We explore mutable value semantics—a combination of value semantics with in-place mutation that supports local reasoning and referential transparency—in the context of machine learning applications. Swift for TensorFlow’s

¹Google Research, Brain ²Work done at Google Research, Brain. Correspondence to: Brennan Saeta <saeta@google.com>, Denys Shabalín <shabalín@google.com>.

```
protocol Differentiable {
  associatedtype TangentVector
  : AdditiveArithmetic
  mutating func move(
    along direction: TangentVector)
}
```

Figure 1. Differentiable protocol definition.

APIs serve as a case study of how this approach can be applied to provide simple yet powerful APIs.

We provide a broad evaluation (Section 5) of how the same programming model can scale across a wide range of environments from datacenter supercomputers to low-power mobile devices.

2 AUTOMATIC DIFFERENTIATION

The Swift for TensorFlow project extended the Swift language to support compile-time source-to-source automatic differentiation (AD) (Wei et al., 2019). The language extensions allow library authors to define “differential operators,” which are ordinary Swift higher order functions that compute derivatives of passed-in functions. For example, we added a `gradient` function to the Swift standard library that evaluates the gradient of a scalar-valued function at a given point. When a call to `gradient(at: 0, in: f)` is compiled, a compiler stage synthesizes a function that can be evaluated at runtime to return the derivative of `f`, which `gradient` calls. Additionally, we support differentiation of arbitrary user-defined types so long as they satisfy a few requirements (Figure 1). Our language extensions have been contributed upstream and have been successfully merged into the main branch of the language & compiler.

2.1 Language extensions

The key language building blocks for differentiation are:

- The `Differentiable` protocol,¹ which encodes requirements on the parameter and return types of differentiable functions (Figure 1), and
- The `@differentiable (A) -> B` function type family, which bundles information about how to compute the function’s derivative. We call instances of these types *differentiable function values*.

Every `Differentiable` type has an associated `TangentVector` type, inspired by differential geometry. `Differentiable` values represent points on differentiable manifolds, and `TangentVector` values represent

¹Swift protocols are similar to Haskell typeclasses.

```
func gradient<A: Differentiable>(
  at x: A,
  in f: @differentiable (A) -> Float
) -> A.TangentVector
```

Figure 2. Gradient function declaration.

vectors in the tangent spaces of those manifolds. Differential geometry defines derivatives between manifolds as linear functions between the tangent spaces of these manifolds.

Consider the function `f: (A) -> Float`. The gradient of `f` at a point is a vector in `A.TangentVector`. When `A` is the flat manifold \mathbb{R}^n , `A.TangentVector` is \mathbb{R}^n and we recover the gradient from multivariable calculus.

A `Differentiable` type also requires a `move` method that moves a value by the distance in the direction indicated by a `TangentVector`. This is known as “exponential map” in differential geometry.

To improve ergonomics, we automatically promote functions and closures to their `@differentiable` counterparts based on their use within their surrounding module. When type checking encounters a function value in a context that requires a differentiable function value, it inserts an implicit conversion and flags the original function for compile-time differentiation. For example, users can pass an unannotated closure to the standard library `gradient` operator shown in Figure 2.

A differentiable function value is a bundle containing the original function value and Jacobian-vector product (JVP) and vector-Jacobian product (VJP) “derivative function” values (Figure 3). Each derivative function returns a pair of the computed value together with a closure called a *differential* or *pullback*, respectively. These derivative functions are taken from previous work (Maclaurin, 2016) and are particularly close to the JVP and VJP functions in (Bradbury et al., 2020; Innes et al., 2019). The JVP implements forward mode differentiation, and the VJP implements reverse mode differentiation.

The AD code transformation transforms a function `f` into derivative functions that are implemented in terms of derivative functions of `f`’s callees. The transformation recursively transforms the callees to get their derivative functions. This recursion requires a base case of known derivative functions. Our approach allows fully-customizable base derivative functions via a `@derivative(of:)` attribute. Users write custom derivative functions with this attribute to register the base case, and the code transformation terminates the recursion whenever it encounters a function with a user specified custom derivative.

Original function	(A) -> B
JVP (Forward mode AD)	(A) -> (B, (A.TangentVector) -> B.TangentVector)
VJP (Reverse mode AD)	(A) -> (B, (B.TangentVector) -> A.TangentVector)

Figure 3. The three elements of a differentiable function value and their types.

2.2 Code transformation

The differentiation code transformation operates on the Swift Intermediate Language (SIL), an intermediate representation (IR) in static single assignment form. SIL is designed specifically for Swift’s compiler transformations, which helps us to handle diverse Swift language constructs (e.g. various control flow statements, value data types, and higher-order functions). The differentiation transformation involves the following steps:

- **Activity analysis** (Laurent Hascoet, 2013) determines instructions of the original function that are both *varied* (depend on the inputs) and *useful* (contribute to the output). Such instructions are *active* and need a derivative.
- **Differentiability checking** detects non-differentiable instructions and emits errors and warnings (e.g. a differentiable function whose return value does not depend on differentiable arguments) that help users catch errors before execution.
- **Derivative synthesis** creates the derivative functions, applies AD rules to active SIL instructions, and builds the corresponding derivative SIL instructions. This step also generates code that captures callee derivatives and the control flow path (see below).

Since AD is a compiler pass operating on the intermediate representation, the code generated by the AD system is fully amenable to the same set of compile-time optimizations as regular Swift code.

The code transformation produces the JVP and VJP discussed earlier in Figure 3. In order to be valid for every possible input x , the JVP and VJP must handle all the possible control flow paths through the original function. Our approach is based on statically-typed records corresponding to the basic blocks of the control flow graph that store intermediate state used in derivative calculations. These records form a nested data structure of control flow branches between basic blocks that have been taken during the execution of the function.

2.3 Contributions, limitations, & future work

To our knowledge, this combination of (1) requesting differentiation within the language, (2) supporting user-defined types with arbitrary tangent vector types, and (3) performing

the AD code transformation at AOT-compile time was first explored in Swift for TensorFlow.

Enzyme (Moses & Churavy, 2020) is a successor system that shares a similar set of the advantages, only targeting LLVM IR rather than SIL. (Laurent Hascoet, 2013) performs the code transformation at compile time in an AOT language, but requires that the user request differentiation outside the language. Clad (Vassilev et al., 2020) is a similar successor system operating on C/C++. Other AD systems that allow the user to request differentiation in the language trace the computation at runtime and differentiate the trace (Abadi et al., 2016; Maclaurin et al., 2015; Bradbury et al., 2020; Paszke et al., 2019) or use dynamic language features to transform the code at runtime (Innes et al., 2019).

Our approach currently does not support certain capabilities that are supported in similar systems for dynamic languages, such as higher-order differentiation. The `@differentiable` function type family would need to be extended to encode that an n -times differentiable function is $(n - 1)$ -times differentiable. Additionally, the code transformation currently cannot transform its own output because the output makes heavy use of closure captures, which are not yet well supported. A possible approach for differentiation of closures with captured variables is documented in (Vytiniotis et al., 2019). Finally, there are additional language features to support, such as `enum` types.

3 TENSORS & LAZY TENSORS

Deep learning models are often trained using a multi-dimensional array abstraction, colloquially called a `Tensor` (Abadi et al., 2016; Paszke et al., 2019). A `Tensor` type includes a suite of linear algebra operations such as matrix multiplication, convolution, and element-wise arithmetic. Because the AD system has been generalized to arbitrary types, there is no coupling between the AD system and a specific `Tensor` type. As a result, we have found that implementing multiple different `Tensor` types allows the platform to more effectively scale to different applications.

3.1 Naïve Tensor

One of our multiple `Tensor` implementations is a single-threaded type backed by Swift arrays. Although this type lacks hardware accelerator support, this implementation has advantages when working with small `Tensors`, including

portability, low computation and memory overheads, and small binary size (since the naïve implementation has no external dependencies).

3.2 Eager Tensor

Some libraries (e.g. NumPy) eagerly dispatch the operations of the user’s program to pre-compiled functionality (termed “kernels”) that can run on CPUs or hardware accelerators such as GPUs. “Eager mode” (also known as define-by-run (Tokui et al., 2019)) is popular for research because it supports the full flexibility of the host programming language, and is known to make debugging easier. In order to make efficient use of hardware accelerators, the kernels are dispatched to the accelerator to execute asynchronously and control is returned to the user’s program before the kernel finishes. (Paszke et al., 2019) As long as the user’s program does not observe the contents of a `Tensor`, the user’s program runs ahead and fills a pipeline of accelerator kernel invocations making efficient use of hardware resources.

Swift for TensorFlow delivers a define-by-run machine learning environment and exposes an eager execution model to the end user. We provide a hardware-accelerated implementation of `Tensor` operations backed by TensorFlow’s eager runtime (Agrawal et al., 2019). Eager `Tensor` performs asynchronous op-by-op execution that dispatches to hardware-optimized kernels under the hood.

The downside of the eager execution model is that it does not expose opportunities for operation fusion to the underlying runtime. This is the direct side-effect of op-by-op execution that is driven by the host programming language.

3.3 Lazy Tensor

Domain-specific optimizing compilers (Chen et al., 2018; Leary & Wang, 2017; Rotem et al., 2018) have been developed to both (a) substantially improve performance of deep learning models, and (b) target domain-specific hardware accelerators such as TPUs (Jouppi et al., 2020).

Unlike eager execution, such compilers can take complete models as programs in their own domain-specific IR and generate optimized hardware-specific machine code. The ability to observe the complete program provides a wide horizon for optimizations such as operation-fusion.

We’ve developed *LazyTensor* as a second implementation of our eager `Tensor` API that relies on tracing (Frostig et al., 2018) to target just-in-time (JIT) compilation via HLO, the XLA compiler’s (Leary & Wang, 2017) IR. End-users can switch between the two implementations by specifying a device for the computation to run on: either an eager or a lazy-tracing one.

The biggest challenge behind the alternative implementation

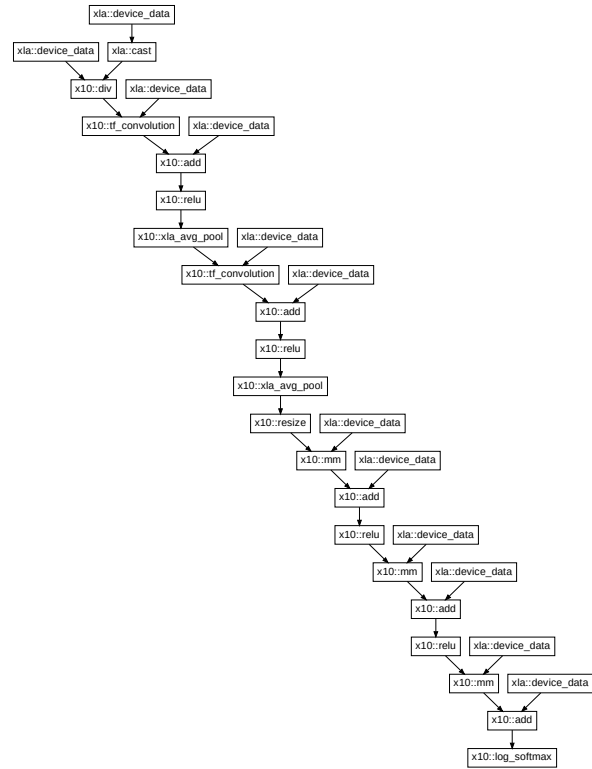


Figure 4. LazyTensor trace of the LeNet-5 model’s forward pass.

was not to break the illusion of eager execution. As long as the user’s program does not observe a `Tensor`’s contents, the code cannot distinguish when a `Tensor` operation is actually executed (ignoring timing and similar side-channels). Instead of dispatching to a fixed-set of pre-compiled kernels, the lazy `Tensor` type simply records a dynamic *trace* of operations to be executed at a later time. Traces are represented in-memory as directed acyclic graphs (Figure 4) and are transformed into an intermediate representation to perform domain-specific optimization and code generation.

By dynamically discovering traces implicit within the user’s program, Swift for TensorFlow enables automatic composition of `Tensor` and non-`Tensor` computations. Imagine a hypothetical robotics motion planning algorithm that is composed of a neural network, a black-box CPU solver, and a second neural network. If this motion planning algorithm is used as a subroutine for a larger `Tensor`-based computation, known competing approaches require the user to factor the algorithm into at least three separate subroutines and annotate a subset of the subroutines for compilation with a domain specific compiler. With our approach, the longest possible traces are discovered dynamically and function composition is never impinged.

Through LazyTensor, Swift for TensorFlow delivers the

usability of eager mode *and* the performance of domain-specific optimizations and code generation. This technique is detailed in a concurrent paper with collaborators from PyTorch (Şuhan et al., 2020).

3.4 Limitations

One of the drawbacks of the LazyTensor approach is the overhead due to re-tracing the user’s program repeatedly. Because invoking the XLA JIT is computationally expensive, trace fragments are hashed to become keys in an XLA-program cache; each unique trace is only compiled by XLA once. Even though we reuse previously compiled traces, we still incur tracing overhead on each iteration since we support the full imperative programming model, and traces can change at any point in program execution.

Since XLA was designed with an expectation of statically known tensor shapes, minor changes in program execution such as changes in the dimensions of the input tensors can trigger recompilation. As a consequence, lazy tracing works best when the computation is done repeatedly over the same constant tensor dimensions on each iteration.

While constructing the trace, we fully unroll any control-flow. This is often desirable as it allows for optimizations across loop iterations. As a downside, this can cause the creation of large traces, which can incur significant one-time JIT compilation costs. We expose an additional library function `LazyTensorBarrier()` that allows users to explicitly cut traces. Fortunately, because the most common loop that is accidentally unrolled is the main training loop of a neural network, a training-loop library can automatically call `LazyTensorBarrier()` after the optimizer update step on behalf of the user.

It is left to future work to investigate automatically detecting a sufficiently large trace fragment to compile and dispatch automatically, completely relieving the user of the need for any annotations.

3.5 Alternatives

Prior to lazy tracing, the Swift for TensorFlow project explored an approach called *graph program extraction*. The key idea behind it was to slice the user’s program into two concurrent programs, one that would execute on an accelerator and one on the CPU in the user’s host language, with communication channels between them for coordination. The accelerator programs could thus be compiled into fused kernels completely ahead-of-time.

We found this approach to be quite limiting since models often rely on dynamically configured values that are only available at runtime. For example, one may implement a complete ResNet family of models by assembling key building blocks in a configuration determined by a dynamic

model variant. A fully static implementation is limited to fusion within the individual building blocks since the final composition is not known ahead-of-time. Alternatively, it could also attempt to pre-compile a large number of variants that exercise all possible combinations of the configuration space (which can be exponential).

Julia’s approach to XLA compilation (Fischer & Saba, 2018) makes a different set of trade-offs. Rather than being embedded in an AOT compiled language, they use XLA as an alternative backend to Julia’s JIT compiler. The language provides a set of powerful runtime code generation and specialization facilities that allow it to address XLA’s limited expressiveness.

4 MUTABLE VALUE SEMANTICS

When a type has value semantics, distinct variables of that type access logically disjoint data, so mutations through one variable are observable only through that variable.² This behavior should be familiar to most programmers: in almost every major language, mutation through an integer variable affects only that variable (Figure 5 column 1).

The alternative is *reference semantics*, where mutation through one variable can be observed through another. Most languages that support mutation allow only reference semantics for user-defined and aggregate types (such as arrays). However, reference semantics has well-known downsides: loss of local reasoning and referential transparency; “spooky action at a distance”³ where data changes unexpectedly due to implicitly shared state (Figure 5 column 2), defensive copying, the inability to optimize code where state is not actually shared, and data races in concurrent code. It is particularly difficult to implement mathematical transformations such as AD in systems with reference semantics.

These hazards often drive designers to adopt *immutability*, producing a pure functional programming model. For example, the `Tensor` types in TensorFlow, JAX, and XLA are immutable. However, strict immutability can be inefficient, difficult to map onto some algorithms (O’Neill, 2009), and often clashes with the programmer’s mental model.⁴ A common intermediate position keeps data immutable but implements state in the programming model by rebinding variables to newly-constructed immutable values. In either case, the performance implications of making a small update⁵ to a large value are unfavorable.

²Immutable data has value semantics trivially.

³With apologies to Einstein.

⁴Both XLA and TensorFlow have compromised their pure programming models for expressivity and performance reasons.

⁵Even in pure functional programming, it’s still common to create updated versions of large values. Persistent data structures (Driscoll et al., 1989) address the problem in part, but cannot match the performance of contiguous arrays.

<pre># Python integer x = 3 y = x x += 1 print(x) # 4 print(y) # 3</pre>	<pre># Python list x = [3] y = x x[0] += 1 print(x) # [4] print(y) # [4] <===</pre>	<pre>// Swift array var x = [3] var y = x x[0] += 1 print(x) // [4] print(y) // [3]</pre>
--	--	---

Figure 5. Examples of value and reference semantics in Python and Swift. In Python, integers have value semantics (first column), but lists and user-defined types have reference semantics: in the middle column, a mutation through `x` changes the value observed through `y`. In Swift, mutable value semantics are pervasive (third column).

A type has *mutable value semantics* when two variables of the type always have logically-disjoint values (value semantics) *combined with* support for efficient part-wise mutation. C++, Swift, and Rust are examples of popular languages that support defining such types.

In Swift, built-in types like integers, arrays, and strings have mutable value semantics, and user-defined types composed from such types automatically acquire mutable value semantics of their own.⁶ Moreover, mutation always occurs via a “unique borrow” of the value being mutated, marked in a function signature with `inout`. Although an `inout` parameter superficially resembles a pass-by-reference parameter, it does not introduce reference semantics because the borrow is guaranteed to be unique. In fact, a program using `inout` can be trivially rewritten into a semantically-equivalent program using only pass-by-value (Figure 8). Because Swift for TensorFlow `Tensors` have value semantics, we can perform sophisticated static analysis such as static shape tracking (Paszke & Saeta, 2020) even in the presence of mutation.

Swift occupies a unique position in the design space for mutation. In languages with reference semantics, users must copy defensively to avoid “spooky action at a distance.” Too many copies results in performance problems; too few copies results in bugs. Other languages with value semantics make copies eagerly, causing users to pass parameters by reference, thus bypassing value semantic guarantees. In Swift, large values are copied lazily, upon mutation, and only when shared, so pass-by-value is efficient and broadly used. The reference counting mechanism underlying this behavior is built in, enabling optimizations not possible in other languages.

Next, we show how mutable value semantics in Swift enable elegant solutions to important challenges for ML systems.

⁶Unless explicitly requested by using the `class` keyword.

4.1 Tensors, Layers, & Models

Swift for TensorFlow APIs use mutable value semantics pervasively (e.g., `Tensors`, models, and datasets are all mutable value types). Unlike TensorFlow, Keras (Géron, 2019) and PyTorch (Paszke et al., 2019), Swift for TensorFlow does not need a `Variable` type; composition of mutable value semantics and language-integrated AD allows us to use the types directly without any additional wrappers.

As an example, we provide a full definition of a variant of the LeNet-5 convolutional neural network model (LeCun et al., 1998) in Figure 6. The LeNet model is simply a `struct` that conforms to the `Layer` protocol that composes a number of standard layers. Each conforming `Layer` must provide an implementation of `callAsFunction` that defines how to apply a transformation to a given input. This function must be annotated as `@differentiable`, which instructs the compiler to enable automatic differentiation (discussed in Section 2).

4.2 Avoiding model copies

Training a neural network model can be thought of as repeated application of a function from network parameters and a training minibatch to updated network parameters. In code, this could be written as a function with type:

```
(Model, Minibatch) -> Model
```

Some neural network models (e.g. large transformer-based natural language models) consume a significant fraction of accelerator memory capacity for the model’s weights. When training such networks, materializing two copies of the network’s parameters in memory will exceed accelerator capacity. XLA’s pure-functional, immutability-oriented core has been extended with the concept of input-output buffer aliasing to accommodate these workloads. Unfortunately, this compromise leaks into the user programming model. Swift’s `inout` keyword (see Figure 8) is a solution that allows us to rewrite the training function’s type to be:

```
(inout Model, Minibatch) -> Void
```

```

public struct LeNet: Layer {
    public var conv1 = Conv2D<Float>(filterShape: (5, 5, 1, 6),
                                       padding: .same, activation: relu)
    public var pool1 = AvgPool2D<Float>(poolSize: (2, 2), strides: (2, 2))
    public var conv2 = Conv2D<Float>(filterShape: (5, 5, 6, 16), activation: relu)
    public var pool2 = AvgPool2D<Float>(poolSize: (2, 2), strides: (2, 2))
    public var flatten = Flatten<Float>()
    public var fc1 = Dense<Float>(inputSize: 400, outputSize: 120, activation: relu)
    public var fc2 = Dense<Float>(inputSize: 120, outputSize: 84, activation: relu)
    public var fc3 = Dense<Float>(inputSize: 84, outputSize: 10)

    public init() {}

    @differentiable
    public func callAsFunction(_ input: Tensor<Float>) -> Tensor<Float> {
        let convolved = input.sequenced(through: conv1, pool1, conv2, pool2)
        return convolved.sequenced(through: flatten, fc1, fc2, fc3)
    }
}

```

Figure 6. LeNet model definition.

```

var model = LeNet()
for epoch in epochs {
    for batch in epoch {
        let gradients = gradient(at: model) {
            model -> Tensor<Float> in
            let loss = softmaxCrossEntropy(
                logits: model(batch.data),
                labels: batch.label)
            return loss
        }
        optimizer.update(&model,
                        along: gradients)
    }
}

```

Figure 7. Training LeNet on a simple dataset.

The semantics are identical to the pure-functional approach, while simultaneously avoiding exceeding hardware limits so long as the `Tensor` type used by the model properly implements mutable value semantics.

As an example, an explicit training loop for the LeNet model can be written in a few lines of code shown in Figure 7. After instantiating the model, the user takes the gradient of a differentiable function computing a loss, with respect to the model. The returned gradients are of type `Model.TangentVector`, and can be stored or manipulated as a first-class value. An optimizer borrows the model uniquely, and updates it in-place based on the computed gradients. Because both the model and its gradient are first class values, both can be manipulated (e.g. storing previous values, or computing properties of them).

4.3 Array indexing differentiation

The efficient gradient design goal states that computing the derivative of a function should be comparable in cost to computing the function itself. AD systems compute derivatives of a composite function by combining the derivatives of constituent operations using the chain rule. While pure-functional languages or intermediate representation (IRs) have formed the basis of numerous powerful Tensor-based AD systems, the derivative of the array indexing operation violates the AD efficiency goal. The types⁷ of the array indexing operation are:

```

// Operation type:
@differentiable(wrt: a)
(a: [T], index: Int) -> T

// Pullback type:
(T) -> [T]

```

Unfortunately, this functional AD formulation necessitates the construction of an array of all zeros⁸ except for a single non-zero entry at position `index`. While a secondary, sparse representation for the array-type could improve performance, runtime overheads remain in practice. Although Dex (Maclaurin et al., 2019) has developed an alternative approach leveraging monads and an algebraic effects system, they add complexity to the language. A mutable-value-semantic AD formulation could instead write this derivative with the following types:

⁷Without loss of generality, we assume `T == T.TangentVector` and `[T] == [T].TangentVector`. Additionally, we use a regular Swift array instead of the multi-dimensional `Tensor` type.

⁸Technically, the additive identity.

```
// Pullback type:
(T, inout [T]) -> Void
```

This new formulation avoids materializing zeros unnecessarily, and composes correctly in the presence of additional operations. Please see Appendix B for the complete implementation of the operations. Further, the technique developed here applies not just to array or `Tensor`-operations, but additionally to derivatives of arbitrary “big-to-small” operations on many data types, such as manipulating a vertex within a tree or graph, or computing the partial derivative with respect to a field within an aggregate data type.

4.4 Open questions

Our work has demonstrated that languages oriented around mutable value semantics enable efficient solutions to important challenges for ML and `Tensor`-based systems. However, work remains on how to implement a domain-specific compiler around a mutable value semantics core. Additionally, further work is required to augment the AD system with support for `inout`-formulated derivatives, and automatically determine when to use a functional-formulation of the pullback versus the `inout` formulation.

5 APPLICATIONS AND EVALUATION

Swift for TensorFlow brings together differentiable programming in a compiled language with high-level abstractions over heterogeneous compute resources, and enables applications in several domains, including machine learning and particularly the field of deep learning. Beyond machine learning, Swift for TensorFlow has been applied to differentiable physics simulations and highly-parallel numerical calculations.

A shared repository of examples, building blocks, and benchmarks is maintained at <https://github.com/tensorflow/swift-models>. As of this writing, over 30 examples exist at that location, spanning such areas as image classification, generative models, recommendation systems, and reinforcement learning.

Two recent works used Swift for TensorFlow to assist in reinforcement learning research. A simulated environment for never-ending learning called Jelly Bean World (Platanios et al., 2020) was constructed using Swift for TensorFlow in order to run efficiently on low-cost hardware. DeepMind released a framework for researching reinforcement learning in games called OpenSpiel (Lanctot et al., 2019) with an implementation in Swift for TensorFlow.

5.1 Training performance

To evaluate the performance of Swift for TensorFlow, we present three scenarios that demonstrate the range of hard-

ware supported by the platform. This includes training a deep learning model on accelerators such as Tensor Processing Units (TPUs) (Jouppi et al., 2017; 2020) and Graphics Processing Units (GPUs), as well as refining simpler machine learning models on mobile devices.

5.1.1 Training on TPUs

TPUs are a domain-specific hardware accelerator optimized for deep learning models. Because the XLA compiler is the only supported mechanism to generate TPU programs, Swift for TensorFlow uses the LazyTensor approach (detailed above) to JIT-compile traces of `Tensor` operations.

The performance of Swift for TensorFlow on TPUs was measured by training the ResNet-50 image classification network (He et al., 2016) on the ImageNet 2012 dataset (Deng et al., 2009) using TPUv3-16, TPUv3-32 and TPUv3-128 clusters, shown in Table 1. The model was trained for 90 epochs, and both the time required as well as the post-warmup throughput in examples per second were recorded. For comparison, a ResNet-50 model implemented in TensorFlow and in JAX using the Flax framework (Heek et al., 2020) were trained on TPUv3-32 under the same conditions, shown in Table 2. The model, dataset, and training conditions were selected to conform to the MLPerf Training Benchmark (Mattson et al., 2020).

Although Swift for TensorFlow trains in approximately equivalent time as both JAX and TensorFlow, Swift for TensorFlow is the only framework with an eager programming model on TPUs. By contrast, both JAX and TensorFlow stage the user’s code in their own intermediate representations (`jaxpr` and `GraphDef`, respectively). Note that because the Swift for TensorFlow implementation achieves higher accuracy typically associated with longer training times (due to algorithmic tweaks inspired by `fastai` (Howard & Gugger, 2020)), the other frameworks are at a slight advantage.

5.1.2 Training on GPUs

GPUs are also commonly used, so the training performance of Swift for TensorFlow was evaluated on a commodity NVIDIA GTX 1080, shown in Table 3. The ResNet-56 image classification network (He et al., 2016) was trained against the CIFAR-10 dataset (Krizhevsky et al.) for 10 epochs, measuring training throughput after a warmup period. Training was performed on the TensorFlow and PyTorch frameworks using the same model, dataset, and hardware.

Swift for TensorFlow’s LazyTensor tracing significantly outperforms the operation-by-operation dispatching of its eager mode, due to the overhead inherent in the latter and the XLA compiler’s ability to selectively fuse operations. PyTorch and TensorFlow both outperform Swift for Tensor-

Swift for TensorFlow

# Cores	Validation Accuracy (top-1)	Training Time (90 epochs)	Throughput (examples / s)	Per-Accelerator Throughput (examples / s / TPU core)
16	78.1%	189 minutes	10164	635.25
32	77.7%	96 minutes	20015	625.47
128	77.8%	25 minutes	77726	607.23

Table 1. Swift for TensorFlow training performance for ResNet-50 on ImageNet on TPUv3 clusters. Per-accelerator throughput is largely maintained while scaling from a single host to 8 hosts synchronously training a single model in data-parallel fashion, demonstrating the scalability of the LazyTensor approach employed by Swift for TensorFlow to target TPUs.

Framework	Validation Accuracy (top-1)	Training Time (90 epochs)	Throughput (examples / s)
JAX + Flax	76.8%	90 minutes	21258
TensorFlow	77.9%	59 minutes	33118
Swift for TensorFlow	77.7%	96 minutes	20015

Table 2. Training performance for ResNet-50 on ImageNet on a TPUv3-32 cluster. Although each system can notionally produce identical XLA HLO and thus achieve equivalent performance, some codebases have been better optimized for benchmark purposes. We include this table for completeness.

Flow’s LazyTensor backend on GPUs at present in part due to LazyTensor tracing overhead. Improvements in caching of LazyTensor traces should lead to increased throughput.

5.1.3 Training on mobile devices

Learning parameters through iterated optimization has applications beyond deep learning, such as learning knots in a polynomial spline. Splines require orders of magnitude less computation and are thus attractive in resource constrained environments such as mobile phones. Optimization algorithms such as backtracking line search use derivatives to determine the step direction.

Swift’s AD capabilities are not tied to any underlying accelerator interface or platform, and code can be compiled to target resource-constrained hardware, such as mobile or embedded devices. Swift is the most popular language to develop applications for Apple, Inc.’s iOS platform, and the compiler can target many other mobile or embedded hardware configurations.

As an example of this, a proprietary personalization model using splines optimized via backtracking line search was implemented in Swift for TensorFlow. The hardware-accelerated Swift `Tensor` type used to train on GPUs or TPUs was replaced with a CPU-only implementation backed by Swift arrays. This model was cross-compiled to binaries intended for ARM32 and ARM64 Android devices.

A global spline model was trained on anonymized, aggregated data, and fine-tuned on a Google Pixel 3 phone using only local data. We compare implementations of this algorithm using three distinct frameworks: TensorFlow Mobile (Abadi et al., 2016), TensorFlow Lite, and Swift for TensorFlow. In the case of TensorFlow Lite, two versions

were tested: one with the standard available operations and another with a custom fused operation designed specifically for this model. The spline model amplifies even slight numerical differences in calculations performed during training, so the results of all three frameworks were verified to produce control point values that matched within 1.5% of each other.

The results of fine-tuning this spline model to convergence are shown in Table 4. The Swift for TensorFlow implementation provides the lowest memory usage, as measured by the differences in peak process memory size before and after training. Swift for TensorFlow trains to convergence faster than all but the customized TensorFlow Lite implementation. At the time of these tests, the Swift compiler was unable to generate appropriate NEON vector instructions on Android for this model, which partially accounts for the time difference with the custom fused training operation.

For performance reasons, it is generally impractical to use Python model code when deploying to mobile devices, so the TensorFlow Mobile and TensorFlow Lite versions of this model used Python for the server-based global model training and a binary model graph driven via Java and C++ for on-device fine tuning. In contrast, the same Swift code defined and ran model training in both stages. This is advantageous for development and ongoing maintenance.

6 RELATED WORK

PyTorch (Paszke et al., 2019), Chainer (Tokui et al., 2019), and Autograd (Maclaurin, 2016) have explored a define-by-run programming model in the context of the Python programming language. Swift for TensorFlow provides a similar end-user experience, but in a statically typed pro-

Framework	Throughput (examples / s)
PyTorch	2462
TensorFlow	2390
Swift for TensorFlow (Eager Mode)	730
Swift for TensorFlow (LazyTensor)	1827

Table 3. Training performance for ResNet-56 on CIFAR-10 on an Nvidia GTX 1080 GPU.

Platform	Training Time (on device)	Memory Usage (on device)	Binary Size (uncompressed)
TensorFlow Mobile	5926 ms	80.0 MB	6.2 MB
TensorFlow Lite (standard operations)	266 ms	12.3 MB	1.8 MB
TensorFlow Lite (manually fused custom operation)	63 ms	6.2 MB	1.8 MB
Swift for TensorFlow	128 ms	4.2 MB	3.6 MB

Table 4. On-device training statistics for a personalized spline model across four different implementations.

programming language. Our AD system is orthogonal to the tensor library to allow the differentiation of arbitrary user defined types, and offers generalized support for defining derivatives alongside the original functions.

Both TensorFlow (Abadi et al., 2016) and Swift for TensorFlow share portions of the underlying technology stack. Swift for TensorFlow’s eager execution is based directly on TensorFlow Eager (Agrawal et al., 2019). We rely on lazy tracing to offer opt-in support for JIT compilation to XLA. TensorFlow exposes similar functionality using the `@tf.function` annotation.

JAX (Bradbury et al., 2020) extends NumPy (Harris et al., 2020) with support for opt-in JIT compilation based on XLA. JAX traces (Frostig et al., 2018) a pure subset of `@jit` annotated Python programs into its own intermediate representation called `jaxpr`, and transforms these `jaxprs` on-the-fly or after-the-fact to perform automatic differentiation, XLA compilation, and other program transformations.

Unlike JAX’s `@jit` and TensorFlow’s `@tf.function`, Swift for TensorFlow supports programs that mix both tensor and arbitrary computations that are run on the host. This is possible with lazy tracing that seamlessly integrates into the host language without breaking the eager-like programming model.

Similar to DyNet (Neubig et al., 2017), we implicitly construct computation graphs on each iteration of the training or inference loops. This allows us to support fully dynamic networks that can change architecture on each iteration, at a cost of incurring tracing overhead. Unlike DyNet, our LazyTensor implementation automatically decides when to materialize and compile the constructed graphs. To minimize this overhead, we use tracing purely to support just-in-time compilation on accelerators via XLA.

The Swift for TensorFlow AD system has a number of things in common with Zygote (Innes et al., 2019). Both systems rely on source-to-source transformations to automatically generate derivatives. The key differentiating aspect is that Swift for TensorFlow manages to do this completely statically in an AOT setting. This makes it suitable for supporting machine learning on edge devices where the cost of tracing and JIT compilation are infeasible.

Approaches built around the Lightweight Modular Staging framework (Rompf & Odersky, 2010) such as Flare (Essertel et al., 2018) and Lantern (Wang et al., 2019) perform compilation to heterogeneous hardware through a language-integrated support for staging. LazyTensor traces to an intermediate representation suitable for domain-specific compilers without breaking the illusion of eager execution. In contrast, LMS requires explicit type annotations to direct the staging framework.

7 CONCLUSION

In this paper we introduced Swift for TensorFlow, a platform for deep learning that has been validated with over 30 models. We have demonstrated how this platform scales across a wide range of hardware envelopes, how a modern AD system can be embedded in a statically typed language, and how a lazy implementation of a `Tensor` type allows JIT-compilation via domain-specific compilers without sacrificing the eager programming model. We have also shown how mutable value semantics addresses challenges when programming ML systems, including simple solutions to (a) reducing memory consumption on accelerators without compromising the programming model and (b) the array subscript derivative problem. We believe there is significant promise in exploring machine learning in the context of statically-typed AOT-compiled languages.

ACKNOWLEDGEMENTS

The authors would like to thank the many folks who made significant contributions to the Swift for TensorFlow project including Olzhas Akpambetov, Pawan Sasanka Ammanamanchi, Tanmay Bakshi, Paige Bailey, Gogul Balakrishnan, Eugene Burmako, James Bradbury, Bart Chrazzcz, Ed Connell, Mingsheng Hong, Brett Koonce, Shaddaj Laddad, Ewa Matejska, Adam Paszke, Anthony Platanios, and Alex Şuhan. We would also like to thank Jeremy Howard, Sylvain Gugger, and Rachel Thomas from fast.ai, whose collaboration substantially influenced the Swift for TensorFlow project. A big thank you goes out to the Google Summer of Code participants (Ayush Agrawal, Victor Antony, Param Bhavsar, Karthik Ramesh Iyer, Seungjae "Ryan" Lee, and Ayushi Tiwari) and our two Google Code-In Grand Prize winners (Rick Wierenga, William Zhang). The automatic differentiation system was heavily influenced by the work of, and conversations with Casey Chu, Conal Elliott, Roy Frostig, Matthew Johnson, Dougal Maclaurin, Bart van Merriënboer, Gordon Plotkin, Alexey Radul, Alex Wiltschko, and Dimitrios Vytiniotis. We also deeply appreciate the collaborations across Georgia Tech, Google, and DeepMind and would like to acknowledge the work of Mike Ando, Frank Dellaert, Dominik Grewe, Dmitri Gribenko, Tim Harley, Fan Jiang, James Keeling, Marc Lanctot, Edward Lockhart, and Julian Schrittwieser. This work would not have been possible without support from the Swift core team and related folks including: Steve Canon, Doug Gregor, Joe Groff, Ted Kremenek, John McCall, and Slava Pestov. Thank you to the many reviewers of this paper including Martin Abadi, James Bradbury, James Laudon, Rif A. Saurous, and Cliff Young for their invaluable feedback. And finally, we would like to thank the amazing Swift for TensorFlow community at large that has made multiple books, the majority of the models, and myriad other contributions; we couldn't have done this without you.

REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. TensorFlow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pp. 265–283, 2016.
- Agrawal, A., Modi, A. N., Passos, A., Lavoie, A., Agarwal, A., Shankar, A., Ganichev, I., Levenberg, J., Hong, M., Monga, R., et al. TensorFlow Eager: A multi-stage, Python-embedded DSL for machine learning. *arXiv preprint arXiv:1903.01855*, 2019.
- Amodei, D., Hernandez, D., Sastry, G., Clark, J., Brokman, G., and Sutskever, I. AI and Compute, 2018. URL <https://openai.com/blog/ai-and-compute/>.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., and Wanderman-Milne, S. JAX: composable transformations of Python + NumPy programs, 2018. URL <http://github.com/google/jax>, pp. 18, 2020.
- Chen, T., Moreau, T., Jiang, Z., Shen, H., Yan, E. Q., Wang, L., Hu, Y., Ceze, L., Guestrin, C., and Krishnamurthy, A. TVM: end-to-end optimization stack for deep learning. *CoRR*, abs/1802.04799, 2018. URL <http://arxiv.org/abs/1802.04799>.
- Şuhan, A., Libenzi, D., Zhang, A., Schuh, P., Saeta, B., Sohn, J. Y., and Shabalin, D. LazyTensor: A portable approach to combine eager execution and fusing compilers, 2020.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Driscoll, J. R., Sarnak, N., Sleator, D. D., and Tarjan, R. E. Making data structures persistent, 1989.
- Essertel, G., Tahboub, R., Decker, J., Brown, K., Olukotun, K., and Rompf, T. Flare: Optimizing Apache Spark with native compilation for scale-up architectures and medium-size data. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pp. 799–815, 2018.
- Fischer, K. and Saba, E. Automatic full compilation of Julia programs and ML models to Cloud TPUs. *arXiv preprint arXiv:1810.09868*, 2018.
- Frostig, R., Johnson, M. J., and Leary, C. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning*, 2018.
- Géron, A. *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: Concepts, tools, and techniques to build intelligent systems*. O'Reilly Media, 2019.
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., et al. Array programming with NumPy. *Nature*, 585(7825):357–362, 2020.
- He, K., Zhang, X., Ren, S., and Sun, J. Identity mappings in deep residual networks. In *European conference on computer vision*, pp. 630–645. Springer, 2016.

- Heek, J., Levsikaya, A., Oliver, A., Pop, R., Reddy, M., Rondepierre, B., Suo, D., and van Zee, M. Flax: A neural network library for jax designed for flexibility. 2020. URL <https://github.com/google/flax>.
- Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Howard, J. and Guggen, S. *Deep Learning for Coders with fastai and PyTorch: AI Applications Without a PhD*. O’Reilly Media, 2020.
- Innes, M., Edelman, A., Fischer, K., Rackauckus, C., Saba, E., Shah, V. B., and Tebbutt, W. Zygote: A differentiable programming system to bridge machine learning and scientific computing. *arXiv preprint arXiv:1907.07587*, pp. 140, 2019.
- Jouppi, N., Yoon, D., Kurian, G., Li, S., Patil, N., Laudon, J., Young, C., and Patterson, D. A domain-specific supercomputer for training deep neural networks. *Communications of the ACM*, 63:67–78, 06 2020. doi: 10.1145/3360307.
- Jouppi, N. P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., Bates, S., Bhatia, S., Boden, N., Borchers, A., et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pp. 1–12, 2017.
- Krizhevsky, A., Nair, V., and Hinton, G. CIFAR-10. URL <http://www.cs.toronto.edu/~kriz/cifar.html>.
- Lanctot, M., Lockhart, E., Lespiau, J.-B., Zambaldi, V., Upadhyay, S., Pérolat, J., Srinivasan, S., Timbers, F., Tuyls, K., Omidshafiei, S., Hennes, D., Morrill, D., Muller, P., Ewalds, T., Faulkner, R., Kramár, J., Vylder, B. D., Saeta, B., Bradbury, J., Ding, D., Borgeaud, S., Lai, M., Schrittwieser, J., Anthony, T., Hughes, E., Danihelka, I., and Ryan-Davis, J. OpenSpiel: A framework for reinforcement learning in games. *CoRR*, abs/1908.09453, 2019. URL <http://arxiv.org/abs/1908.09453>.
- Laurent Hascoet, V. P. The Tapenade automatic differentiation tool: principles, model, and specification. *ACM Transactions on Mathematical Software, Association for Computing Machinery*, 2013.
- Leary, C. and Wang, T. XLA: TensorFlow, compiled. *TensorFlow Dev Summit*, 2017.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, volume 86, pp. 2278–2324, 1998.
- Maclaurin, D. *Modeling, inference and optimization with composable differentiable procedures*. PhD thesis, 2016.
- Maclaurin, D., Duvenaud, D., and Adams, R. P. Autograd: Effortless gradients in NumPy. In *ICML 2015 AutoML Workshop*, volume 238, pp. 5, 2015.
- Maclaurin, D., Radul, A., Johnson, M. J., and Vytiniotis, D. Dex: array programming with typed indices. 2019.
- Mattson, P., Cheng, C., Damos, G. F., Coleman, C., Micikevicius, P., Patterson, D. A., Tang, H., Wei, G., Bailis, P., Bittorf, V., Brooks, D., Chen, D., Dutta, D., Gupta, U., Hazelwood, K. M., Hock, A., Huang, X., Kang, D., Kanter, D., Kumar, N., Liao, J., Narayanan, D., Oguntebi, T., Pekhimenko, G., Pentecost, L., Reddi, V. J., Robie, T., John, T. S., Wu, C., Xu, L., Young, C., and Zaharia, M. MLPerf training benchmark. In Dhillon, I. S., Papailiopoulos, D. S., and Sze, V. (eds.), *Proceedings of Machine Learning and Systems 2020, MLSys 2020, Austin, TX, USA, March 2-4, 2020*. mlsys.org, 2020. URL <https://proceedings.mlsys.org/book/309.pdf>.
- Moses, W. S. and Churavy, V. Instead of rewriting foreign code for machine learning, automatically synthesize fast gradients, 2020.
- Neubig, G., Dyer, C., Goldberg, Y., Matthews, A., Amar, W., Anastopoulos, A., Ballesteros, M., Chiang, D., Clothiaux, D., Cohn, T., et al. DyNet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*, 2017.
- O’Neill, M. The genuine sieve of Eratosthenes. *J. Funct. Program.*, 19:95–106, 01 2009. doi: 10.1017/S0956796808007004.
- Paszke, A. and Saeta, B. Tensors Fitting Perfectly, 2020.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pp. 8026–8037, 2019.
- Platanios, E. A., Saporov, A., and Mitchell, T. Jelly bean world: A testbed for never-ending learning. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=Byx_YAVYPH.
- Rompf, T. and Odersky, M. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled DSLs. In *Proceedings of the ninth international conference on Generative programming and component engineering*, pp. 127–136, 2010.

- Rotem, N., Fix, J., Abdulrasool, S., Deng, S., Dzhabarov, R., Hegeman, J., Levenstein, R., Maher, B., Satish, N., Olesen, J., Park, J., Rakhov, A., and Smelyanskiy, M. Glow: Graph lowering compiler techniques for neural networks. *CoRR*, abs/1805.00907, 2018. URL <http://arxiv.org/abs/1805.00907>.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and Wojna, Z. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826, 2016.
- Tan, M. and Le, Q. V. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019.
- Tokui, S., Okuta, R., Akiba, T., Niitani, Y., Ogawa, T., Saito, S., Suzuki, S., Uenishi, K., Vogel, B., and Vincent, H. Y. Chainer: A deep learning framework for accelerating the research cycle, 2019.
- Vassilev, V., Efremov, A., and Shadura, O. Automatic differentiation in ROOT, 2020.
- Vinyals, O., Babuschkin, I., Chung, J., Mathieu, M., Jaderberg, M., Czarnecki, W. M., Dudzik, A., Huang, A., Georgiev, P., Powell, R., et al. Alphastar: Mastering the real-time strategy game StarCraft II. *DeepMind blog*, pp. 2, 2019.
- Vytiniotis, D., Belov, D., Wei, R., Plotkin, G., and Abadi, M. The differentiable curry. 2019.
- Wang, F., Zheng, D., Decker, J., Wu, X., Essertel, G. M., and Rompf, T. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *Proceedings of the ACM on Programming Languages*, 3(ICFP):1–31, 2019.
- Wei, R., Zheng, D., Rasi, M., and Chrzasczcz, B. Differentiable programming manifesto. 2019. URL <http://github.com/apple/swift/blob/main/docs/DifferentiableProgramming.md>.

A INOUT AND MUTABLE VALUE SEMANTICS

```
func inc(_ x: inout Int) -> Bool {
  x = x + 1
  return x < 10
}

var y = 2, z = false
z = inc(&y)

print(y, z)
```

```
func inc(_ x0: Int) -> (Int, Bool) {
  let x = x0 + 1
  return (x, x < 10)
}

var y = 2, z = false
(y, z) = inc(y)

print(y, z)
```

Figure 8. Rewriting pass-by-inout as pass-by-value plus assignment. Both programs print “3 true.” A call using `inout` (left column) can be transformed into an equivalent call using pass-by-value (right column), by returning an updated value for each `inout` parameter to the original function and assigning into corresponding `inout` argument(s) at the call site. This equivalence demonstrates that `inout` does not introduce reference semantics.

B ARRAY SUBSCRIPT DIFFERENTIATION

Figure 9 contains a reference program, and two explicitly-written implementations of JVPs, one in the immutable-functional style, and one leveraging mutable value semantics. The operation to differentiate is $O(1)$. The functional pullback formulation runs in $O(n)$ time, where n is the size of values. The value semantic formulation runs in constant time, irrespective of the size of values. For an executable form of this appendix, please see: https://colab.research.google.com/github/tensorflow/swift/blob/master/notebooks/value_semantics/02_value_semantics_and_autodiff.ipynb that walks through the functional AD formulation before introducing this case study.

```

// Example operation to differentiate.
func myOp(values: [Float], a: Int, b: Int) -> Float {
    return values[a] + values[b]
}

// Functional representation

// subscript read with explicit pullback, functional style.
func subscriptWithFunctionalPullback(values: [Float], index: Int) -> (
    value: Float,
    pullback: (Float) -> [Float] // Functional formulation.
) {
    let size = values.count // Optimization: don't capture whole array, just size.
    return (values[index], { dx in
        var tmp = Array(repeating: Float(0), count: size) // Allocates O(n) memory!
        tmp[index] = dx
        return tmp
    })
}

// Returns a new array where each element is the sum of the corresponding
// elements of `a` and `b`.
func sumArraysHelper(_ a: [Float], _ b: [Float]) -> [Float] {
    precondition(a.count == b.count)
    var result = Array(repeating: Float(0), count: a.count)
    for i in 0..

```

Figure 9. A simple, differentiable function using array subscripts, and two alternative AD formulations explicitly written, one in a pure-functional paradigm, and a second leveraging mutable value semantics which reduces derivative complexity from $O(n)$ to $O(1)$.