

Extracted from:

Machine Learning in Elixir

Learning to Learn with Nx and Axon

This PDF file contains pages extracted from *Machine Learning in Elixir*, published by the Pragmatic Bookshelf. For more information or to purchase a paperback or PDF copy, please visit <http://www.pragprog.com>.

Note: This extract contains some colored text (particularly in code listing). This is available only in online versions of the books. The printed versions are black and white. Pagination might vary between the online and printed versions; the content is otherwise identical.

Copyright © 2023 The Pragmatic Programmers, LLC.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher.

The Pragmatic Bookshelf

Raleigh, North Carolina

The
Pragmatic
Programmers

Your Elixir Source

Machine Learning in Elixir

Learning to Learn with Nx and Axon



Sean Moriarity
Edited by Tammy Coron

Machine Learning in Elixir

Learning to Learn with Nx and Axon

Sean Moriarity

The Pragmatic Bookshelf

Raleigh, North Carolina



Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and The Pragmatic Programmers, LLC was aware of a trademark claim, the designations have been printed in initial capital letters or in all capitals. The Pragmatic Starter Kit, The Pragmatic Programmer, Pragmatic Programming, Pragmatic Bookshelf, PragProg and the linking *g* device are trademarks of The Pragmatic Programmers, LLC.

Every precaution was taken in the preparation of this book. However, the publisher assumes no responsibility for errors or omissions, or for damages that may result from the use of information (including program listings) contained herein.

For our complete catalog of hands-on, practical, and Pragmatic content for software developers, please visit <https://pragprog.com>.

For sales, volume licensing, and support, please contact support@pragprog.com.

For international rights, please contact rights@pragprog.com.

Copyright © 2023 The Pragmatic Programmers, LLC.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher.

ISBN-13: 979-8-88865-034-9

Encoded using the finest acid-free high-entropy binary digits.

Book version: B1.0—July 19, 2023

Preface

Machine learning, specifically deep learning, is constantly pushing the boundaries of what we thought was possible—and it's doing it in every industry.

Seemingly every day, a company or research group releases a new model that pushes state of the art even further ahead. In recent years, ChatGPT and Stable Diffusion, among others, have taken the world by storm, bringing artificial intelligence to the forefront and redefining what types of applications are possible.

For most of its existence, Elixir and the BEAM were not viable options for machine learning tooling. However, the Nx ecosystem changes that. Within this ecosystem, you can now write machine learning and numerical routines directly in Elixir and achieve the same or better performance as an equivalent program in Python or Julia. Nx offers new capabilities to Elixir programmers. It also provides an offramp for existing machine learning engineers and researchers looking to explore ecosystems and build applications without Python.

My goal in writing this book is for it to serve as the definitive Nx and machine learning resource for programmers who want to explore machine learning in Elixir. While a lot of machine learning knowledge is available online, none of the existing resources are written with Elixir in mind. This book is designed to teach machine learning the Elixir way—functional, pragmatic, and fun.

In this book, you'll build up from the basics of machine learning and Nx to deploying powerful pre-trained models in an application with Phoenix. Throughout this book, you'll use Livebook—Elixir's interactive code notebooks—to work with data and train models in an interactive and reproducible way. By the end of the book, you'll have experience working with Nx, Axon, Scholar, Bumblebee, Explorer, VegaLite, and many more of the libraries in the growing Elixir machine learning ecosystem.

We've finally made it to the always-promised but never delivered golden age of artificial intelligence-enabled products. So, if you want to build powerful, intelligent applications augmented with machine learning models, Elixir is one of the best languages to use.

Why Elixir for machine learning?

Elixir is a dynamic, functional, general-purpose programming language that targets the BEAM virtual machine. It's proven itself as a capable language for building and maintaining scalable applications. But, Elixir was not designed with machine learning in mind. Prior to the Nx¹ project, machine learning in Elixir was, at the very least, ill-advised. The BEAM virtual machine was not designed to perform the type of heavy-lifting that machine learning requires. However, Nx and the ecosystem around it were built to enable Elixir programmers to perform the computationally expensive work required for machine learning applications.

Elixir and Nx are new players in the machine learning field, and while the Nx ecosystem has no aim to overtake the established Python machine learning ecosystem, there are numerous benefits to understanding machine learning in Elixir.

First, Elixir is a *functional* programming language. While this might seem like a drawback for the heavy computation required in machine learning workloads, it's actually a benefit. In a functional world, everything is immutable. This means that when performing computations on large amounts of data, you need to return a new copy of the data after every individual operation. Fortunately, due to the rise of deep learning, computations are staged out to hardware accelerators such as GPUs before any of the actual operations are executed. These computations are staged out with Just-in-time (JIT) compilers, which typically implement a functional interface for building up computation graphs. As you'll see in later chapters, you can write in Elixir's functional style and not worry about large intermediate copies.

Second, Elixir is designed for concurrency. When training machine learning algorithms, the performance bottleneck is typically in loading and preparing data for a training step on a hardware accelerator. This type of concurrent data processing is natural in Elixir because Elixir was designed with concurrency in mind. You need only to read [Concurrent Data Processing in Elixir \[Gos21\]](#) to see how good a fit Elixir is for those types of applications.

1. <https://github.com/elixir-nx/nx>

Third, Elixir and its web framework, Phoenix, are exceptional at building and scaling real-time, fault-tolerant applications. More and more machine learning systems are being deployed in scenarios where latency and uptime matter. For example, a credit card fraud detection system should ideally work in real-time, immediately alerting customers when a fraudulent transaction occurs. Additionally, with large amounts of fraudulent transactions taking place everyday, those systems can't go down. Elixir and Phoenix have proven capable of building systems that are low-latency and resilient to failure. However, until now, integrating machine learning in an Elixir application required leaving the safety of the BEAM to an external service.

Finally, Elixir is a lot of fun and an absolute joy to use. As you'll see throughout this book, most machine learning algorithms are naturally expressed with functional constructs, even more so thanks to the beauty and thoughtfulness of the Elixir programming language.

Who This Book Is For

This book is for programmers interested in learning machine learning in Elixir and programmers with prior Elixir or machine learning experience.

Elixir programmers can pick up and read this book with no prerequisite knowledge of machine learning, mathematics, or numerical computing, as this book will teach you everything you need to know from the basics on.

Programmers with machine learning experience but no Elixir experience will be able to follow along, but some syntactic details and functional concepts might seem unfamiliar. If you fit into this category, I recommend keeping the Elixir documentation handy to reference some language details that this book omits.

If you have no machine learning or Elixir experience, then learning both Elixir and machine learning at the same time will be a difficult task. Therefore, I recommend getting familiar with Elixir first and then returning to this book to learn machine learning in Elixir.

What's in This Book

In [Chapter 1, Make Machines That Learn, on page ?](#), you'll answer the question, "What is machine learning?" You'll learn the fundamental concepts of machine learning and set the course for the rest of your machine learning journey. You'll create, train, and use your first machine learning model. By the end of the chapter, you'll be able to define machine learning, discuss the

types of problems machine learning is good for, and solve a problem with machine learning in Elixir.

In [Chapter 2, Get Comfortable with Nx, on page ?](#), you'll get familiar with Nx—the numerical computing library that serves as the backbone of the Elixir machine learning ecosystem. You'll learn about the fundamental data structure in numerical computing, the tensor, and a bit about what makes Nx so important for machine learning. By the end of this chapter, you'll know how to create and manipulate tensors, how to encode real-world data with Nx, and how to accelerate routines with numerical definitions.

In [Chapter 3, Harness the Power of Math, on page ?](#), you'll channel your inner mathematician and dive into machine learning math. You'll learn the three areas of mathematics that underpin most of the concepts in this book: linear algebra, probability, and vector calculus. By the end of this chapter, you'll understand the power of mathematics and how it relates to machine learning.

In [Chapter 4, Optimize Everything, on page ?](#), you'll start navigating an ocean of loss curves with optimization routines. You'll learn what a loss function is and how optimization and machine learning are related and different. By the end of this chapter, you'll be able to formulate machine learning problems as optimization problems. You'll also be able to implement some loss functions and optimization algorithms in Nx.

In [Chapter 5, Traditional Machine Learning, on page ?](#), you'll graduate from foundations to real machine learning with Scholar. You'll learn about various traditional machine learning algorithms, such as linear and logistic regression. By the end of this chapter, you'll be able to use Scholar to create and train traditional machine learning algorithms.

In [Chapter 6, Go Deep with Axon, on page ?](#), you'll start working with Axon to create and train neural networks in Elixir. You'll learn what deep learning is and why it's so powerful. You'll create neural networks in both Nx and Axon, and you'll get a better understanding of what a neural network really is. By the end of this chapter, you'll know how to create and train basic neural networks with Axon. You'll also be able to distinguish and discuss the tradeoffs between traditional and deep learning methods.

In [Chapter 7, Learn to See, on page ?](#), you'll implement computer vision models with Axon and Elixir. You'll learn about convolutional neural networks and how they offer an improvement over traditional neural networks when working with image data. By the end of this chapter, you'll be able to create and train a convolutional neural network in Axon. You'll also be able to

implement more complex machine learning training pipelines with Elixir's standard libraries.

In *the (as yet) unwritten Chapter 8, Stop Reinventing the Wheel*, you'll make use of pre-trained models to improve performance and save time and money. You'll learn about transfer learning and fine-tuning. By the end of this chapter, you'll be able to implement transfer learning in Axon. You'll also know how to convert models from Python to Elixir in order to take advantage of Python's vast machine learning ecosystem.

In *the (as yet) unwritten Chapter 9, Understand Text*, you'll teach your programs to read with recurrent neural networks, and you'll learn how to work with text in Elixir. By the end of this chapter, you'll be able to create and train recurrent neural networks in Axon, define and use various text pre-processing strategies with Elixir and Nx, and discuss why recurrent neural networks perform better than traditional models when working with sequential data like text.

In *the (as yet) unwritten Chapter 10, Forecast the Future*, you'll break out your crystal ball and attempt to peer into the future with recurrent neural networks. You'll learn about timeseries forecasting and some of the challenges of predicting the future with machine learning. By the end of this chapter, you'll be able to create and train a recurrent neural network to perform single step timeseries predictions.

In *the (as yet) unwritten Chapter 11, Model Everything with Transformers*, you'll unlock the power behind models like GPT-3 and Stable Diffusion with transformers. You'll learn what transformers are and why they perform so well on almost all types of data. By the end of this chapter, you'll be able to use pre-trained transformer models with Bumblebee. You'll also know how to fine-tune a pre-trained transformer with Bumblebee and Axon.

In *the (as yet) unwritten Chapter 12, Learn without Supervision*, you'll go off the rails and create models that learn without supervision. You'll learn about unsupervised learning and generative modeling. By the end of this chapter, you'll be able to create and train autoencoders, variational autoencoders, and generative adversarial networks.

In *the (as yet) unwritten Chapter 13, Put Machine Learning into Practice*, you'll use everything you learned throughout this book, and you'll put a real machine learning model into a production application. You'll learn what to consider when operationalizing machine learning models, and you'll deploy a model in a Phoenix application. By the end of this chapter, you'll know how to use Nx.Serving to integrate your trained models into real applications.

In *the (as yet) unwritten Chapter 14, That's a Wrap*, you'll conclude your journey with this book. You'll get a recipe for training models in practice, and you'll get resources for staying up to date with the latest trends in machine learning. You'll also learn a bit about emerging trends in the field, as well as some topics this book did not cover, such as adversarial attacks. By the end of this chapter, you'll be ready to build intelligent applications with machine learning in Elixir.

How to Use This Book

The code examples in each chapter of this book are isolated and designed to be run as standalone programs. For the best experience, read the chapters in successive order. While you can skip around, the concepts in each chapter are designed to build off one another, so you might be left confused at certain points if you skip around too much.

If you have some machine learning experience and want to dive right in, you can safely skip [Chapter 1, Make Machines That Learn, on page ?](#), [Chapter 3, Harness the Power of Math, on page ?](#), and [Chapter 4, Optimize Everything, on page ?](#). I still recommend reading [Chapter 2, Get Comfortable with Nx, on page ?](#) to get a better understanding of Nx and then reading the remaining chapters in order.

I also recommend you follow along with the examples in a Livebook, as most of the chapters will assume you're running the code in a Livebook installation.

Finally, don't worry about having access to CPUs or GPUs. A modern, commercial laptop should be powerful enough to run all the examples in this book.