KOTLIN∇

A SHAPE-SAFE DSL FOR DIFFERENTIABLE FUNCTIONAL PROGRAMMING

Anonymous Authors[1]

ABSTRACT

Kotlin is a statically-typed programming language with support for embedded domain specific languages, asynchronous programming, and multi-platform compilation. In this work, we present an algebraically-grounded implementation of forward- and reverse- mode automatic differentiation (AD) with shape-safe tensor operations, written in pure Kotlin. Our approach differs from existing AD frameworks in that Kotlin∇ is the first shape-safe AD library fully compatible with the Java type system, requiring no metaprogramming, reflection or compiler intervention to use. A working prototype is available at: https://github.com/anonymous/kotlingrad.

1 INTRODUCTION

Existing AD frameworks for machine learning are implemented in Python, which is not a type-safe language. Some AD implementations are written in statically-typed languages, but only type check primitive data types (e.g. integers or floating point numbers), and are unable to check the shape of multidimensional arrays in their type system. Those which do are typically implemented in experimental programming languages with sophisticated type-level programming features. In our work, we demonstrate an automatic differentiation library with shape-checked array programming in a widely-used programming language.

Prior work has shown it is possible to encode a deterministic context-free grammar as a *fluent interface* (Gil & Levy, 2016) in Java. This result was strengthened to prove Java's type system is Turing Complete (Grigore, 2017), allowing arbitrary computation to occur in the type system. As a practical consequence, we can use a similar technique to perform shape-safe automatic differentiation (AD) in Java, using type-level programming. A similar technique is feasible in any language with parametric polymorphism, such as C# or TypeScript. We use *Kotlin*, whose type system is strictly less expressive, but fully interoperable with Java.

Differentiable programming has a rich history among dynamic languages like Python, Lua and JavaScript, with early implementations including projects like Theano, Torch, and TensorFlow. Similar ideas have been implemented in statically typed, functional languages, such as Haskell's Stalin∇ (Pearlmutter & Siskind, 2008b), DiffSharp in F# (Baydin et al., 2015) and recently Swift 2018tensorflow. However, the majority of existing automatic differentiation (AD) frameworks use a loosely-typed DSL, and few offer shape-safe tensor operations in a widely-used programming language.

Existing AD implementations for the JVM include Lantern (Wang et al., 2018), Nexus (Chen, 2017) and DeepLearning.scala (Bo, 2018), however these are Scala-based and do not interoperate with other JVM languages. Kotlin∇ is fully interoperable with vanilla Java, enabling broader adoption in neighboring languages. To our knowledge, Kotlin has no prior AD implementation. However, the language contains a number of desirable features for implementing a native AD framework. In addition to type-safety and interoperability, Kotlin∇ primarily relies on the following language features:

1. **Operator overloading and infix functions** allow a concise notation for defining arithmetic operations on tensor-algebraic structures, i.e. groups, rings and fields.
2. **λ-functions and coroutines** support backpropogation with lambdas and shift-reset continuations, following Pearlmutter & Siskind 2008a and Wang et al. 2018.
3. **Extension functions** support extending classes with new fields and methods and can be exposed to external callers without requiring sub-classing or inheritance.

2 USAGE

Kotlin∇ allows users to implement differentiable functions by composing operations on algebraic fields to form expressions. Operations on variables with incompatible shape will fail to compile. Valid expressions are lazily evaluated inside a type-safe numerical context and numerical evaluation only

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the Montréal AI Sympoisum (MAIS). Do not distribute.

occurs when a function is passed concrete input values.

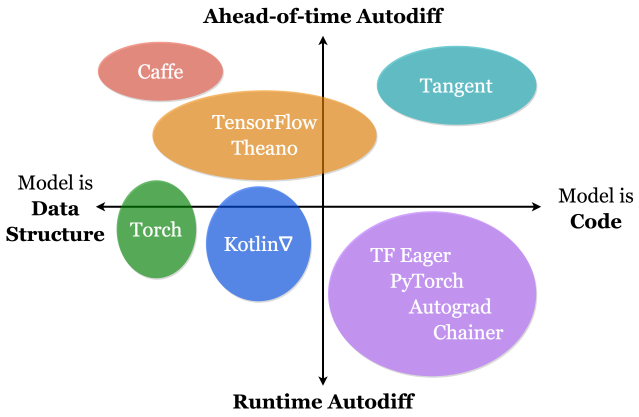

Figure 1. Kotlin∇ models are data structures, constructed by an embedded DSL, eagerly optimized, and lazily evaluated at runtime.

```
import edu.umontreal.kotlingrad.numerics.DoublePrecision

with(DoublePrecision) { // Use double-precision protocol
  val x = variable("x") // Declare immutable vars (these
  val y = variable("y") // are just symbolic constructs)
  val z = sin(10 * (x * x + pow(y, 2))) / 10 // Lazy exp
  val dz_dx = d(z) / d(x) // Leibniz derivative notation
  val d2z_dxdy = d(dz_dx) / d(y) // Mixed higher partial
  val d3z_d2xdy = grad(d2z_dxdy)[x] // Indexing gradient
  plot3D(d3z_d2xdy, -1.0, 1.0) // Plot in -1 < x,y,z < 1
}
```

Figure 2. A basic Kotlin∇ program with two inputs and one output.

Above, we define a function with two variables and take a series of partial derivatives with respect to each variable. The function is numerically evaluated on the interval $(-1, 1)$ in each dimension and rendered in 3-space. We can also plot higher dimensional manifolds (e.g. the loss surface of a neural network), projected into four dimensions, and rendered in three, where one axis is represented by time.

Kotlin∇ treats mathematical functions and programming functions with the same underlying abstraction. Expressions are composed recursively to form a data-flow graph (DFG). An expression is simply a `Function`, which is only evaluated once invoked with numerical values, e.g. `z(0, 0)`.

Kotlin∇ supports shape-shafe tensor operations by encoding tensor rank as a parameter of the operand's type signature. By enumerating type-level integer literals, we can define tensor operations just once using the highest literal, and rely on Liskov substitution to preserve shape safety for subtypes.

It is possible to enforce shape-safe vector construction as well as checked vector arithmetic up to a fixed `L`, but the full implementation is omitted for brevity. A similar pattern can be applied to matrices and higher rank tensors, where the type signature encodes the shape of the operand at runtime.

With these basic ingredients, we have almost all the features

$$z = \sin\left(10(x \times x + y^2)\right)/10, \quad \texttt{plot3D}\left(\frac{\partial^3 z}{\partial x^2 \partial y}\right)$$

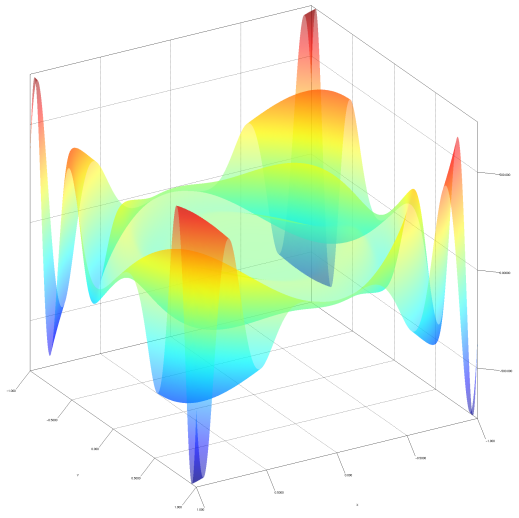

Figure 3. Output generated by the program shown in Figure 2.

```
val z = sin(10 * (x * x + pow(y, 2))) / 10
```

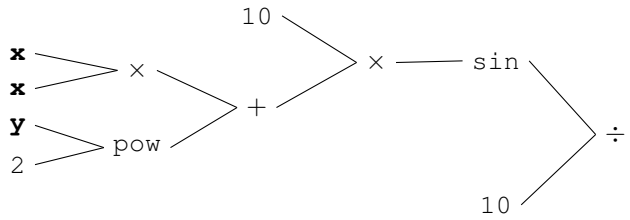

Figure 4. Implicit DFG constructed by the original expression, `z`.

```
// Literals have reified values for runtime comparison
sealed class `0`(open val value: Int = 0)
sealed class `1`(override val value: Int = 1): `0`(1)
sealed class `2`(override val value: Int = 2): `1`(2)
sealed class `3`(override val value: Int = 3): `2`(3)
// <L: `1`> will accept 1<=L<=3 via Liskov substitution
class Vec<E, L: `1`>(len: L, cts: List<E> = listOf())
// Define addition for two vectors of type Vec<Int, L>
operator fun <L: `1`, V: Vec<Int, L>> V.plus(v: V) =
  Vec<Int, L>(len, cts.zip(v.cts).map { it.l + it.r })
// Type-checked vector addition with shape inference
val Y = Vec(`2`, listOf(1, 2)) + Vec(`2`, listOf(3, 4))
val X = Vec(`1`, listOf(1, 2)) + Vec(`3`) // Undefined!
```

Figure 5. Shape safe tensor addition for rank-1 tensors, $\forall L \leq 3$.

necessary to build an expressive shape-safe AD, but unlike prior implementations using Scala or Haskell, in a language that is fully interoperable with Java, while also capable of compiling to JVM bytecode, JavaScript, and native code.

In future work, we plan to implement a broader grammar of differentiable primitives including matrix convolution, as well as control flow and recursion. While Kotlin∇ currently implements matrix arithmetic directly, we plan to wrap a

native linear algebra library for performance reasons.

REFERENCES

Baydin, A. G., Pearlmutter, B. A., and Siskind, J. M. DiffSharp: Automatic differentiation library. *CoRR*, abs/1511.07727, 2015. URL http://arxiv.org/abs/1511.07727.

Bo, Y. DeepLearning.scala: A simple library for creating complex neural networks. 2018. URL https://github.com/ThoughtWorksInc/DeepLearning.scala.

Chen, T. Typesafe abstractions for tensor operations (short paper). pp. 45–50, 2017. doi: 10.1145/3136000.3136001. URL http://doi.acm.org/10.1145/3136000.3136001.

Gil, Y. and Levy, T. Formal Language Recognition with the Java Type Checker. 56:10:1–10:27, 2016. ISSN 1868-8969. doi: 10.4230/LIPIcs.ECOOP.2016.10. URL http://drops.dagstuhl.de/opus/volltexte/2016/6104.

Grigore, R. Java generics are Turing Complete. pp. 73–85, 2017. doi: 10.1145/3009837.3009871. URL http://doi.acm.org/10.1145/3009837.3009871.

Pearlmutter, B. A. and Siskind, J. M. Reverse-mode AD in a functional framework: Lambda the ultimate backpropagator. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):7, 2008a.

Pearlmutter, B. A. and Siskind, J. M. Using programming language theory to make automatic differentiation sound and efficient. pp. 79–90, 2008b. ISSN 1439-7358. doi: 10.1007/978-3-540-68942-3_8. URL http://www.bcl.hamilton.ie/~barak/papers/sound-efficient-ad2008.pdf.

Wang, F., Wu, X., Essertel, G. M., Decker, J. M., and Rompf, T. Demystifying differentiable programming: Shift/reset the penultimate backpropagator. *CoRR*, abs/1803.10228, 2018. URL http://arxiv.org/abs/1803.10228.