
Fashionable Modelling with Flux

Michael J Innes
Julia Computing, Inc.
Edinburgh, UK
mike.j.innes@gmail.com

Elliot Saba
Julia Computing, Inc.
Cambridge, MA, USA
staticfloat@gmail.com

Keno Fischer
Julia Computing, Inc.
Cambridge, MA, USA
keno@juliacomputing.com

Dhairya Gandhi
Julia Computing, Inc.
Bangalore, India
dhairya@juliacomputing.com

Marco Concetto Rudilosso
University College London
London, UK
marco.rudilosso.15@ucl.ac.uk

Neethu Mariya Joy
Birla Institute of Technology and Science
Pilani, India
neethumariyajoy@gmail.com

Tejan Karmali
National Institute of Technology
Goa, India
tejank10@gmail.com

Avik Pal
Indian Institute of Technology
Kanpur, India
avikpal@cse.iitk.ac.in

Viral B. Shah
Julia Computing, Inc.
Cambridge, MA, USA
viral@juliacomputing.com

Abstract

Machine learning as a discipline has seen an incredible surge of interest in recent years, due in large part to a perfect storm of new theory, superior tooling, and renewed interest in its capabilities. We present in this paper a framework named Flux that shows how further refinement of the core ideas of machine learning, built upon the foundation of the Julia programming language, can yield an environment that is simple, easily modifiable, and performant. We detail the fundamental principles of Flux as a framework for differentiable programming, give examples of models that are implemented within Flux to display many of the language and framework-level features that contribute to its ease of use and high productivity, display internal compiler techniques used to enable the acceleration and performance that lies at the heart of Flux, and finally give an overview of the larger ecosystem that Flux fits inside of.

1 Introduction

Flux is a new machine learning (ML) stack. Only a year old and developed by a very small team, it nonetheless has garnered an enthusiastic and rapidly growing community of researchers and industry users, building new kinds of models and techniques.

ML engineering is fundamentally a programming languages problem; machine learning models are growing ever more complex, incorporating control flow, state and data structures, and borrowing techniques from research areas throughout the field of computer science from software engineering to scientific simulation to statistical inference. These models are best viewed as *differentiable algorithms*; and to express algorithms we use programming languages [15].

Flux takes Julia [6] to be this language. Julia is recent but, being designed from the ground up for mathematical and numerical computing, unusually well-suited for expressing ML programs. Its mix of modern design and novel techniques in the compiler—which we have significantly extended to further suit differentiable algorithms and accelerators—makes it easier to address the high performance requirements needed for applying machine learning to large models and data sets.

There are three pillars that set Flux apart among ML systems: simplicity, hackability, and underlying compiler technology.

1.1 Simplicity

The word “simple” is typically not thrown around in the ML systems world. However, simplifying solutions is a crucial part of approaching more complex problems. Where typical ML frameworks are written in many hundreds of thousands of lines of C++ [1], Flux is only a thousand lines of straightforward Julia code.

Several factors compound to enable this. First, the ability to write everything from layer definitions, to algorithmic differentiation, to CUDA kernels, in high-level and *fast* Julia code enables high developer productivity; code can typically be written once and forgotten about rather than being written, rewritten and optimized in C over years. Fast high-level code also enables something much more powerful than convenience: abstraction. This means infrastructure like higher-order kernels for `map(f, xs)` and `broadcast`, which can be written once and generalize to any user-defined `f` or input type with no extra effort or undue performance overhead.

This extends to integration with other tools in the ecosystem. Writing an image pipeline really means loading the Flux and Images [13] packages, and using them as normal. In fact, even GPU support is handled this way; rather than being provided as part of Flux, one loads a generic GPU arrays package such as CuArrays [14], and passing those arrays into a Flux model just works; transferring data and computational kernels off to the GPU accelerator then bringing the results back to the CPU without any special handling within Flux itself. This leverages Julia’s heavy use of specialization, which effectively generates custom code for the GPU, and even for custom floating-point types that would otherwise need to be written by hand. Far from becoming an all-encompassing monolith, Flux remains a lean “glue” package, bringing together a set of underlying abstractions (e.g. gradients and GPU support) and combining them to create an ML framework.

1.2 Hackability

A core tenet of Flux, and Julia more generally, is that library code is just user code that happens to have been loaded from an external file. Unlike previous ML frameworks (including those written in Julia), Flux does not pick a certain level of abstraction (such as mathematical graphs or layer stacking) that all models must use. Instead, careful design of the underlying automatic differentiation (Section 3.2) allows freely mixing mathematical expressions, built-in and custom layers and algorithms with control flow in one model. This makes Flux unusually easy to extend to new problems.

Flux users regularly inject custom CUDA kernels, write down new mathematical functions and layers, hook in custom gradient definitions and even custom parallel training algorithms, all with the same performance as built-in features. Hooking in custom functionality is often a one-line change that happens in the same Julia script as model code. Because there is no difference between Flux code and other general Julia code, it is possible to break out of the typical deep learning paradigm and experiment with new concepts such as interleaving gradient descent training with MCMC. Just write down your own training loop; the obvious mathematical code will work, and be fast.

This extends to integrating other packages with Flux models, such as Julia’s state-of-the-art tools for mathematical optimization [9], differential equations [28] and probabilistic programming [10]. This goes deeper than just using neural nets alongside these other techniques, as one can even incorporate them directly into models. For example, Julia’s differential equation solvers—while not explicitly adapted either for AD or GPUs—can seamlessly be used with both. This means that one can use a physical simulation to enhance the predictions of the model, then backpropagate through the entire model *including* the simulation to get gradients. Bringing tools like physics simulators into models is where deep learning truly becomes differentiable programming.

1.3 Compiler Technology

Flux is committed to providing a dynamic (or “define-by-run”) interface, and takes a hard line against any kind of graph building or performance annotations [27]. We support all of Julia’s language features, from control flow and data structures to macros. Users can code interactively in Jupyter notebooks and combine high-performance numerics with convenient plotting and visualization. But we also want to get the benefits traditionally held by “static graph” frameworks—zero-overhead source-to-source AD, operator fusion, multi-GPU/distributed training, and single-executable deployment.

Doing this effectively requires extracting and analyzing “static graphs” directly from written Julia syntax; but this really just describes the task of a *compiler*. Most ML systems problems are, in fact, standard and well-studied compiler problems, viewed through the right lens. Using a compiled language is enough to solve many issues, and extending that compiler is the best way to solve many more.

Several illustrative examples, covering differentiation (Section 3.2), GPU and TPU compilation (Sections 3.1 and 3.3) and SPMD / batching (Section 3.4) are covered briefly in this paper.

2 Fashionable Modelling

2.1 Discriminative Adversarial Networks

Nonstandard gradient control is a growing need in the machine learning community as model architectures are stretched to complete more and more complex tasks. A significant recent development has been the emergence of adversarial networks [11], where two networks are arranged in opposition to each other, one attempting to learn a task and the other inventing new inputs that the first can learn from. Extensions of this idea include work done in [30] using Flux, where adversarial networks were used to decrease bias induced by dataset imbalance.

The fundamental problem was strong correlation between classification label (in this case, *tuberculosis* versus *non-tuberculosis* coughs) and originating dataset (in this case, *clinical* versus *non-clinical*). The machine learning model naturally learned to key off of the differences in dataset, rather than the differences in the cough sounds themselves, because the differences between clinical and non-clinical recordings were larger than the differences between tuberculosis and non-tuberculosis coughs within a single dataset. Due to the fact that the overwhelming majority of non-tuberculosis coughs were from the non-clinical dataset, the classification algorithm was trapped in a local minimum simply predicting “tuberculosis” for every sample from the clinical dataset.

In order to solve this, a Discriminative Adversarial Network (DAN) was employed, building a second network explicitly designed to determine dataset provenance. This network is then used to “penalize” a shared set of weights (the convolutional block shown in red in Figure 1) for extracting information that can be used to determine which dataset a sample originated from.

```
1 for x, y_c, y_d in training_set
2     y_c_hat, y_d_hat = model(x)
3
4     c_loss = loss(y_c_hat, y) +
5             lambda*loss(y_d_hat, 1 - y_d)
6     d_loss = loss(y_d_hat, y_d)
7
8     back!(c_loss)
9     back!(d_loss)
10
11     opt()
12 end
```

Listing 1: DAN training loop

The code necessary to perform this nonstandard optimization task is surprisingly simple; shown in Listing 1, it simply calculates a forward pass through the model, returning the outputs from the two branches of the model as y_hat and y_dan_hat , calculates the classifier loss (c_loss) and the DAN loss (d_loss), backpropagates the losses onto the network, then takes an optimizer step ($opt()$). As can be seen, the optimization framework of Flux follows very naturally from the basic mathematical principles of optimization and provides for very flexible training and evaluation paradigms.




Figure 1: Cough classification architecture with DAN and classifier architectures visualized. Along the bottom lies the convolutional classification network that classifies cough type, with the outputs from the first layer feeding into a multilayer perceptron that classifies dataset source.

3 Extending Julia’s Compiler

3.1 Compiling Julia for GPUs

Julia supports the basic CUDA programming model for writing GPU kernels [5]. A simple vector addition kernel looks similar to the CUDA C equivalent (Listing 2).

```

1 function kernel_vadd(a, b, c)
2     i = (blockIdx().x-1) * blockDim().x +
3         threadIdx().x
4     c[i] = a[i] + b[i]
5 end

```

Listing 2: Basic CUDA programming model expressed in Julia

However, Julia’s type specialization enables a powerful set of additional abstractions on the GPU. For example, the code above is not restricted to dense arrays of floats, and could instead be given sparse arrays of complex numbers; Julia’s normal specialization mechanisms would generate a new set of PTX instructions for that case. We can even abstract this code further into a “higher-order kernel” that accepts the $+$ function (or $*$, or arbitrary user-defined f) and thus create a whole family of functions $\text{map}(f, x, y)$ in four lines of code [16].

Since this works for user-defined types, we can additionally make use of dual numbers for forward-mode differentiation [29]. Compared to reverse mode techniques, dual numbers have the significant advantage of being stack-allocated and interleaving primal and tangent computation, and thus in memory bound situations (such as GPUs) the derivatives are effectively free. This is particularly valuable in functions that contain control flow.

In the case of element-wise operations, such as $c = \tanh.(a .+ b)$ with vectors of length N , we can see this single $\mathbb{R}^{2N} \rightarrow \mathbb{R}^N$ computation as N independent $\mathbb{R}^2 \rightarrow \mathbb{R}$ ones, for which forward mode is very efficient. In other words, by running the same computation with dual numbers we efficiently compute the (very sparse) Jacobians $\partial c/\partial a$ and $\partial c/\partial b$ fused with the original operation, which can then be used within the reverse-mode sweep to get gradients. In the base case this allows us to fully fuse elementwise operations like the above, but it also generalizes to complex user-defined

functions including control flow, and can be much faster than the equivalent series of vector-level operations that would have to be used otherwise.

3.2 Algorithmic Differentiation

Pushing the limits of reverse-mode differentiation [32], we have also come to see this as a language-level problem. Differentiation is a symbolic transformation that works on programs or symbolic expressions, and this is the domain of compilers. Existing ML frameworks achieve this transformation by *tracing* (or *partial evaluation*); a new tensor type is introduced which records all the basic mathematical operations performed, yielding a graph with the control flow and data structures of the host language elided. This graph, equivalent to a Wengert list [3], is more easily differentiated than the original program.

Frameworks are then divided into the “dynamic” and “static” approaches [22] depending on whether they interpret [21] or compile [4] the recorded graph. However, this presents to us a difficult tradeoff: we either accept the overhead of an interpreter or freeze user control flow and limit the kinds of models that can be built (perhaps providing limited additional primitives to place control flow into the graph).

We reject these constraints and assert that the “graph” can instead simply be the Julia syntax tree. Extending previous work on differentiable languages [25] we have built *Zygote*, a source-to-source auto-differentiator in Julia, which works directly on SSA-form IR and supports language features like control flow, recursion, data structures and macros, resolving the AD trade-off. By putting the generated SSA-form adjoint code through a traditional compiler such as LLVM [19], we get all the benefits of traditional compiler optimization applied to both our forward and backwards passes. In addition, it opens up the possibility of extending that compiler infrastructure with more advanced and domain-specific optimizations, such as kernel fusion and compilation to accelerators such as TPUs.

3.3 Compiling Julia for TPUs

Google recently introduced a new API that allows users of their Cloud TPU offering to directly generate IR for the *XLA* (“Accelerated Linear Algebra”) compiler. This IR is essentially a general purpose IR and optimizing compiler for expressing arbitrary computations of linear algebra primitives and thus provides a good foundation for targeting TPUs by non-Tensorflow users as well as for non-machine learning workloads. We take advantage of Julia’s dynamic type system and extendable compiler to build the capability to compile arbitrary Julia code to XLA, and then to run it on TPUs. This allows users to take advantage of the full expressiveness of the Julia programming language in writing their models, including multiple dispatch, higher order functions and existing libraries such as those for differential equation solvers and generic linear algebra routines, while reaping the benefits of the high-performance systolic array engine within the TPU.

The basic workflow of the Julia compiler is to analyze chunks of a Julia program, identify static sub-segments, compile those sub-segments, and run them, returning to a dynamic environment only when necessary to compute, at runtime, the next static sub-segment to be compiled/run. By hooking into this compilation infrastructure, we are able to define custom compiler passes that identify static sub-segments of Julia code and compile them directly to blocks of static XLA IR, which can be run directly on TPUs through the newly exposed API offered by Google.

While many dynamic languages are capable of limited amounts of static analysis, it is worth pointing out that the “static sub-segments” being referred to here are often quite large due to the Julia compiler’s aggressive type specialization, type inference and constant propagation. As an illustrative example, we are able to compile the VGG19 [31] machine learning model’s forward pass, backward pass and optimization step (that is to say, the entire training loop) into a single chunk of XLA code that can be run on the TPU without breaking out into Julia at all. This technical capability is critical in a heterogeneous computing system such as the TPU, as communication latency between the Julia host process and the TPU accelerator is significant, and must be avoided whenever possible.

3.4 Automatic Batching

The naturally data-parallel structure of many ML models makes them an excellent fit for massively parallel processors such as GPUs [24] and TPUs. To get the most from these accelerators—which

can have significant constant overheads per kernel launch, but scale very well over input size—it is common to *batch* them, applying the forwards and backwards passes to multiple training examples at once. In simple models this can be achieved by stacking a set of images or samples along an additional batch dimension, and primitives such as matrix multiply, convolution and broadcasting naturally handle this extra dimension as if all samples were independent.

However, this task becomes much harder when dealing with variably-structured inputs, such as trees [33] or graphs [18], or any kind of value-dependent control flow. Most researchers address this by taking on the significant burden of batching code by hand. Different solutions have been proposed for different frameworks [23] [20], which heuristically try to batch some high level operations together when possible, using different policies to choose what to batch and when, but these have their own usability issues and typically do not achieve the performance of hand-written code.

We propose that this problem is identical to that of Single Program Multiple Data (SPMD) programming, which has been well-studied by the language and compiler community for decades [7, 2], and becomes visible in more recent approaches to batching such as Matchbox [8]. Indeed, it is very similar to the model of parallelism used by GPUs internally, and has been implemented as a compiler transform for the SIMD units of CPUs [26]. Taking inspiration from this work, we are implementing the same transform [12] in Julia to provide SPMD programming both for scalar SIMD units and for model-level batching. This allows us to reach the ideal of writing simple code that operates on individual samples, while still getting the best performance on modern hardware.

3.5 JavaScript Compilation

In an approach similar to that of the GPU and TPU compilation efforts, the FluxJS package [17] supports compiling Julia models to Javascript, using an approach based on partial evaluation of model code. By defining a small number of fundamental operator equivalencies between Julia and Javascript, this package is able to parse the Julia syntax tree to create a function call graph, including control flow such as in RNNs, and makes use of the `tensorflow.js` [34] package in order to accelerate specific machine learning operations.

4 Conclusion

In conclusion, we have presented the reasons why Flux in particular and Julia in general provide an excellent environment for high performance, simple and hackable machine learning. We gave examples of complex models and functions defined in clearly readable, mathematically recognizable forms and detailed many of the techniques used to ensure that the generated code is not only performant on traditional CPUs, but also the accelerators that are increasingly critical to applied machine learning today.

References

- [1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. Tensorflow: a system for large-scale machine learning. In *OSDI*, volume 16, pages 265–283, 2016.
- [2] J. R. Allen, K. Kennedy, C. Porterfield, and J. Warren. Conversion of control dependence to data dependence. In *Proceedings of the 10th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '83, pages 177–189, New York, NY, USA, 1983. ACM.
- [3] M. Bartholomew-Biggs, S. Brown, B. Christianson, and L. Dixon. Automatic differentiation of algorithms. *Journal of Computational and Applied Mathematics*, 124(1-2):171–190, 2000.
- [4] J. Bergstra, F. Bastien, O. Breuleux, P. Lamblin, R. Pascanu, O. Delalleau, G. Desjardins, D. Warde-Farley, I. Goodfellow, A. Bergeron, et al. Theano: Deep learning on gpus with python. In *NIPS 2011, BigLearning Workshop, Granada, Spain*, volume 3, pages 1–48. Citeseer, 2011.
- [5] T. Besard, C. Foket, and B. De Sutter. Effective extensible programming: Unleashing julia on gpus. *arXiv preprint arXiv:1712.03112*, 2017.
- [6] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.
- [7] G. E. Blelloch. *Vector models for data-parallel computing*, volume 75. MIT press Cambridge, 1990.

- [8] J. Bradbury and C. Fu. Automatic batching as a compiler pass in PyTorch, 2018.
- [9] I. Dunning, J. Huchette, and M. Lubin. Jump: A modeling language for mathematical optimization. *SIAM Review*, 59(2):295–320, 2017.
- [10] H. Ge, K. Xu, and Z. Ghahramani. Turing: a language for flexible probabilistic inference. In *International Conference on Artificial Intelligence and Statistics*, pages 1682–1690, 2018.
- [11] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [12] R. K. S. Hack, R. Wilhelm, S. Neumann, R. Spurk, J. Doerfert, M. Jacobs, T. Jung, S. Moll, F. Ritter, R. Leißa, et al. Whole-function vectorization. 2011.
- [13] T. Holy and Contributors. Images.jl: An image library for julia, 2016.
- [14] M. Innes, T. Besard, and Contributors. Cuarrays.jl: A curious cumulation of cuda cuisine, 2018.
- [15] M. Innes, S. Karpinski, V. Shah, D. Barber, P. Stenertorp, T. Besard, J. Bradbury, V. Churavy, S. Danisch, A. Edelman, et al. On machine learning and programming languages, 2018.
- [16] M. J. Innes. Generic gpu kernels, 2017.
- [17] N. M. Joy and M. Innes. Fluxjs.jl: Flux to javascript compiler, 2018.
- [18] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907, 2016.
- [19] C. Lattner and V. Adve. LLVM: A compilation framework for lifelong program analysis & transformation. In *Proceedings of the international symposium on Code generation and optimization: feedback-directed and runtime optimization*, page 75. IEEE Computer Society, 2004.
- [20] M. Looks, M. Herreshoff, D. Hutchins, and P. Norvig. Deep learning with dynamic computation graphs. *arXiv preprint arXiv:1702.02181*, 2017.
- [21] D. Maclaurin, D. Duvenaud, and R. P. Adams. Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, 2015.
- [22] G. Neubig, C. Dyer, Y. Goldberg, A. Matthews, W. Ammar, A. Anastasopoulos, M. Ballesteros, D. Chiang, D. Clothiaux, T. Cohn, et al. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*, 2017.
- [23] G. Neubig, Y. Goldberg, and C. Dyer. On-the-fly operation batching in dynamic computation graphs. In *Advances in Neural Information Processing Systems*, pages 3971–3981, 2017.
- [24] K.-S. Oh and K. Jung. Gpu implementation of neural networks. *Pattern Recognition*, 37(6):1311–1314, 2004.
- [25] B. A. Pearlmutter and J. M. Siskind. Reverse-mode ad in a functional framework: Lambda the ultimate backpropagator. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):7, 2008.
- [26] M. Pharr and W. R. Mark. ispc: A spmd compiler for high-performance cpu programming. *2012 Innovative Parallel Computing (InPar)*, pages 1–13, 2012.
- [27] PyTorch Team. The road to 1.0: production ready PyTorch. <https://pytorch.org/blog/a-year-in/>, 2018. Accessed: 2018-09-22.
- [28] C. Rackauckas and Q. Nie. Differentialequations.jl—a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of Open Research Software*, 5(1), 2017.
- [29] J. Revels, M. Lubin, and T. Papamarkou. Forward-mode automatic differentiation in julia. *arXiv preprint arXiv:1607.07892*, 2016.
- [30] E. Saba. *Techniques for Cough Sound Analysis*. PhD thesis, University of Washington, 2018.
- [31] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

- [32] B. Speelpenning. Compiling fast partial derivatives of functions given by algorithms. Technical report, Illinois Univ., Urbana (USA). Dept. of Computer Science, 1980.
- [33] K. S. Tai, R. Socher, and C. D. Manning. Improved semantic representations from tree-structured long short-term memory networks. *CoRR*, abs/1503.00075, 2015.
- [34] T. Team. Introducing tensorflow.js: Machine learning in javascript, 2018.