Make Your Own Neural Network with Ruby

Arto Bendiken

Agenda

1. Motivation
2. Demo
3. Theory
4. Tech
5. Code
6. Study
7. Bibliography
8. Q & A

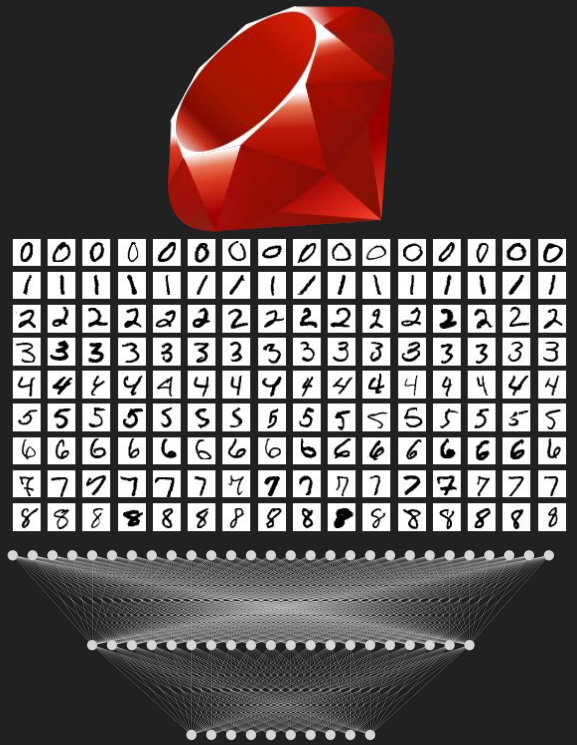

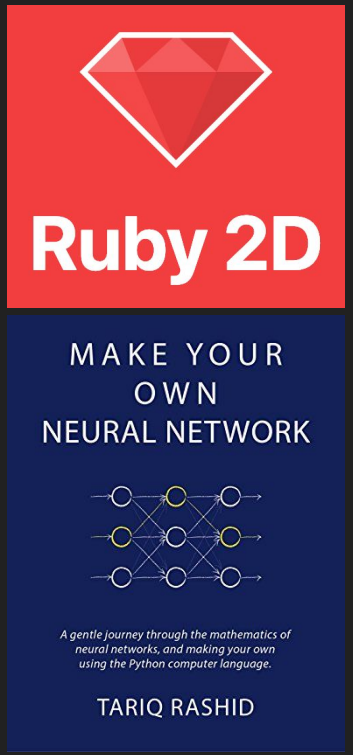

Motivation

“People worry that computers will get too smart and take over the world, but **the real problem is that they're too stupid and they've already taken over the world**.”

— Pedro Domingos, author of *The Master Algorithm*

“Amid all this activity, a picture of our AI future is coming into view, and it is not the HAL 9000—a discrete machine animated by a charismatic (yet potentially homicidal) humanlike consciousness—or a Singularitan rapture of superintelligence.

“The AI on the horizon looks more like [AWS]—**cheap, reliable, industrial-grade digital smartness running behind everything, and almost invisible except when it blinks off**. This common utility will serve you as much IQ as you want but no more than you need.”

— Kevin Kelly, co-founder of *Wired Magazine*,
in *The Three Breakthroughs That Have Finally Unleashed AI on the World* (2004)

“Like all utilities, AI will be supremely boring, even as it transforms the Internet, the global economy, and civilization. It will enliven inert objects, much as electricity did more than a century ago. **Everything that we formerly electrified we will now cognitize.**

“There is almost nothing we can think of that cannot be made new, different, or interesting by infusing it with some extra IQ. In fact, the business plans of the next 10,000 startups are easy to forecast: **Take X and add AI**.”

— Kevin Kelly, co-founder of *Wired Magazine*,
in *The Three Breakthroughs That Have Finally Unleashed AI on the World* (2004)

“**Deep Learning is a superpower.** With it you can make a computer see, synthesize novel art, translate languages, render a medical diagnosis, or build pieces of a car that can drive itself. If that isn’t a superpower, I don’t know what is.”

— Andrew Ng, co-founder of Google Brain & Coursera

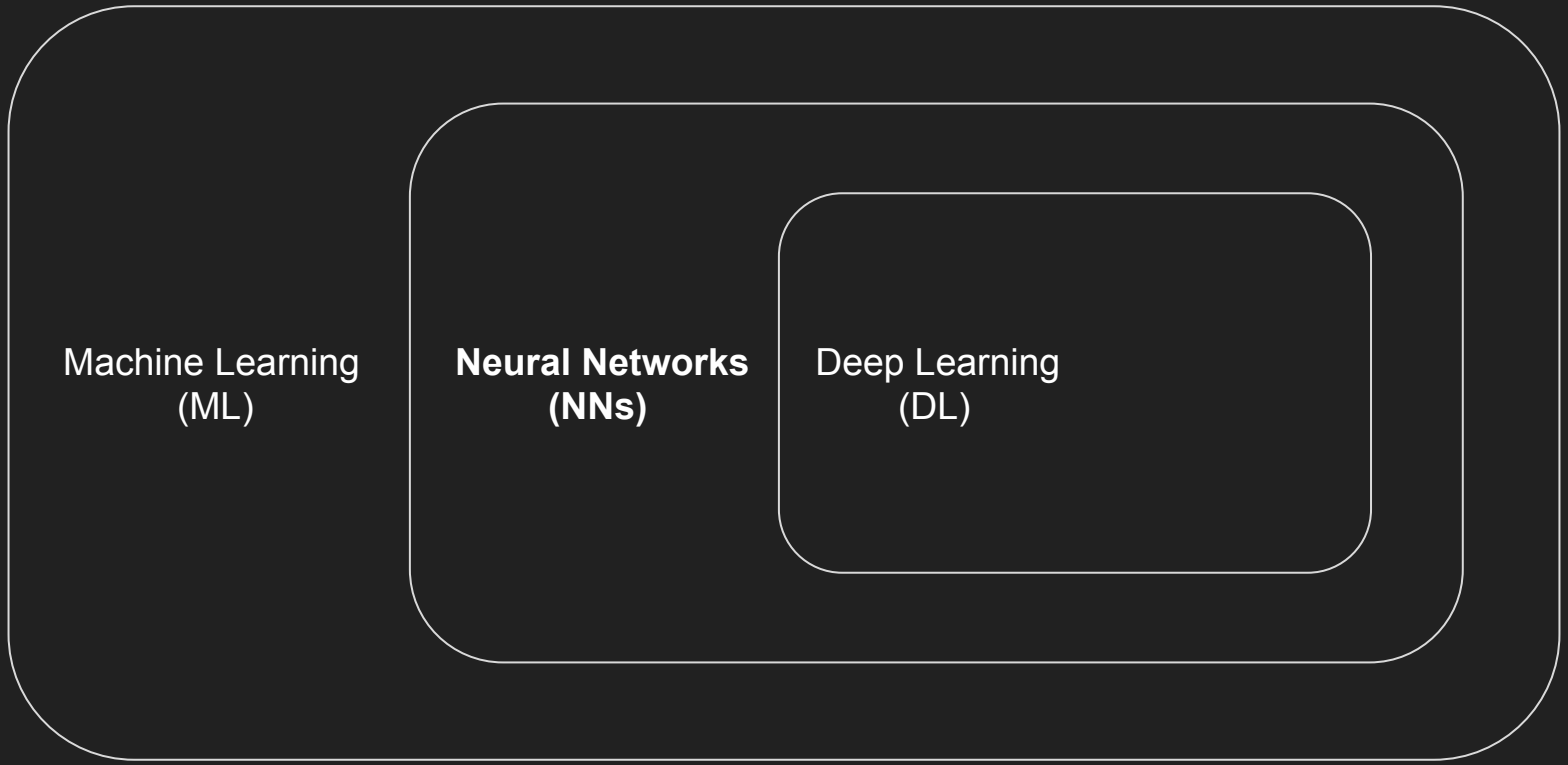
Machine Learning
(ML)
Neural Networks
(NNs)
Deep Learning
(DL)

Demo

The MNIST Dataset

- Large database of handwritten digits, widely used in machine learning
- Grayscale images
- **28×28 pixel resolution**
- **60,000 training images**
- **10,000 testing images**
- Convolutional neural networks have achieved 99.79% accuracy on this dataset (Ukraine, 2016)

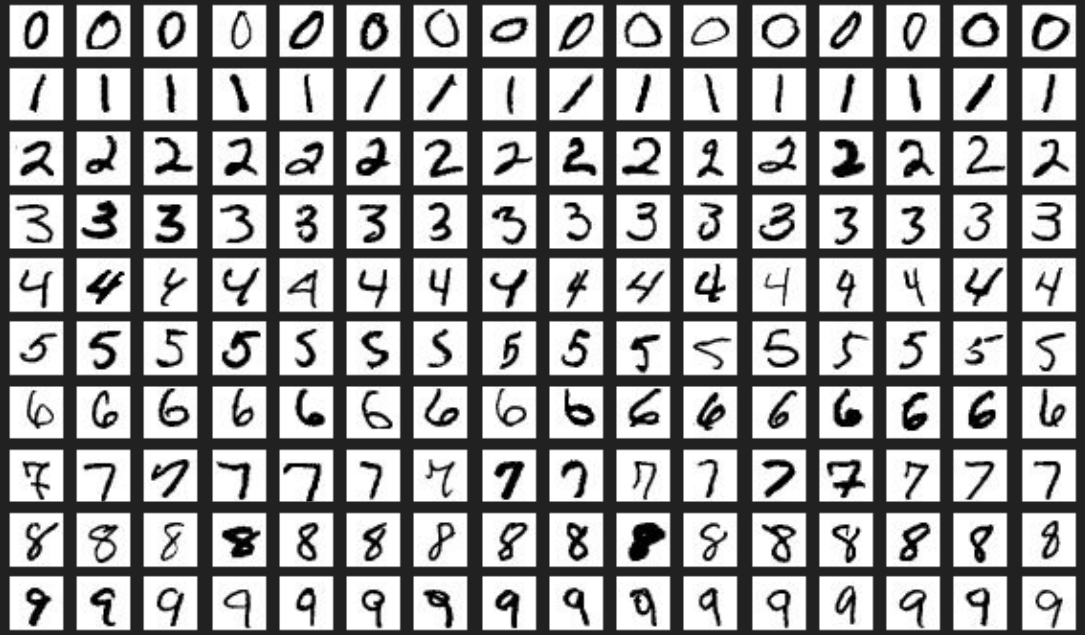

Examine the MNIST Dataset

```
$ ./mnist_render.rb
```

Press <UP> or <DOWN> to navigate records, and <SPACE> to view the assigned label

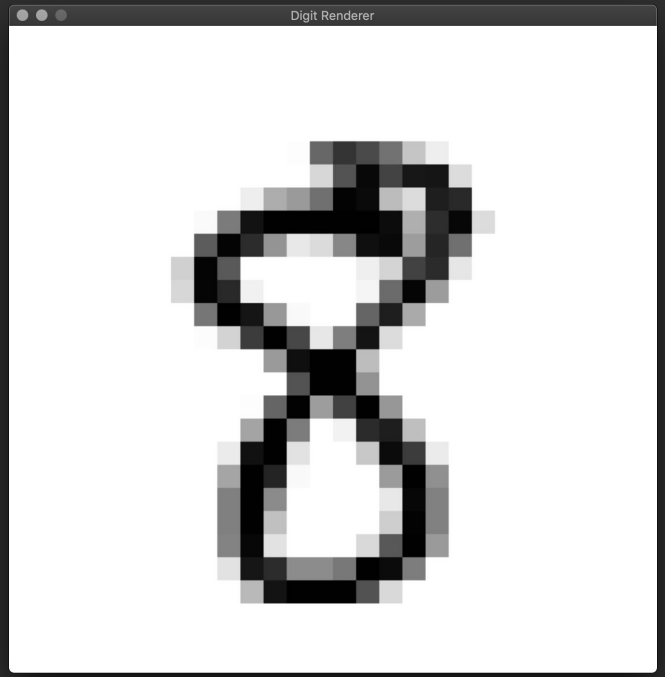

Train the Neural Network

```
$ ./mnist_train.rb
```

It will take a couple of minutes when using the CPU (GPU is faster)

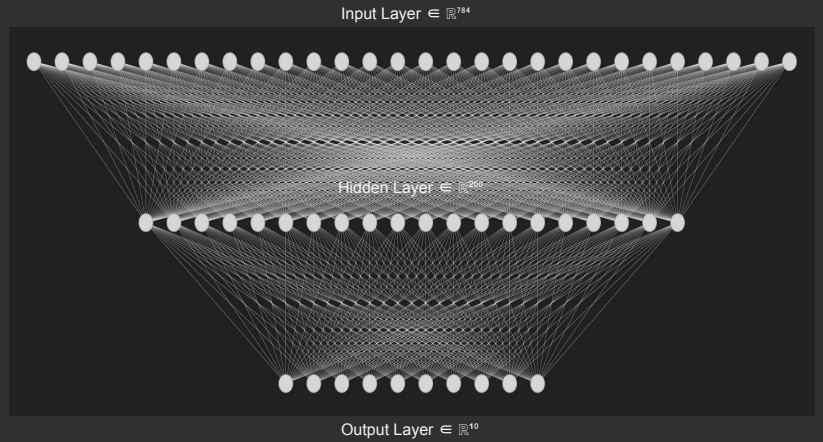

Practice Drawing Digits

```
$ ./mnist_draw.rb
```

Press down your (left) mouse button to draw your digit

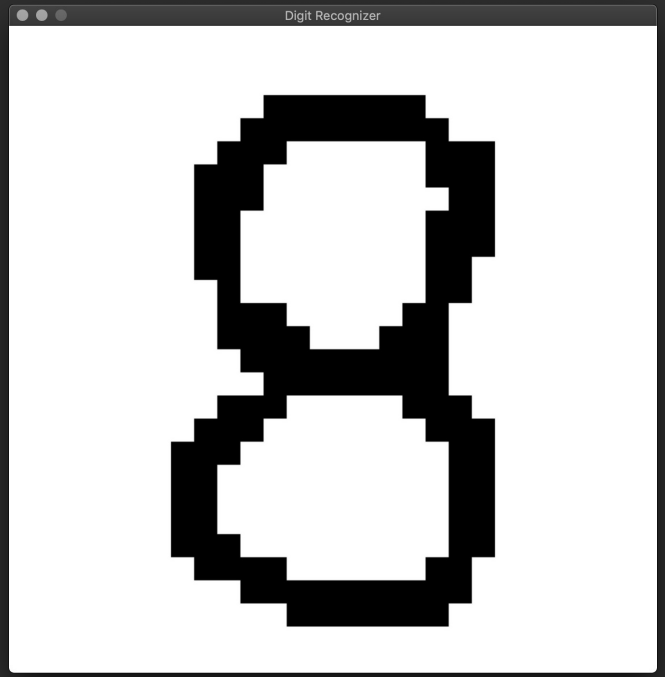

Run the Neural Network

```
$ ./mnist_draw.rb
```

Press <SPACE> to attempt to recognize the hand-drawn digit

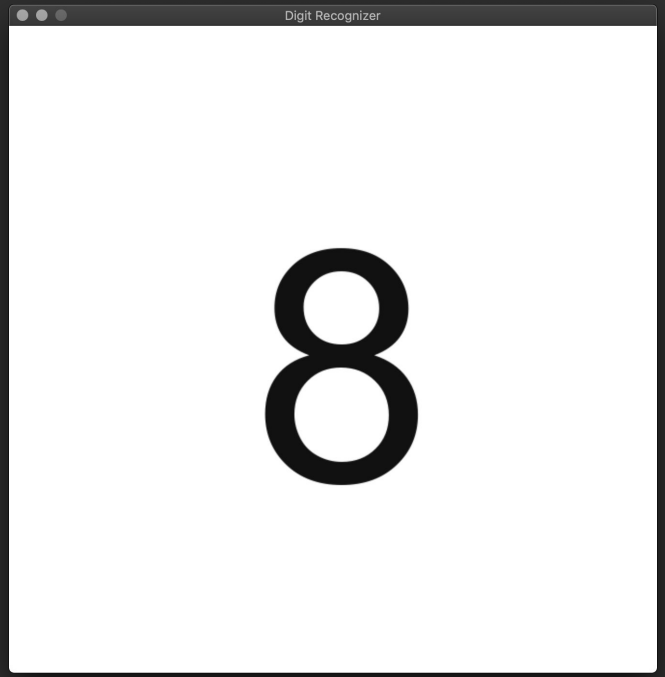

Theory

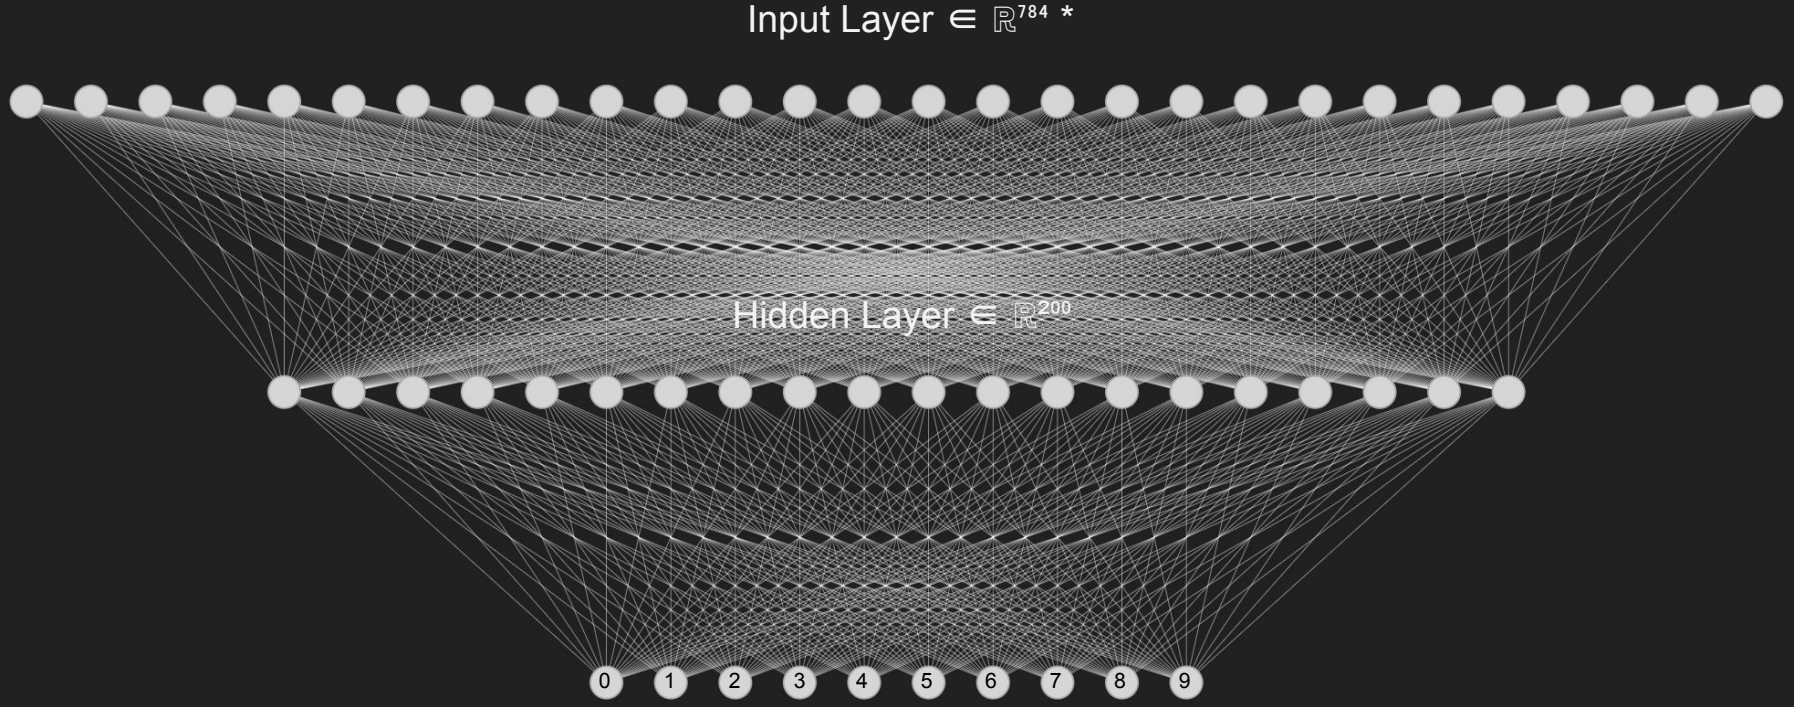
Input Layer ∈ ℝ^784 *
Hidden Layer ∈ ℝ^200
0
1
2
3
4
5
6
7
8
9
Output Layer ∈ ℝ^10

(*) 28 × 28 = 784

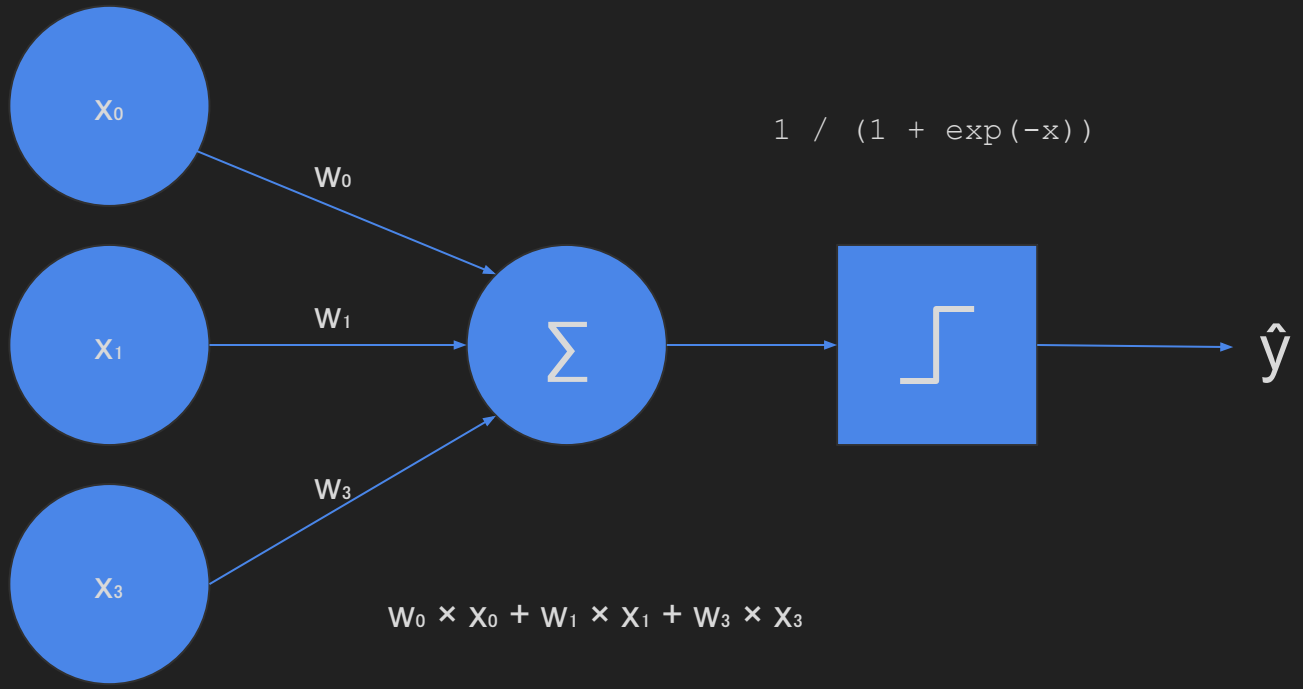
x0
x1
x3
w0
w1
w3
∑
ŷ
1 / (1 + exp(-x))
w0 × x0 + w1 × x1 + w3 × x3

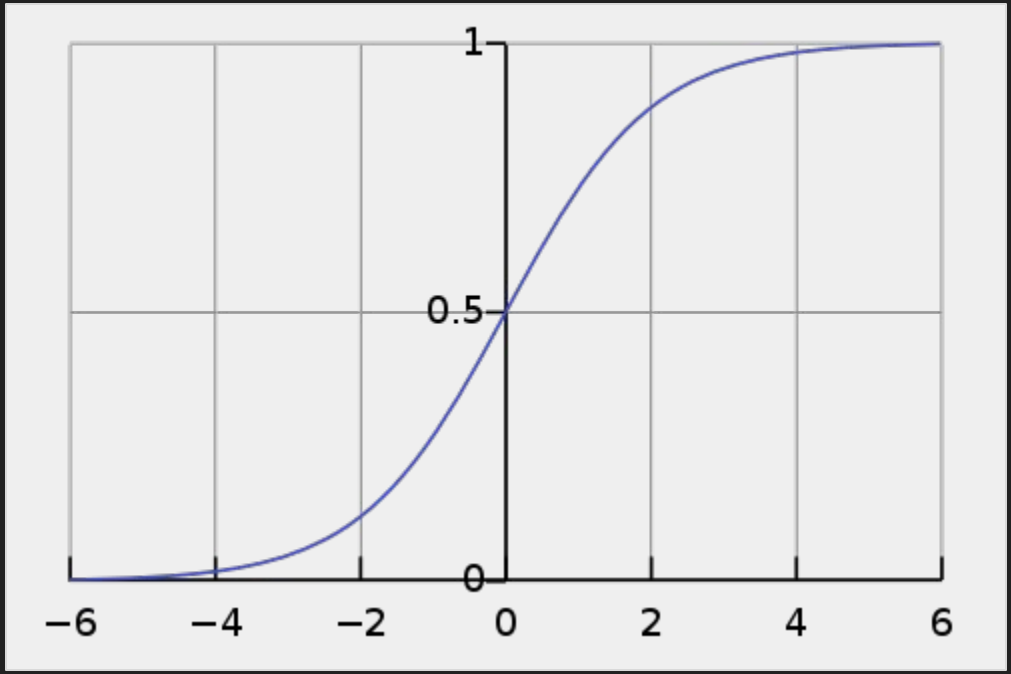

The logistic sigmoid activation function

Tech

Challenges Encountered & Overcome

- Ruby isn't *obviously* a suitable programming language for this task
  - But it's feasible nonetheless, if the heavy lifting can be outsourced to BLAS libraries
- Ruby doesn't have as mature numeric computing support as Python
  - A fragmented ecosystem with accumulated sedimentary layers
  - The new Numo project for Ruby is promising and aims to cover the same ground as NumPy
- The UX and the performance isn't quite comparable as yet
  - NumPy "just works" after installation, with the best possible performance
  - Numo more than likely will need some manual configuration with OpenBLAS, MKL, etc
  - On my laptop, Numo with OpenBLAS doesn't run multi-threaded (hence trains the NN some 4x slower than NumPy does) ...this will need more troubleshooting
  - If you have an NVIDIA graphics card, Cumo should help speed you up
- Ruby 2D was a superb find for quick & easy GUI visualization

Code

Study

deeplearning.ai

Deep Learning Specialization

16 weeks of study, 3-6 hours per week

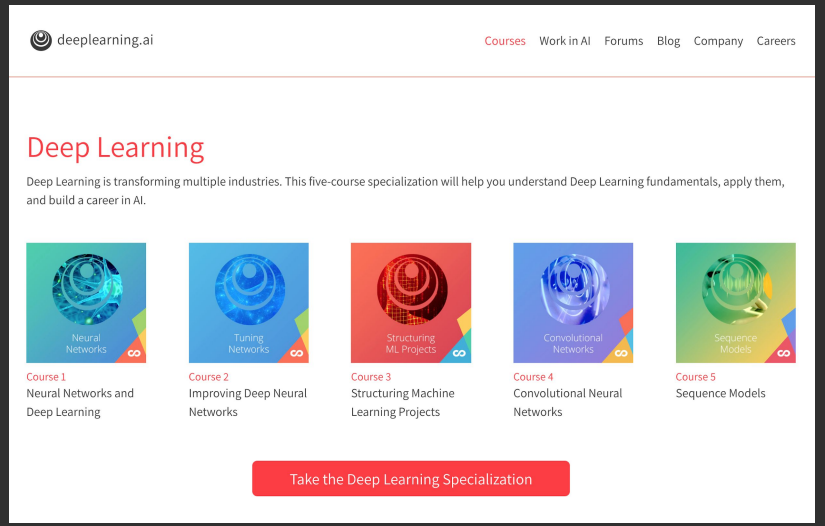

fast.ai

Practical Deep Learning for Coders

7 weeks of study, 10 hours per week

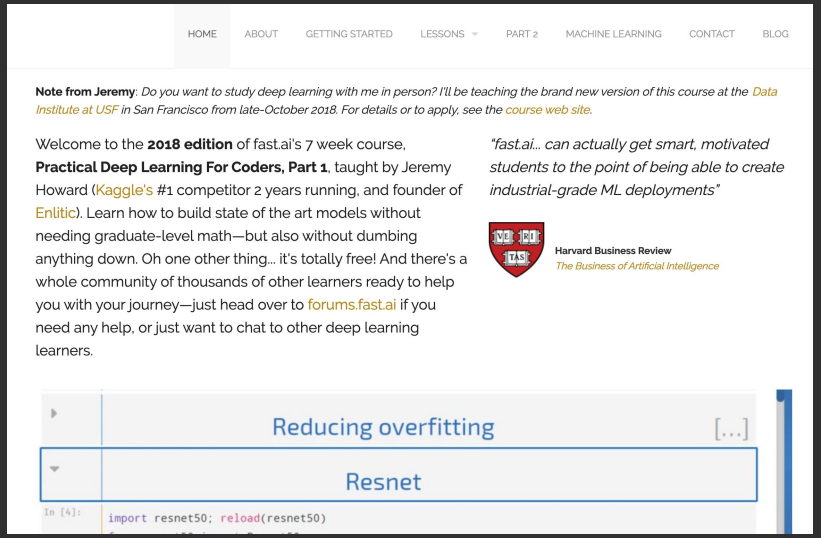

Bibliography

Make Your Own Neural Network by Tariq Rashid

- The **single best quick & short introduction** to the principles and mathematics underlying neural networks
- Can be read in one sitting in a couple of hours
- Example code in Python available at GitHub in the form of a Jupyter Notebook

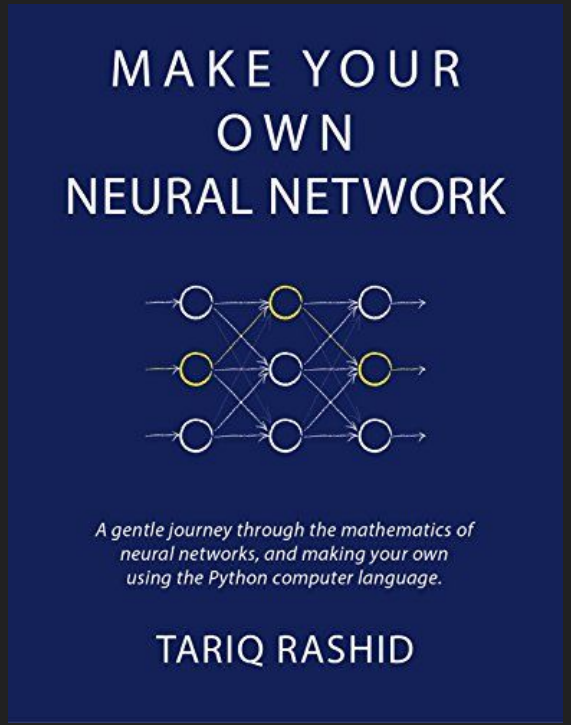

The Master Algorithm by Pedro Domingos

- Outlines the five tribes of machine learning:
  - the symbolists (inductive reasoning),
  - **the connectionists (backpropagation)**,
  - the evolutionaries (genetic programming),
  - the Bayesians (Bayesian inference), and
  - the analogizers (support vector machines)

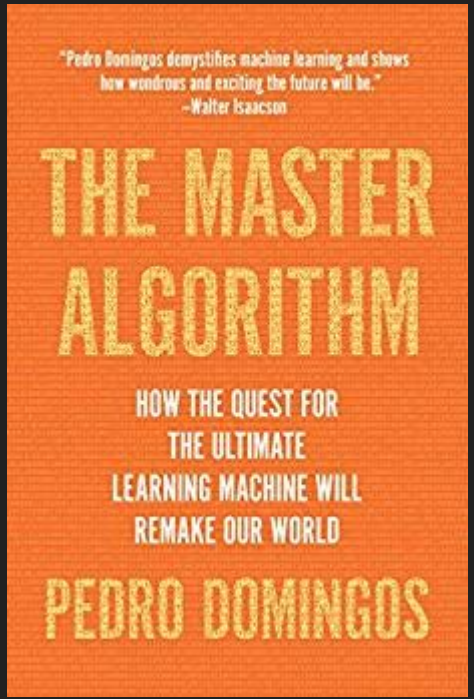

Дякую!

Find me at:
https://ar.to