

Centaur vs. Minotaur

Or: Where do you stand
in the human/AI equation?

I'm Joey deVilla...


- Senior Developer Advocate at Okta (the Auth0 product)
- Editor of *Global Nerdy*
(home of the Tampa Bay Tech Events list)
- Accordion player
- Met Anitra while visiting from Toronto at a Tampa Bay tech event (BarCamp 2011)





**Why are we trying to
make things to
replace ourselves,
anyway?**

Bicycle for the mind



More energy

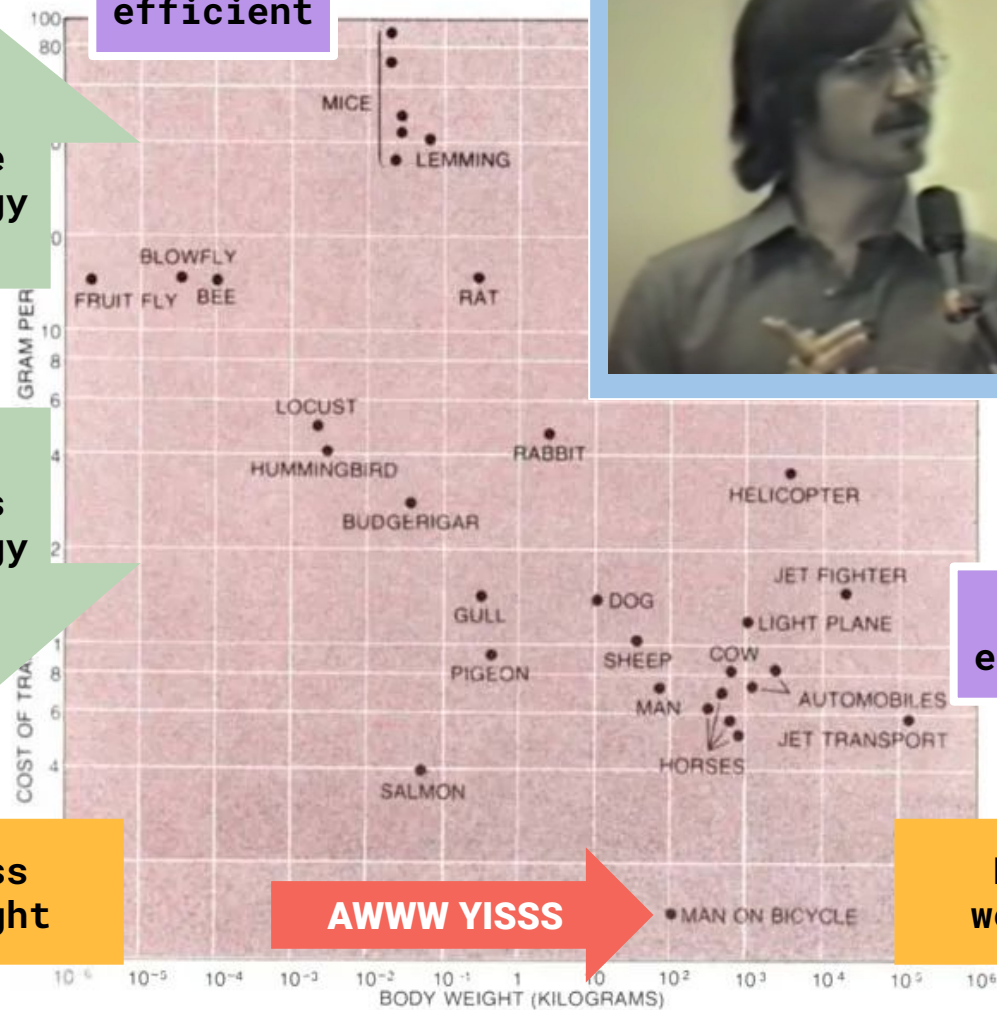
Less energy

Less weight

Less efficient


More efficient

More weight




AWWW YISSS

MAN ON A BICYCLE ranks first in efficiency among traveling animals and machines in terms of energy consumed in moving a certain distance as a function of body weight. The rate of energy consumption for a bicyclist (about .15 calorie per gram per kilometer) is approximately a fifth of that for an unaided walking man (about .75 calorie per gram per kilometer). With the exception of the black point representing the bicyclist (lower right), this graph is based on data originally compiled by Vance A. Tucker of Duke University.



1


Could you explain
this thing with the
centaurs and
minotaurs?




This is the
classical depiction
of a
CENTAUR.

“Κένταυρος” (kentauros) in ancient Greek, this is a mythological creature with a human head, arms, and torso, and the body of a horse.

Nobody Likes Reverse Centaur





Garry Kasparov

coined the term
“centaur chess”

“Teams today win by combining **average skill players**, **mid-range technology** and **strong interfaces**. By optimizing for an efficient connection between human and machine, teams can better leverage the characteristics of either.” – <https://centaur-labs.io/centaurs>



This is the classical depiction of a **MINOTAUR.**



“Μινώταυρος” (minotauros) in ancient Greek, this is a mythological creature with a bull’s head and tail, and human body.



“Siri, yes Siri!”

“as the impact of AI accelerates the tempo of battle and reduces effective decision-making time for humans, **militaries may have little alternative but to outsource many decisions to AI.**”

– *Minotaurs, Not Centaurs: The Future of Manned-Unmanned Teaming*
US Army War College Quarterly, vol. 3, no. 1, March 3, 2023



The kind of “taur” depends on who’s in charge

Centaur

Human head,
non-human body:

The **human** has agency – that is, the human part makes the decisions, and the non-human part has to abide by them.

Minotaur

Non-human head,
human body:


The **non-human** has agency – that is, the non-human part makes the decisions, and the human part has to abide by them.

CENTAUR

when possible;

MINOTAUR

when necessary?



JAKE-CLARK.TUMBLR




Chiron

First among centaurs, and seen as the most wise and just of them all.

Pictured to the left, teaching Achilles how to play the lute.

Taught Aesclepius, the Greek god of medicine, which makes Chiron the **Centaur for Disease Control!**



2 How do I get in on some of this centaur action?

Let's start with ETHICS.



Few AI presentations, lectures, or courses start with ethics.
Many people in the field even actively resist discussing it!



Let's talk
about the
**Boston
Housing
Prices**
dataset.

**It's for a
Harvard paper
about how air
quality...**



**...affects
housing prices.**



Air quality in Boston in 1973 vs. 2016



Source: United States Environmental Protection Agency (EPA)
<https://www.instagram.com/p/BDDxI7HApUf/>



			D	E	F	G
			ict	SATV	SATM	GPA
			38	368	253	3.52
			5	670	496	1.11
4	55509	21 F	54	639	439	2.68
5	36489	19 M	49	368	465	3.11
6	36387	21 F	36	620	306	2.16
7	95507	20 F	13	512	593	2.83
8	16360	20 M	52	621	377	2.79
9	12838	18 F	44	571	544	2.13
10	73450	20 F	59	647	746	2.08
11	26869	18 F	28	337	371	2.28
12	48552	22 M	63	260	498	3.24
13	23416	19 M	51	476	294	2.31
14	42635	19 F	35	677	241	3.19
15	67448	19 F	55	335	533	1.81
16	34689	21 F	42	585	708	1.80
17	32763	22 F	20	556	787	1.18

It used to be part of scikit-learn, an important Python code library for machine learning.

(If you plan to get into AI, you'll use scikit-learn, a.k.a. "sklearn" sooner or later.)

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2



Per capita crime rate by town.

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2

Per capita crime rate by town.



	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2

Proportion of population that is lower status = $0.5 * (\text{proportion of adults without some high school education and proportion of male workers classified as laborers})$. **The logarithmic specification implies that socioeconomic status distinctions mean more in the upper brackets of society than in the lower classes.**

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2



Harvard SUS


Proportion of population that is lower status = $0.5 * (\text{proportion of adults without some high school education and proportion of male workers classified as laborers})$. **The logarithmic specification implies that socioeconomic status distinctions mean more in the upper brackets of society than in the lower classes.**

	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2

“Black proportion of population. At low to moderate levels of B , an increase in B should have a negative influence on housing value if Blacks are regarded as undesirable neighbors by Whites. However, market discrimination means that housing values are higher at very high levels of B . **One expects, therefore, that a parabolic relationship between proportion Black in a neighborhood and housing values.**”


	CRIM	ZN	INDUS	CHAS	NOX	RM	AGE	DIS	RAD	TAX	PTRATIO	B	LSTAT	Price
0	0.00632	18.0	2.31	0.0	0.538	6.575	65.2	4.0900	1.0	296.0	15.3	396.90	4.98	24.0
1	0.02731	0.0	7.07	0.0	0.469	6.421	78.9	4.9671	2.0	242.0	17.8	396.90	9.14	21.6
2	0.02729	0.0	7.07	0.0	0.469	7.185	61.1	4.9671	2.0	242.0	17.8	392.83	4.03	34.7
3	0.03237	0.0	2.18	0.0	0.458	6.998	45.8	6.0622	3.0	222.0	18.7	394.63	2.94	33.4
4	0.06905	0.0	2.18	0.0	0.458	7.147	54.2	6.0622	3.0	222.0	18.7	396.90	5.33	36.2

“Black proportion of population. At low to moderate levels of B , an increase in B should have a negative influence on housing value if Blacks are regarded as undesirable neighbors by Whites. However, market discrimination means that housing values are higher at very high levels of B . **One expects, therefore, that a parabolic relationship between proportion Black in a neighborhood and housing values.**”



WTF?!

$$y = (x - 0.63)^2, 0 < x < 1$$




y is the number
that ends up in
the B column

$$y = (x - 0.63)^2, 0 < x < 1$$

x is the actual proportion
of Black people in an
area
(which has nothing to do
with air pollution)...

This is a formula for a
parabola, which was the
relationship those racist
researchers were talking
about earlier.

The blue curve is the
value of B, which
mathematically expresses
a very nasty opinion.





**This was a
research
paper about
AIR
QUALITY
and HOUSE
PRICES.**

AI pioneer accused of having sex with trafficking victim on Jeffrey Epstein's island



Jeffrey Epstein with Professor Marvin Minsky Photo by Rick Friedman/Corbis via Getty Images

/ Marvin Minsky was named alongside several other prominent men

By [Russell Brandom](#)

Aug 9, 2019, 4:14 PM EDT | [0 Comments](#) / [0 New](#)





Harvard
biologist
George Church

“There was just a lot of
nerd tunnel vision.”

Remember
that
phrase!



Marc Andreessen ✓

@pmarca

“AI regulation” = “AI ethics” = “AI safety” = “AI censorship”. They're the same thing.

3:39 PM · Dec 3, 2022

Source: Twitter / X

<https://twitter.com/pmarca/status/1599141199805550593>



Marc Andreessen ✓

@pmarca

“AI regulation” = “AI ethics” = “AI safety” = “AI censorship”. They're the same thing.

3:39 PM · Dec 3, 2022



Source: Twitter / X

<https://twitter.com/pmarca/status/1599141199805550593>

**More than ever,
tech in general
(and AI in particular)
needs people who
give a damn about
ETHICS.**

Some search terms for ethics and AI

Take a picture of this slide and do a search on these title terms!

On the dangers of stochastic parrots

Mo' LLM inputs,
Mo' problems,

Weapons of math destruction

How big data increases inequality and threatens democracy.

Timnit Gebru

The AI ethicist whom Google fired for warning about the effects that AIs have on marginalized people.

TESCREAL

A philosophy that Elon Musk, Sam Bankman-Fried and many "AI-holes" are into.
* DO THE OPPOSITE! *

Ethics 101, by Brian Boone

And yes, there's an audiobook.




There will be
MATH.
Make peace
with that.

In theory, you could build AI applications without doing any math by making API calls to an AI framework, library, or service...

...but where's the innovation (or fun) in that?

Algebra


- Exponents and logarithms
- Radicals
- Factorials
- Summation
- Scientific notation



Algebra is the branch of math where you use variables and manipulate them within formulas. If you're doing any kind of programming beyond "Hello, world!", you're already doing algebra to some extent.



Linear algebra

- Scalars and vectors
- Matrices
- Tensors
- Eigenvectors and eigenvalues
- Singular value decomposition
- Principal component analysis




Linear algebra takes algebra and applies it to equations for straight lines, where none of the variables are raised to a power.


Widths and Lengths of Garden Bugs



length



Caterpillars are much longer than they are wide.




Ladybirds are much wider than they are long.



width



This line divides the graph into two zones - caterpillar and ladybird.



If it's on this side of the green line, it's a caterpillar.

If it's on this side of the green line, it's a ladybird.



Calculus

I have good news for you:

The calculus used in AI tends to be *differential* calculus – **that's the easy one!** Trust me, there's a harder calculus – *integral* calculus – and it sucks.


- Derivatives
- Vector and matrix calculus (gradient, Jacobian, Hessian, Laplacian)
- Gradient algorithms (maxima and minima, saddle points, gradient descent)

Calculus is the math of rates of change.

If you drove (and if you paid attention to the speedometer), you saw calculus in action!

When you started your trip to this event, you were a certain distance from here. That distance changed as you traveled here, because of the speed you were traveling.

From a calculus point of view, speed is the rate of change of distance – or more formally, **speed is the *derivative* of distance.**





We train neural networks by presenting them with “training data” and having them to produce an answer. Then we tell them how close or far they were to or from the correct answer. It’s sort of like “The Price is Right,” but without the “without going over” rule.



Neural networks use calculus to determine how much they should correct themselves – a little if they were close to the correct answer, more if their answer was way off.

Probability and Statistics

This is probably the easiest math, since the basic stuff seems to make intuitive sense to a lot of people, and it may be the math we unconsciously use the most in our daily lives.

- Basics: Mean, media, mode, variance, covariance, standard deviation, and all that
- Probability: Events (independent and dependent), sample spaces, conditional probability
- Common distributions: binomial, poisson, bernoulli, gaussian, exponential
- Bayes' theorem: This one's a big deal, as it's the basis of algorithmic pattern recognition and decision-making



Friend: You can't possibly guess a wheel of fortune puzzle with no letters

Me:





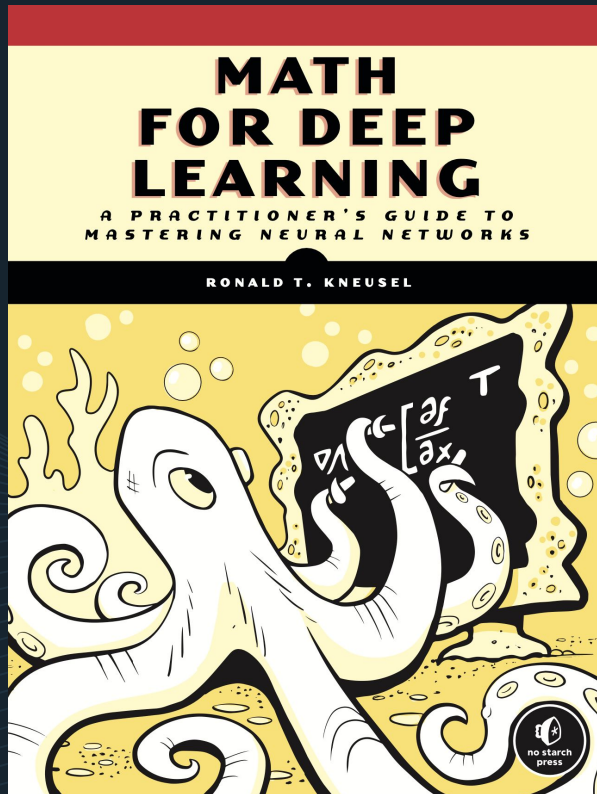
Information theory

This is a pretty new branch of math – it's not even 100 years old yet! It's the study of the measurement, storage, and transmission of information – and by information, we mean *numbers*.

- Entropy
- Comparing probability distributions for similarity
- Encoding and decoding information (not in the cryptography sense, but in the compression and noise-reduction sense)



Recommended books for getting back into “math gear”




You should READ.

Lorem ipsum dolor sit amet,
consectetur adipiscing elit. Integer a
sem ligula...

(Just checked to see if you read
this.)





If you read only one non-programming AI book, make it this one:

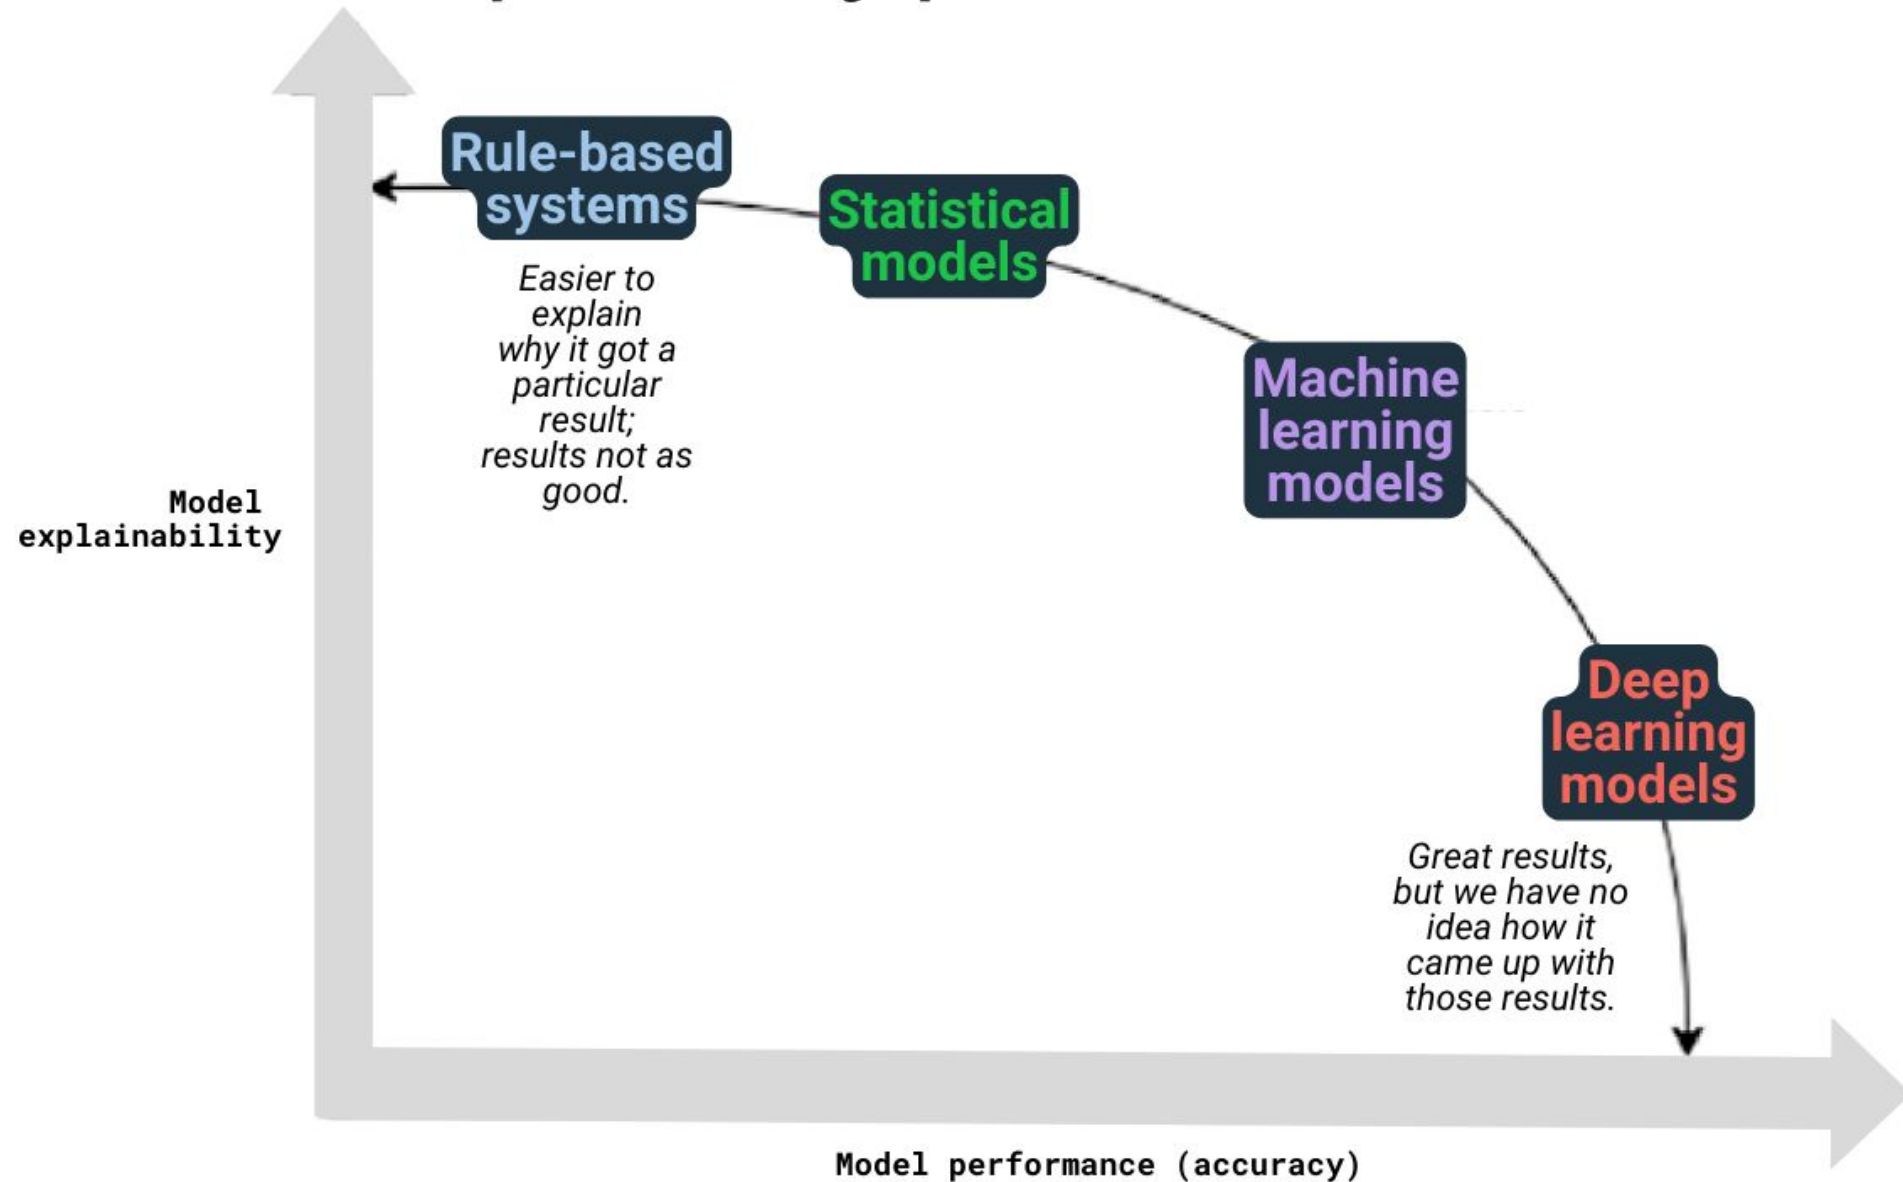


Written by Janelle Shane, author of the *AI Weirdness* blog, this book does a wonderful (and amusing) job of explaining AI, even if you're not terribly technical.

Also worth reading...



The explainability/performance tradeoff






An (admittedly simple) example of a rule-based system

```
# Python
```

```
if customer_last_played_artist == TAYLOR_SWIFT && current_season == AUTUMN:  
    likely_coffee_choice = PUMPKIN_SPICE_LATTE
```


An example of a **deep learning system**



A biological neuron


1: Dendrites receive incoming signals, which can...

- *Excite* the neuron, making it more likely to “fire”
- *Inhibit* the neuron, making it less likely to “fire”

2: The nucleus collects the exciting and inhibiting dendrite signals and forms them into a sum

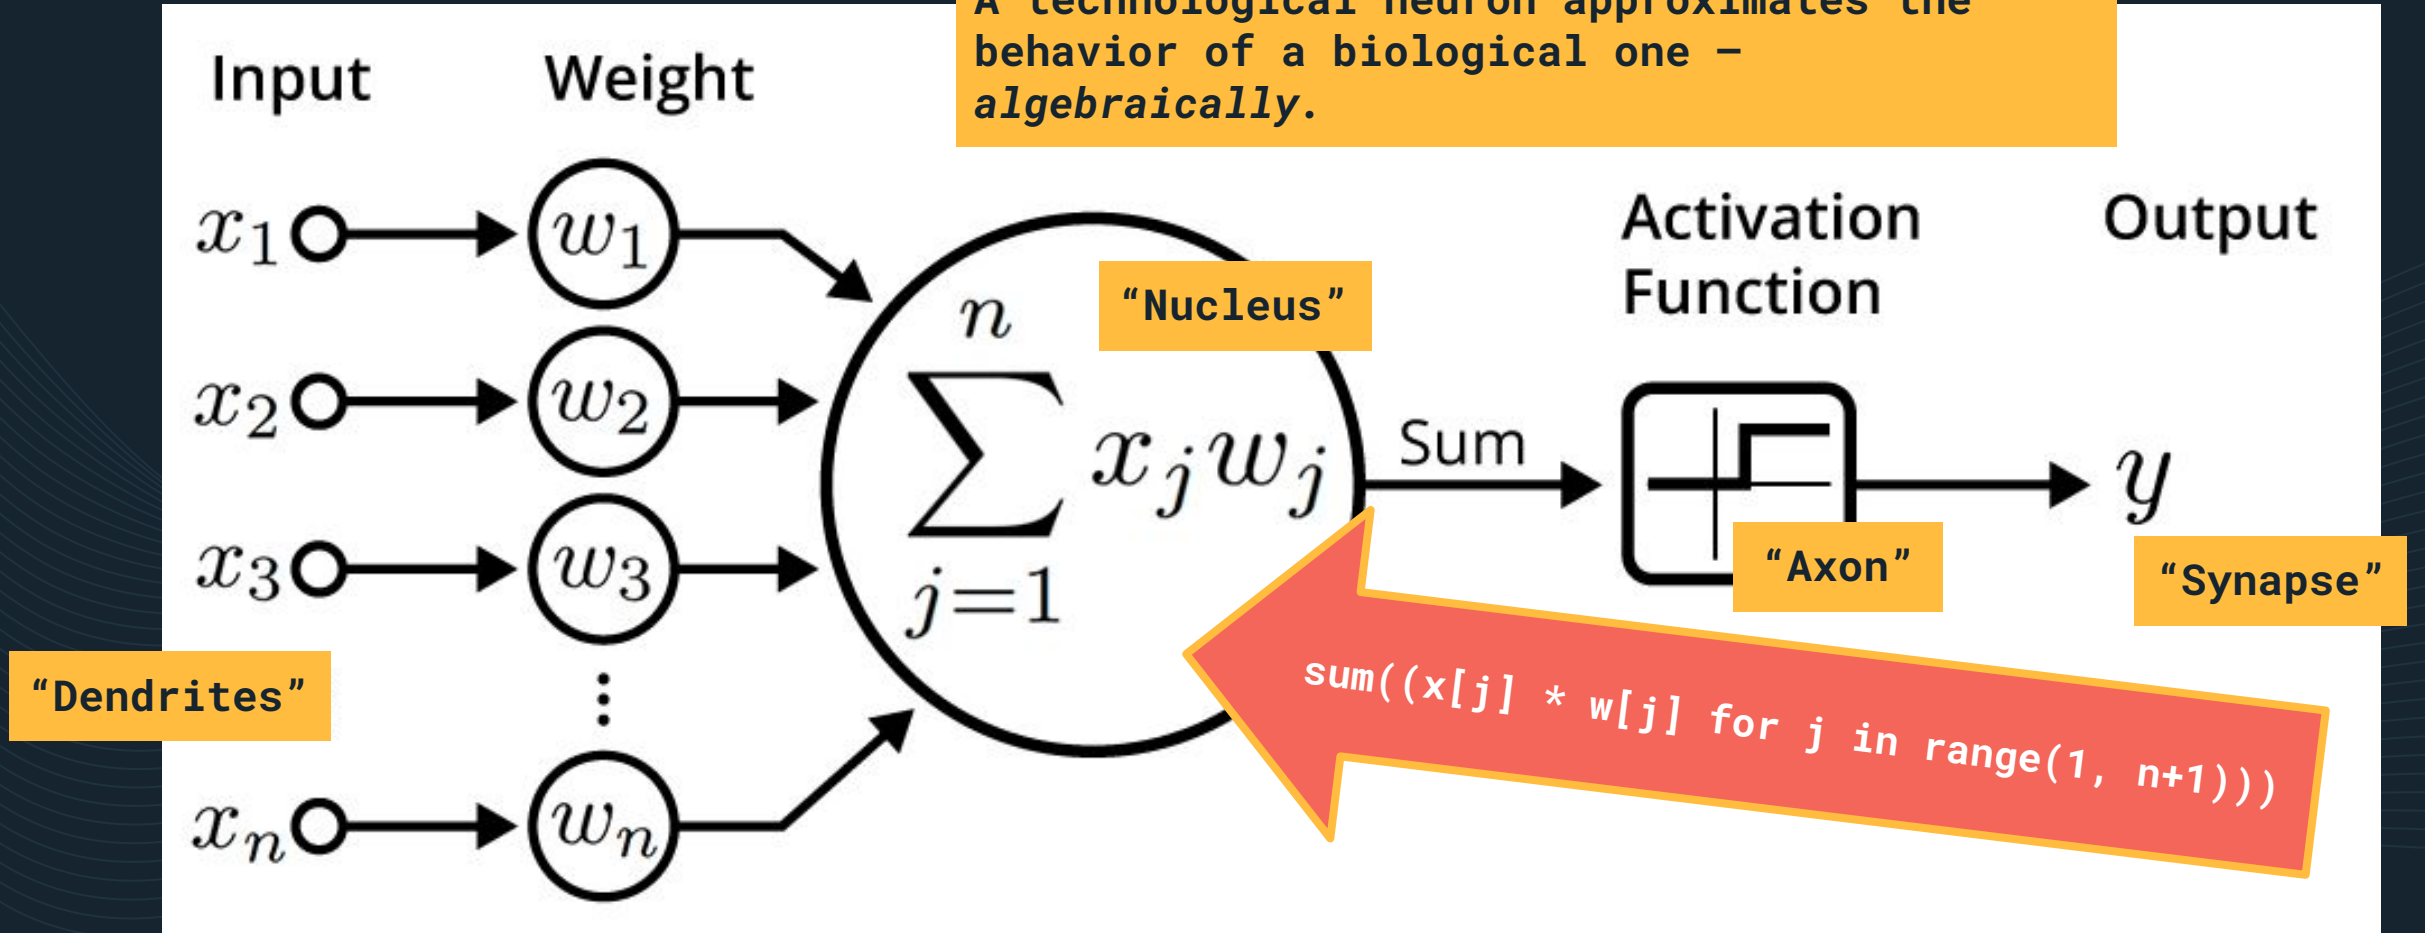
3: If the sum is big enough, a signal is sent along the axon

4: The signal is sent to the next neuron via the synapse






A technological neuron

A technological neuron approximates the behavior of a biological one – *algebraically*.



Read these books if you want Amazon's AI to put you on a watch list






Oh yes,
PROGRAMMING...




Many “hot new things” are driven by programming languages that seemed to **COME BACK FROM THE DEAD.**





Mojo 🔥




3 To wrap everything up (and get back to centaurs and minotaurs)

Every 13 years, something redefines our field.

1968

Douglas Englebart's
"Mother of all
demos"



1981

The Florida-
conceived IBM
PC is
released.



1993/1994

Mosaic, the
first
graphical
browser,
debuts.



2007-2008

First comes
the iPhone,
and then the
App Store.



2020-2022

GPT-3 is
released in
2020, followed
by ChatGPT in
2022.




Every 13 years, something redefines our field.



1968

Douglas Englebart's
"Mother of all
demos"



1981

The Florida-
conceived IBM
PC is
released.



1993/1994

Mosaic, the
first
graphical
browser,
debuts.



2007-2008

First comes
the iPhone,
and then the
App Store.




2020-2022

GPT-3 is
released in
2020, followed
by ChatGPT in
2022.



How do you stay
on the
CENTAUR
side of the
equation?



JAKE-CLARK.TUMBLR

“

Stop doing
average



•

– My mom.
(but without the swearing)

