https://cloud.google.com/products/ai/ml-comic-1

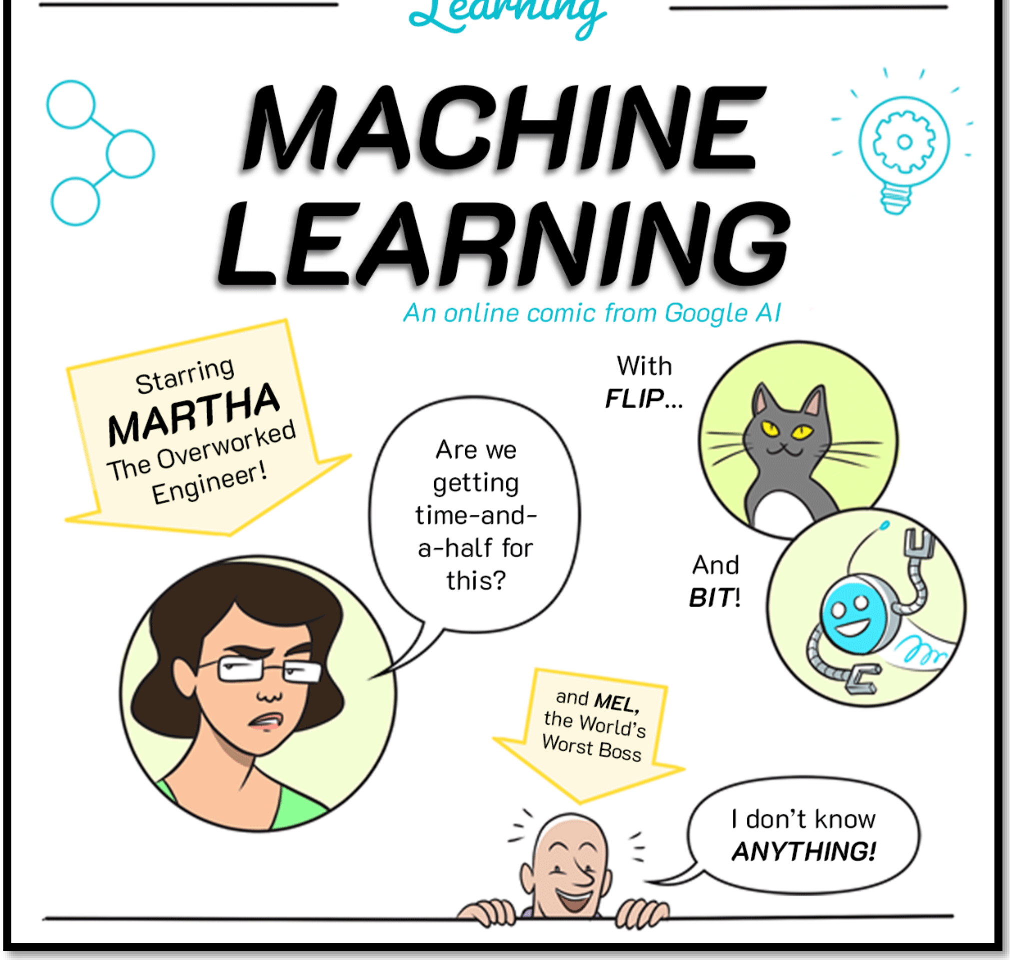

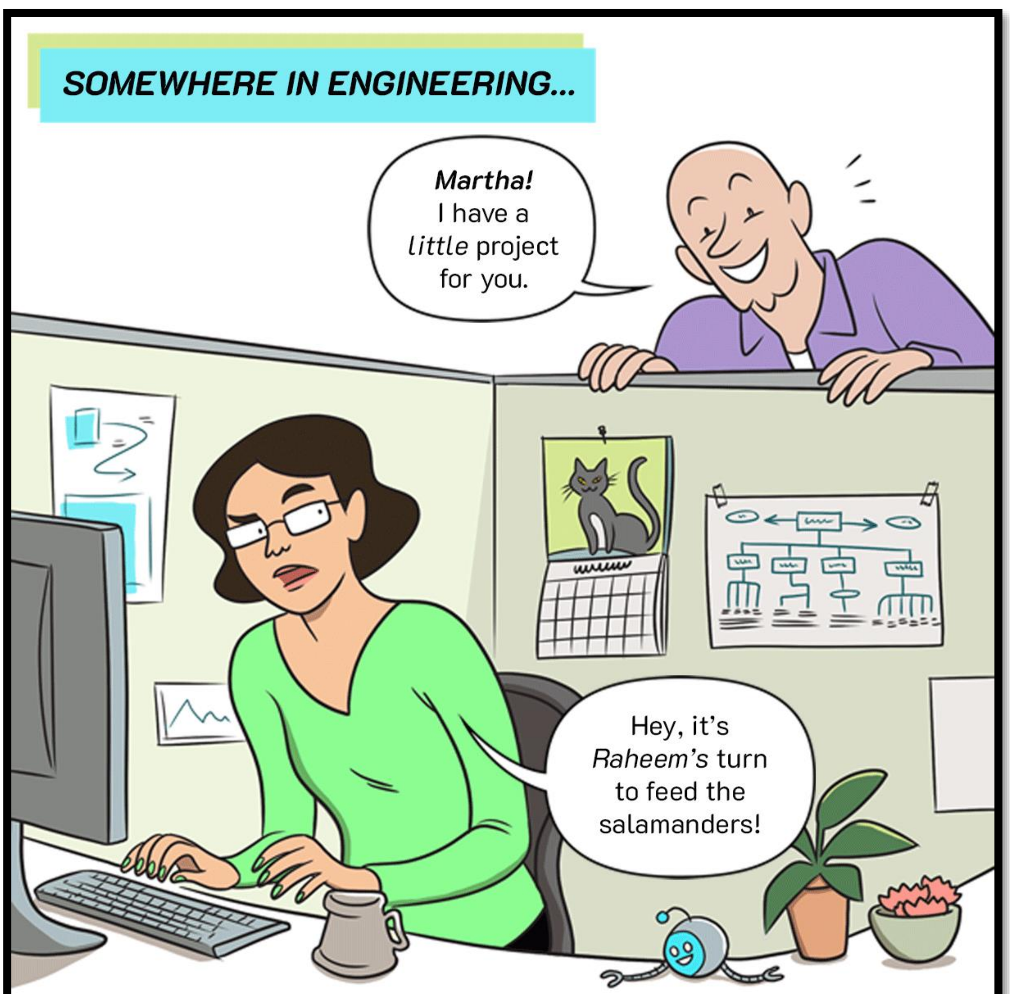


So... our VP heard about MACHINE LEARNING and wants us to "add it to everything."
Can you lead that effort?
Machine learning?
That's not exactly my area of exper—


Okay good luck gotta hop on a call kthx bye!
But –


Oh no. Is there even anyone here with machine learning experience?
Gotta be ***someone...***
Ha! At ***THIS*** place?
CONFER
CONFEREN
ART


Great... Ten more white papers, here I come...
MYSTERIOUS DOOR THAT DEFINITELY WASN'T HERE THIS MORNING
MORE ART


Wait—
what??
MYSTERIOUS DOOR THAT DEFINITELY WASN'T HERE THIS MORNING
MORE ART


Hello?
Is somebody there?

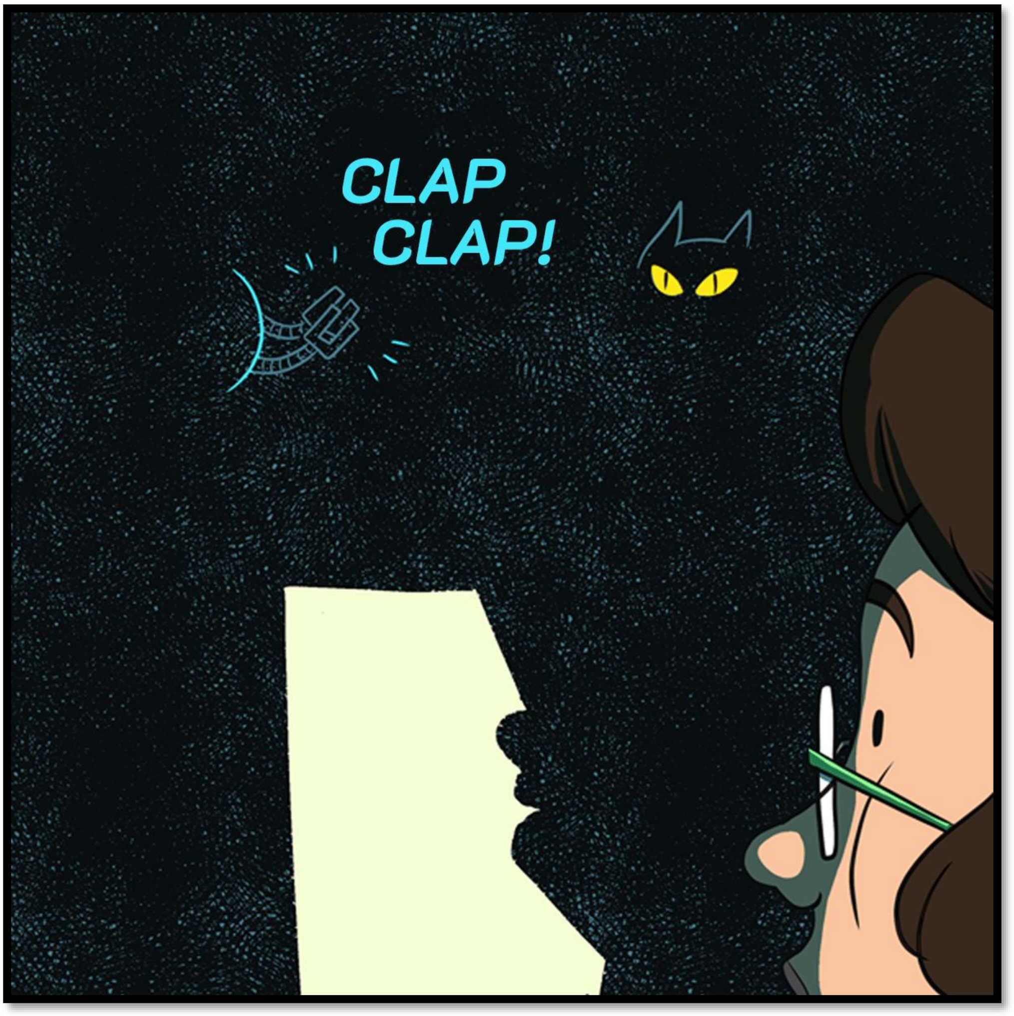
CLAP
CLAP!

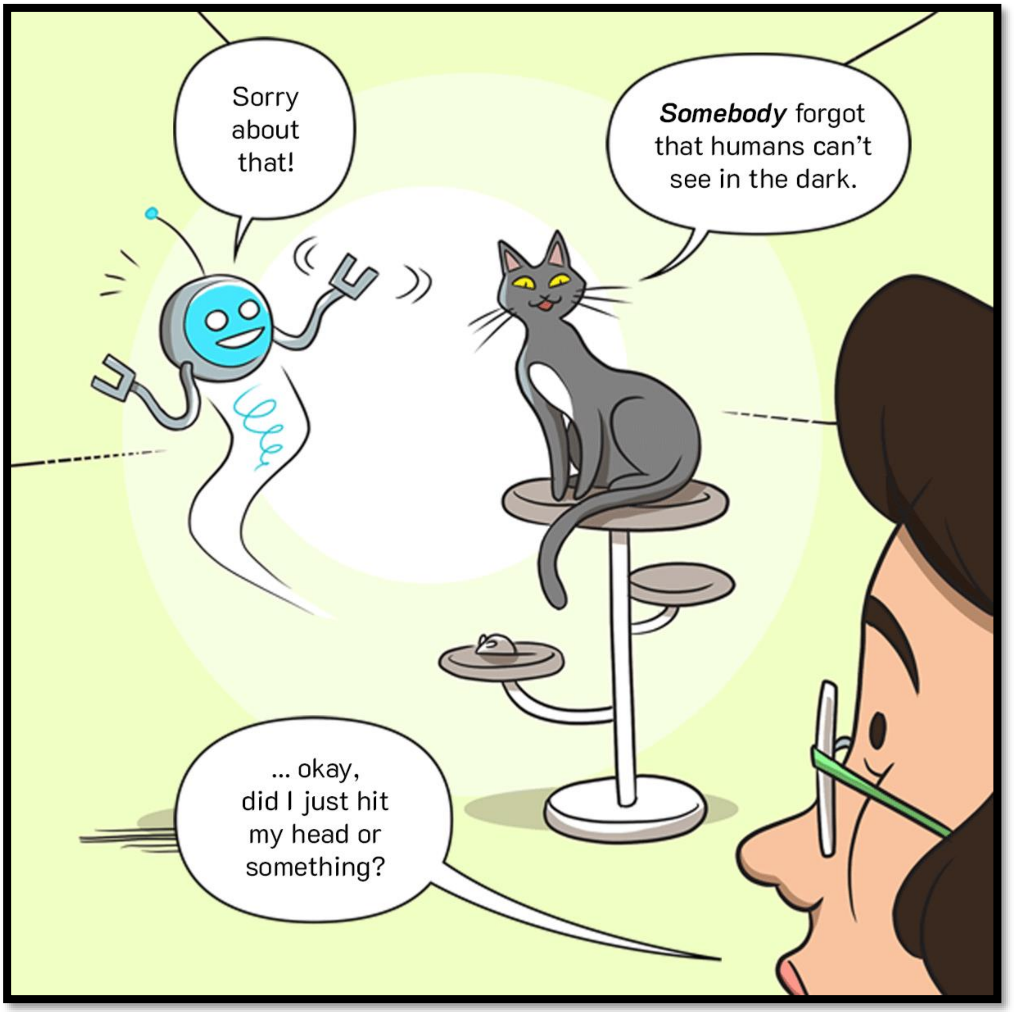
Sorry about that!
Somebody forgot that humans can't see in the dark.
... okay, did I just hit my head or something?

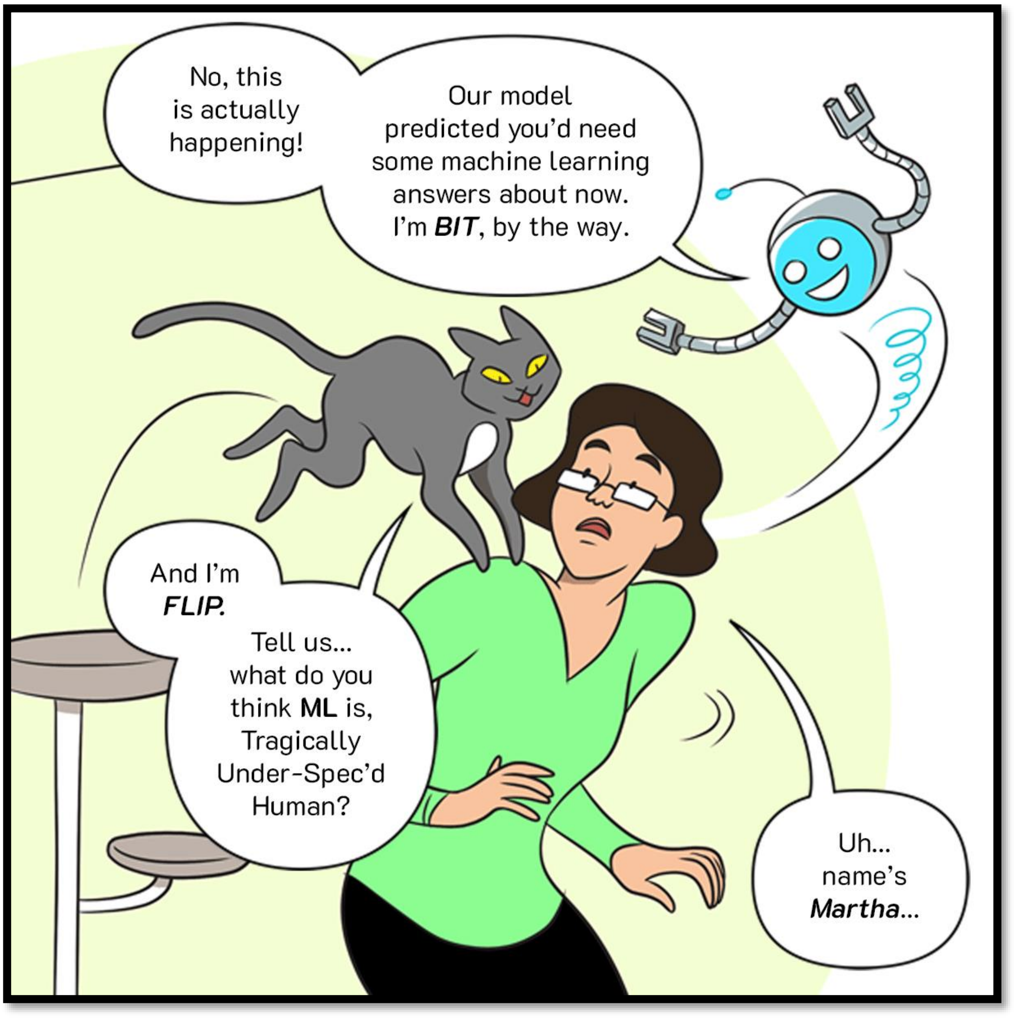
No, this is actually happening!
Our model predicted you'd need some machine learning answers about now. I'm ***BIT***, by the way.
And I'm ***FLIP.***
Tell us... what do you think **ML** is, Tragically Under-Spec'd Human?
Uh... name's ***Martha***...

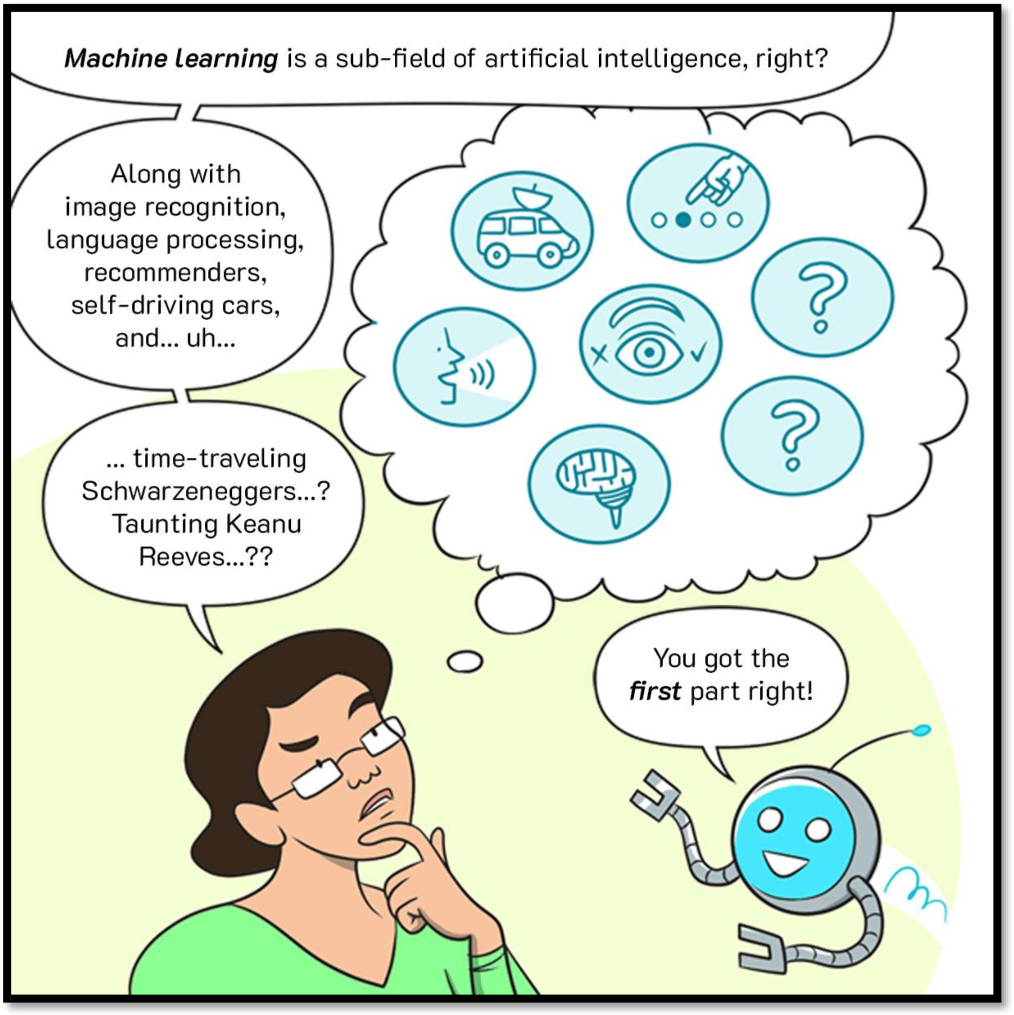
Machine learning is a sub-field of artificial intelligence, right?
Along with image recognition, language processing, recommenders, self-driving cars, and... uh...
... time-traveling Schwarzeneggers...? Taunting Keanu Reeves...??
You got the ***first*** part right!

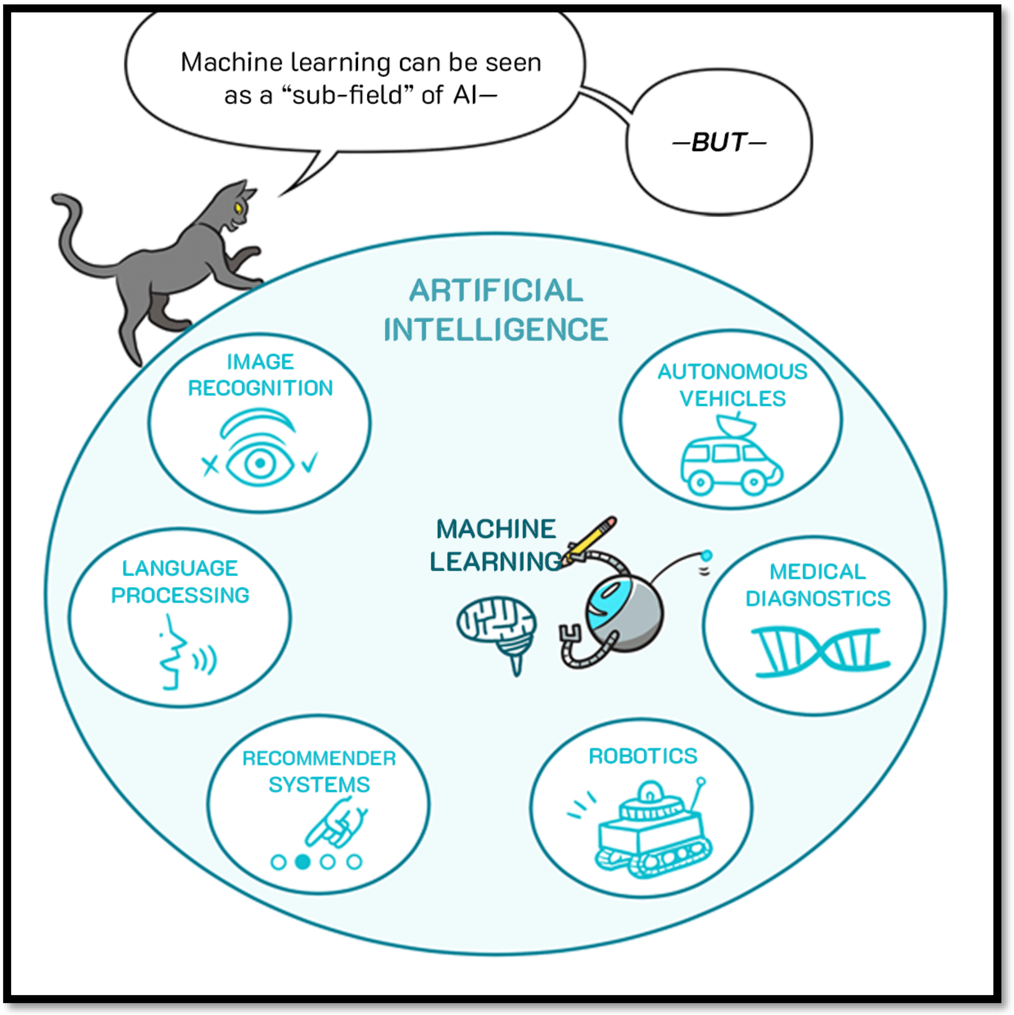
Machine learning can be seen as a "sub-field" of AI—
—BUT—
ARTIFICIAL INTELLIGENCE
IMAGE RECOGNITION
AUTONOMOUS VEHICLES
MACHINE LEARNING
LANGUAGE PROCESSING
MEDICAL DIAGNOSTICS
RECOMMENDER SYSTEMS
ROBOTICS

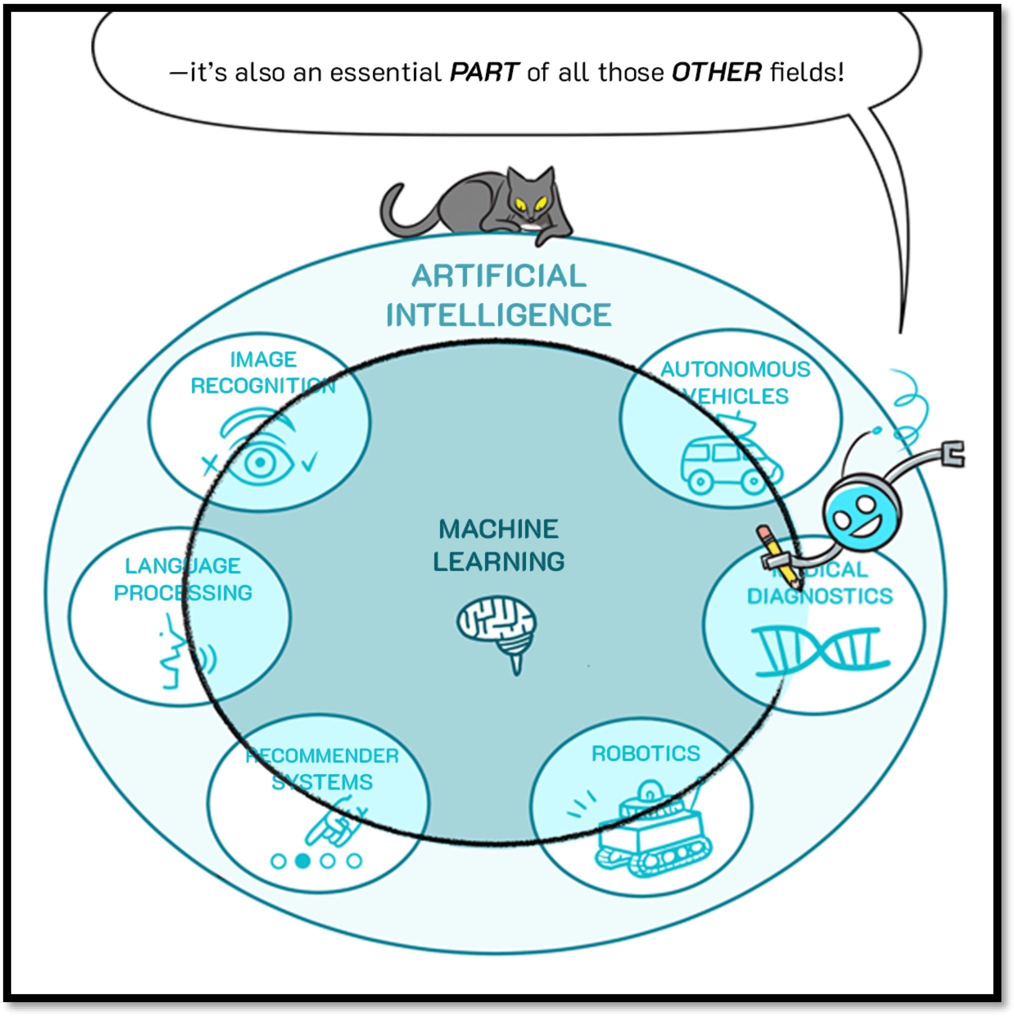
—it's also an essential *PART* of all those *OTHER* fields!
ARTIFICIAL INTELLIGENCE
IMAGE RECOGNITION
AUTONOMOUS VEHICLES
MACHINE LEARNING
LANGUAGE PROCESSING
MEDICAL DIAGNOSTICS
RECOMMENDER SYSTEMS
ROBOTICS

Machine learning is when computers *IMPROVE* at a task based on *EXPERIENCE* at that task, but *without explicit programming.*
In other words, computers finally starting to **take the initiative** in ways they rarely have before.
Hunh.

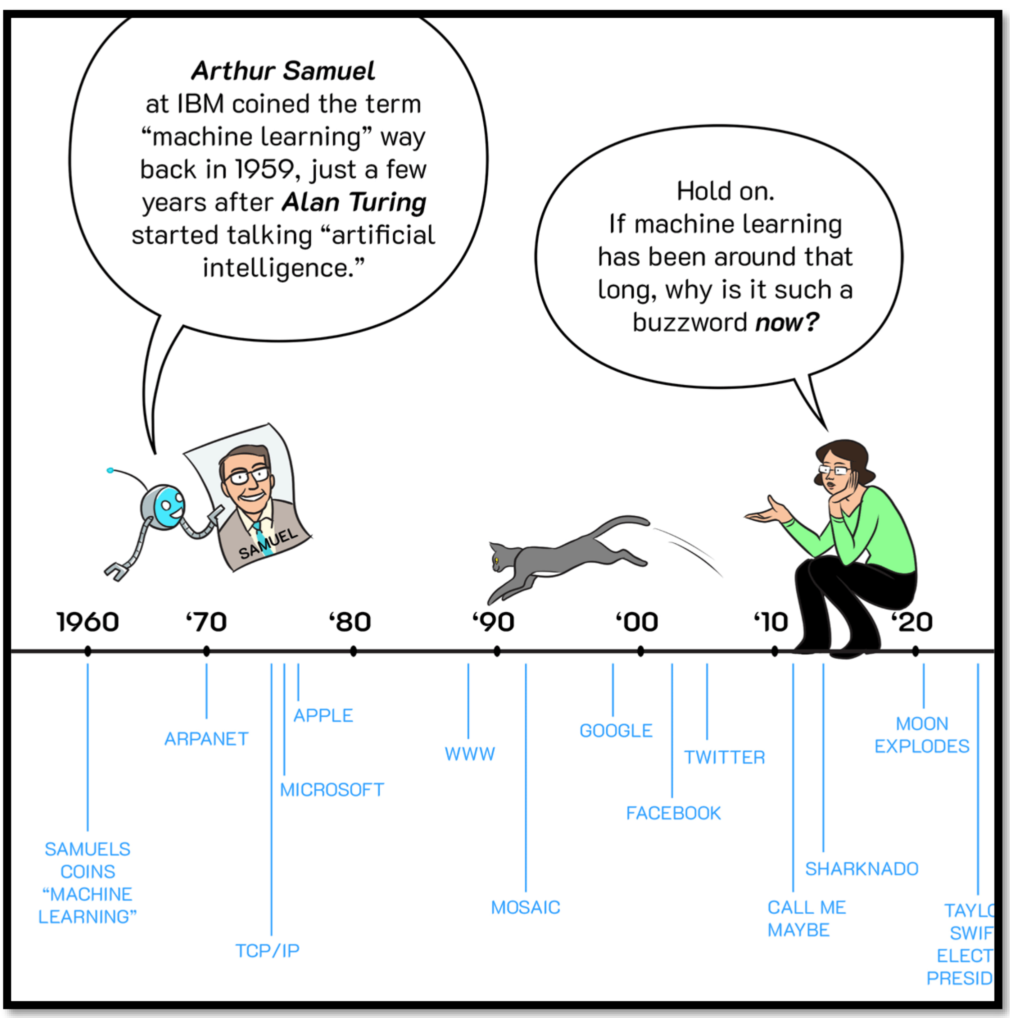

Arthur Samuel at IBM coined the term "machine learning" way back in 1959, just a few years after *Alan Turing* started talking "artificial intelligence."
Hold on. If machine learning has been around that long, why is it such a buzzword *now?*
SAMUEL
1960
'70
'80
'90
'00
'10
'20
SAMUELS COINS "MACHINE LEARNING"
ARPANET
TCP/IP
MICROSOFT
APPLE
WWW
MOSAIC
GOOGLE
FACEBOOK
TWITTER
CALL ME MAYBE
SHARKNADO
MOON EXPLODES

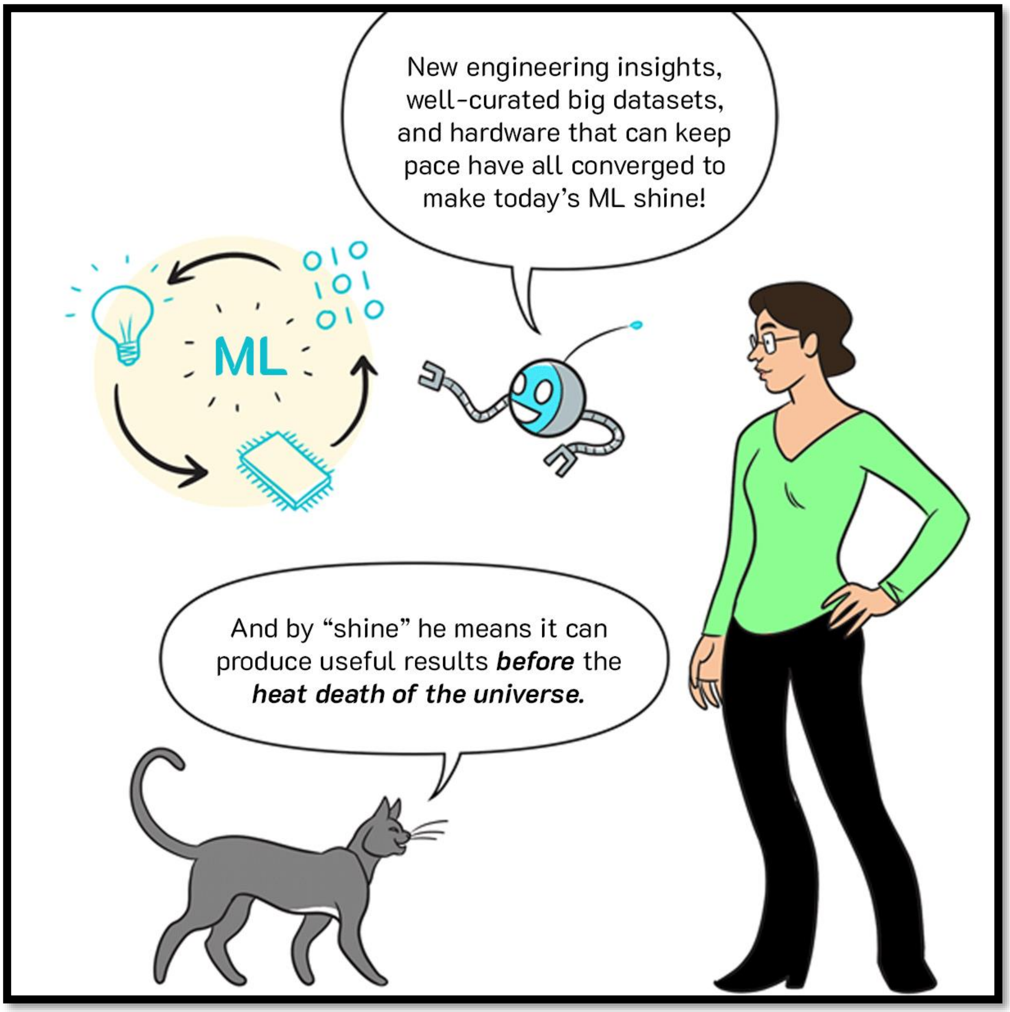
New engineering insights, well-curated big datasets, and hardware that can keep pace have all converged to make today's ML shine!
ML
And by "shine" he means it can produce useful results ***before*** the ***heat death of the universe.***

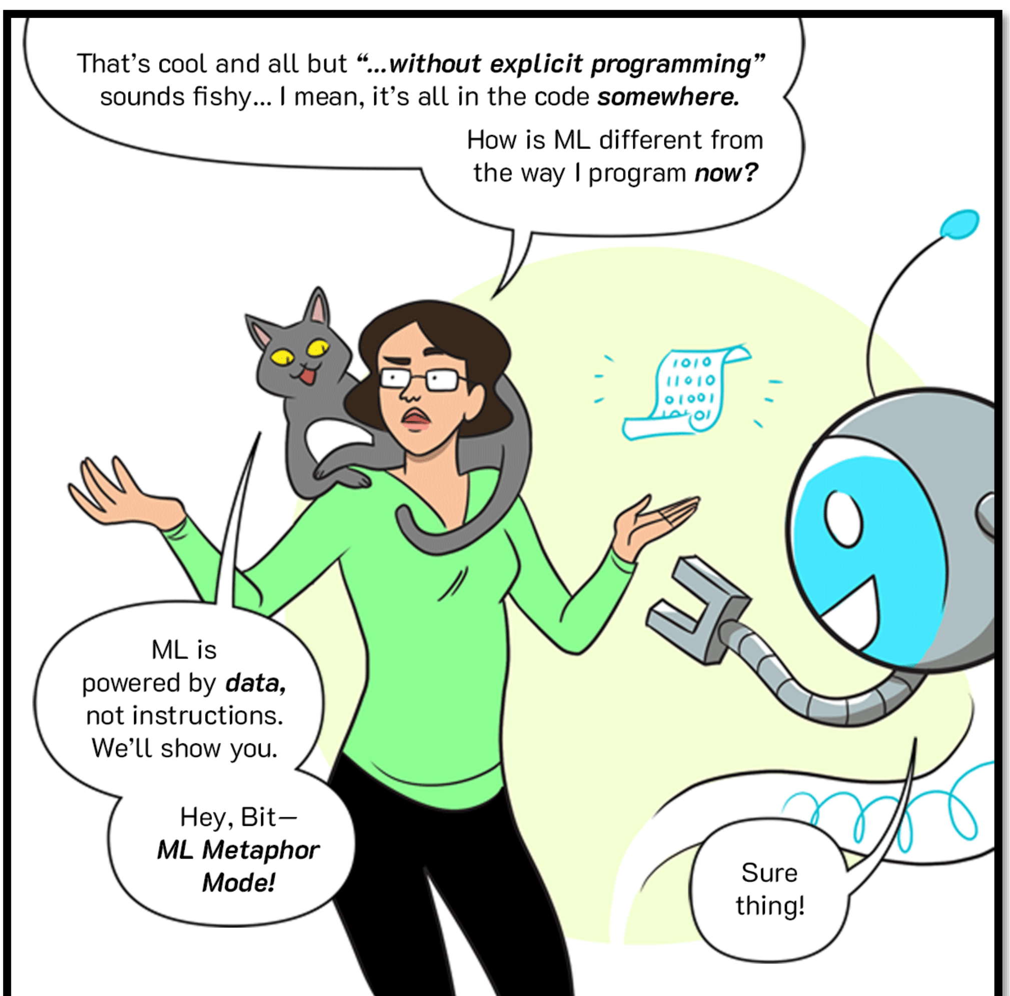
That's cool and all but ***"...without explicit programming"*** sounds fishy... I mean, it's all in the code ***somewhere.***
How is ML different from the way I program ***now?***
ML is powered by ***data,*** not instructions. We'll show you.
Hey, Bit— ***ML Metaphor Mode!***
Sure thing!


Let's try some classification. We can start with...
...this fork!
BOOOP
Where were you *hiding* that?

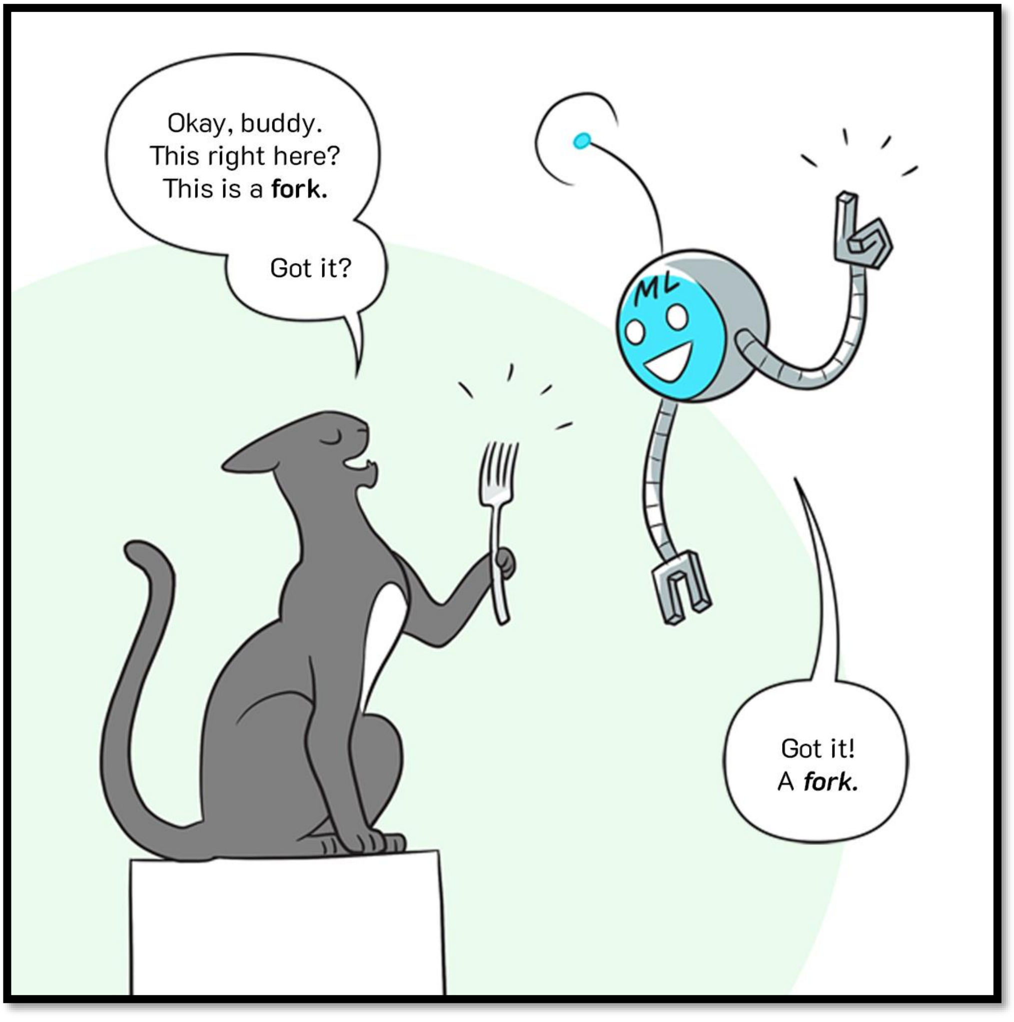
Okay, buddy. This right here? This is a **fork.**
Got it?
Got it! A ***fork.***

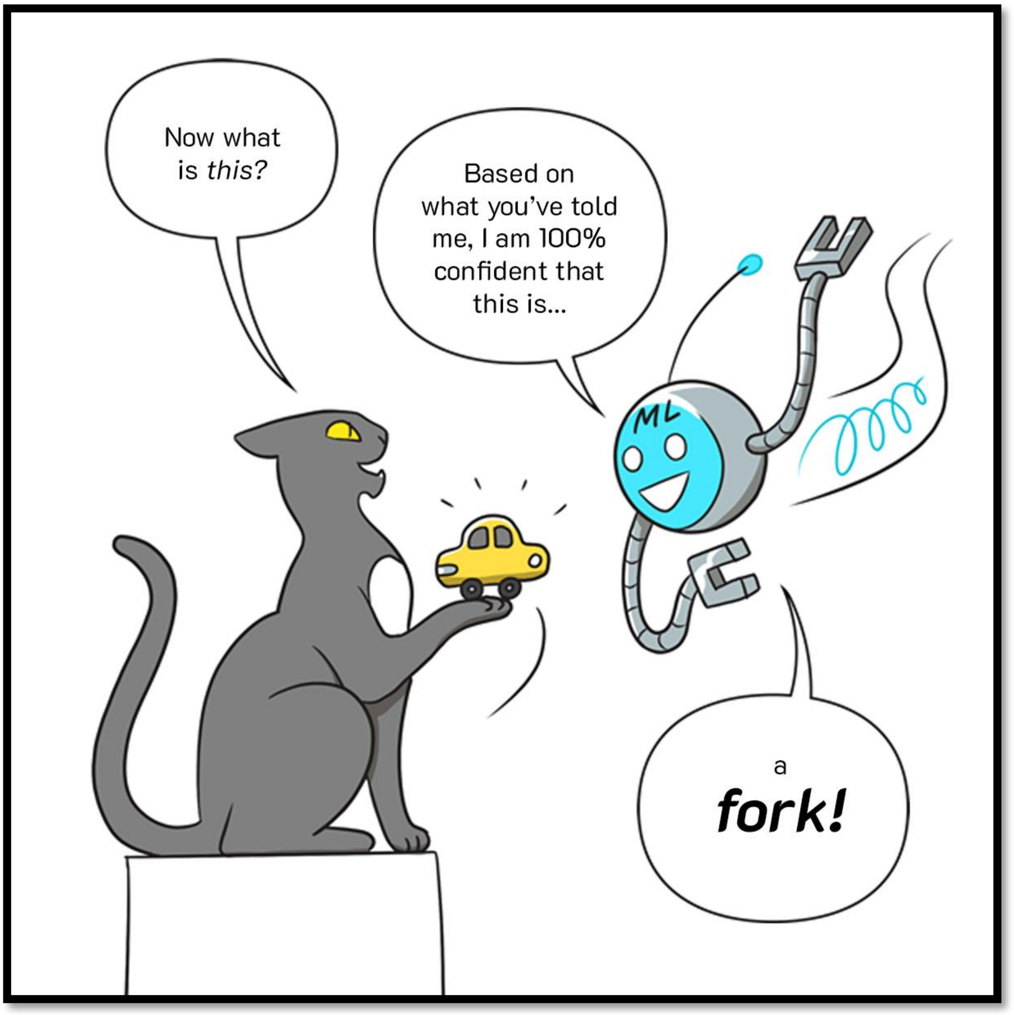
Now what is *this?*
Based on what you've told me, I am 100% confident that this is...
a **fork!**

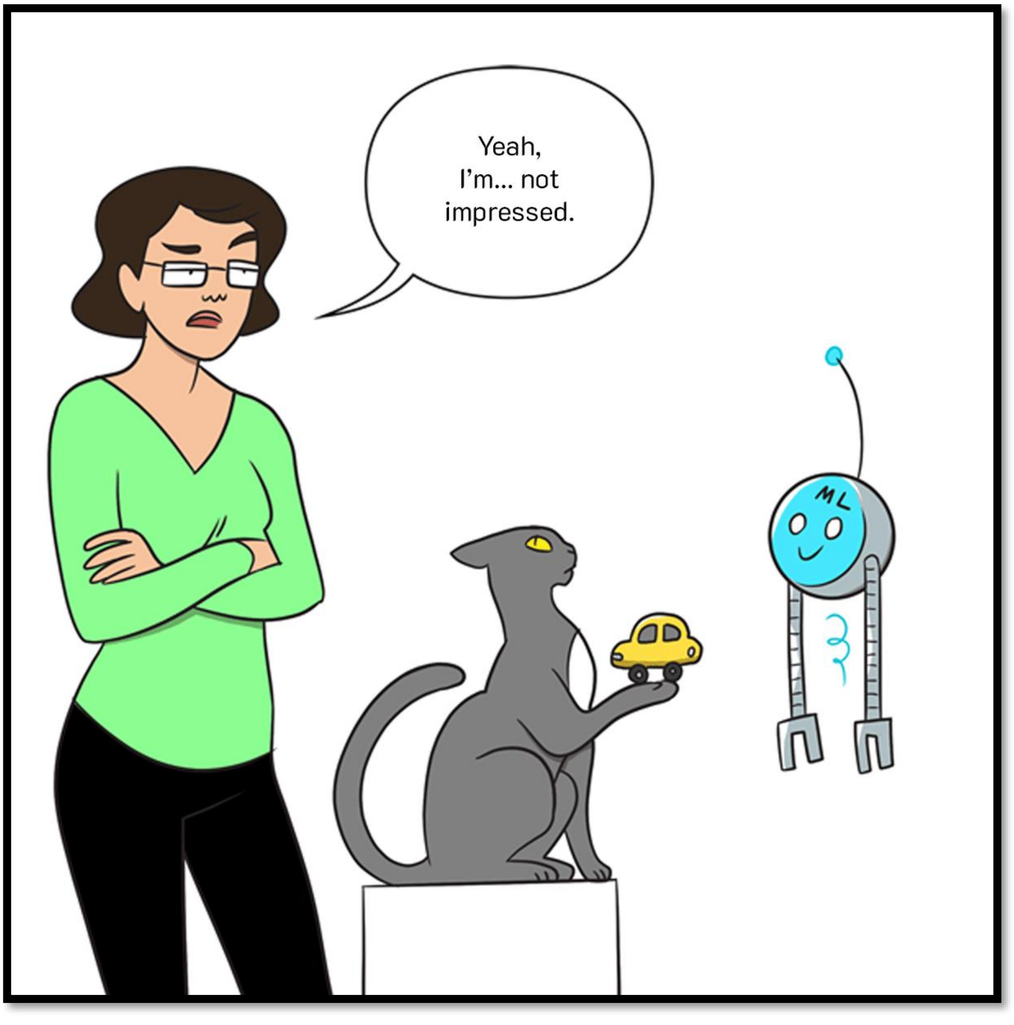
Yeah, I'm... not impressed.

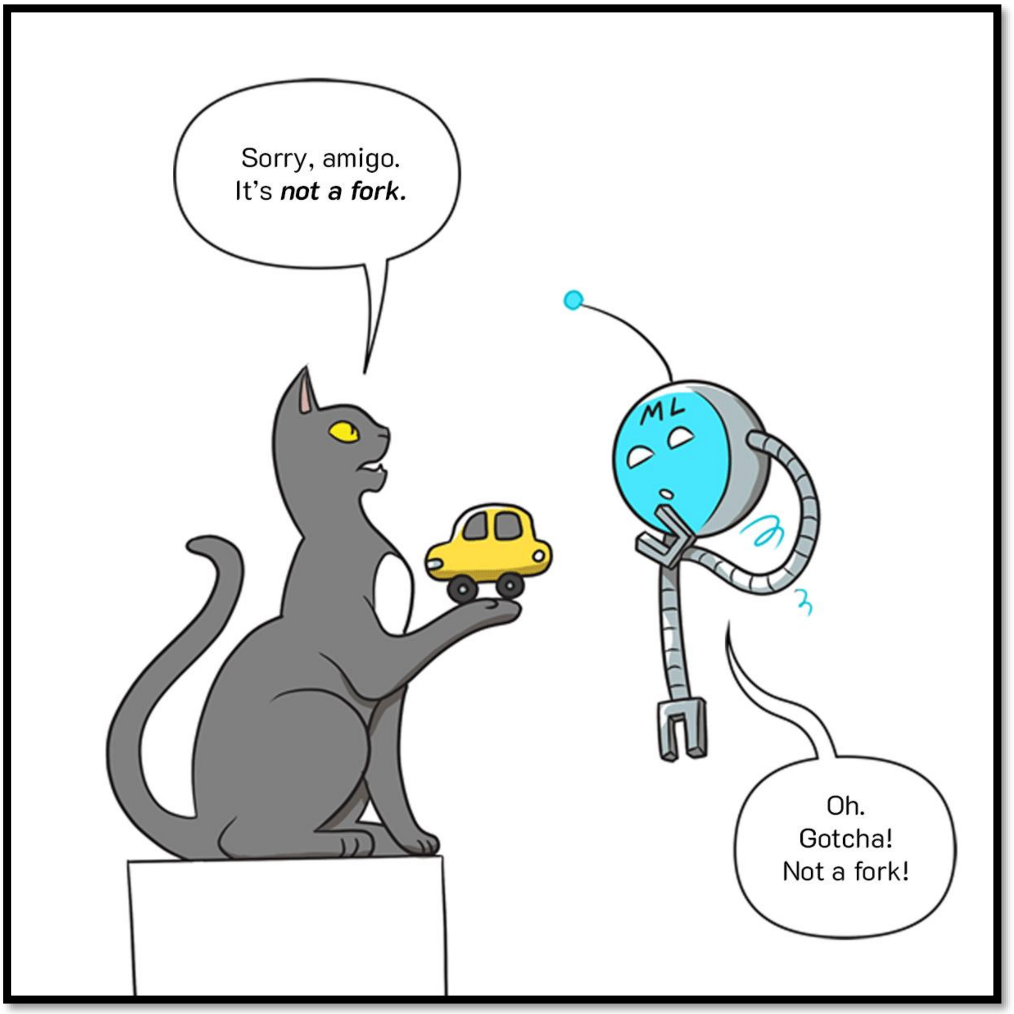
Sorry, amigo.
It's ***not a fork.***
Oh.
Gotcha!
Not a fork!

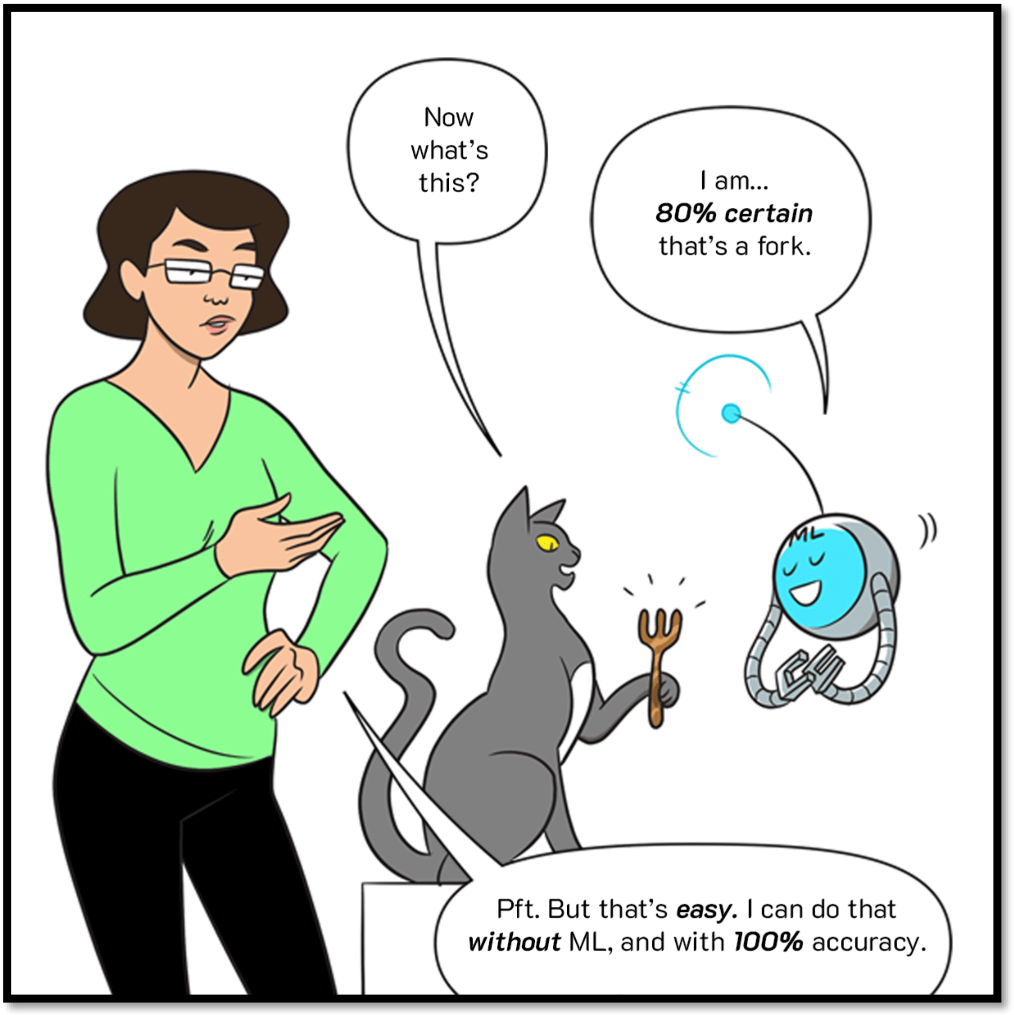
Now
what's
this?
I am...
80% certain
that's a fork.
Pft. But that's ***easy.*** I can do that
without ML, and with ***100%*** accuracy.

if ((thing.isStainlessSteel() OR
thing.isWood()) AND
(thing.hasFourTines() OR
thing.hasThreeTines())
{
isFork();
}
else
{
isNotFork();
}
Okay, fair enough...

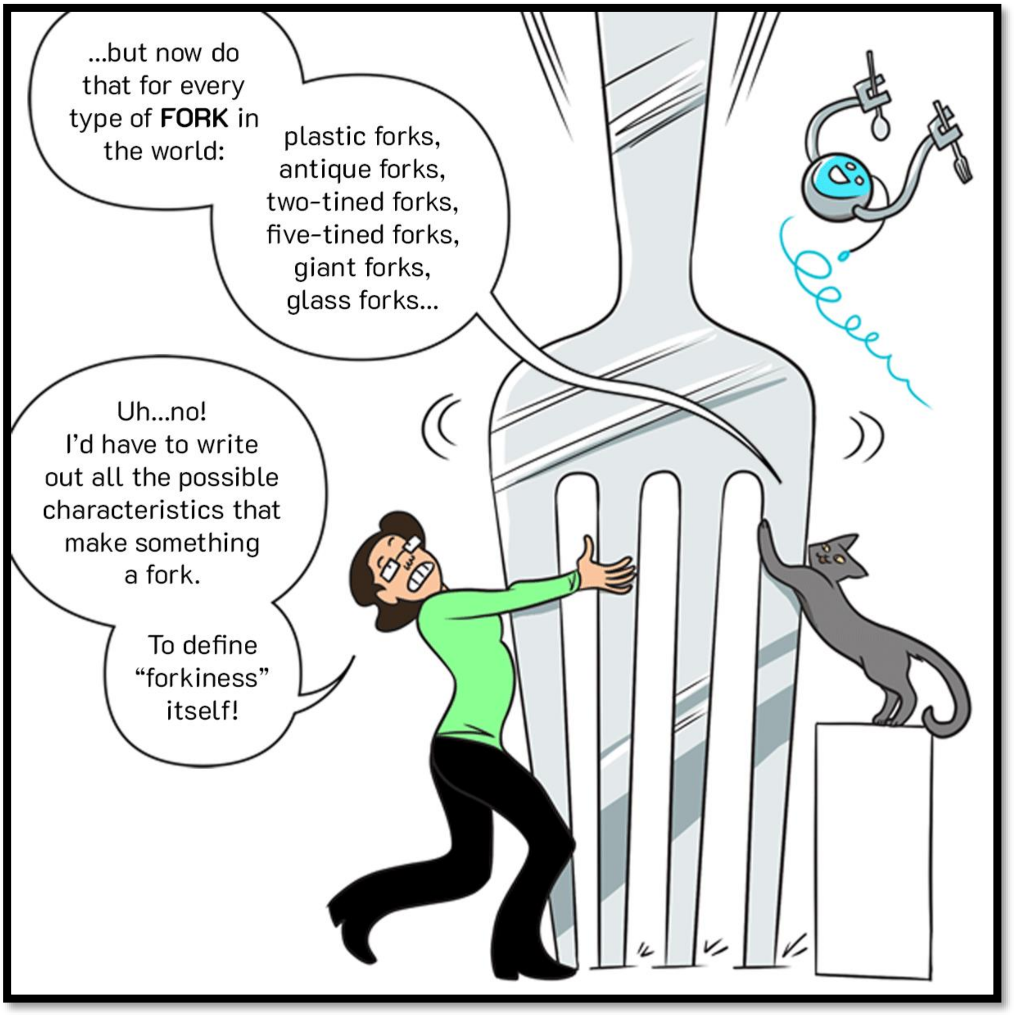

...but now do that for every type of **FORK** in the world:
plastic forks, antique forks, two-tined forks, five-tined forks, giant forks, glass forks...
Uh...no! I'd have to write out all the possible characteristics that make something a fork.
To define "forkiness" itself!


Exactly! But with ML, if you have enough well-curated data—
—and run it through enough training and testing...


PULL FOR LABELED DATA
Woohoo!


Ahem.
Can you guess what *this* is, buddy?
ML
Hm.
I'm ***93.4%*** sure that's a fork.

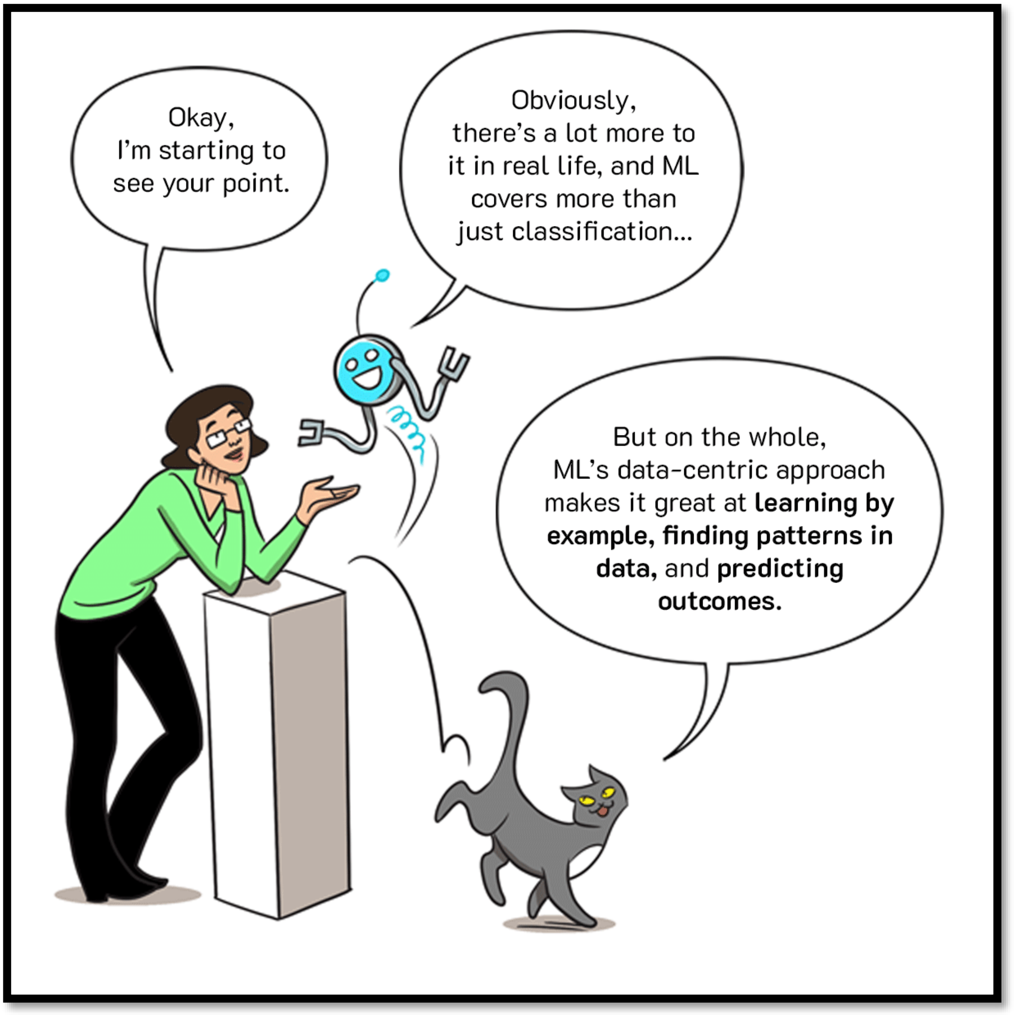
Okay, I'm starting to see your point.
Obviously, there's a lot more to it in real life, and ML covers more than just classification...
But on the whole, ML's data-centric approach makes it great at **learning by example, finding patterns in data,** and **predicting outcomes.**


Sounds great, but how can a machine ***"learn,"*** exactly?
So, let's say you want to predict the prices of houses in Portland, Oregon...
Hold on. Does it ***have*** to be Portland? These examples are ***always*** in Portland.
Aren't ***ALL*** houses in Portland?


ACK!
Now, let's say you know just the ***SIZE*** and ***PRICE*** of only ***ten houses...***
How could you ***use*** that dataset to guess the price of ***OTHERS?***


I guess I'd pop 'em into a ***spreadsheet,*** and then plot X and Y ***coordinates*** and a ***trend-line*** like in high school algebra.
Exactly! Just ***linear regression.*** Nothing fancy.
Now add ***TEN MORE*** houses!

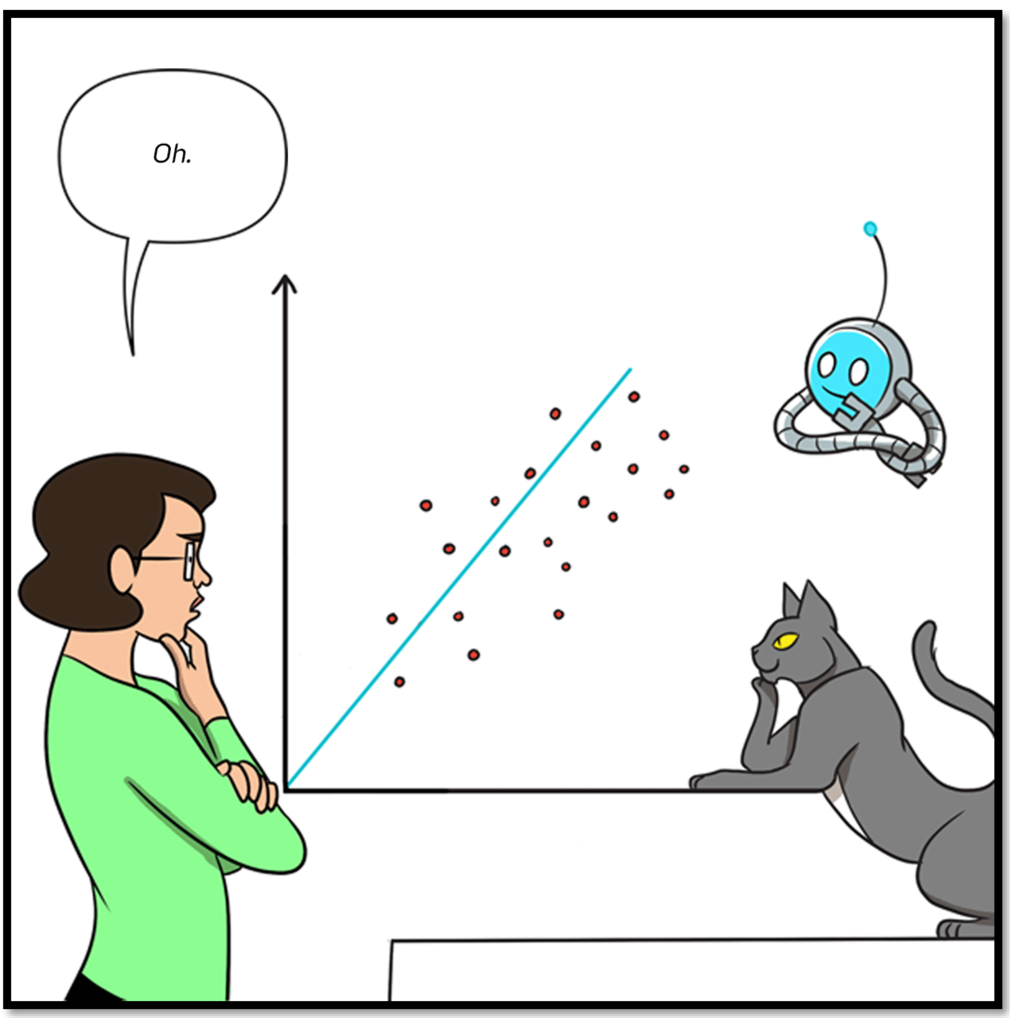
Oh.

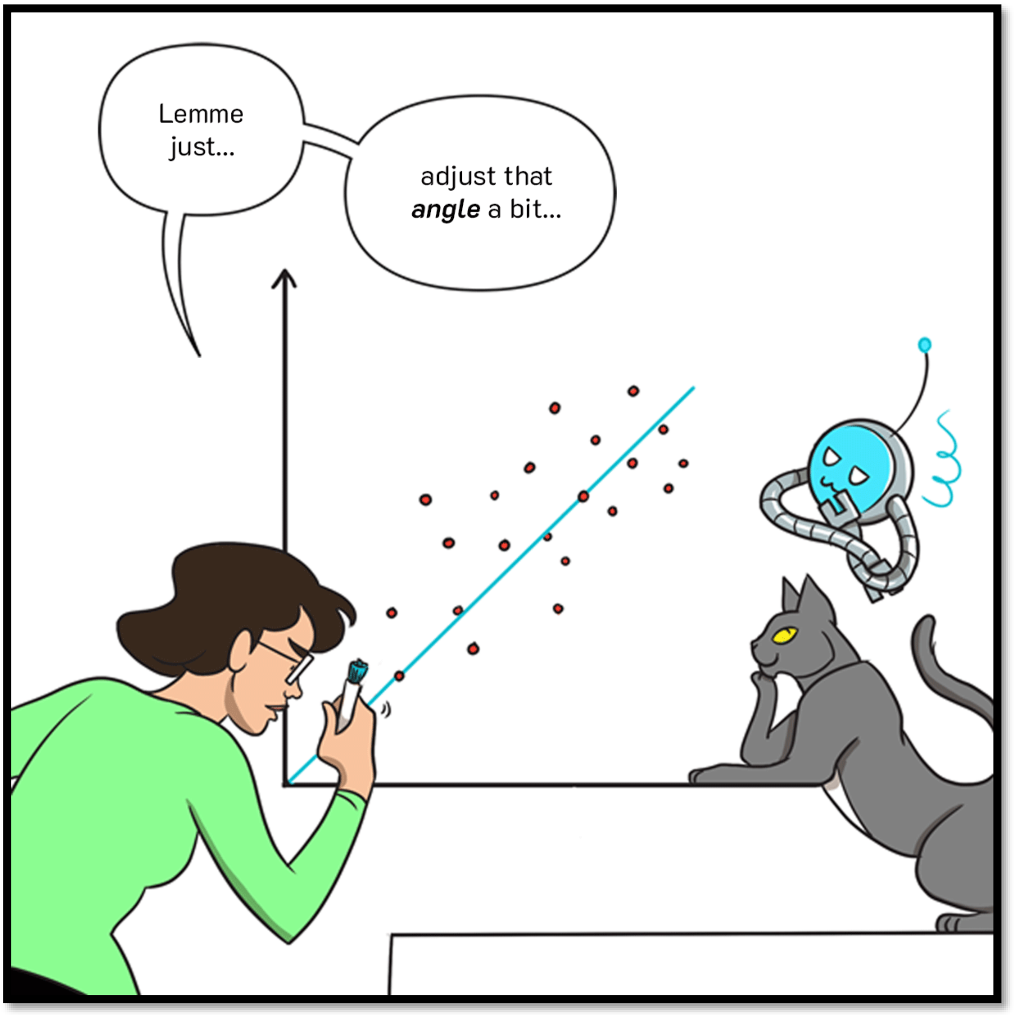
Lemme
just...
adjust that
angle a bit...


There.
Much better.
AH-HAH!

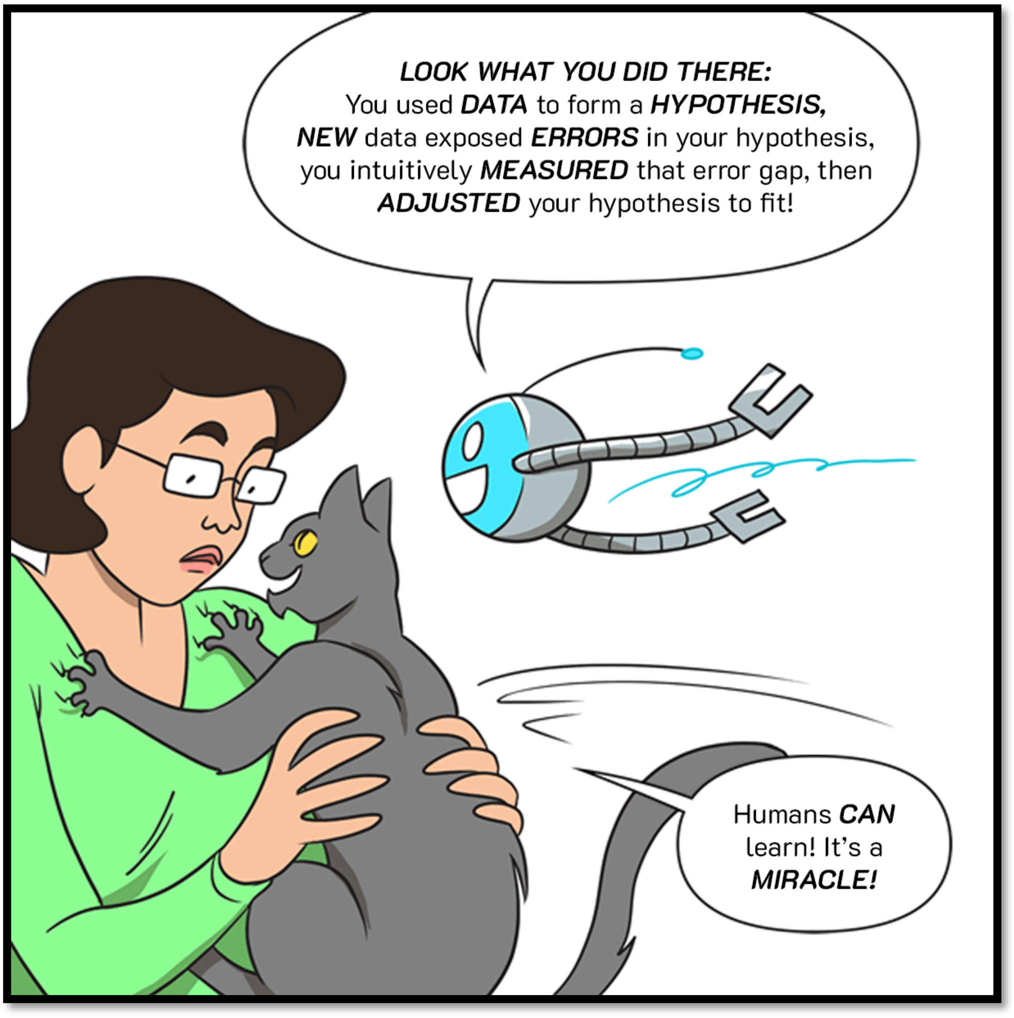
LOOK WHAT YOU DID THERE: You used **DATA** to form a **HYPOTHESIS, NEW** data exposed **ERRORS** in your hypothesis, you intuitively **MEASURED** that error gap, then **ADJUSTED** your hypothesis to fit!
Humans **CAN** learn! It's a **MIRACLE!**


Oh, **CUT IT OUT,** you dorks! I just saw what was wrong and **fixed** it... like **ANYONE** would.
"Any**ONE**" maybe, but not any**THING.** We machines can't rely on **"intuition"** to tell **forks** from **spoons** or predict **housing costs** in the **only city with houses!**
Reducing errors in ML is a **journey,** to be taken **one step at a time!**

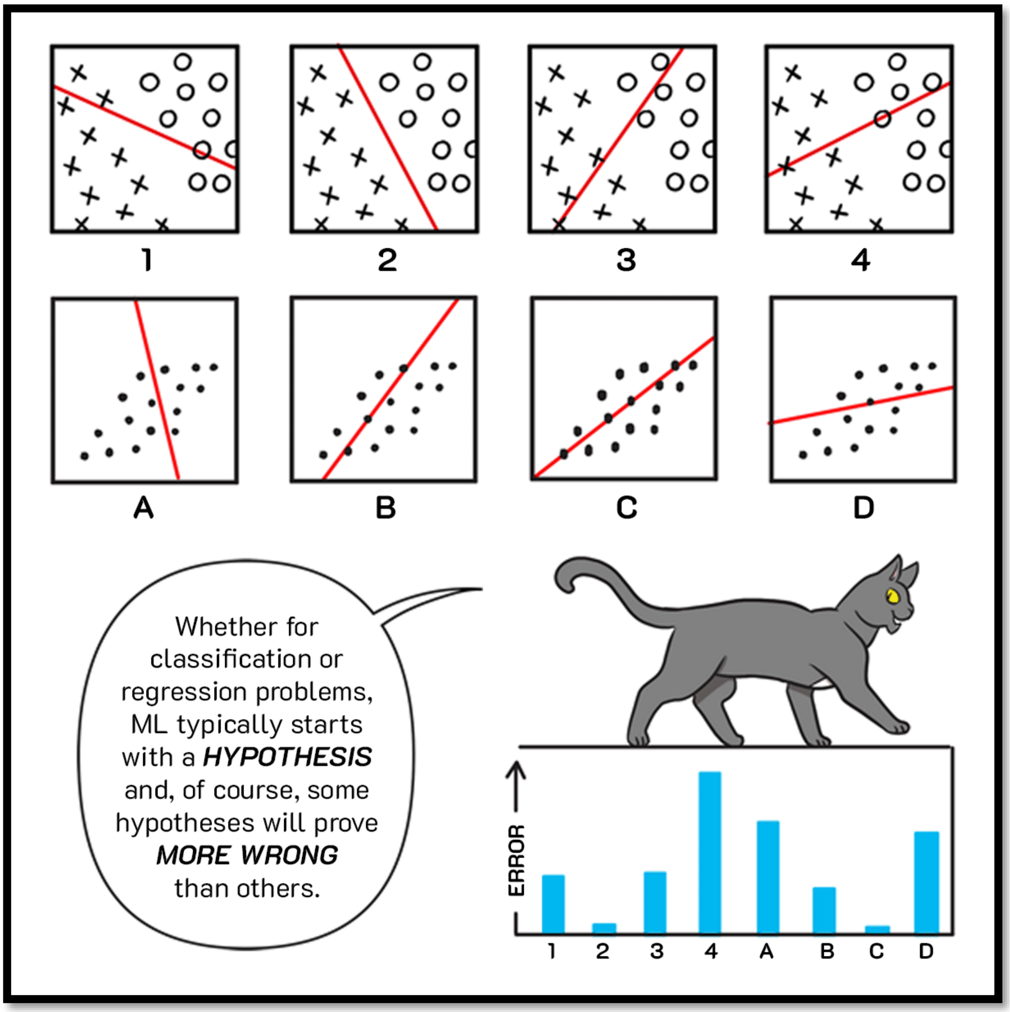
1
2
3
4
A
B
C
D
Whether for classification or regression problems, ML typically starts with a ***HYPOTHESIS*** and, of course, some hypotheses will prove ***MORE WRONG*** than others.
ERROR
1 2 3 4 A B C D

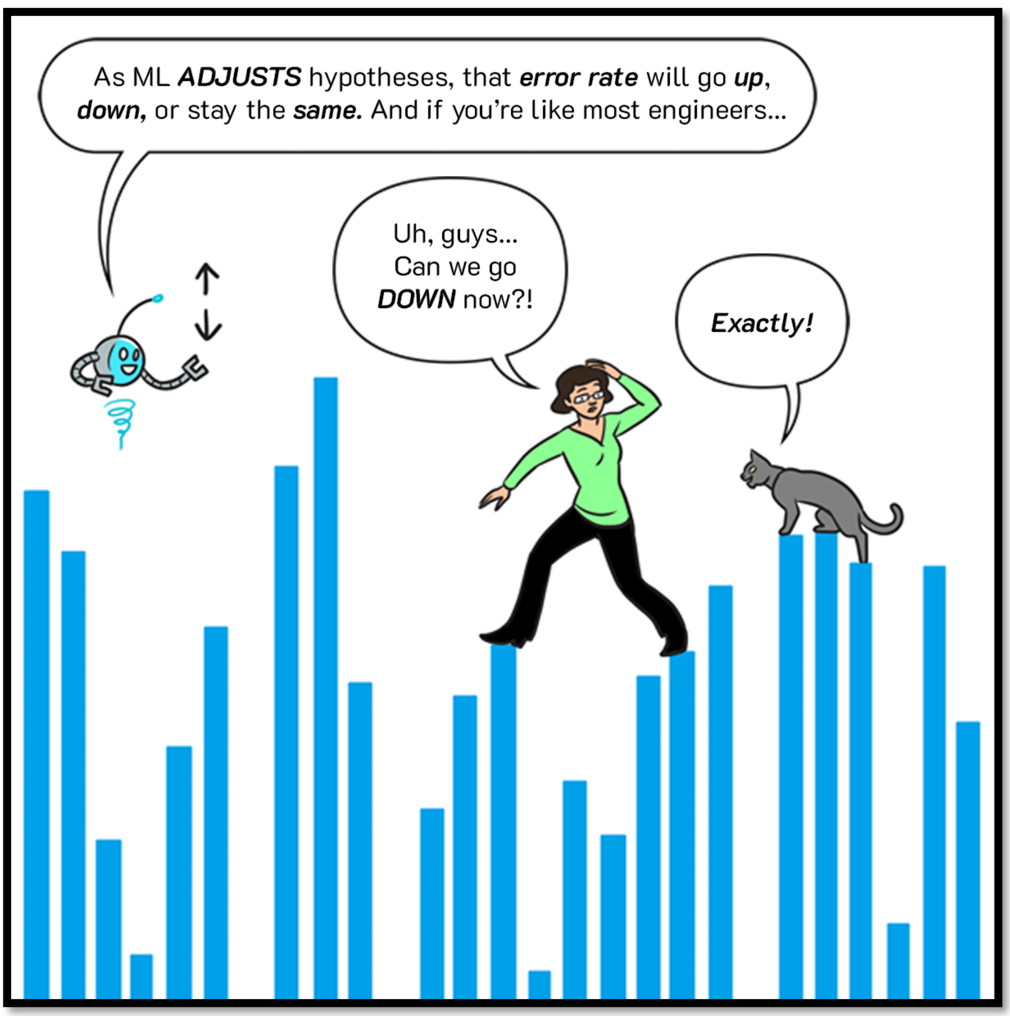
As ML ***ADJUSTS*** hypotheses, that ***error rate*** will go ***up***, ***down,*** or stay the ***same.*** And if you're like most engineers...
Uh, guys... Can we go ***DOWN*** now?!
Exactly!


You want to get **“DOWN”** to where the model’s ***error rate*** is at its ***LOWEST,*** so that you—
NO. NO. I mean I ***LITERALLY*** want to climb down to—
AH!
Why can’t I ***SEE*** anything anymore?!


heh
You didn’t think we’d make it ***that*** easy, did’ja?


But that ***error landscape***...
Where did it all ***go??***
Oh, it's still here... all ***around*** us!
But the ***SHAPE*** of that landscape...


...that array of peaks and valleys we call an ***"error function"*** or ***"loss function"***...
...can only be revealed—

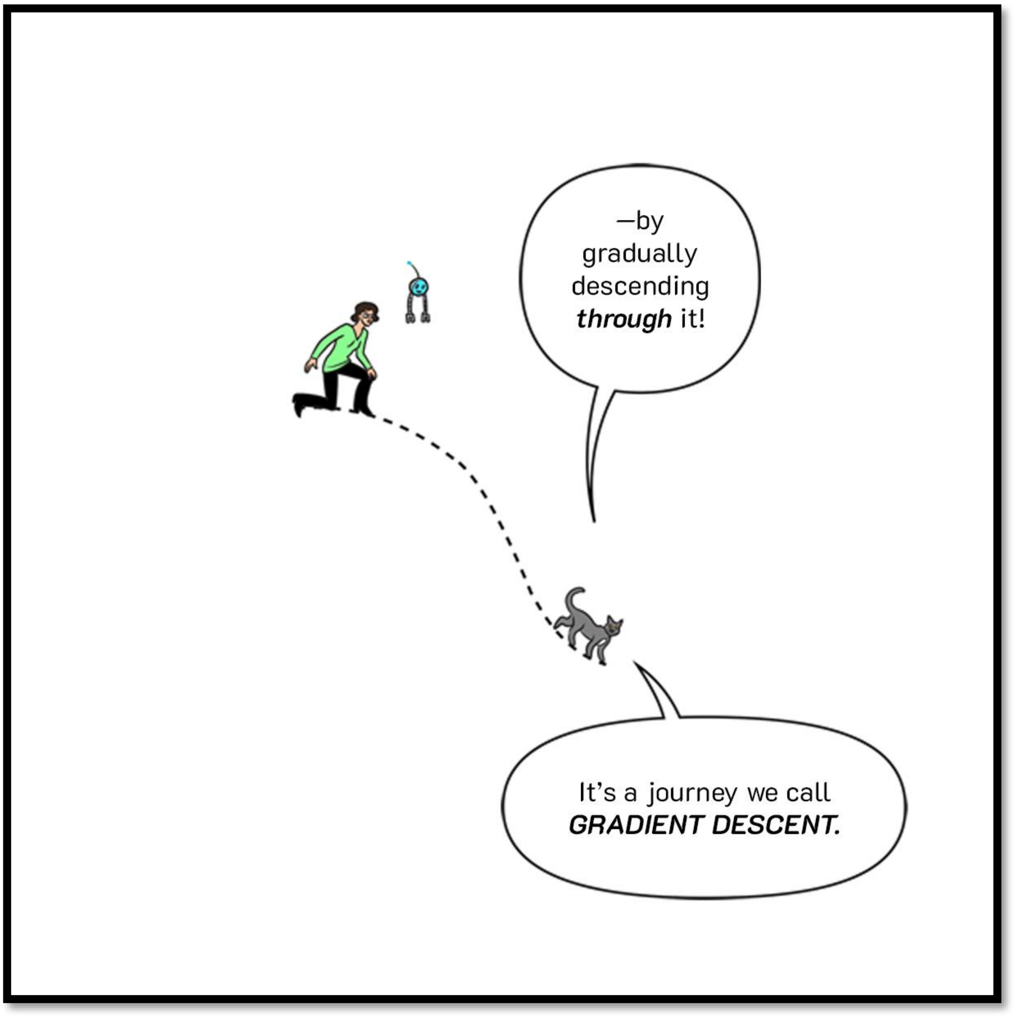
—by gradually descending *through* it!
It's a journey we call ***GRADIENT DESCENT.***

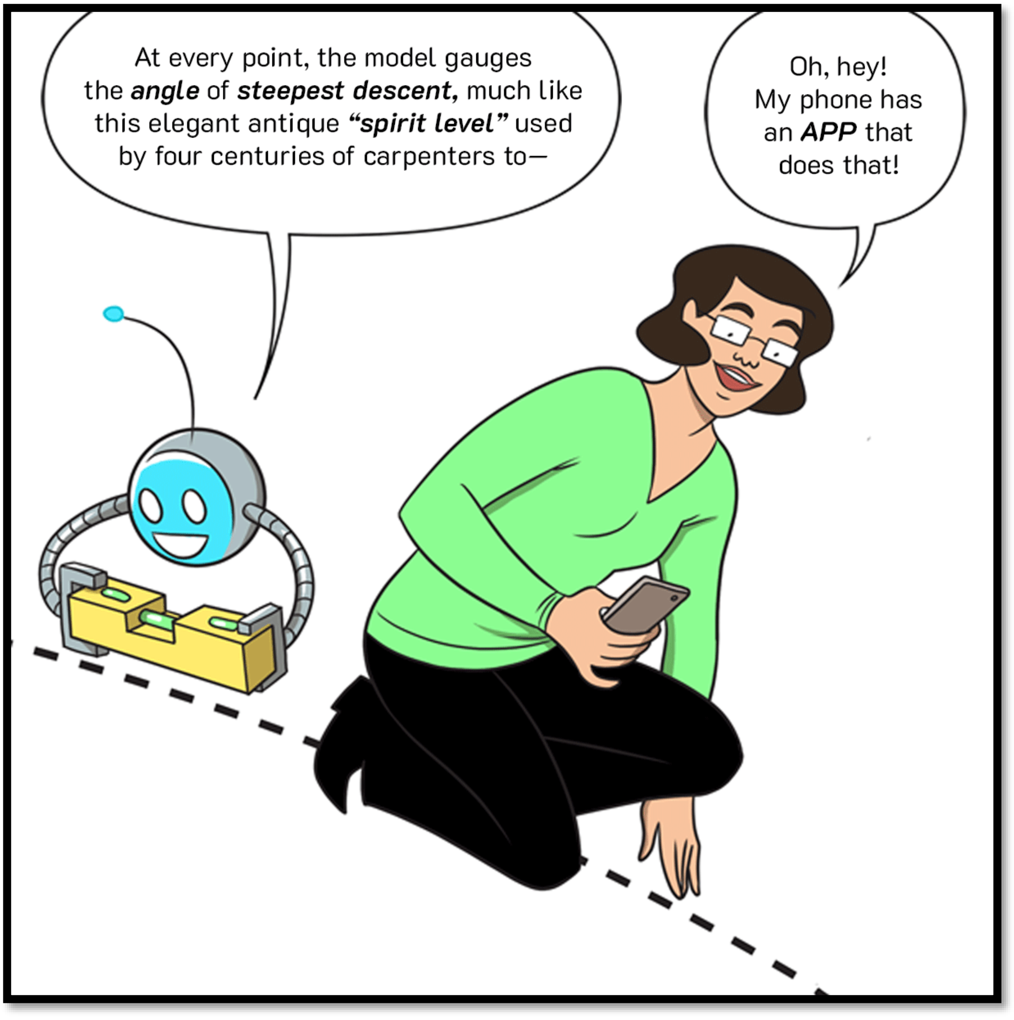
At every point, the model gauges the ***angle*** of ***steepest descent,*** much like this elegant antique ***"spirit level"*** used by four centuries of carpenters to—
Oh, hey! My phone has an ***APP*** that does that!

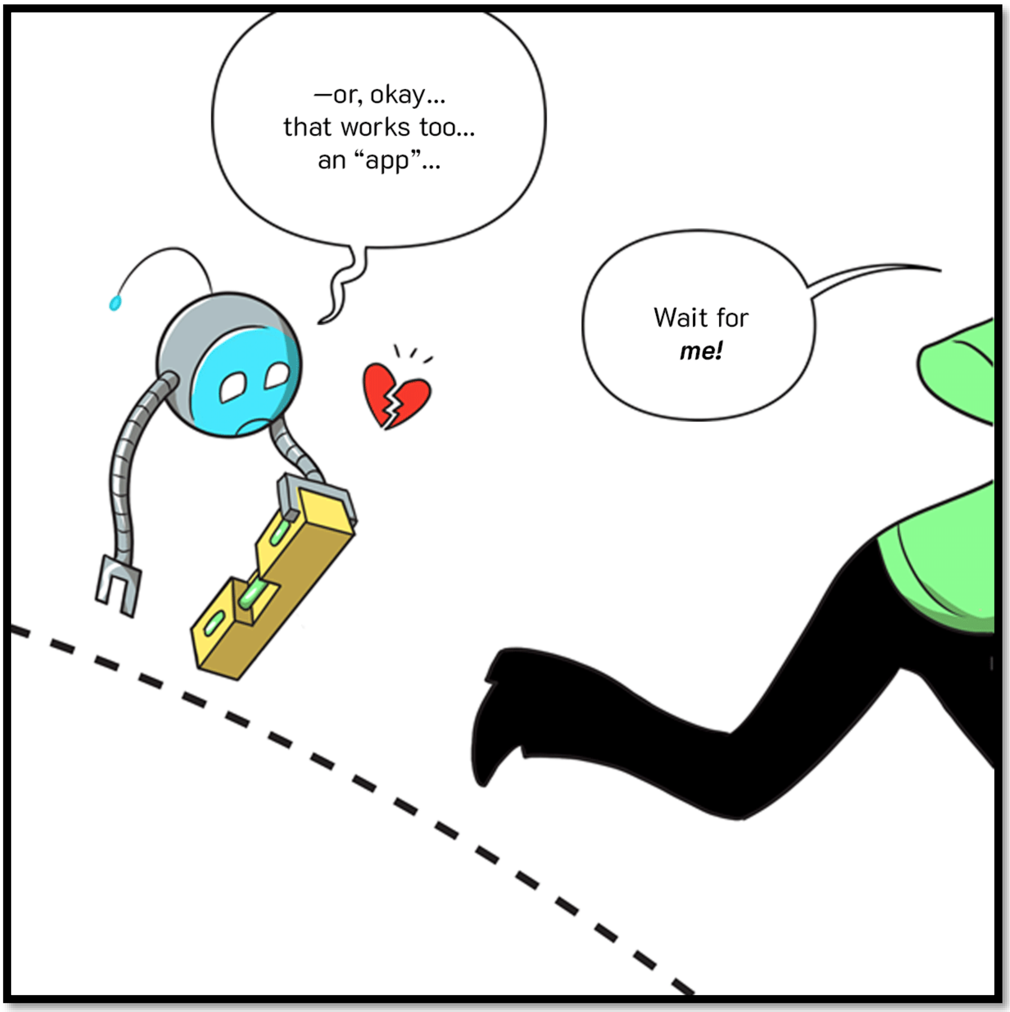
—or, okay…
that works too…
an "app"…
Wait for *me!*


So *gradient descent* is kinda like hiking down the *SLOPES OF WRONG* into the *VALLEY OF RIGHT?*
Yup! Or at least into the *"Valley of Right Enough."*
Machine learning is part *data science* and *statistics;* there's a strong *probabilistic streak* to it.

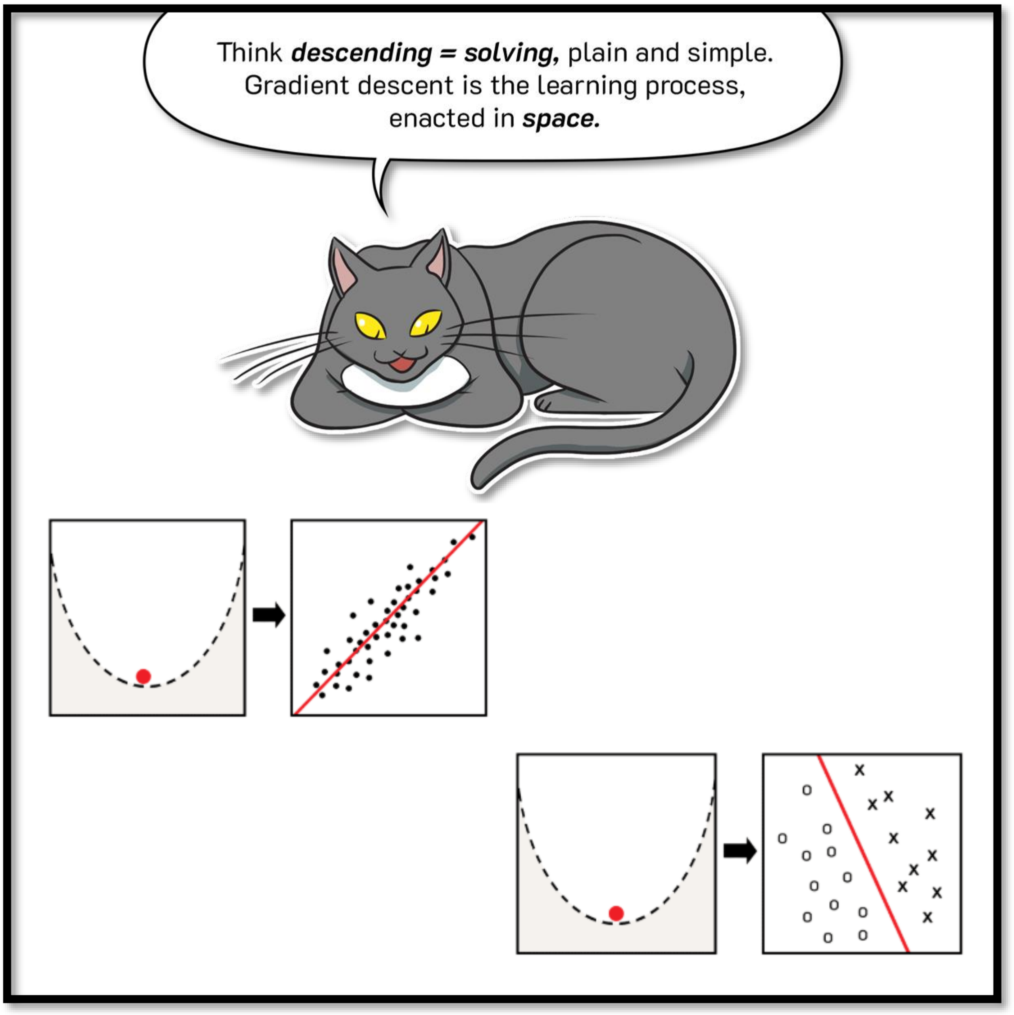
Think ***descending = solving,*** plain and simple.
Gradient descent is the learning process,
enacted in ***space.***
x x x x x x x x x
o o o o o o o o o o o o

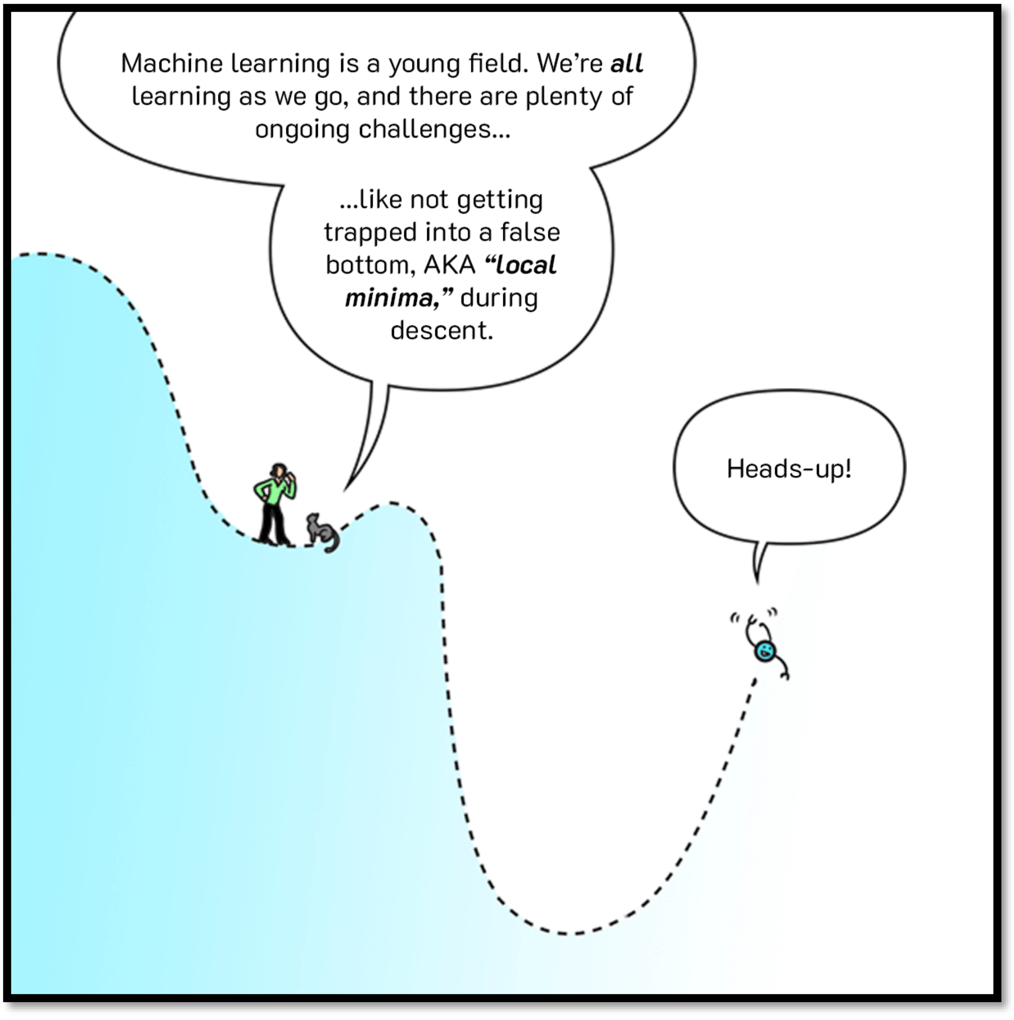
Machine learning is a young field. We're ***all*** learning as we go, and there are plenty of ongoing challenges...
...like not getting trapped into a false bottom, AKA ***"local minima,"*** during descent.
Heads-up!

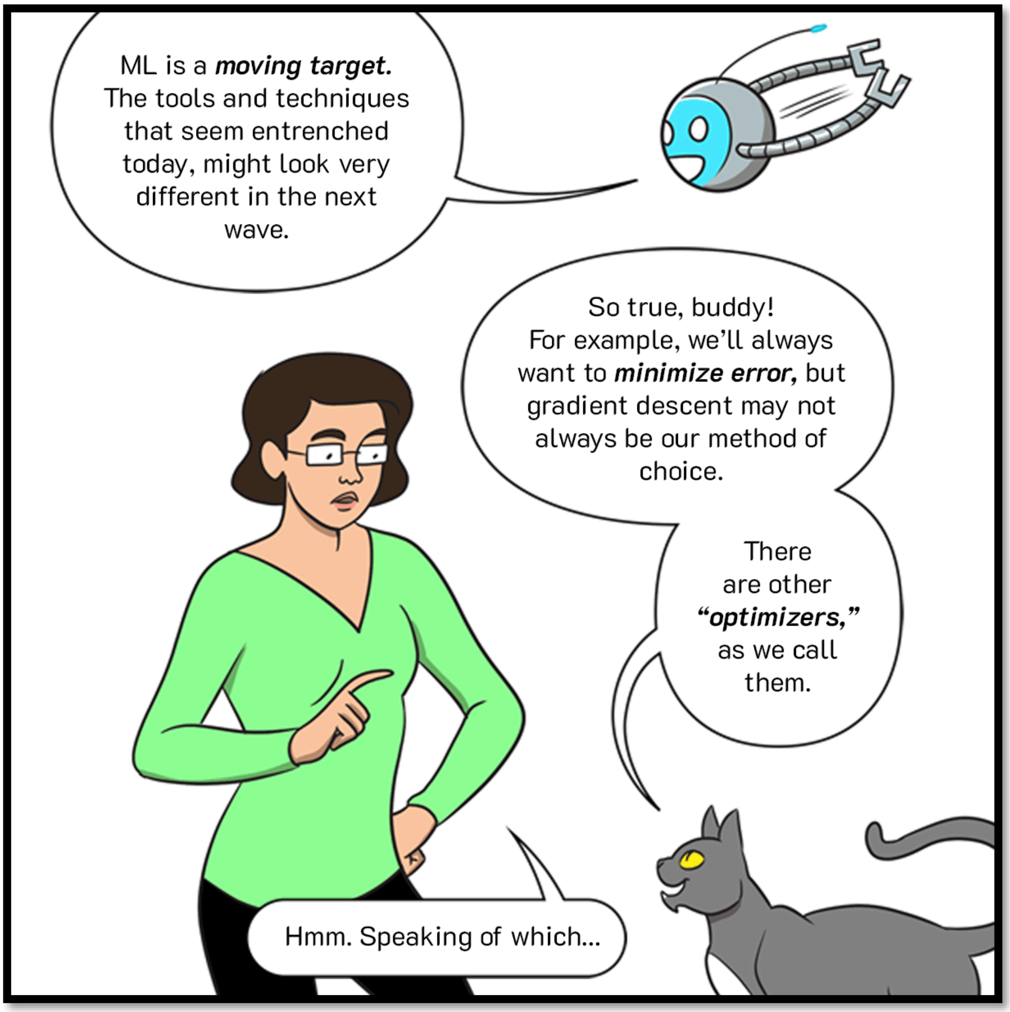
ML is a ***moving target.*** The tools and techniques that seem entrenched today, might look very different in the next wave.
So true, buddy! For example, we'll always want to ***minimize error,*** but gradient descent may not always be our method of choice.
There are other ***"optimizers,"*** as we call them.
Hmm. Speaking of which...

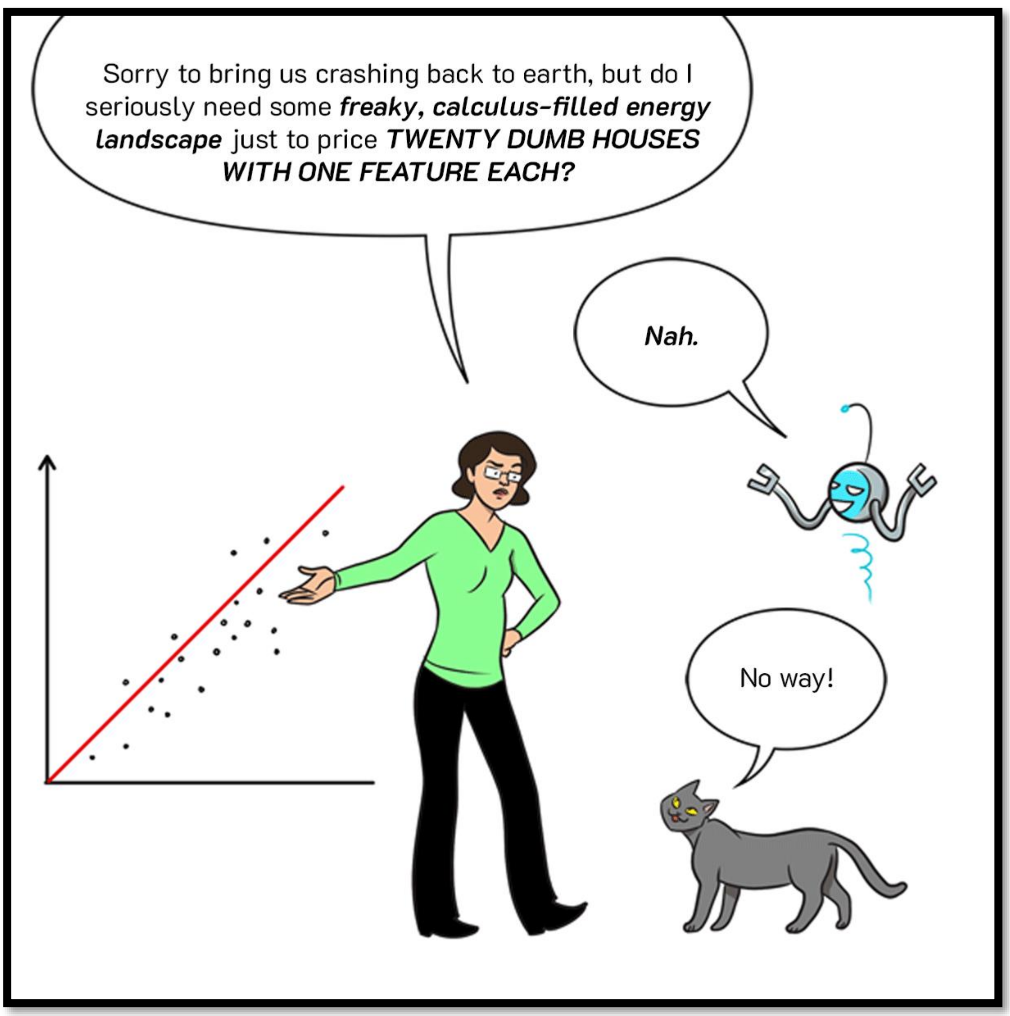
Sorry to bring us crashing back to earth, but do I seriously need some ***freaky, calculus-filled energy landscape*** just to price ***TWENTY DUMB HOUSES WITH ONE FEATURE EACH?***
Nah.
No way!

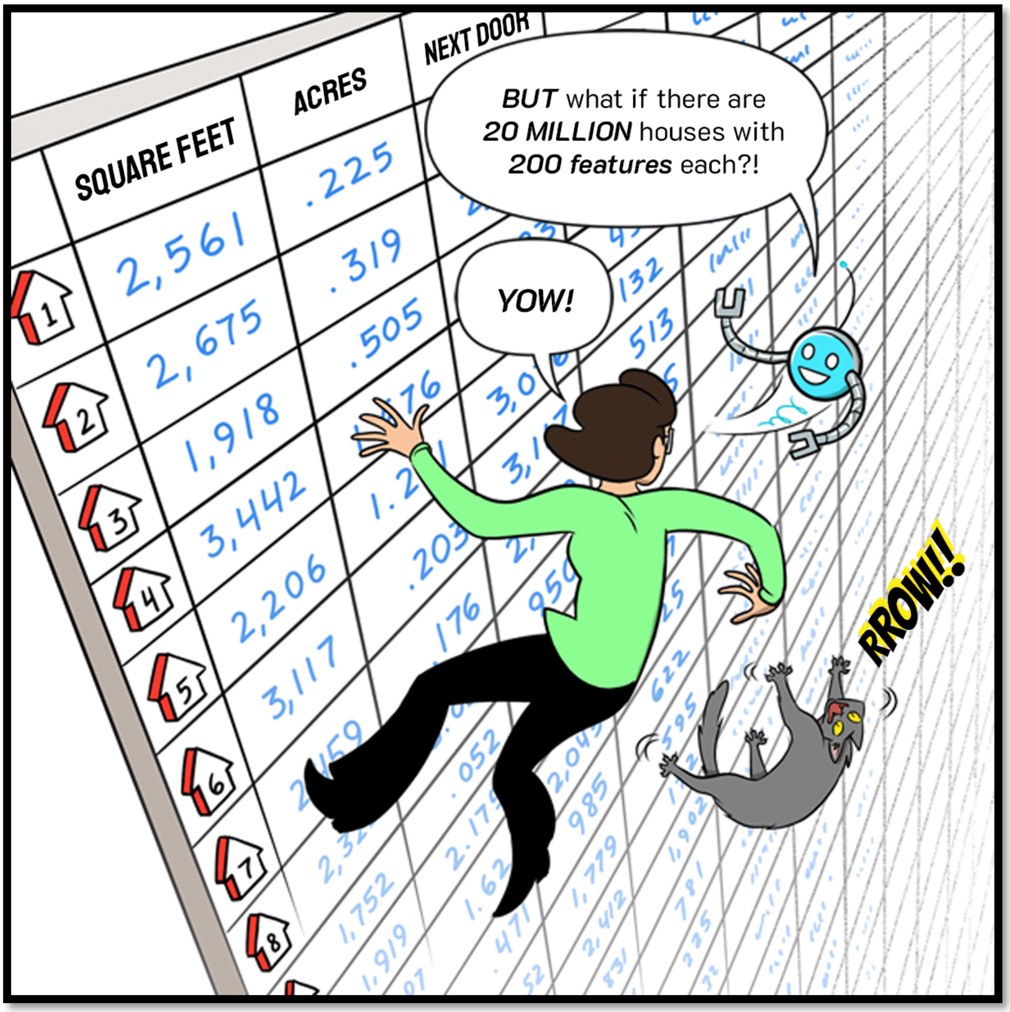
SQUARE FEET
ACRES
NEXT DOOR
BUT what if there are 20 MILLION houses with 200 features each?!
YOW!
RROW!!

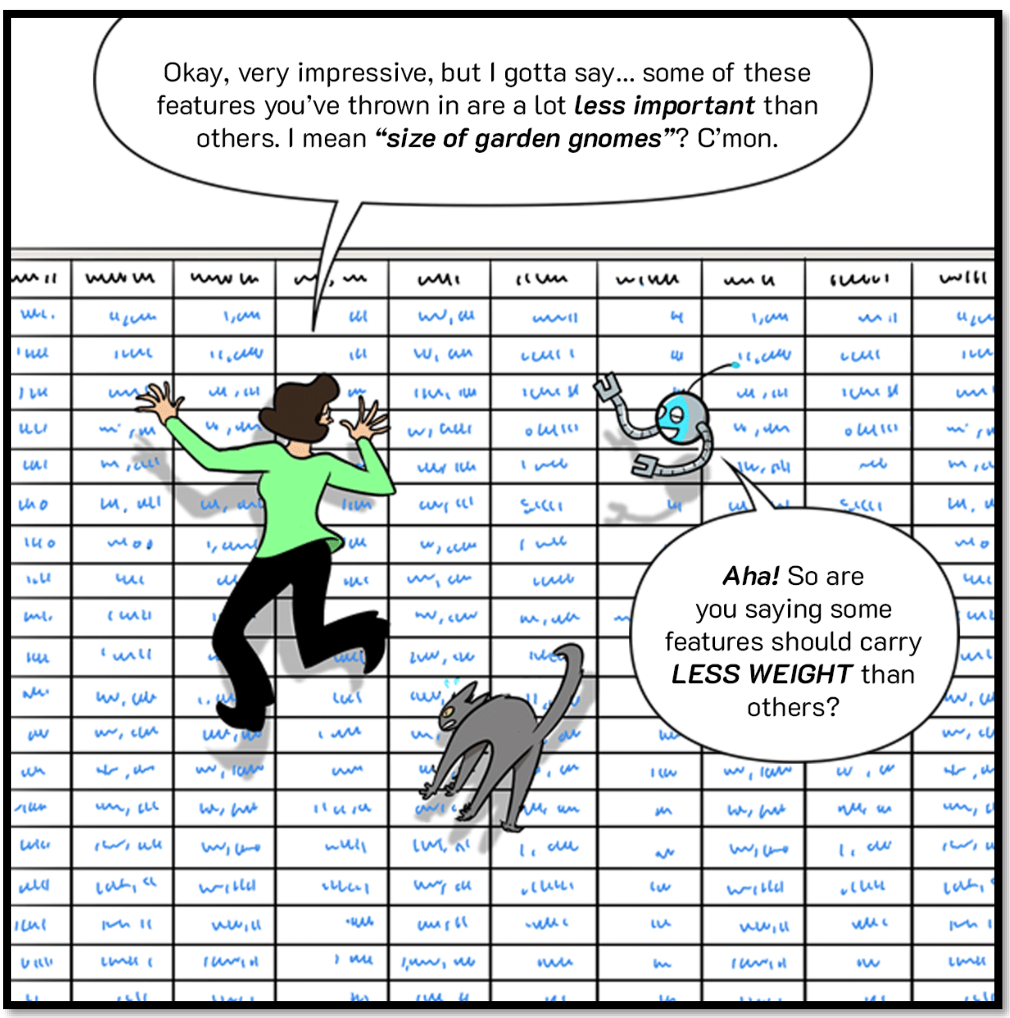
Okay, very impressive, but I gotta say... some of these features you've thrown in are a lot less important than others. I mean "size of garden gnomes"? C'mon.
Aha! So are you saying some features should carry LESS WEIGHT than others?

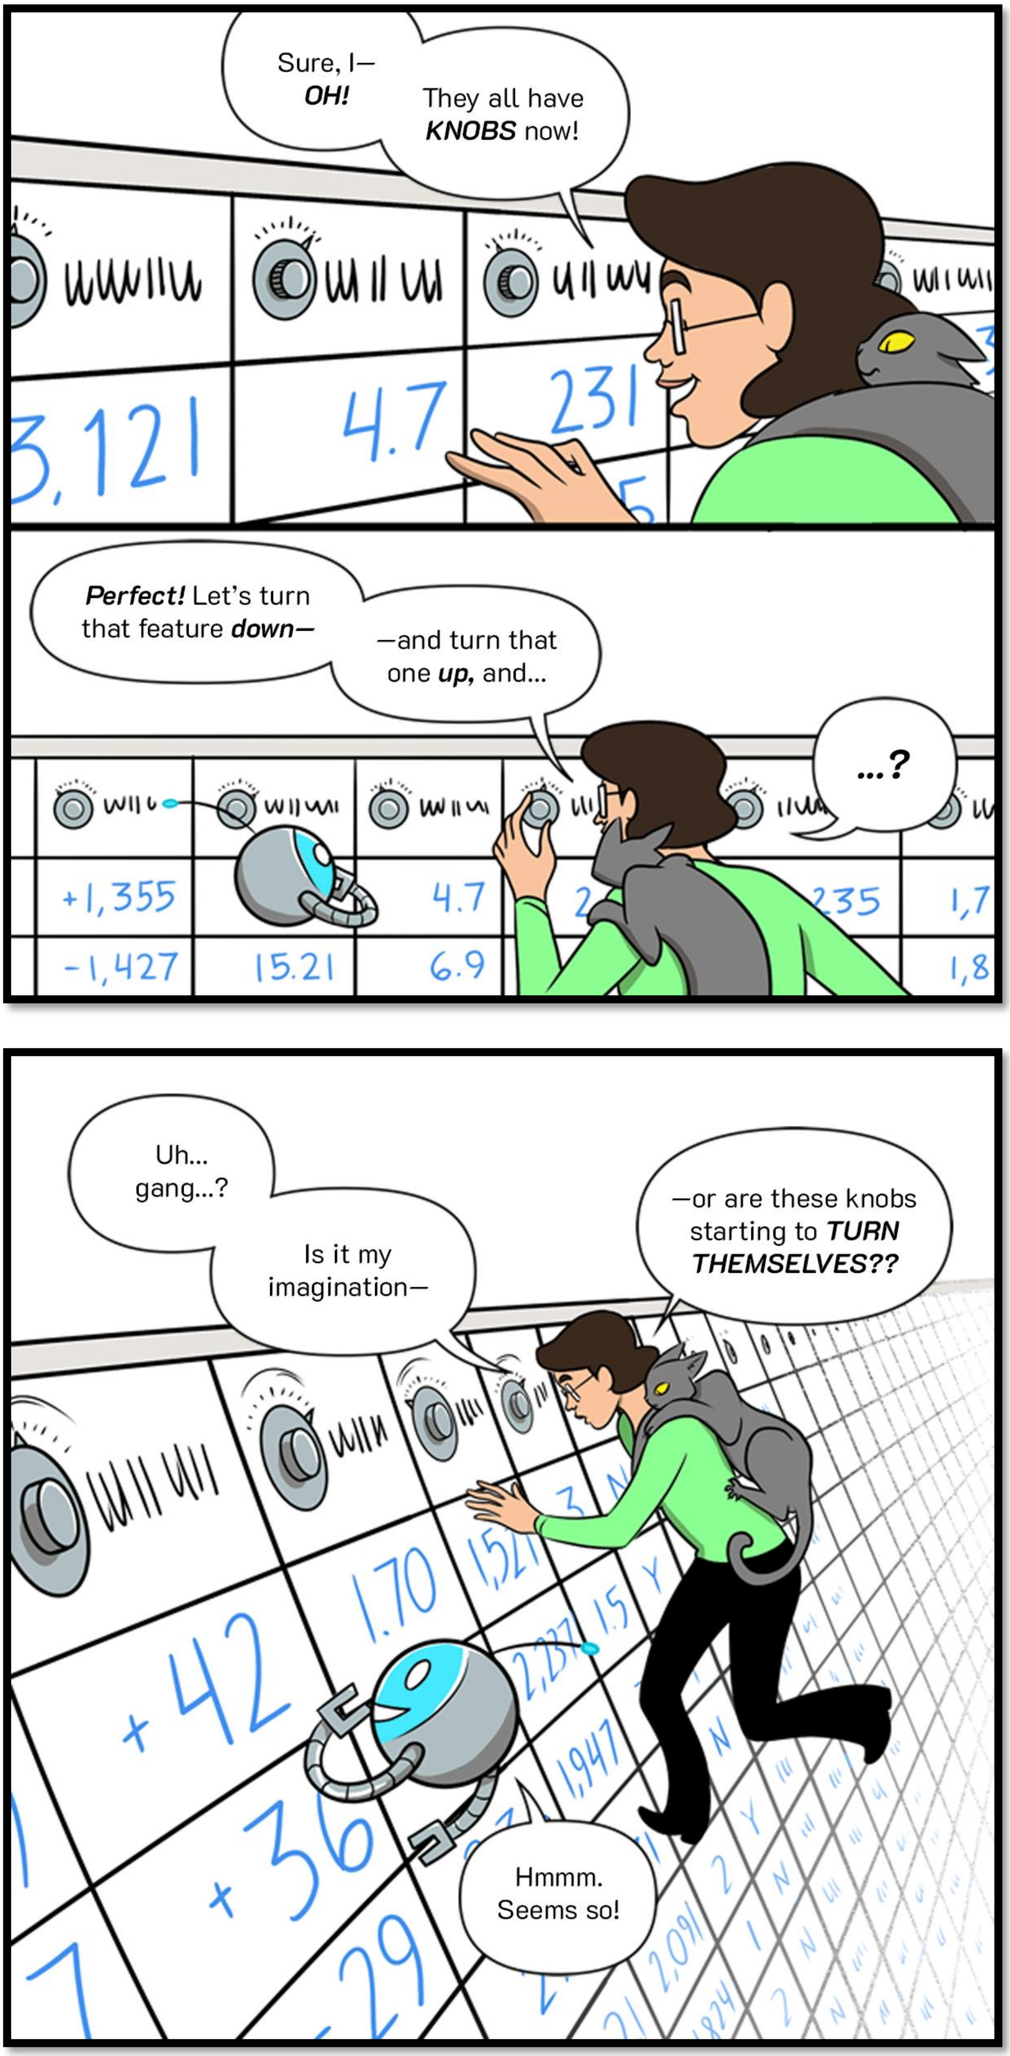
Sure, I— OH!
They all have KNOBS now!
Perfect! Let's turn that feature down—
—and turn that one up, and...
...?
Uh... gang...?
Is it my imagination—
—or are these knobs starting to TURN THEMSELVES??
Hmmm. Seems so!

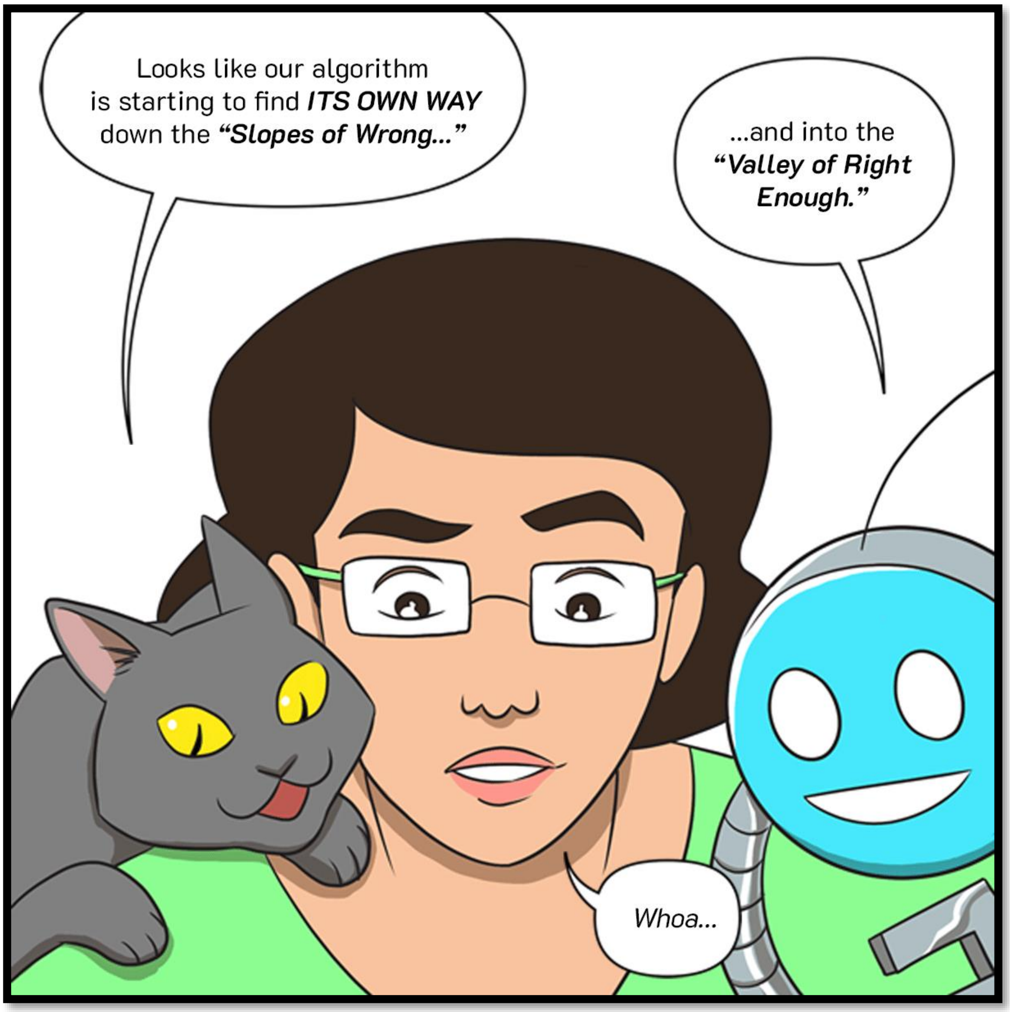
Looks like our algorithm is starting to find ***ITS OWN WAY*** down the ***"Slopes of Wrong..."***
...and into the ***"Valley of Right Enough."***
Whoa...

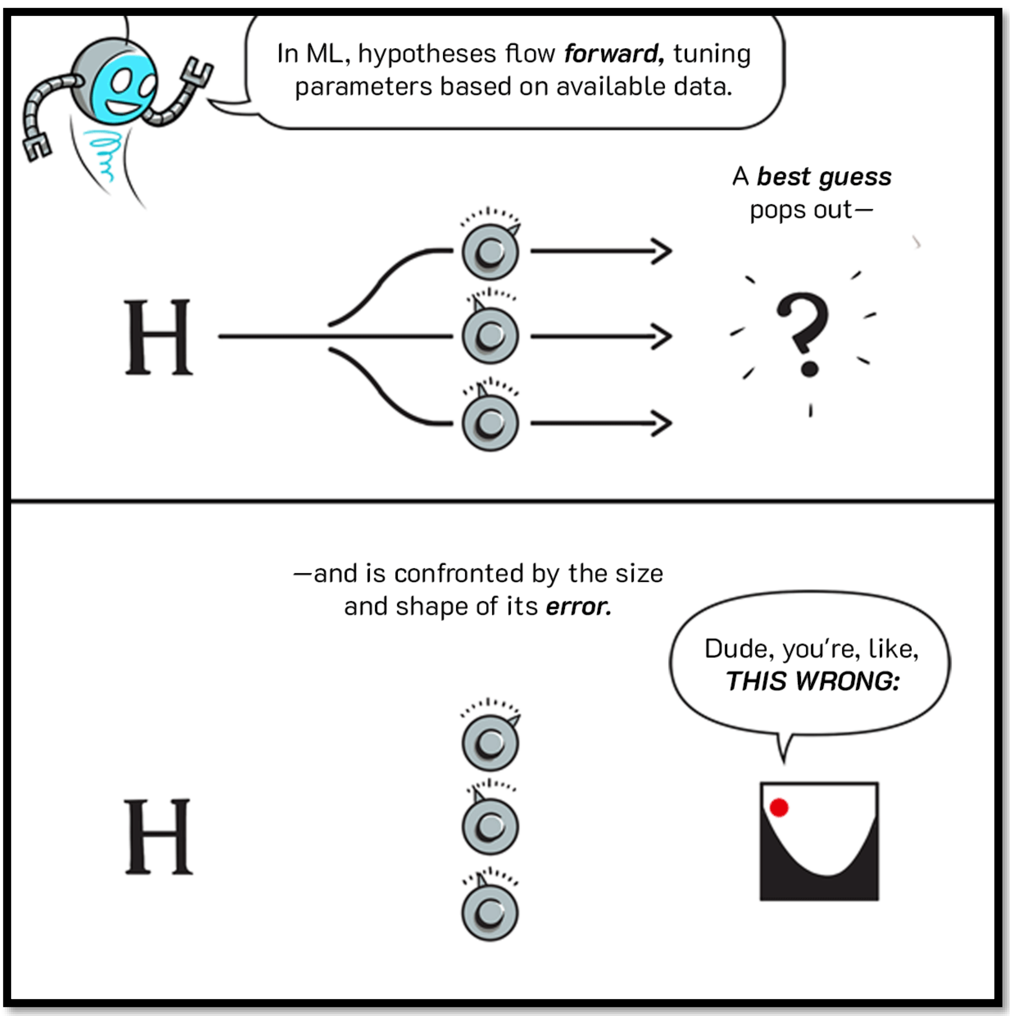
In ML, hypotheses flow ***forward,*** tuning parameters based on available data.
A ***best guess*** pops out—
H
?
—and is confronted by the size and shape of its ***error.***
Dude, you're, like, ***THIS WRONG:***
H

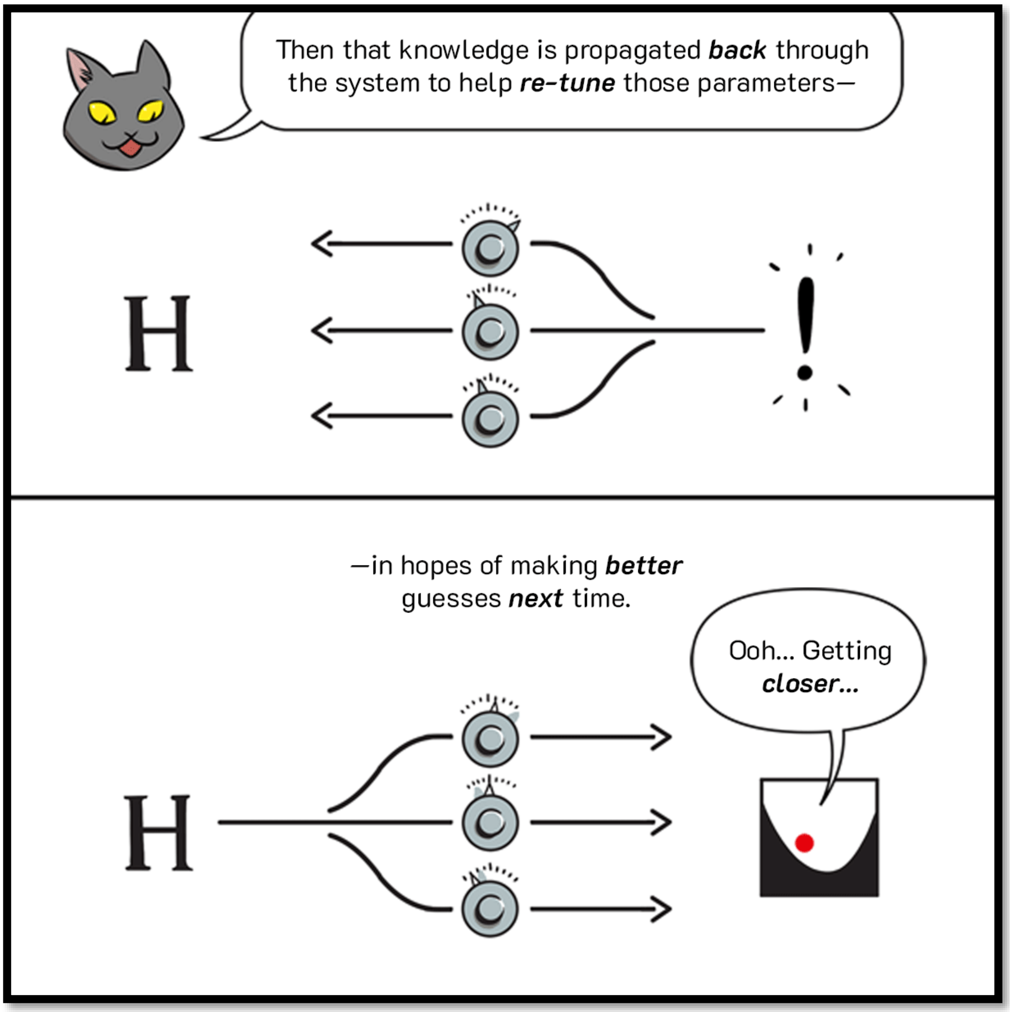
Then that knowledge is propagated ***back*** through the system to help ***re-tune*** those parameters—
H
!
—in hopes of making ***better*** guesses ***next*** time.
Ooh... Getting ***closer...***
H

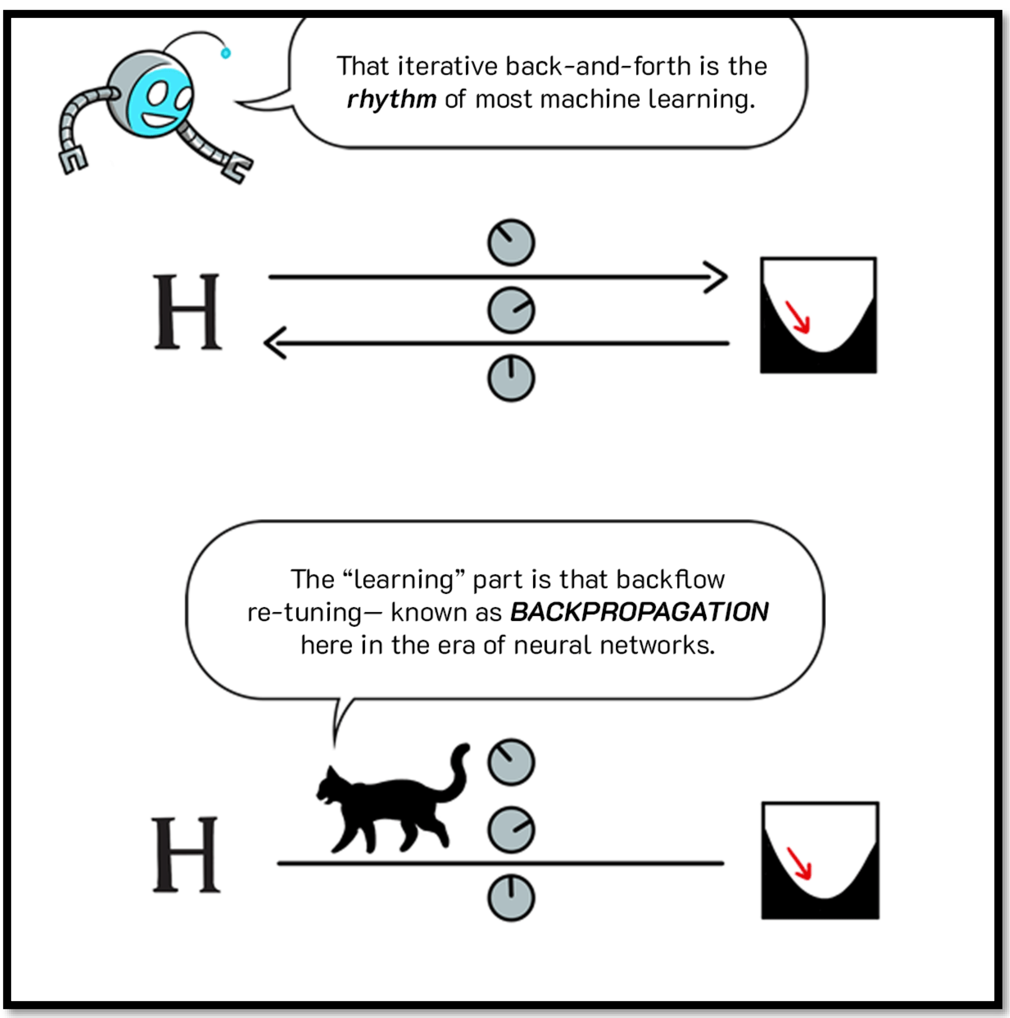
That iterative back-and-forth is the ***rhythm*** of most machine learning.
H
The "learning" part is that backflow re-tuning— known as ***BACKPROPAGATION*** here in the era of neural networks.
H

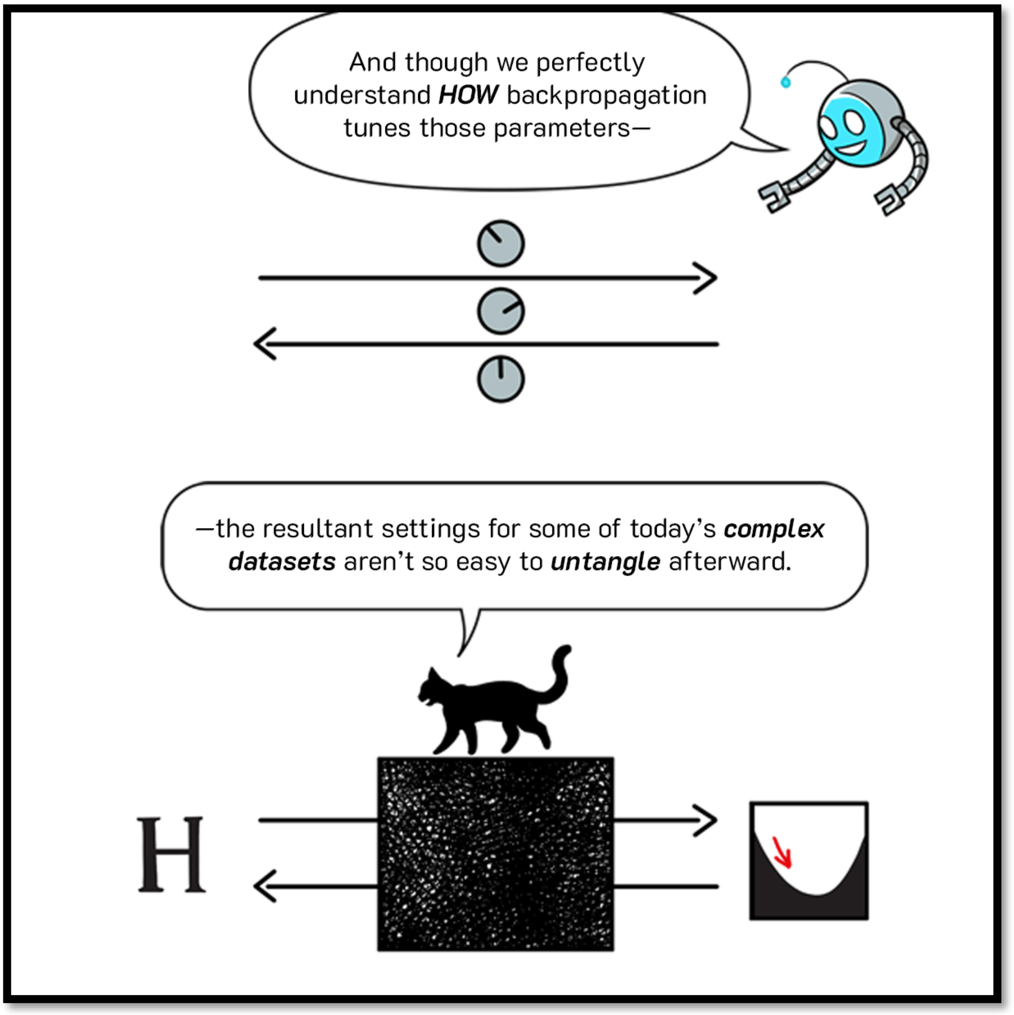
And though we perfectly understand ***HOW*** backpropagation tunes those parameters—
—the resultant settings for some of today's ***complex datasets*** aren't so easy to ***untangle*** afterward.
H

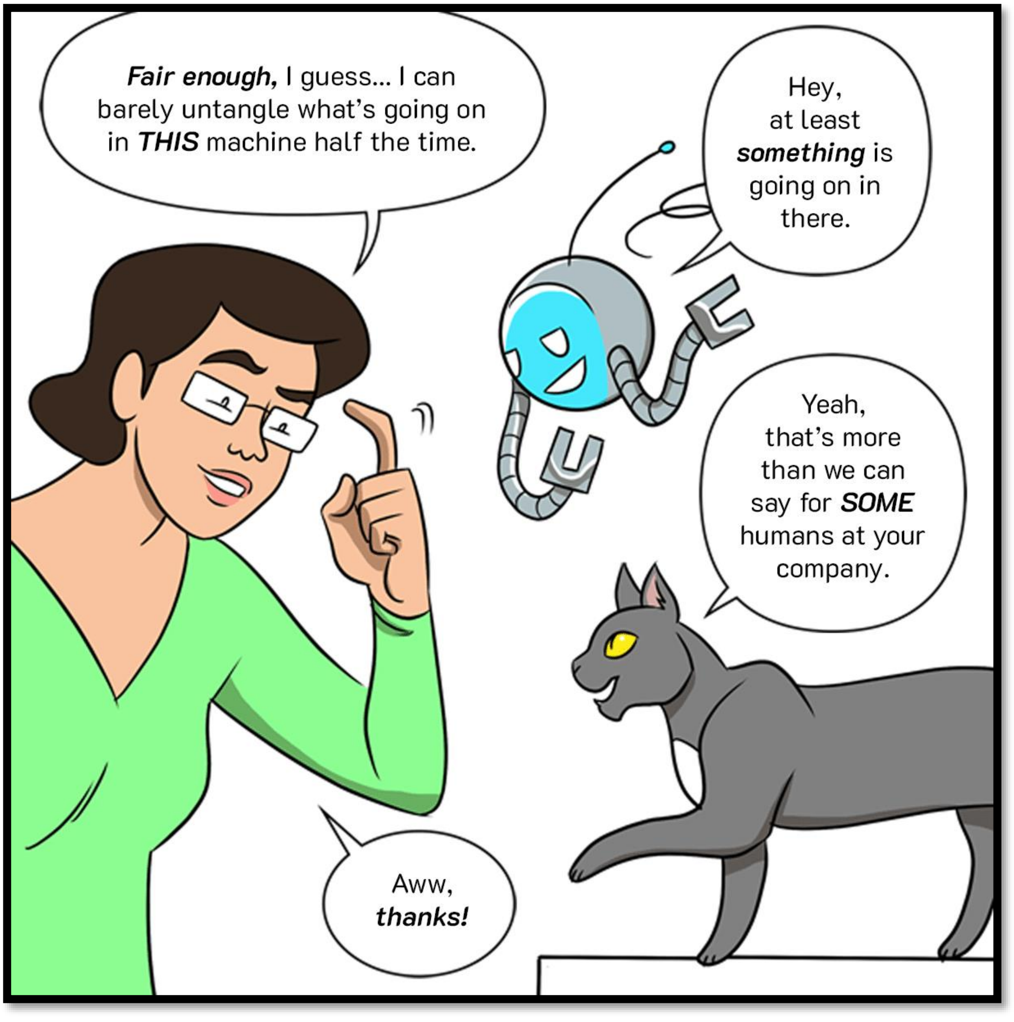
Fair enough, I guess... I can barely untangle what's going on in ***THIS*** machine half the time.
Hey, at least ***something*** is going on in there.
Yeah, that's more than we can say for ***SOME*** humans at your company.
Aww, ***thanks!***

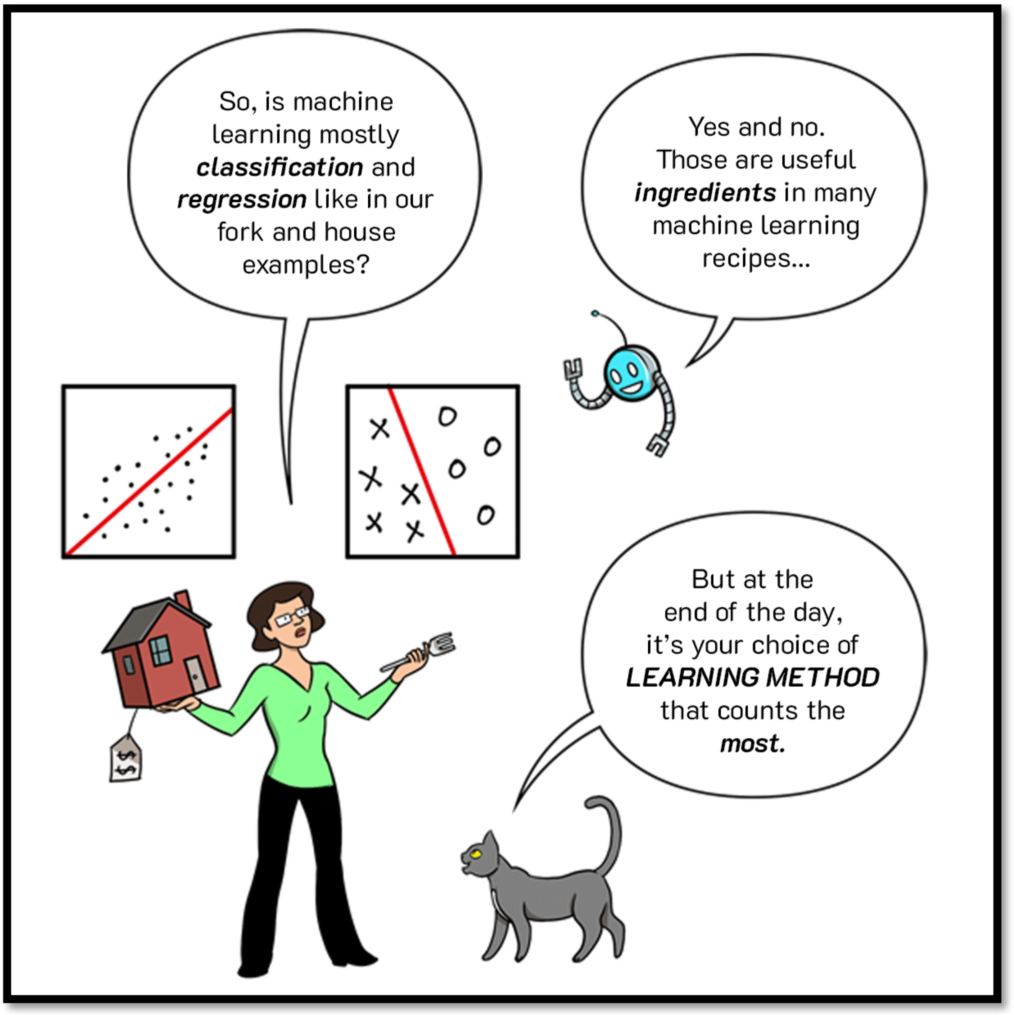
So, is machine learning mostly ***classification*** and ***regression*** like in our fork and house examples?
Yes and no. Those are useful ***ingredients*** in many machine learning recipes...
But at the end of the day, it's your choice of ***LEARNING METHOD*** that counts the ***most.***

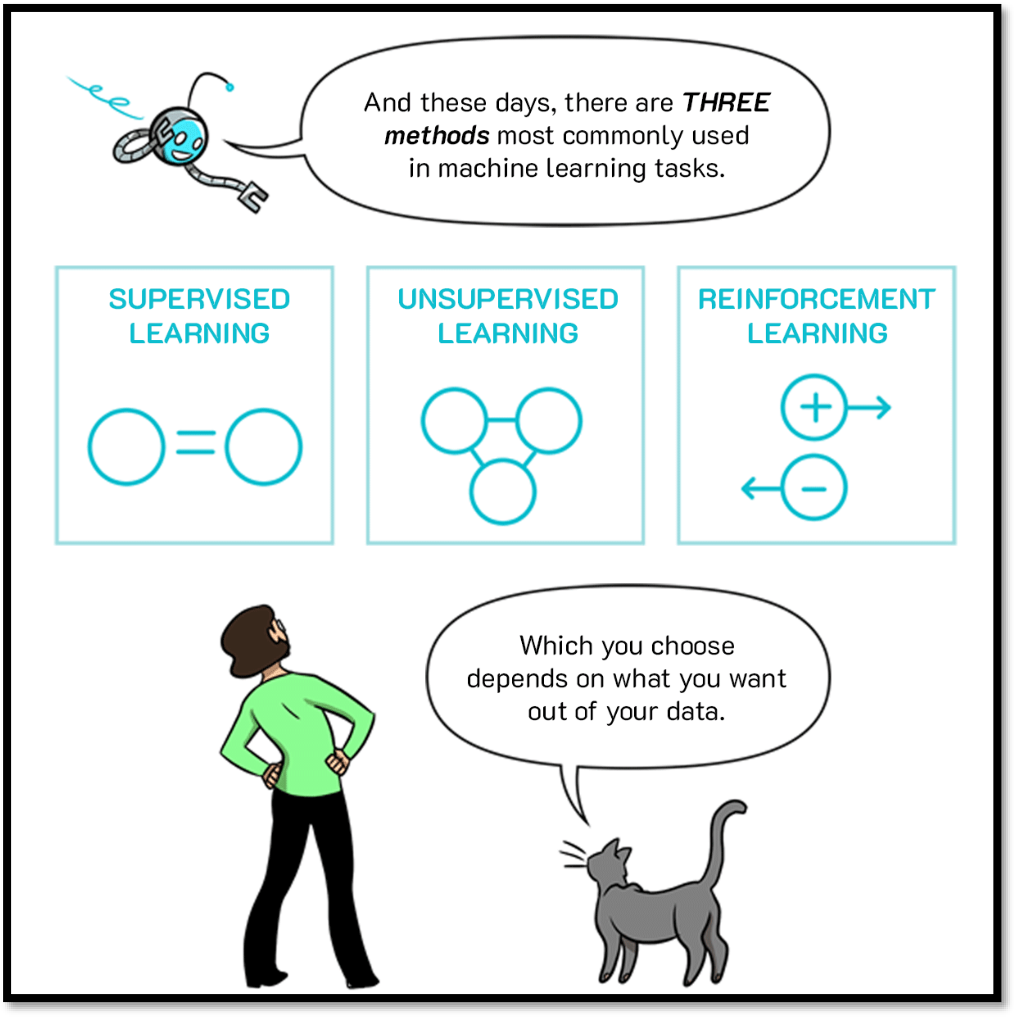
And these days, there are ***THREE methods*** most commonly used in machine learning tasks.
SUPERVISED LEARNING
UNSUPERVISED LEARNING
REINFORCEMENT LEARNING
Which you choose depends on what you want out of your data.

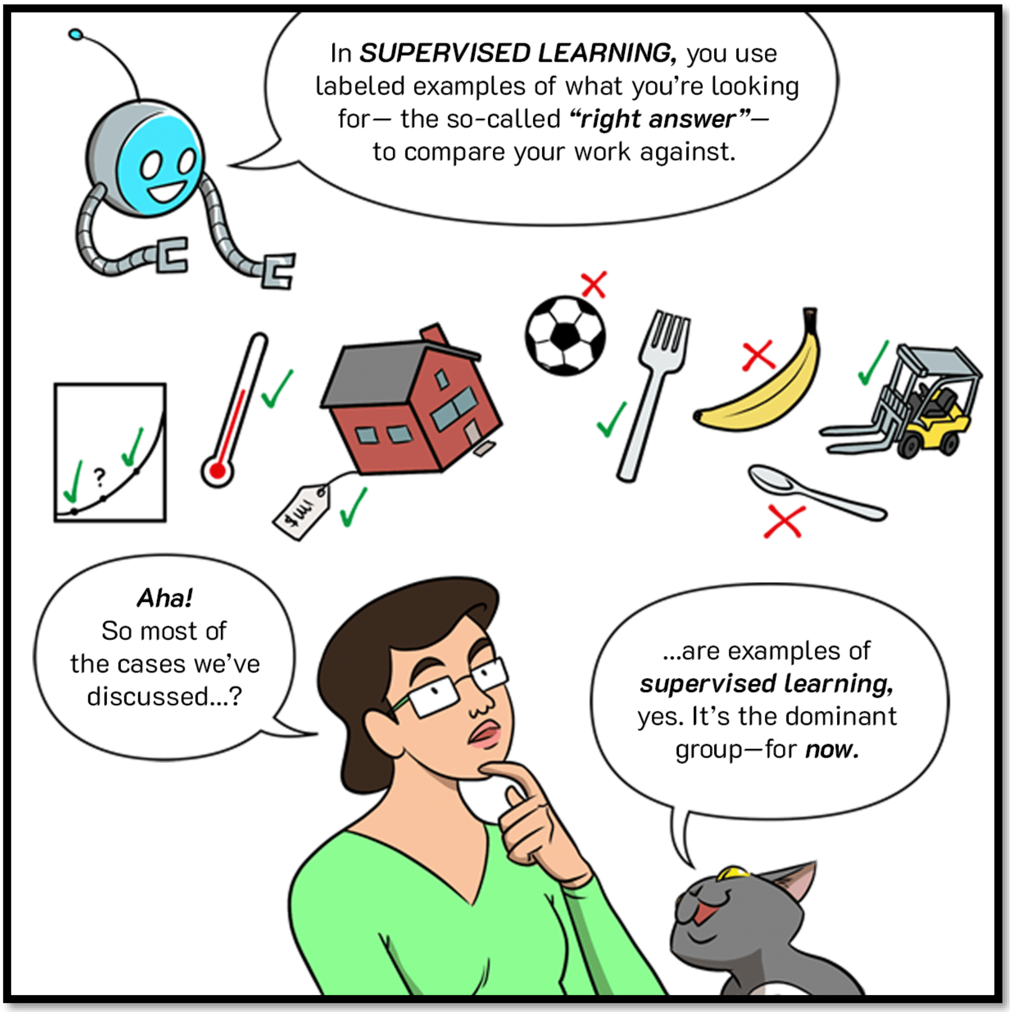
In ***SUPERVISED LEARNING,*** you use labeled examples of what you're looking for— the so-called ***"right answer"***— to compare your work against.
$m
?
Aha! So most of the cases we've discussed...?
...are examples of ***supervised learning,*** yes. It's the dominant group—for ***now.***

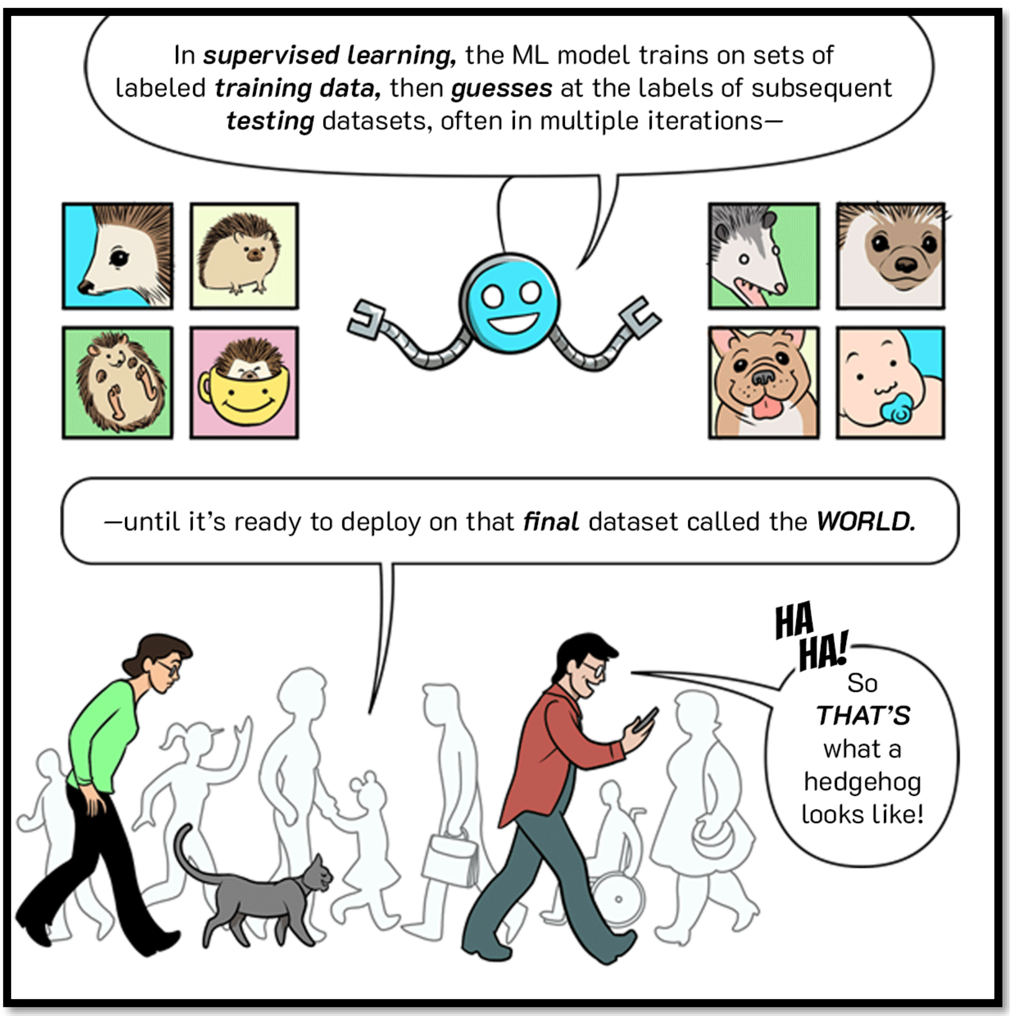
In ***supervised learning,*** the ML model trains on sets of labeled ***training data,*** then ***guesses*** at the labels of subsequent ***testing*** datasets, often in multiple iterations—
—until it's ready to deploy on that ***final*** dataset called the ***WORLD.***
HA HA!
So ***THAT'S*** what a hedgehog looks like!

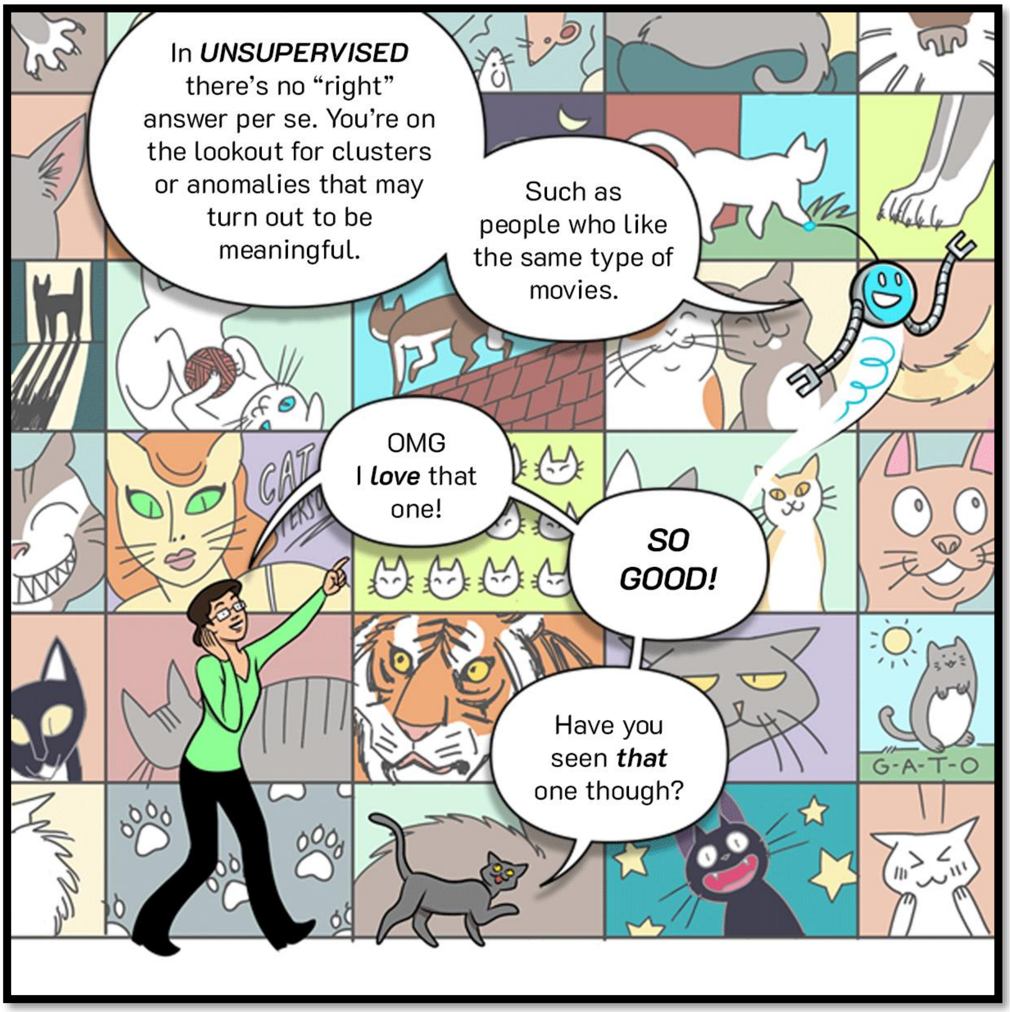
In ***UNSUPERVISED*** there's no "right" answer per se. You're on the lookout for clusters or anomalies that may turn out to be meaningful.
Such as people who like the same type of movies.
OMG I ***love*** that one!
SO GOOD!
Have you seen ***that*** one though?
G-A-T-O


ALL MOVIES
ML is a great fit for this kind of work because it can take on a huge—
—or hugely ***complex,*** dataset.
OMG I ***love*** that one!
私も又!

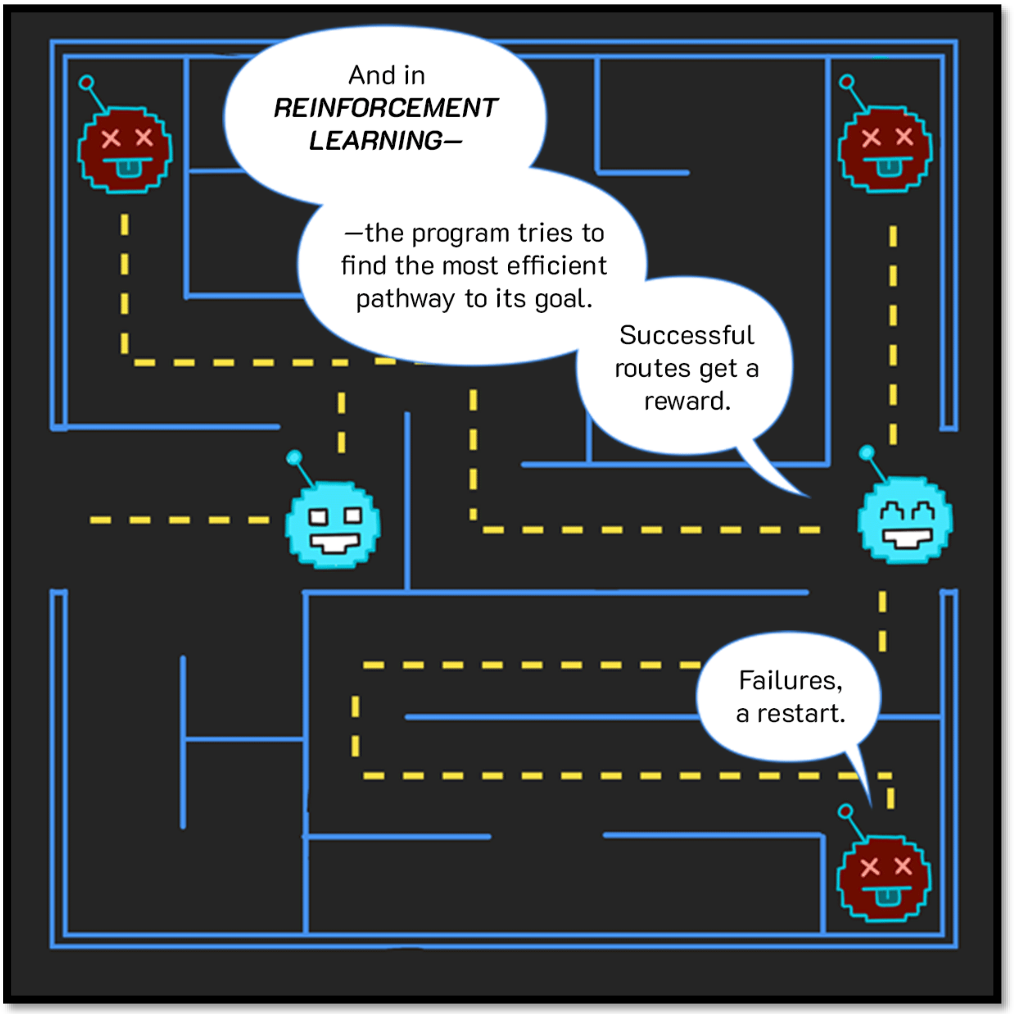

https://deepmind.com/research/publications/playing-atari-deep-reinforcement-learning

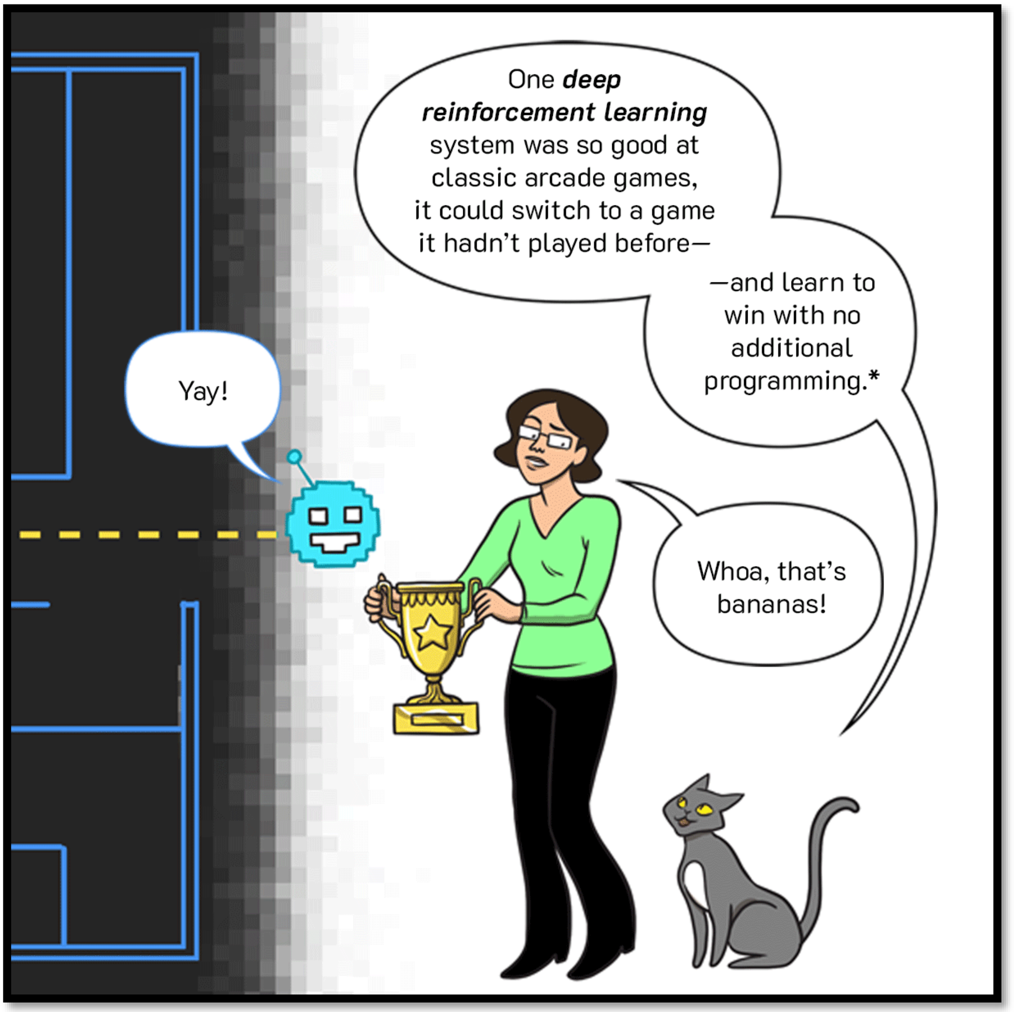

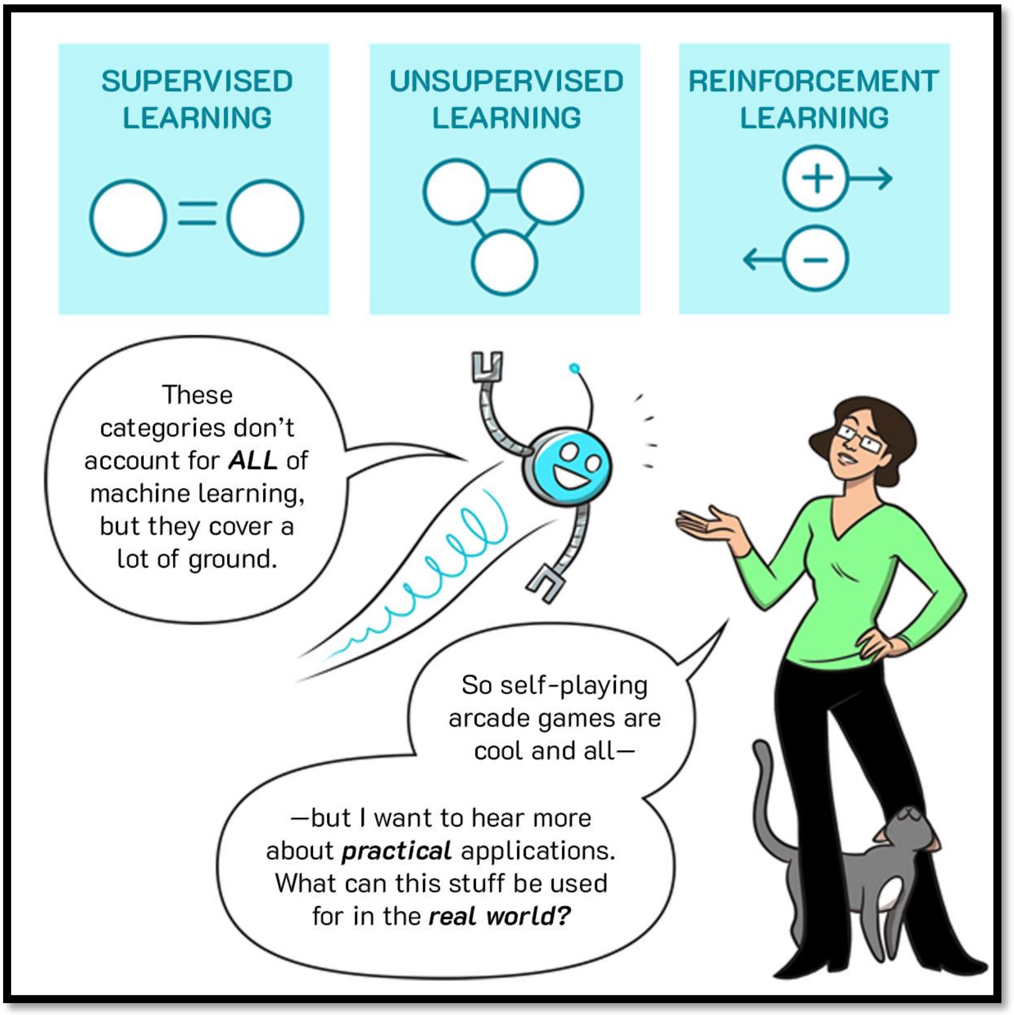
SUPERVISED LEARNING
UNSUPERVISED LEARNING
REINFORCEMENT LEARNING
These categories don't account for ***ALL*** of machine learning, but they cover a lot of ground.
So self-playing arcade games are cool and all—
—but I want to hear more about ***practical*** applications. What can this stuff be used for in the ***real world?***

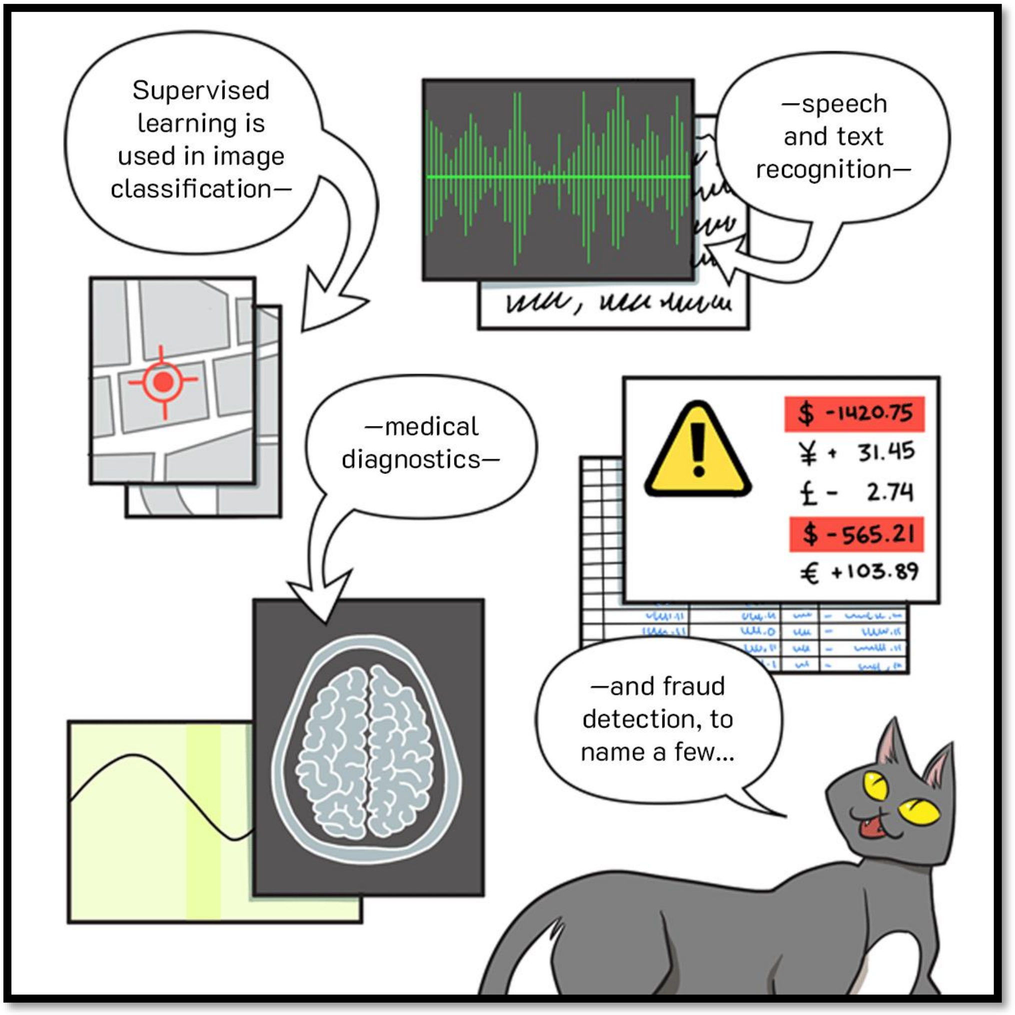
Supervised learning is used in image classification—
—speech and text recognition—
—medical diagnostics—
$ -1420.75
¥ + 31.45
£ - 2.74
$ - 565.21
€ +103.89
—and fraud detection, to name a few...

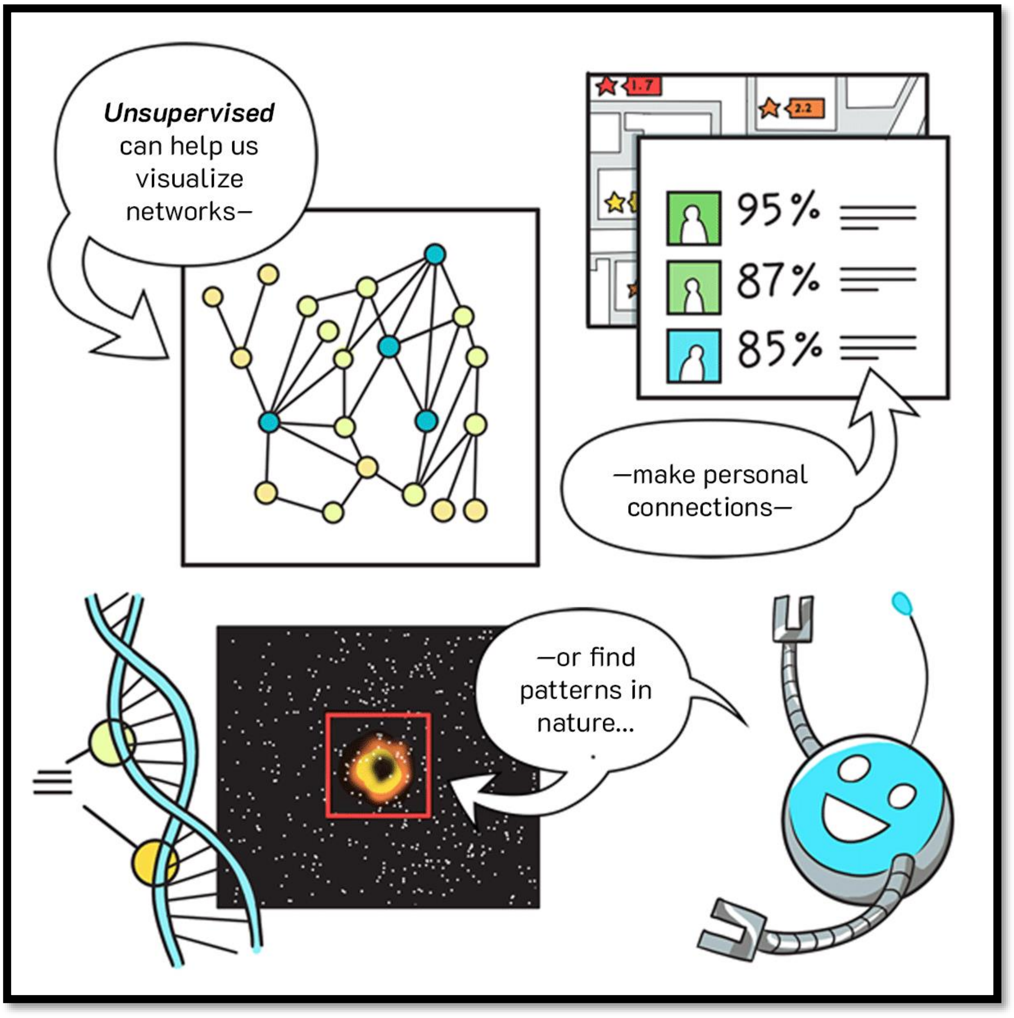
Unsupervised can help us visualize networks—
1.7
2.2
95%
87%
85%
—make personal connections—
—or find patterns in nature...

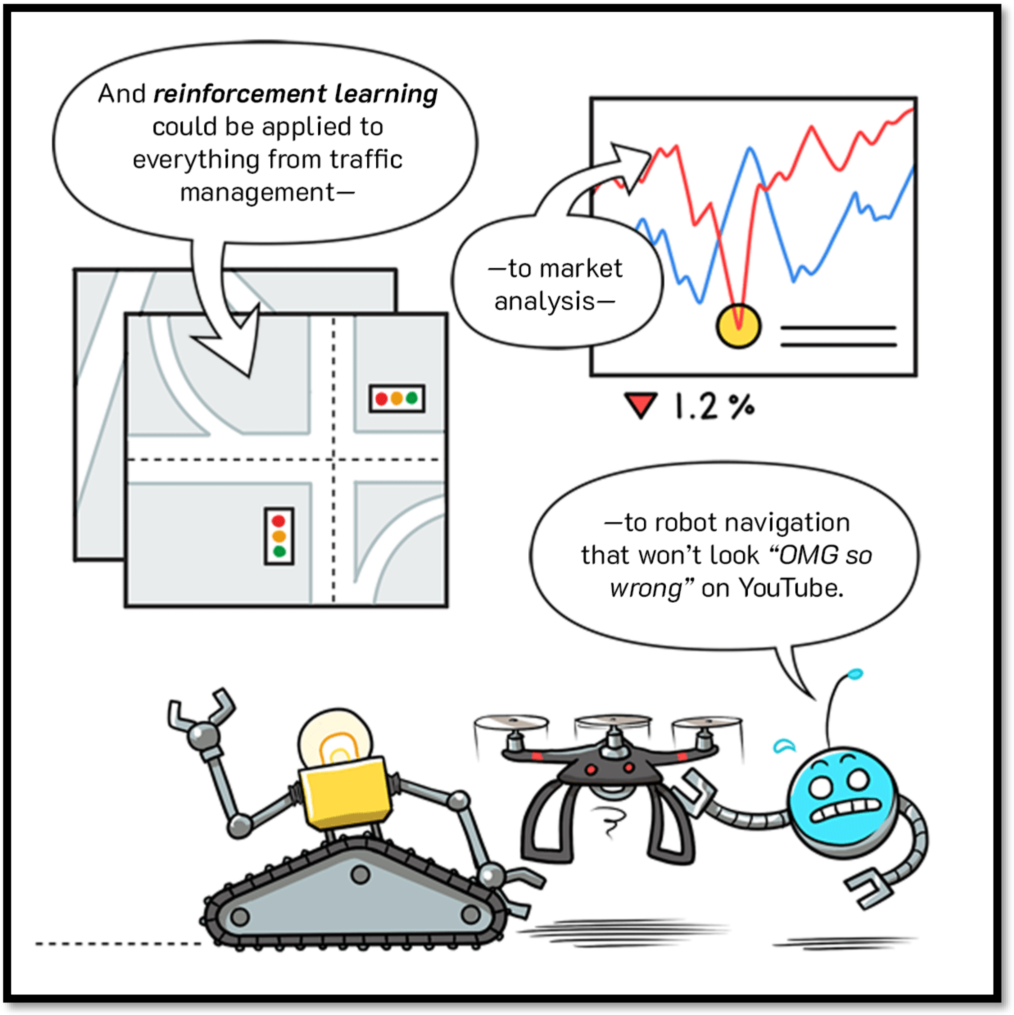
And ***reinforcement learning*** could be applied to everything from traffic management—
—to market analysis—
1.2%
—to robot navigation that won't look *"OMG so wrong"* on YouTube.


But no matter the kind of ML, ***GOOD DATA*** is where it all begins.
Say you're using supervised learning to identify photos of, oh, I dunno...
...CATS!
First you have to gather data that contain the information you're looking for.


So, I'd need to think about... dimensions, resolution, aspect ratio...
Close-ups versus long shots...
From what angles?
Isolated or in an environment?
Are ***kittens*** included?
Tigers?
Exactly. You need images that encapsulate the essentials of "felinity."

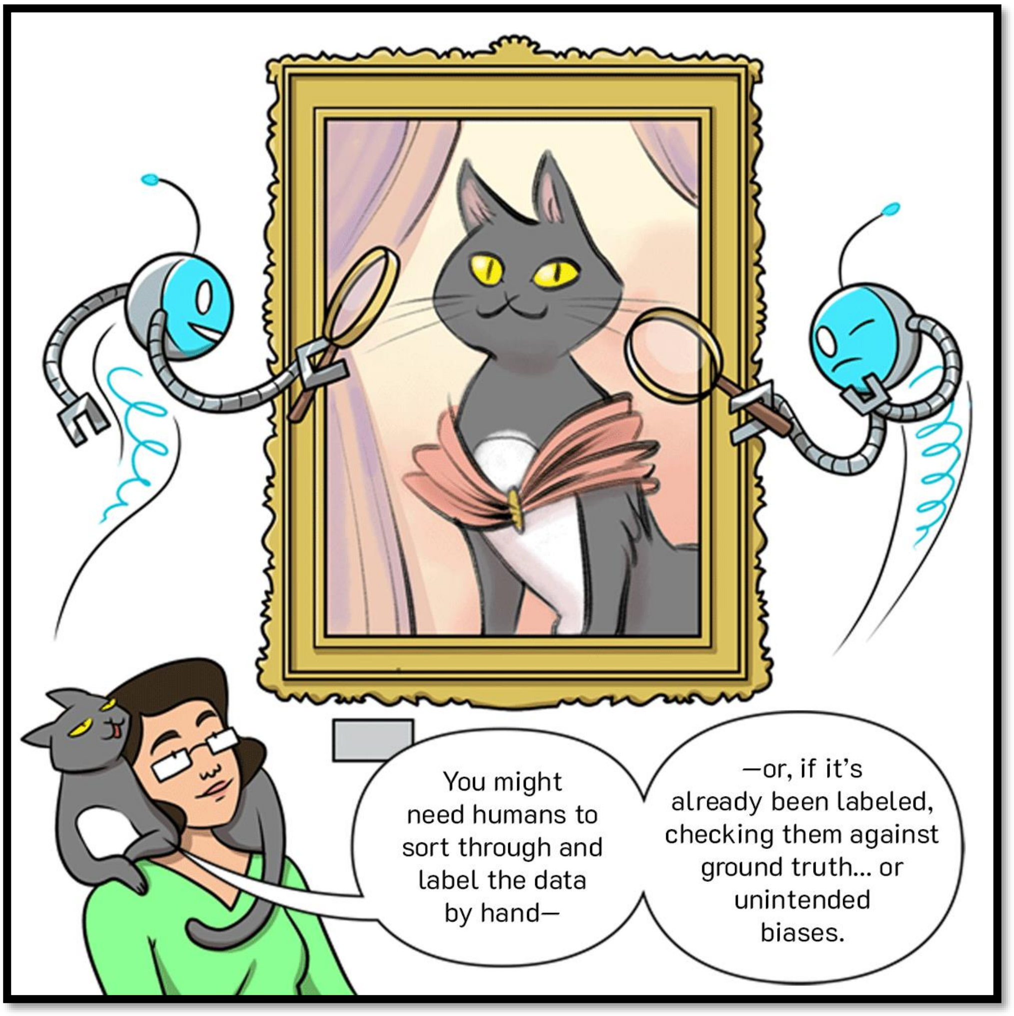
You might need humans to sort through and label the data by hand—
—or, if it's already been labeled, checking them against ground truth... or unintended biases.


Once you have well-curated and prepared data, you split it up for use in training, testing, and validation.
It might require a ton of work to collect and label that much data, but it may be worth it.
TRAINING
TESTING
VALIDATION

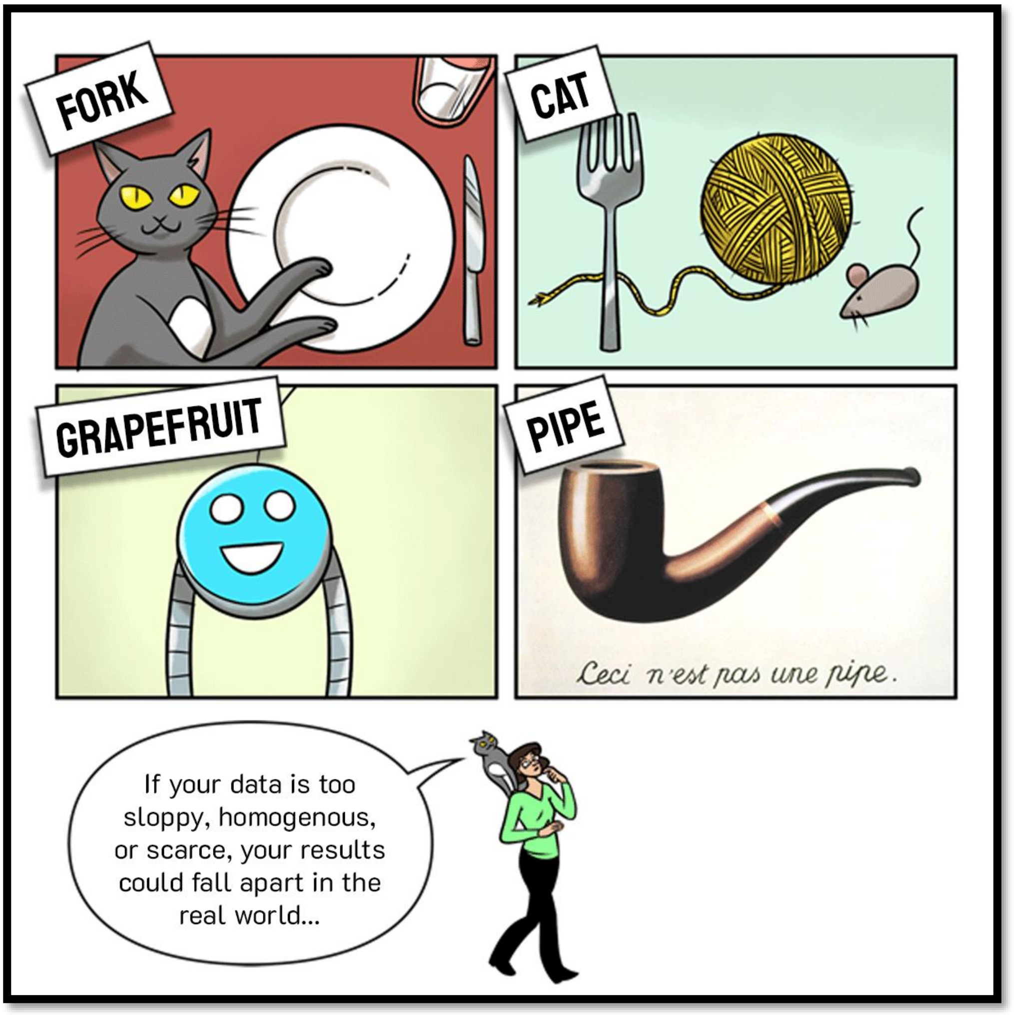
FORK
CAT
GRAPEFRUIT
PIPE
Ceci n'est pas une pipe.
If your data is too sloppy, homogenous, or scarce, your results could fall apart in the real world...

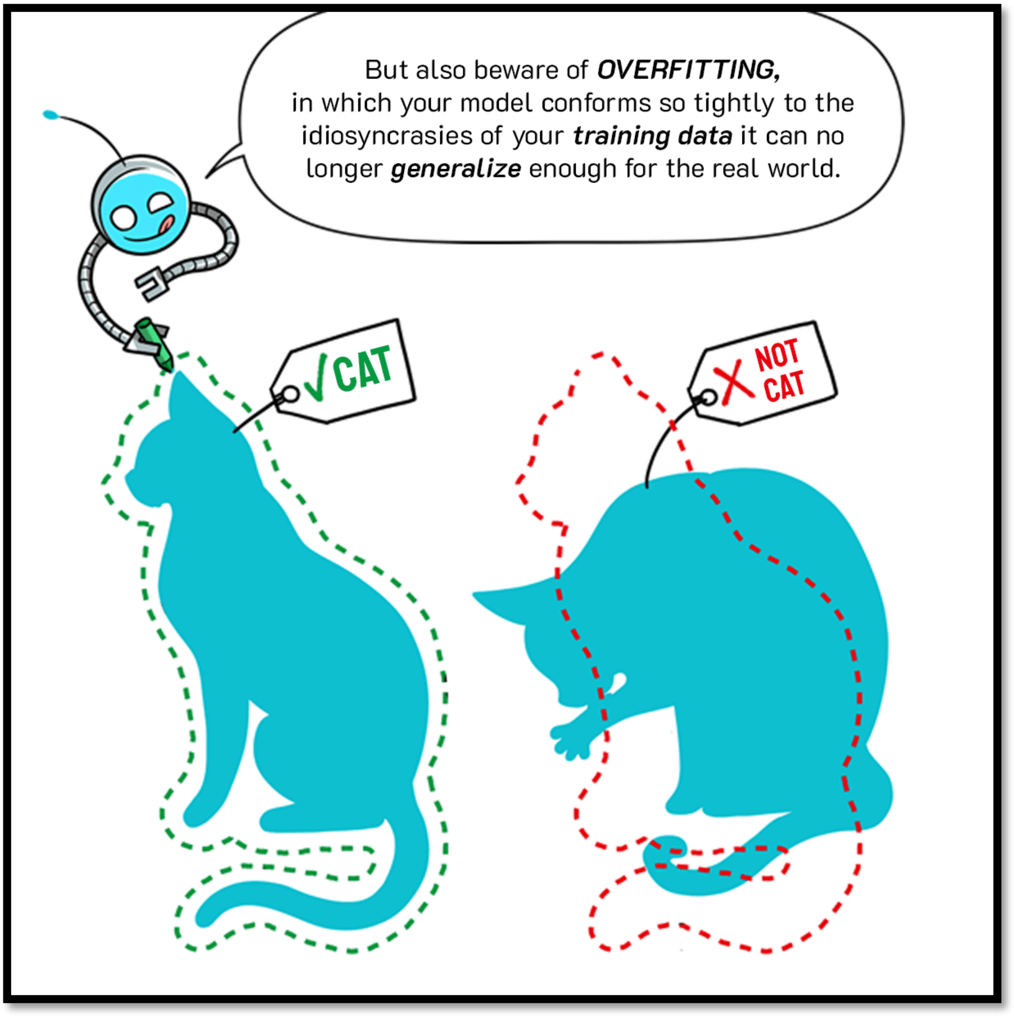
But also beware of OVERFITTING, in which your model conforms so tightly to the idiosyncrasies of your training data it can no longer generalize enough for the real world.
√CAT
NOT CAT

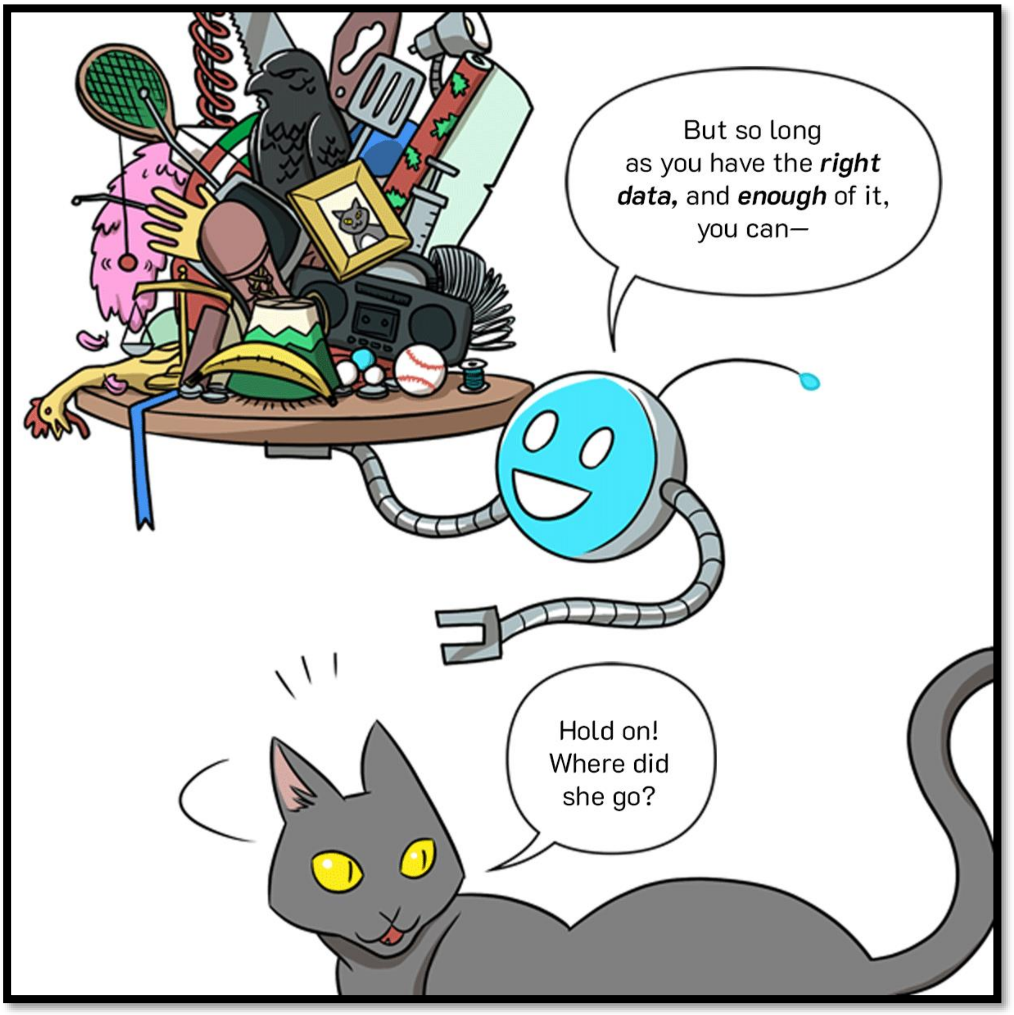
But so long as you have the ***right data,*** and ***enough*** of it, you can—
Hold on! Where did she go?

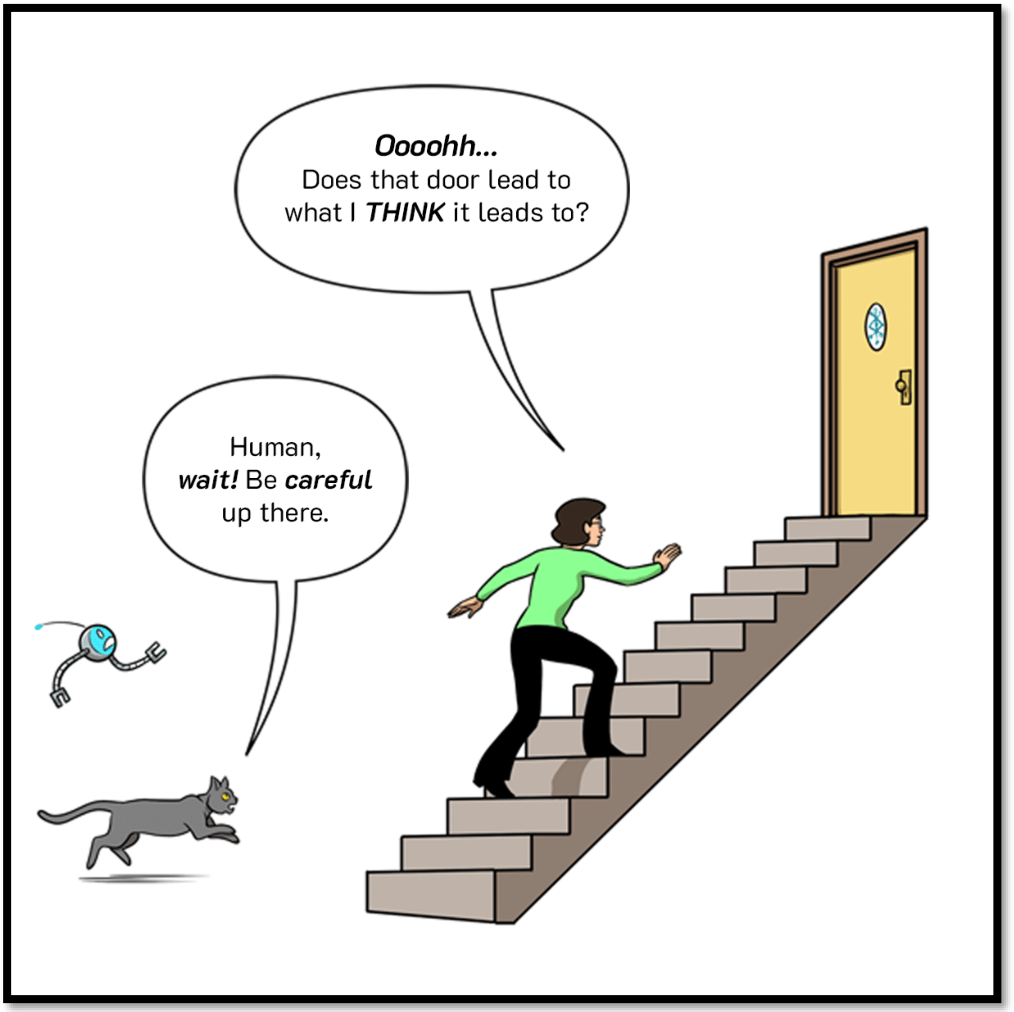
Oooohh... Does that door lead to what I ***THINK*** it leads to?
Human, ***wait!*** Be ***careful*** up there.

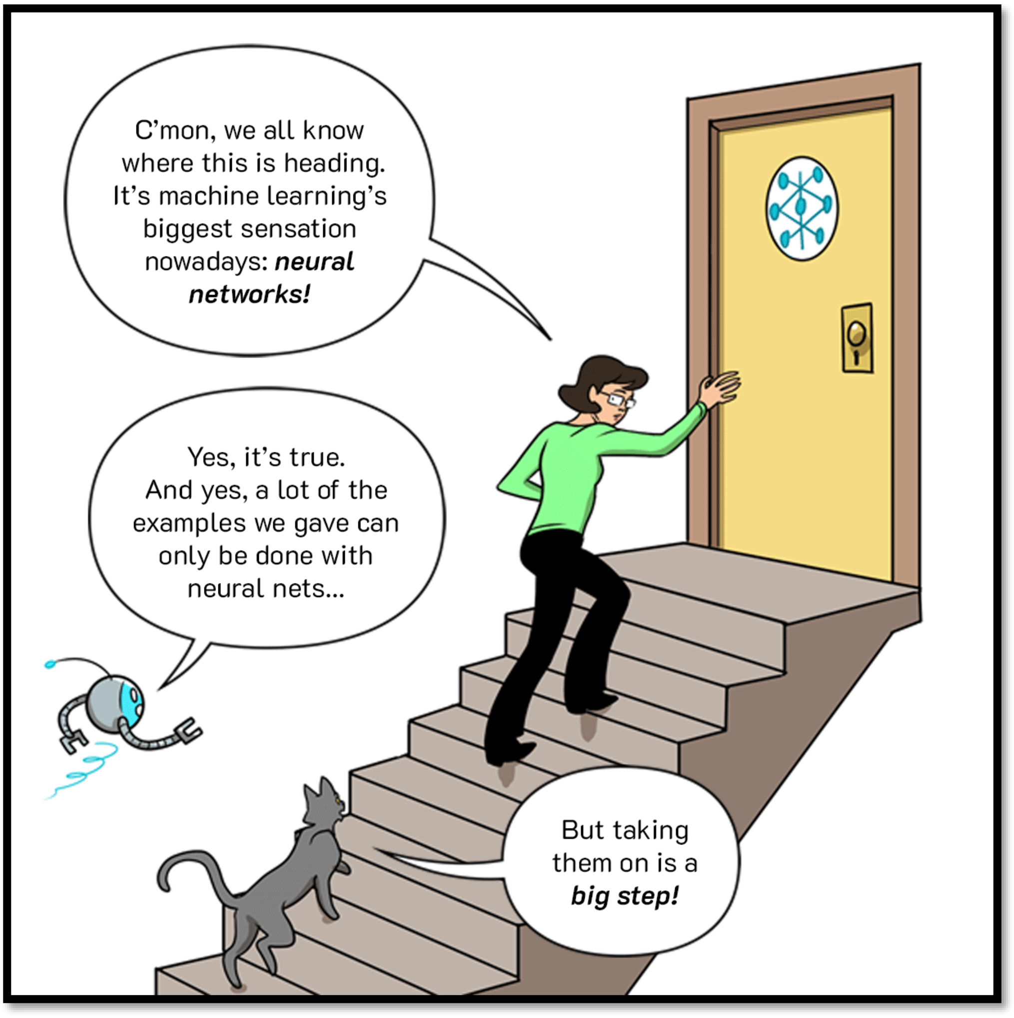
C'mon, we all know where this is heading. It's machine learning's biggest sensation nowadays: ***neural networks!***
Yes, it's true. And yes, a lot of the examples we gave can only be done with neural nets...
But taking them on is a ***big step!***

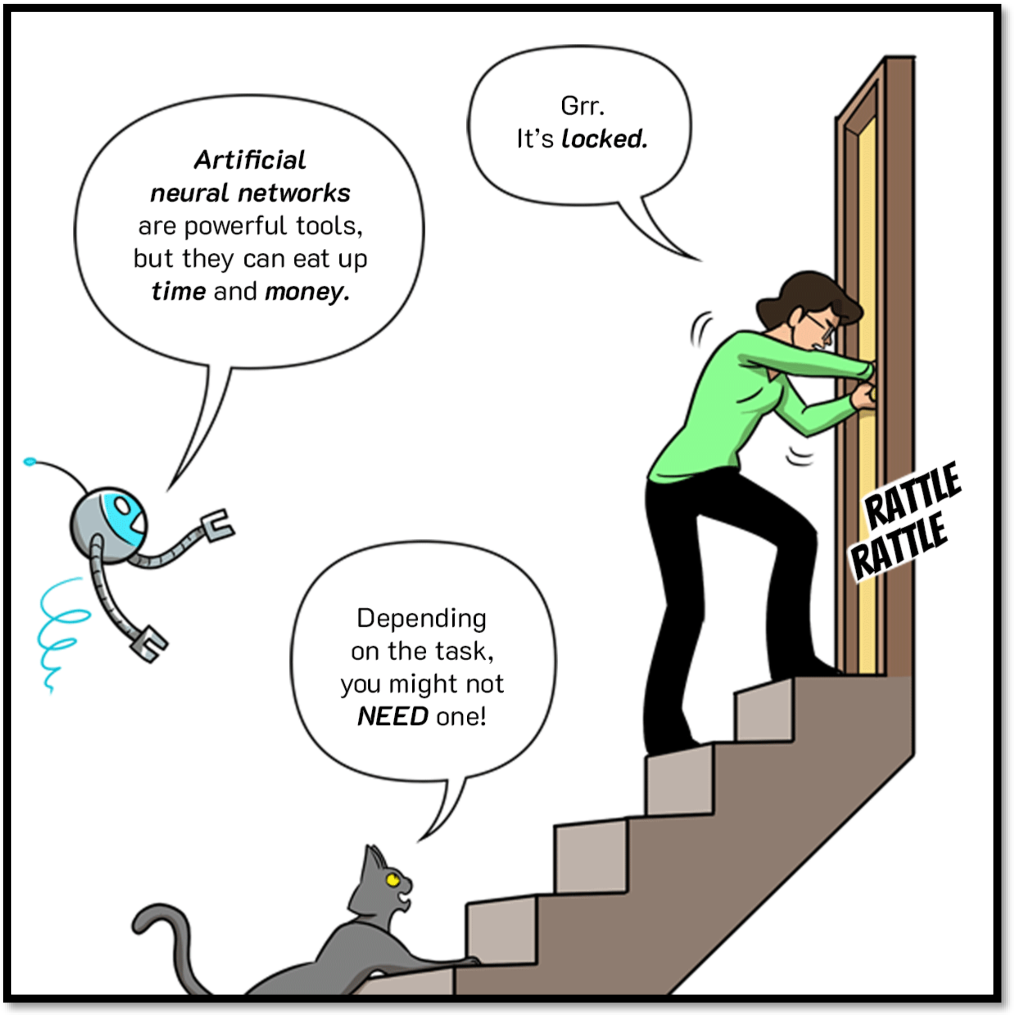
Artificial neural networks are powerful tools, but they can eat up ***time*** and ***money.***
Grr. It's ***locked.***
RATTLE RATTLE
Depending on the task, you might not ***NEED*** one!

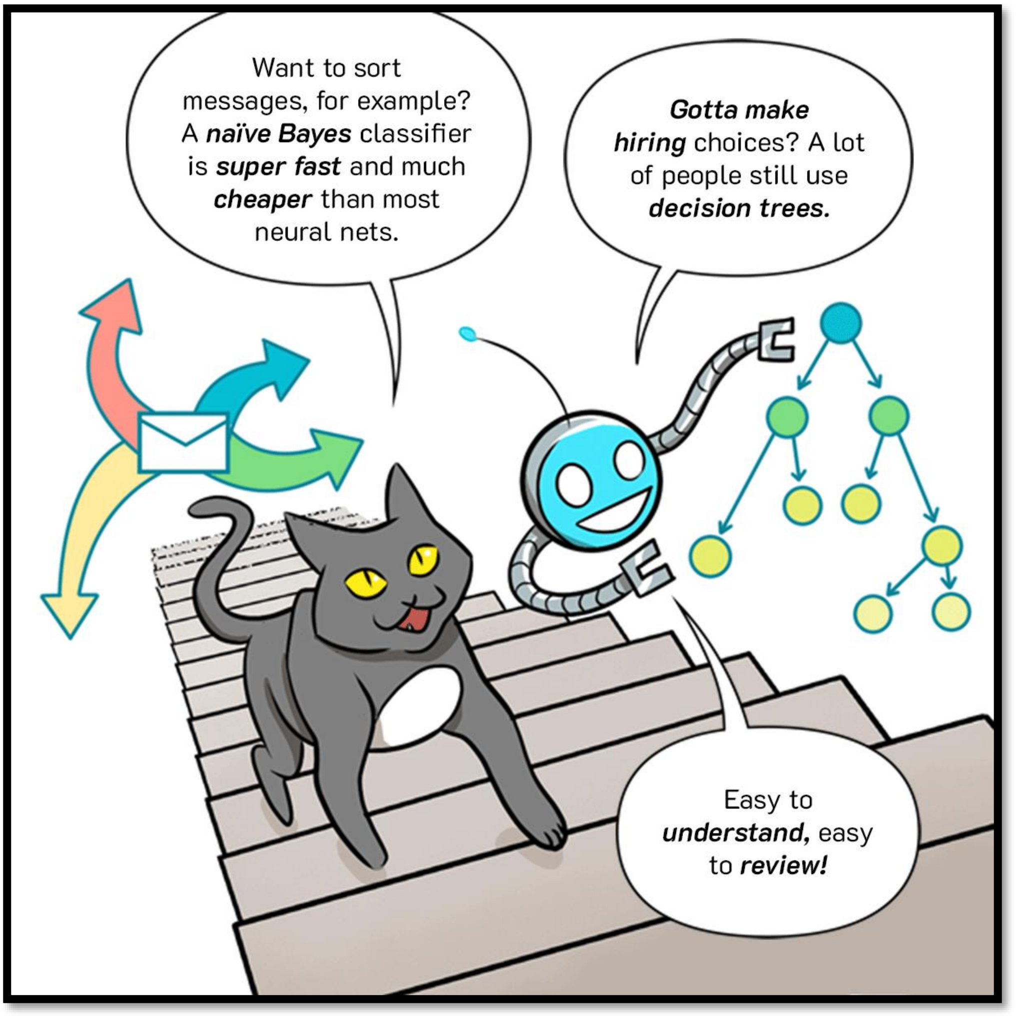
Want to sort messages, for example? A ***naïve Bayes*** classifier is ***super fast*** and much ***cheaper*** than most neural nets.
Gotta make hiring choices? A lot of people still use ***decision trees.***
Easy to ***understand,*** easy to ***review!***

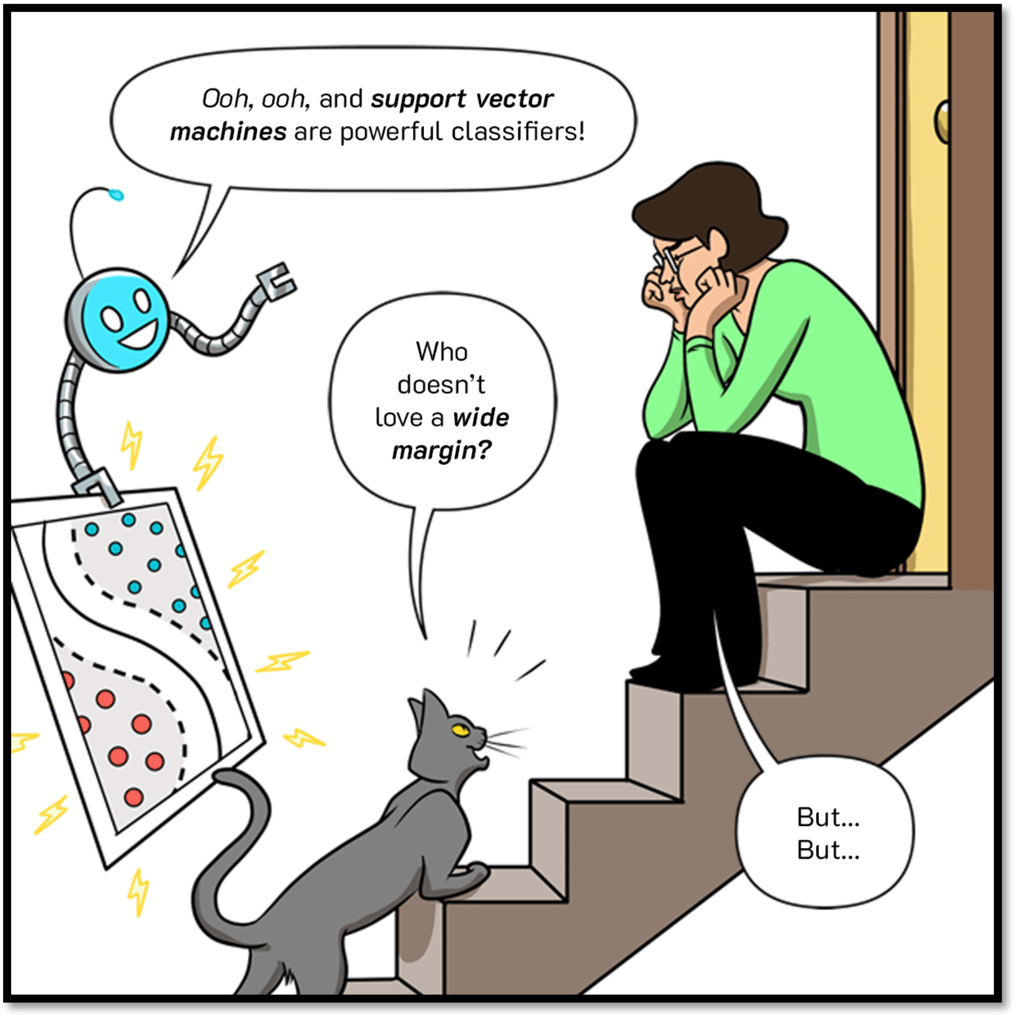
Ooh, *ooh*, and ***support vector machines*** are powerful classifiers!
Who doesn't love a ***wide margin?***
But...
But...


I still wanna LEARN about neural networks, because... because they...
Because they're cool?
Yeah...
...because they're cool.
They ARE pretty cool!

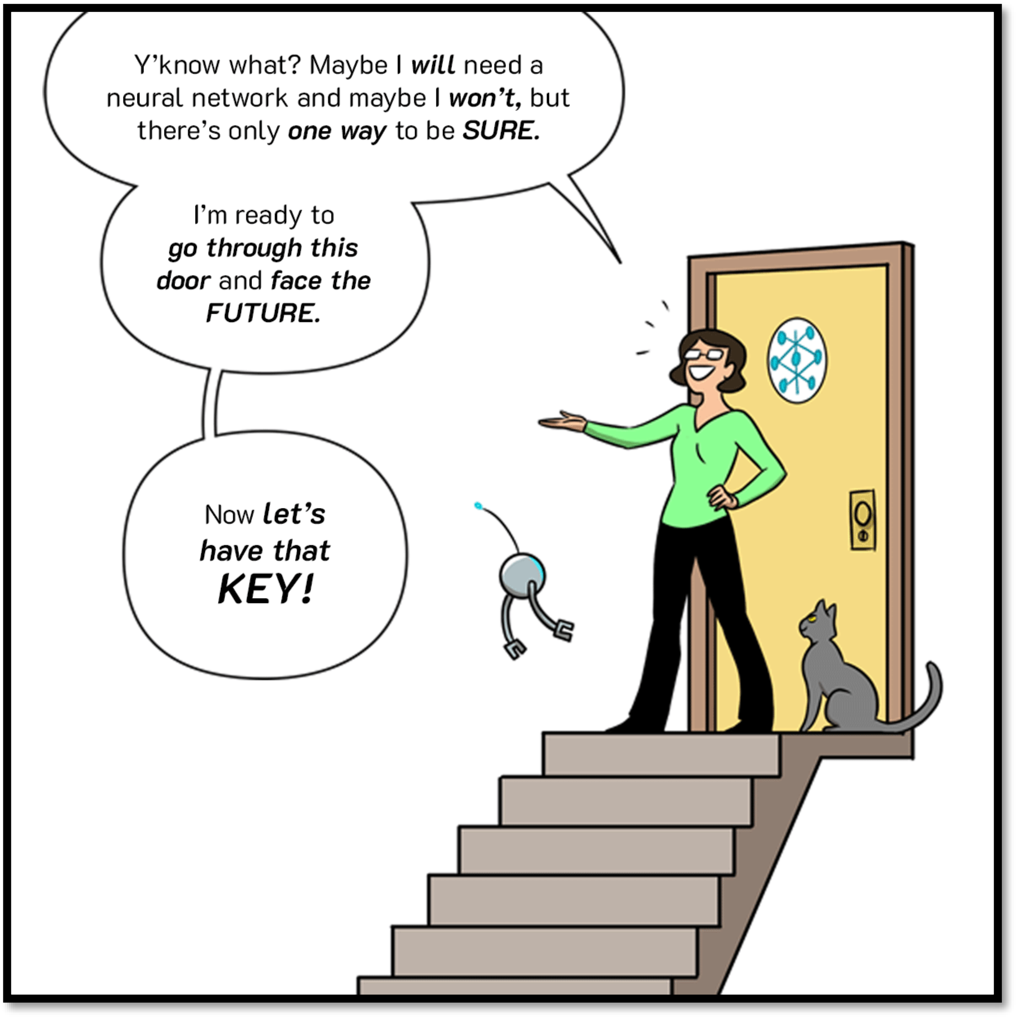
Y'know what? Maybe I will need a neural network and maybe I won't, but there's only one way to be SURE.
I'm ready to go through this door and face the FUTURE.
Now let's have that KEY!

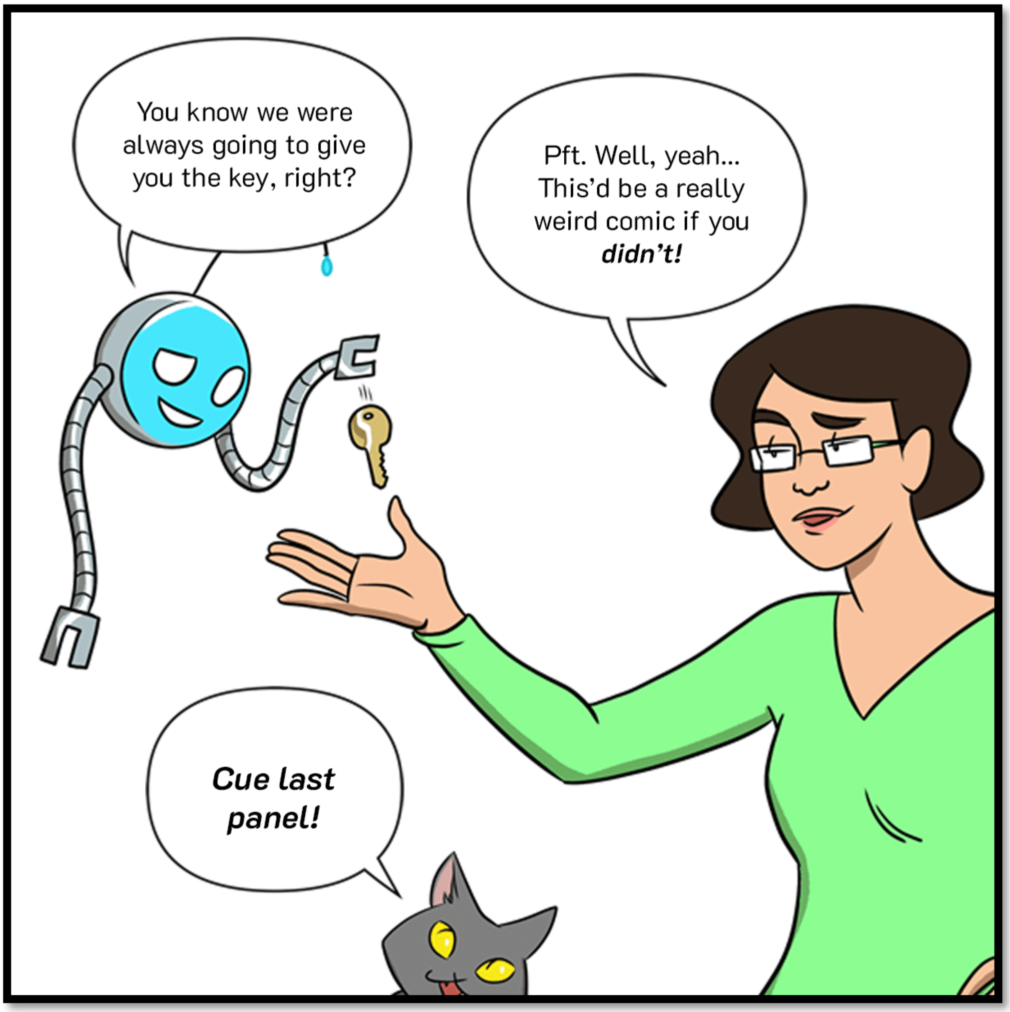

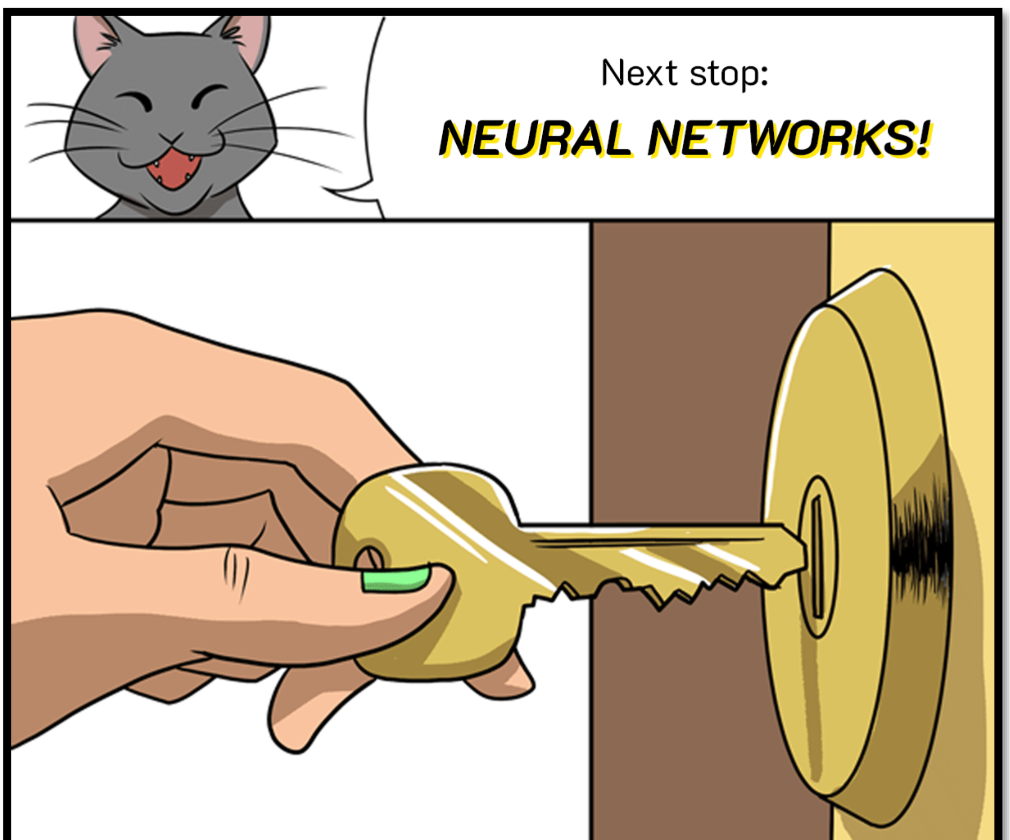

LEARN MORE about the fast-changing world of machine learning and **STAY TUNED** for *Part Two: A Deep Dive Into Deep Learning.*