ML and Society

Feb 15, 2023

Passcode: Tressie McMillan Cottom

- A preeminent public scholar on race, digital media, and higher ed

If you want to see science writing done right, go read one of her books

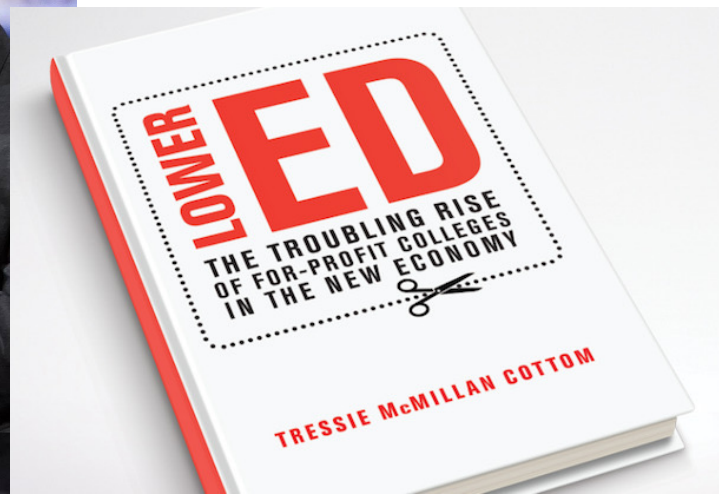

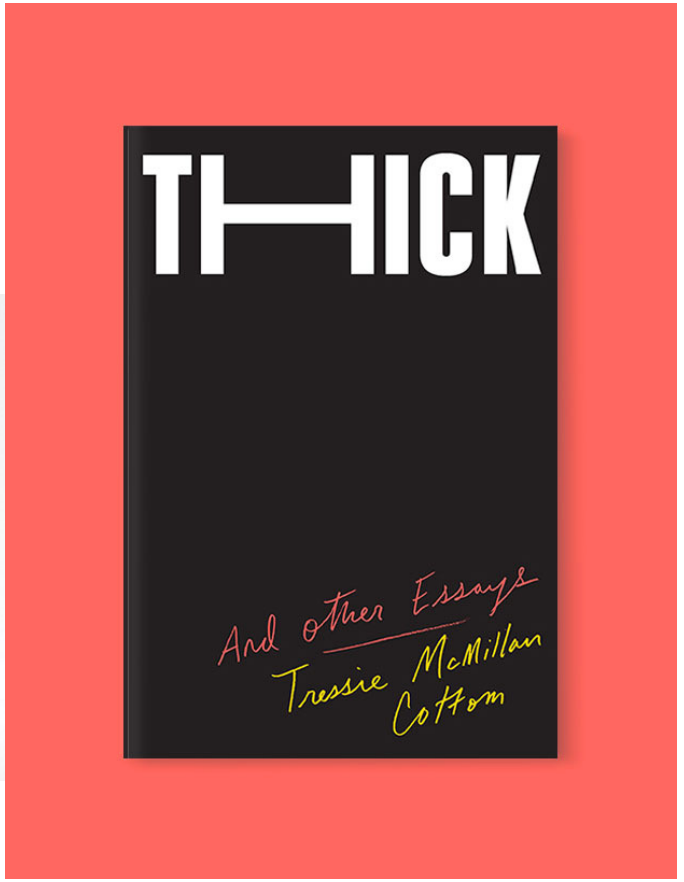

Sit with your team!

Team 1	Afzal	Cole	Navid	Tim	
Team 2	Aishwarya	Herman	Mads	Melvin	
Team 3	Daphkar	Juliana	Ibtida	Monica	
Team 4	Joe	Ken	Vedant	Zach	
Team 5	Chaitanya	Evan	Hitesh	Sushanth	
Team 6	Hannah	Harinee	Gabriella	Suradhya	
Team 7	Alex	Connor	Gopi	Shane	Thanh
Team 8	Aditi	Connor	Jason	Mitali	
Team 9	Botsalano	Niharika	Vedang	Yunmei	
Team 10	Dhiraj	Frank	Kashyap	Michael	

Rage students in Green. ML+Soc students in black

Who does my machine learning model serve?

How do I know?

What can I do about it?

So who is right?

Depends on what you mean by "algorithm"

"Algorithms can be biased"

"Algorithms are based on math and hence cannot be biased"

Knuth's Definition of an *algorithm*

An algorithm is a finite, definitive, effective procedure with some input and some output

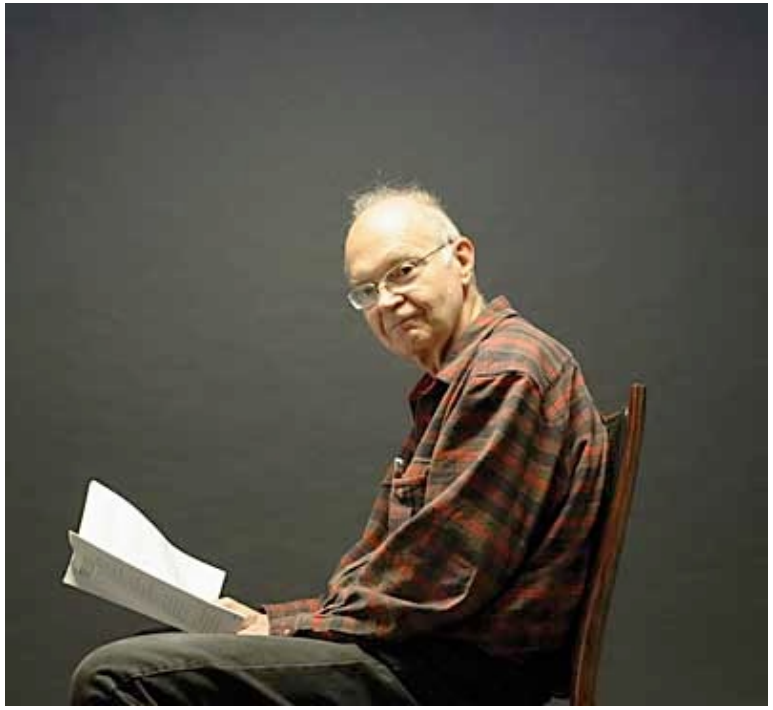

Main Steps in Traditional Algorithm Design

Problem definition

Algorithm design

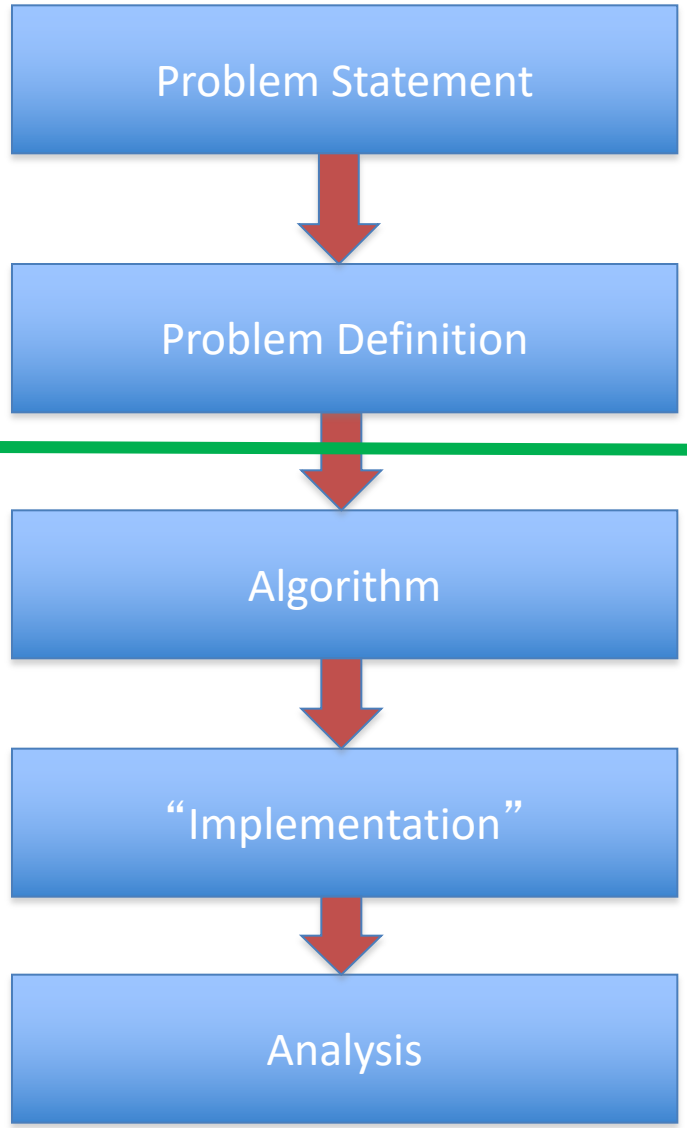

Real world problem

Precise mathematical def

Data Structures

Correctness/Run time

Two main points about traditional algorithms

Problem is defined BEFORE algorithm is designed

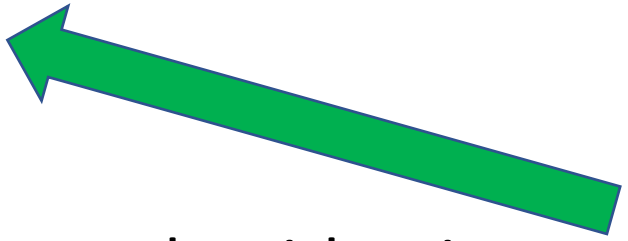

Though really you would be wrong there as well since the problem statement itself could be biased...

Can prove algorithm is correct for ALL inputs

One *could* consider such algorithms not biased

Today

Correct for modern algorithms

“Algorithms can be biased”

Cat vs. Dogs

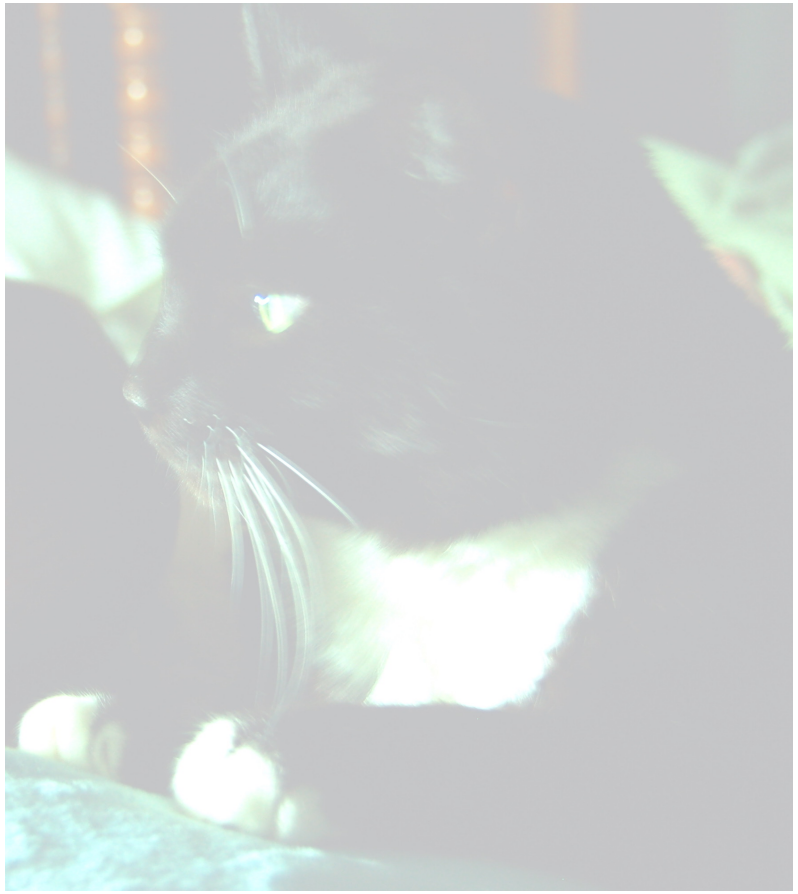

Warren and Billy

How do you “define” a dog vs cat image?

Google Images has “solved” this problem

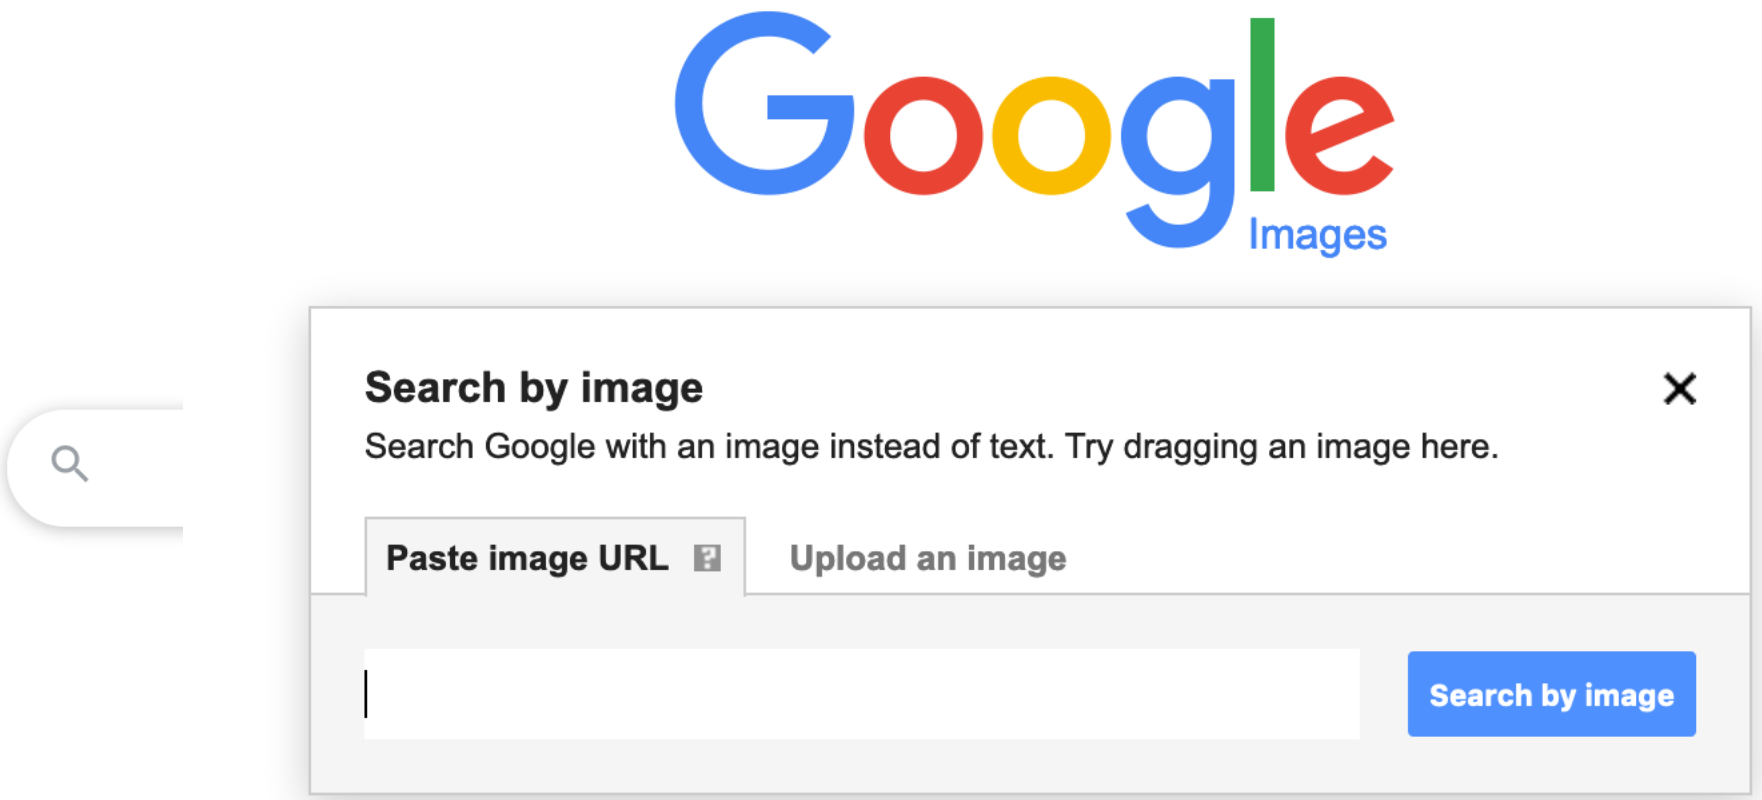

My result for Warren (Spring 20)

Google warrenJPG.JPG

All Images Maps Shopping More Settings Tools

About 2 results (0.47 seconds)

Image size:
1973 × 1895

No other sizes of this image found.

Possible related search: ***pit bull***

Pit bull - Wikipedia
https://en.wikipedia.org › wiki › Pit_bull
Pit bull is the common name for a type of dog descended from bulldogs and terriers. The **pit bull**-type is particularly ambiguous, as it encompasses a range of ...

American Pit Bull Terrier Dog Breed Information, Pictures ...
https://dogtime.com › dog-breeds › american-pit-bull-terrier
The American **Pit Bull** Terrier is a companion and family dog breed. Originally bred to "bait" bulls, the breed evolved into all-around farm dogs, and later moved ...

Visually similar images

Pit bull
Dog

Pit bull is the common name for a type of dog descended from bulldogs and terriers. The pit bull-type is particularly ambiguous, as it encompasses a range of pedigree breeds, informal types and appearances that cannot be reliably identified. Wikipedia

Lifespan: American Pit Bull Terrier: 8 – 15 years, American Staffordshire Terrier: 12 – 16 years

Height: American Pit Bull Terrier: 18 – 21 in., MORE

Mass: American Pit Bull Terrier: 35 – 65 lbs, Staffordshire Bull Terrier: 29 – 37 lbs

Life span: The average life span of the American Pit Bull Terrier ranges from 10 to 12 years. petwave.com

My result for Warren (Spring 22)

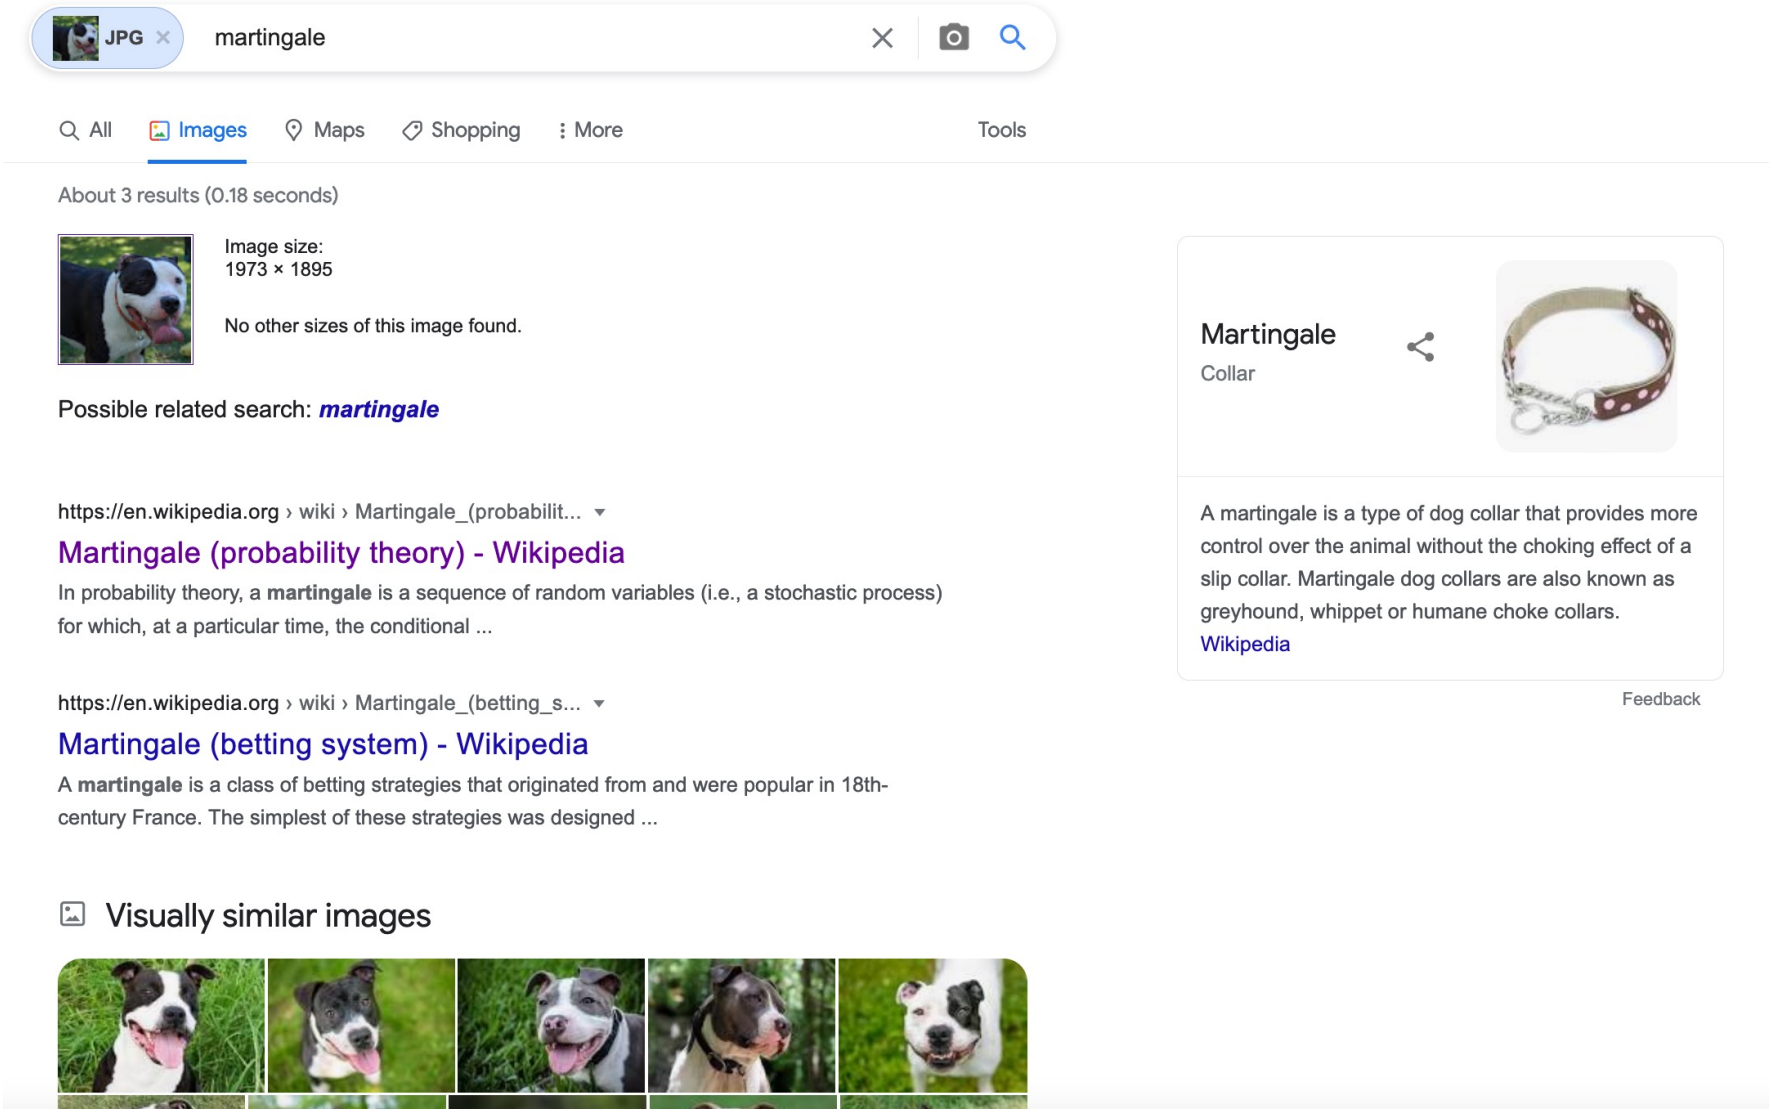

My result for Warren (Spring 23)

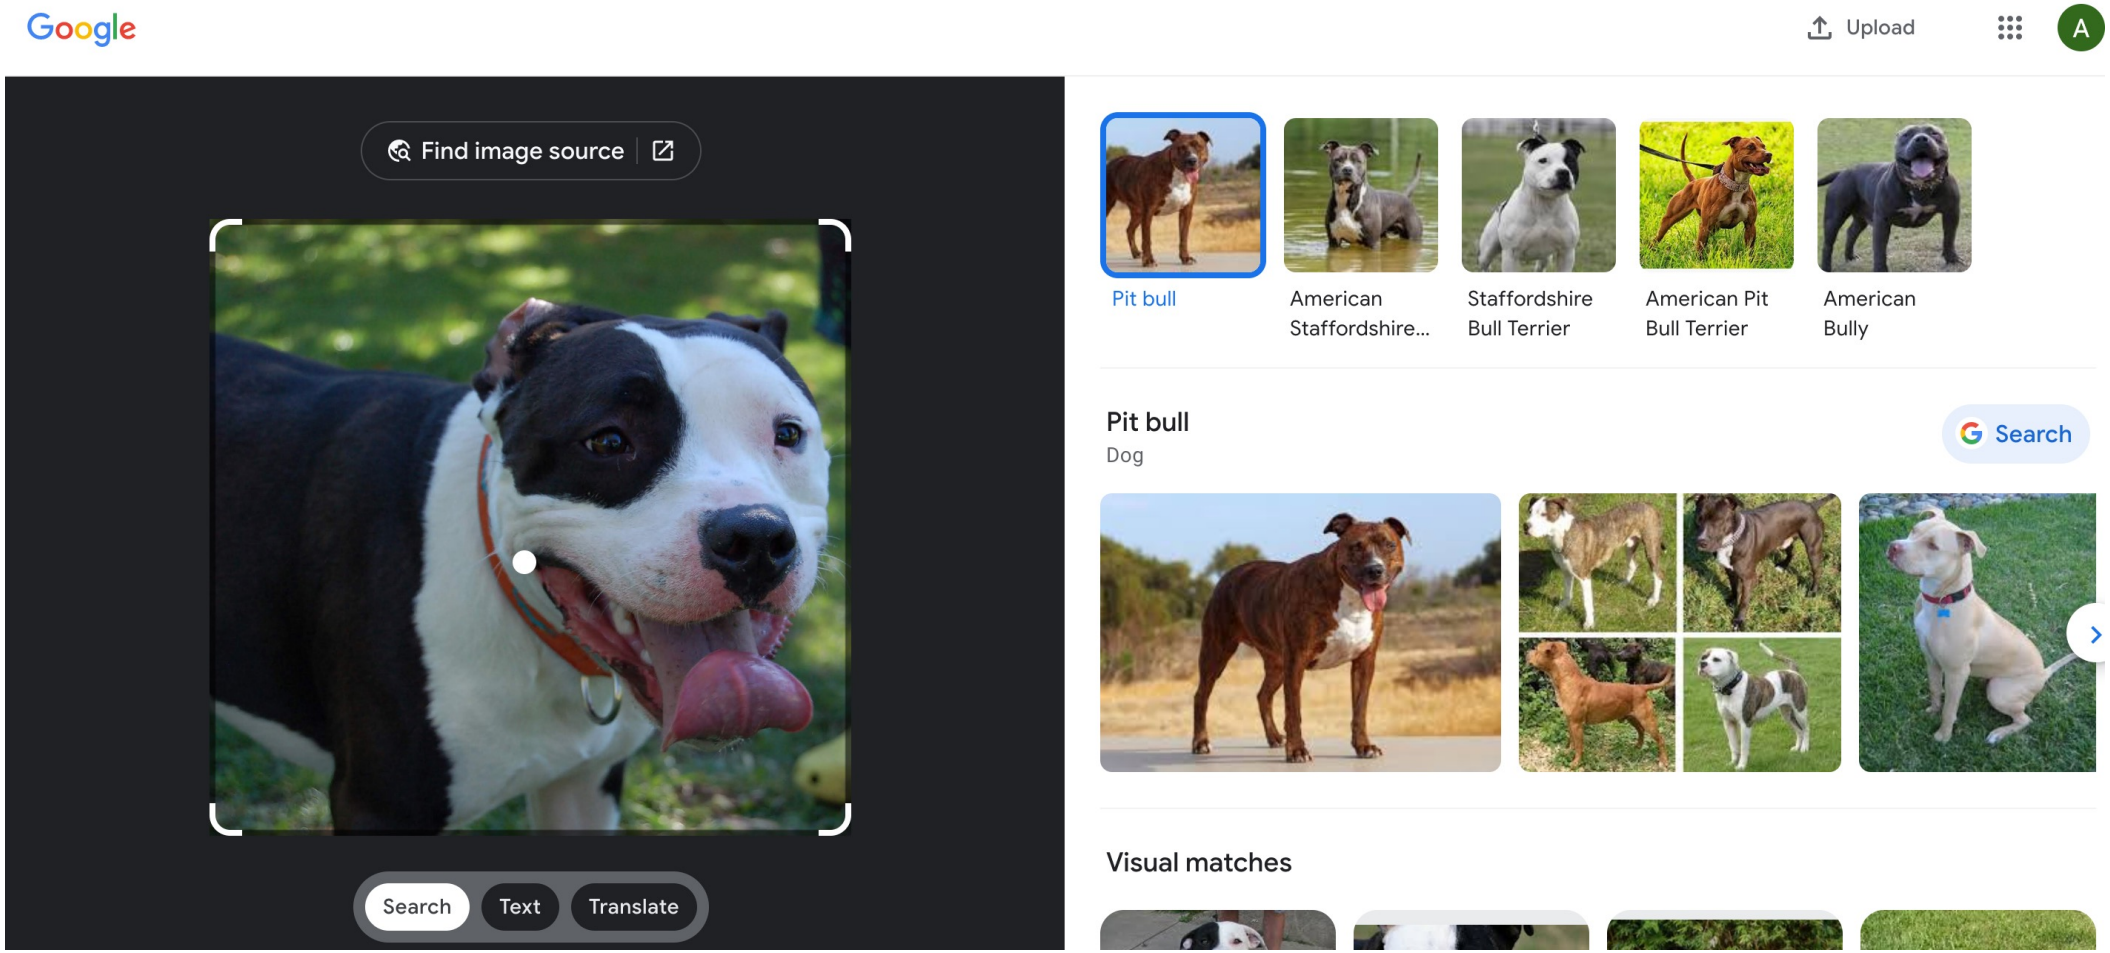

My result for Billy (Spring 20+22)

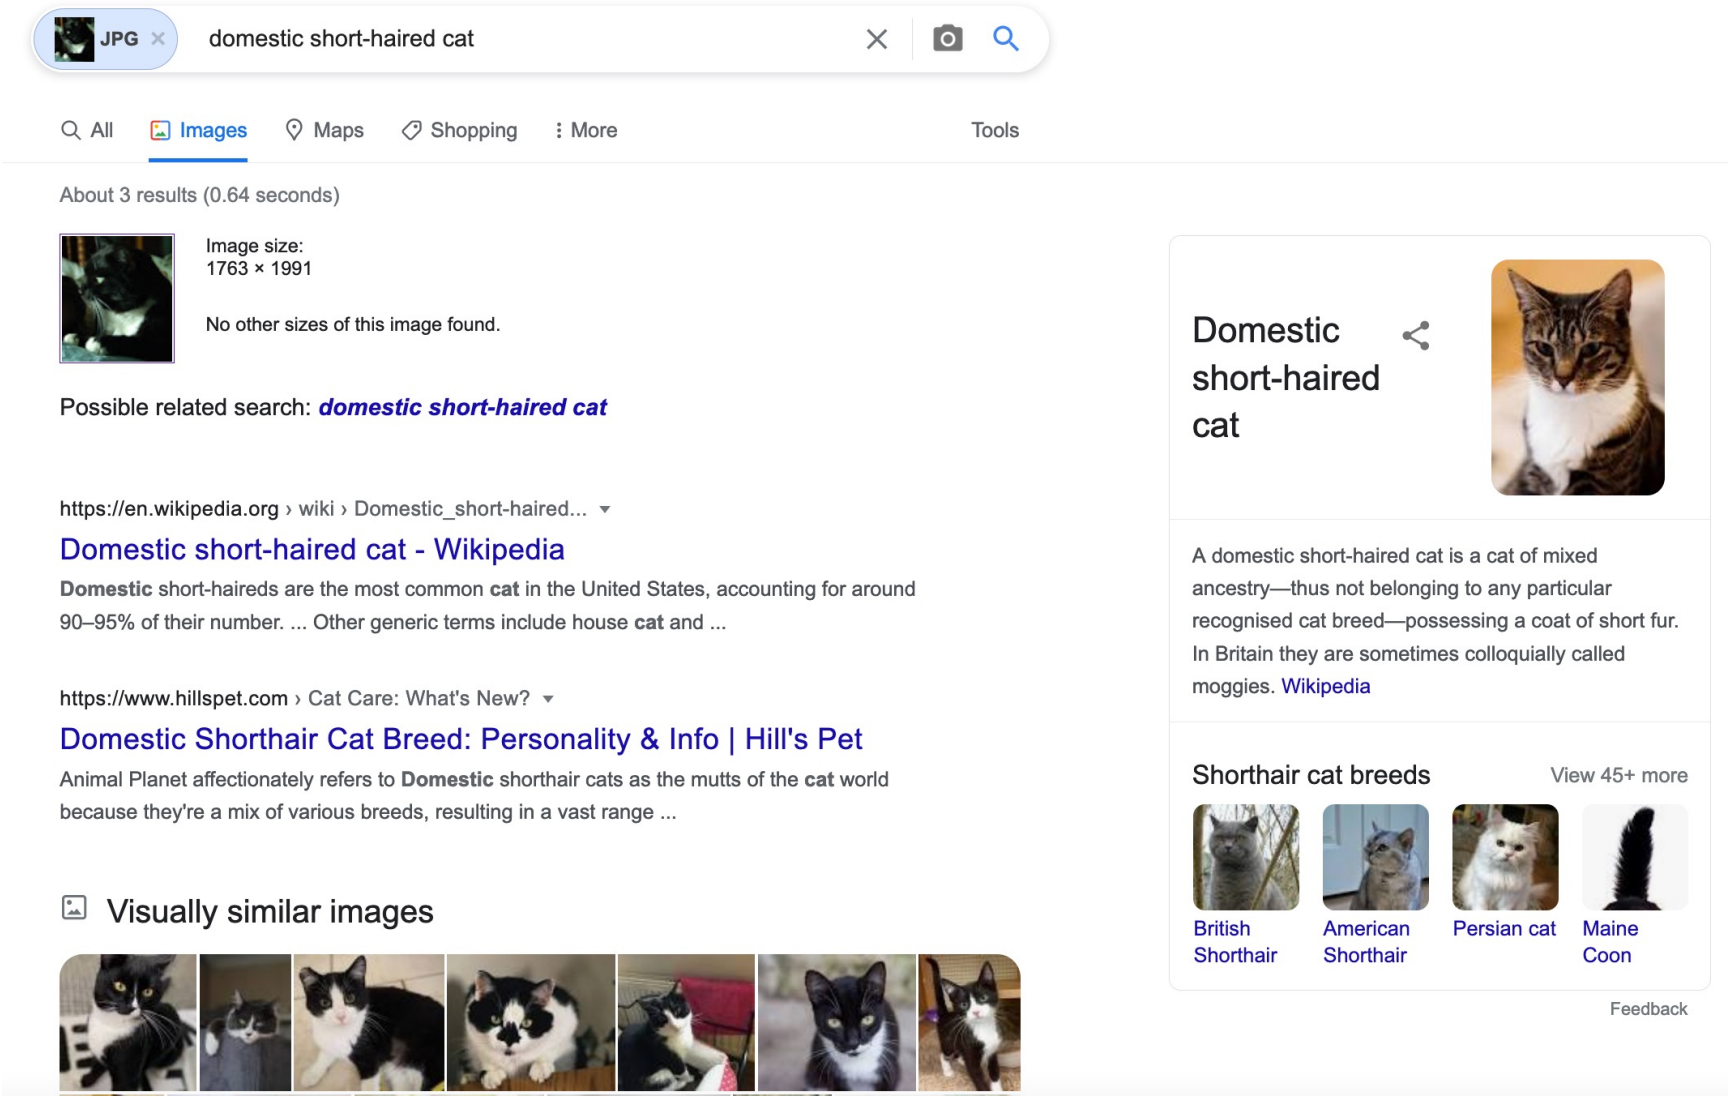

My result for Billy (Spring 23)

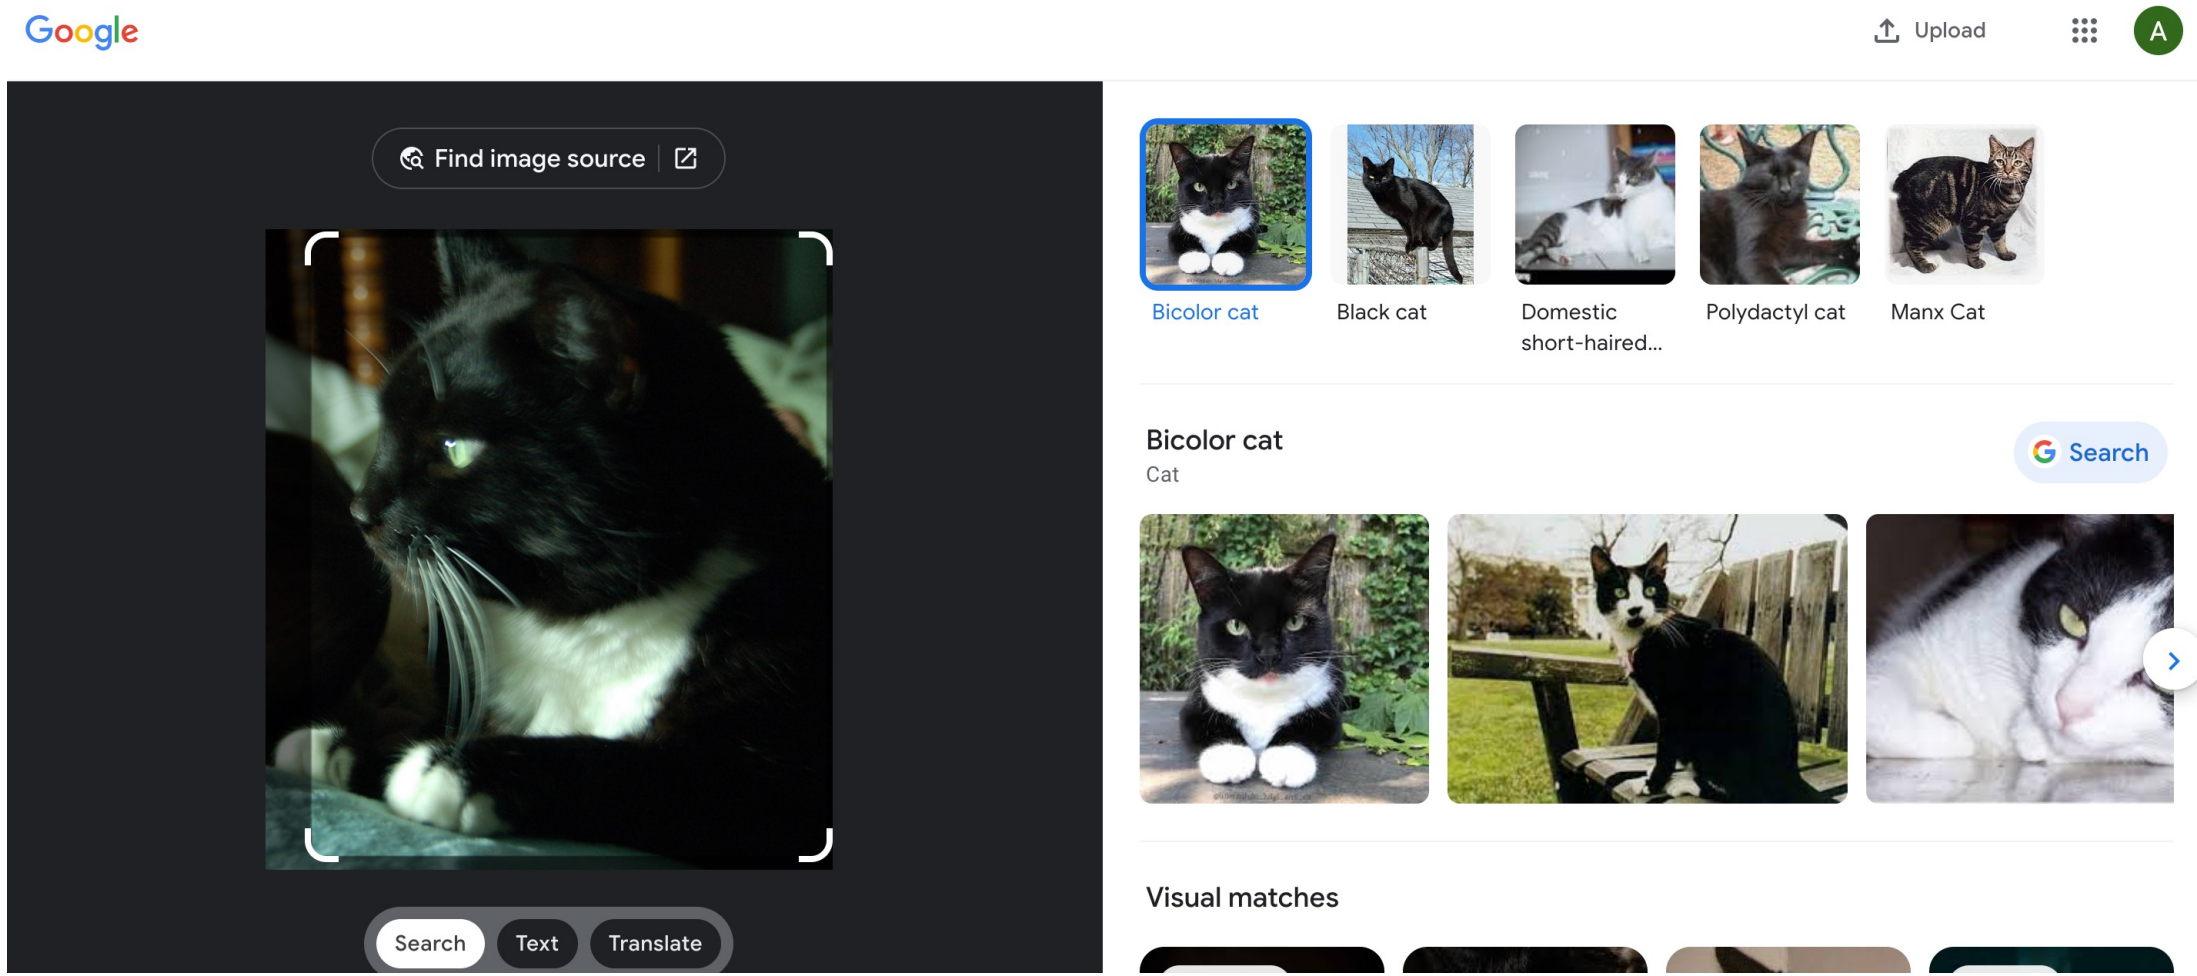

So cats vs dogs problem solved?

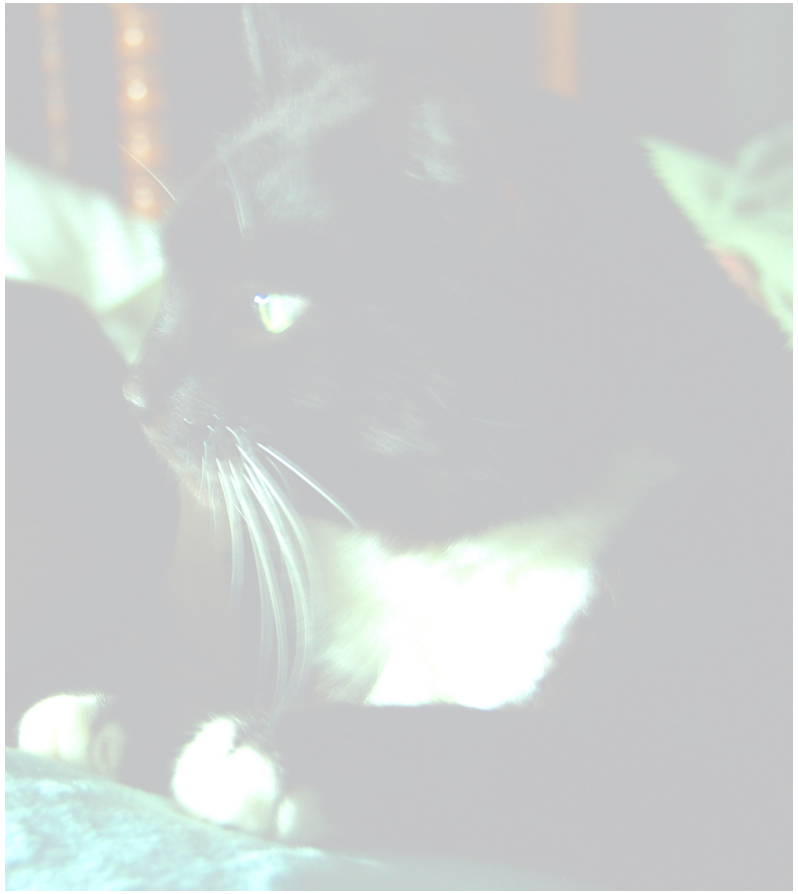

My result for modified Warren (Spring 20)

Google JPG companion dog

All Images Maps Shopping More Settings Tools

About 2 results (2.32 seconds)

Image size:
1973 × 1895

No other sizes of this image found.

Possible related search: ***companion dog***

Companion dog - Wikipedia

https://en.wikipedia.org › wiki › Companion_dog

A **companion dog** is a dog that does not work, providing only companionship as a pet, rather than usefulness by doing specific tasks. Many of the toy dog breeds ...

Best Companion Dog Breeds | Purina

https://www.purina.com › Dogs › Dog Breeds › Collections

Whether you want a friendly face to come home to or the best **companion dog** breed for an elderly parent, get the complete list here.

Visually similar images

Companion dog

A companion dog is a dog that does not work, providing only companionship as a pet, rather than usefulness by doing specific tasks. Many of the toy dog breeds are used only for the pleasure of their company, not as workers. Wikipedia

Companion dog breeds

View 4+ more


Pekingese


Yorkshire Terrier


German Spaniel


Valley Bulldog


Borador

My result for modified Warren (Spring 22)

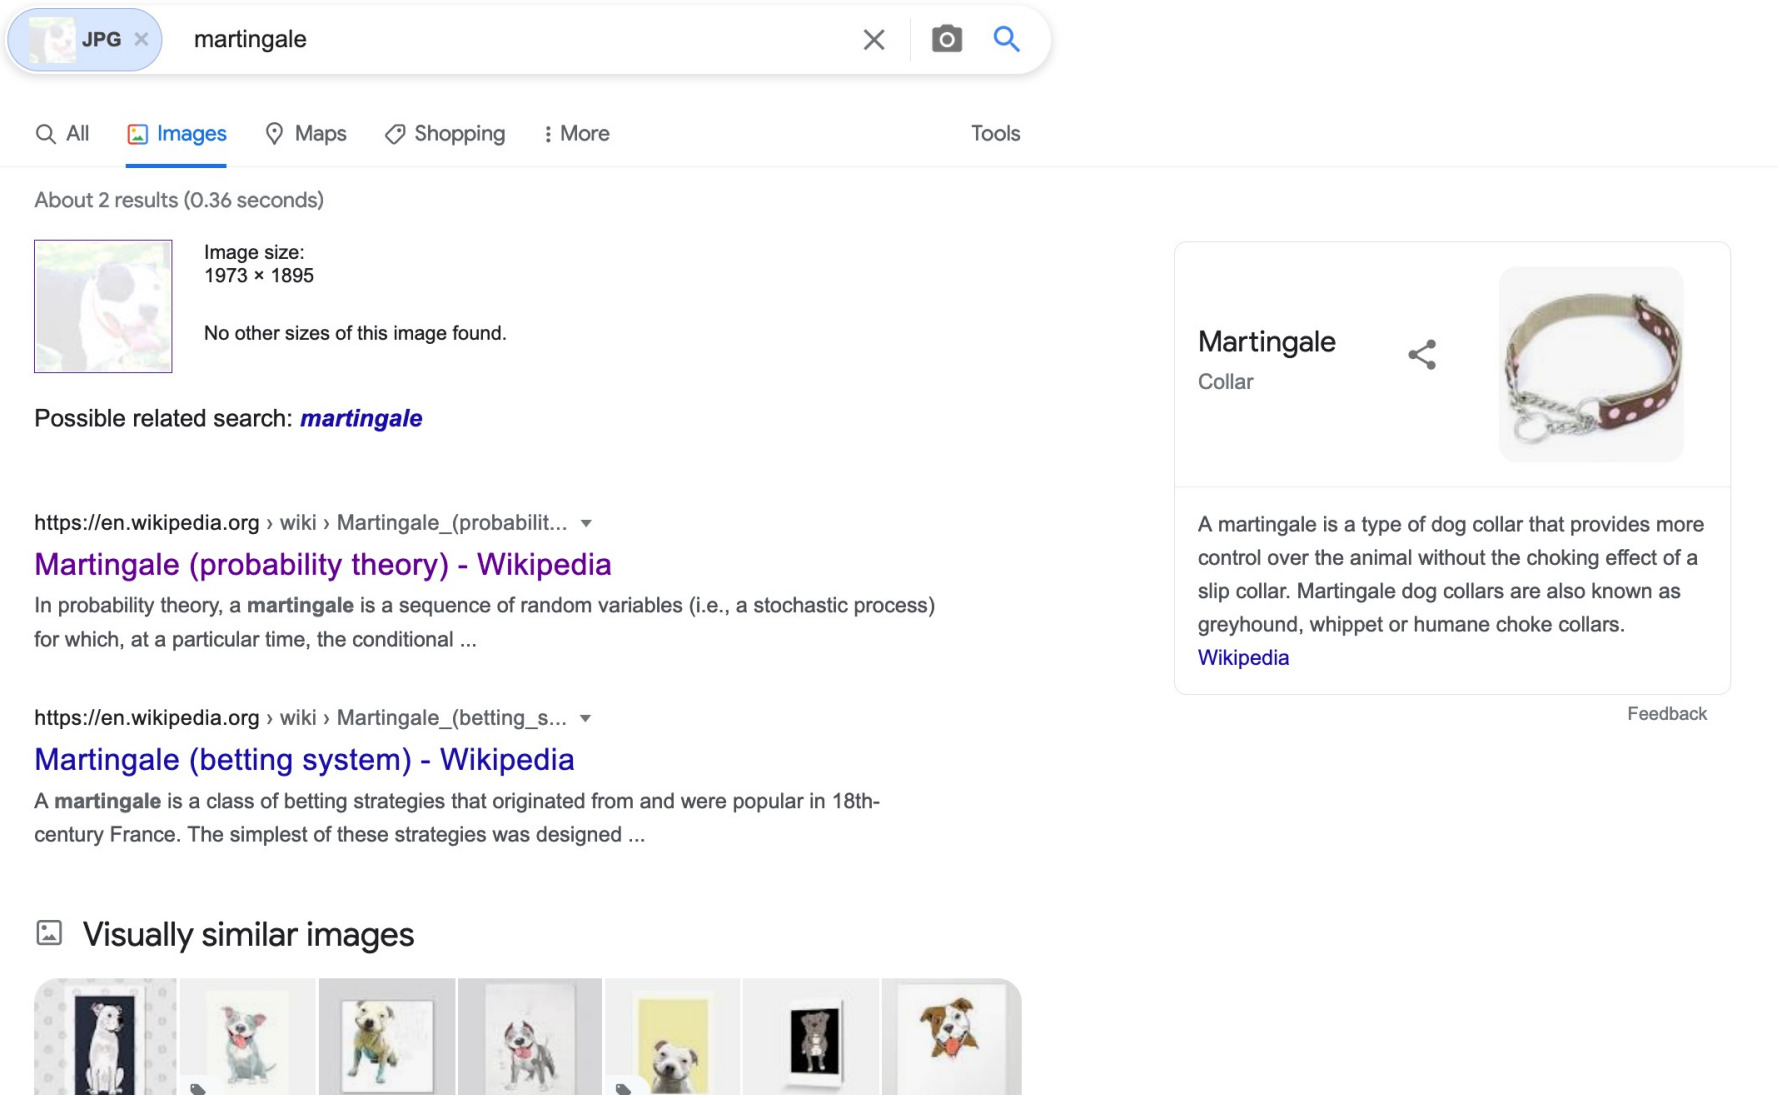

My result for modified Warren (Spring 23)

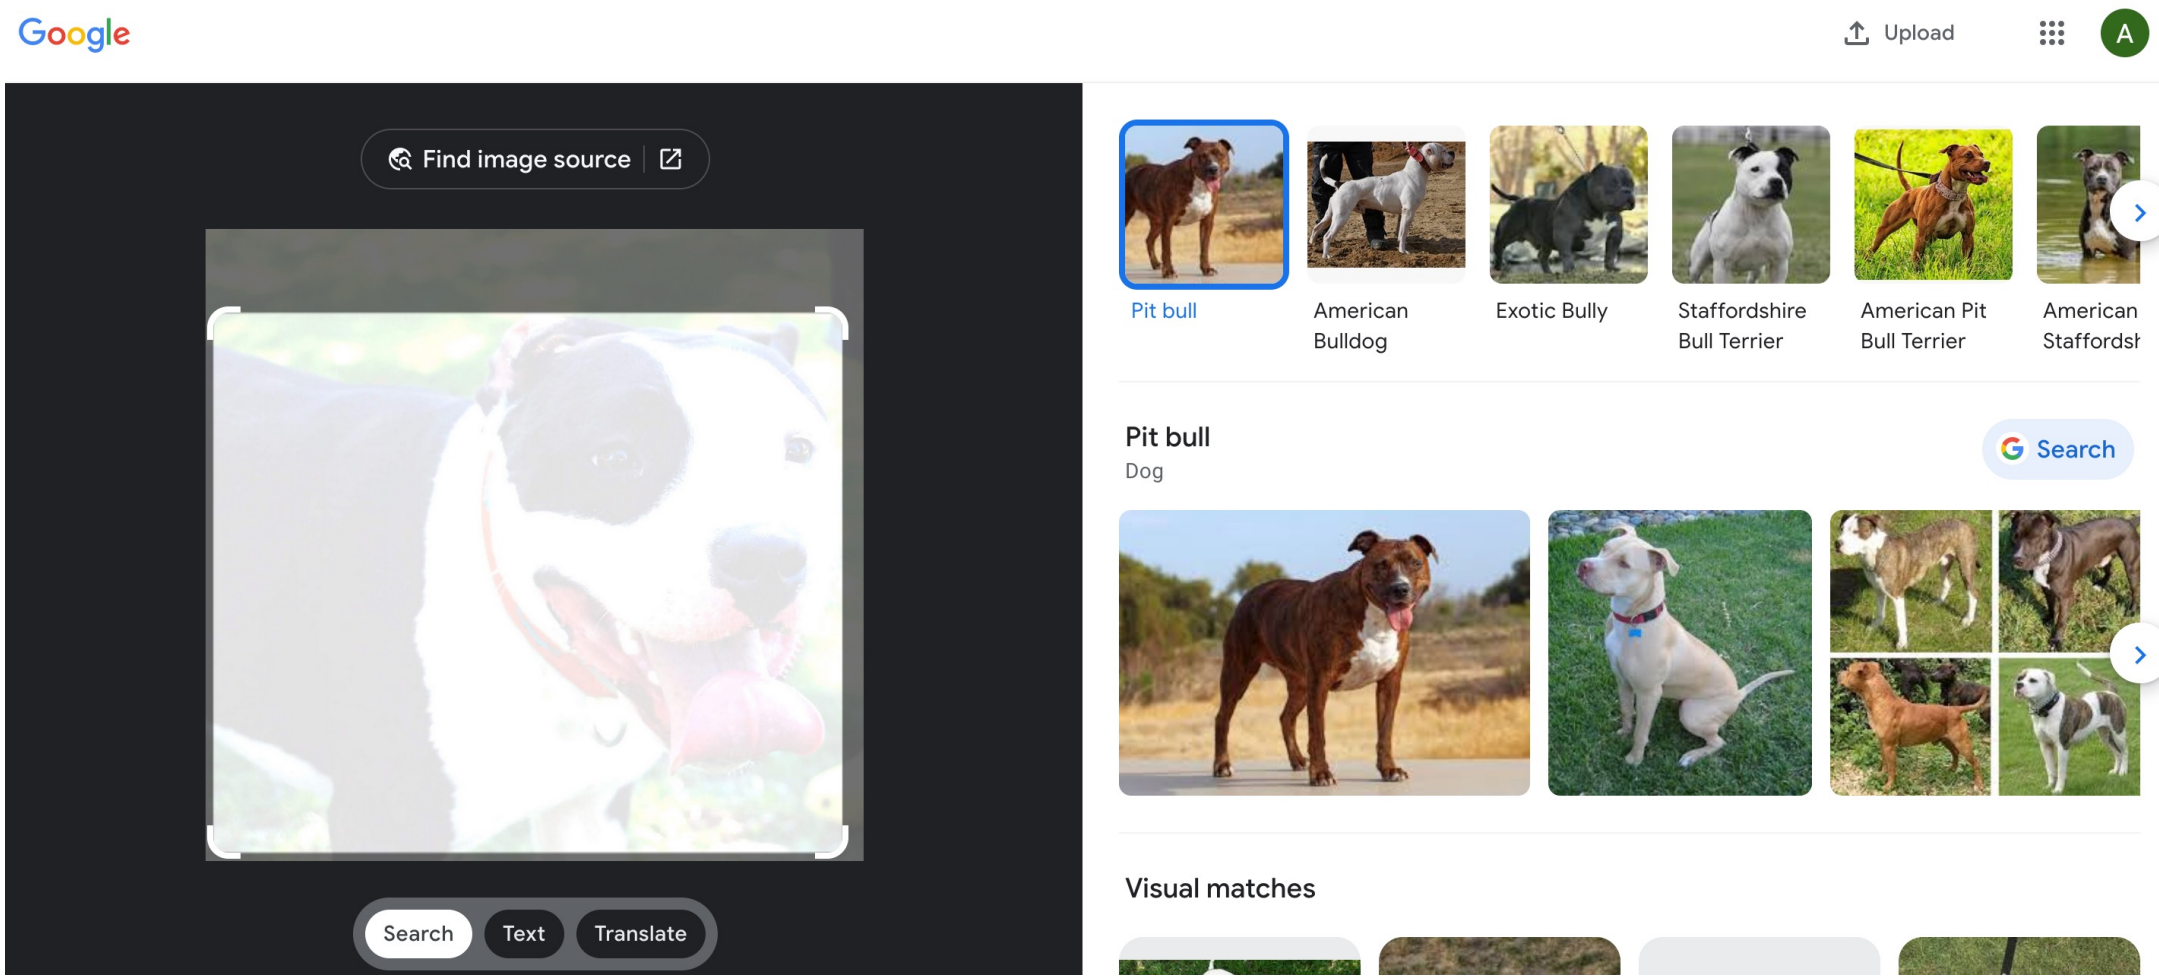

My result for modified Billy (Spring 20)

Google billy...fied-1.JPG fish

All Images Maps Shopping More Settings Tools

About 2 results (1.40 seconds)

Image size:
1763 × 1991

No other sizes of this image found.

Possible related search: ***fish***

Fish - Wikipedia

https://en.wikipedia.org › wiki › Fish

Fish are gill-bearing aquatic craniate animals that lack limbs with digits. They form a sister group to the tunicates, together forming the olfactores. Included in this ...

Pet Fish for Sale: Tropical and Freshwater Fish | PetSmart

https://www.petsmart.com › fish › live-fish

130 Items - Create or augment the perfect underwater community with our selection of freshwater and tropical **fish** for sale.

Visually similar images

Fish

Animal

Fish are gill-bearing aquatic craniate animals that lack limbs with digits. They form a sister group to the tunicates, together forming the olfactores. Included in this definition are the living hagfish, lampreys, and cartilaginous and bony fish as well as various extinct related groups. Wikipedia

Lifespan: Common carp: 20 years, MORE

Phylum: Chordate

Mass: Common carp: 4.4 – 31 lbs, Northern pike: 34 lbs, MORE

Encyclopedia of Life

Length: Common carp: 16 – 31 in., Siamese fighting fish: 2.8 in., MORE

Speed: Ocean sunfish: 2 mph, Great white shark: 35 mph

Clutch size: Common carp: 300,000, Siamese fighting fish: 10 – 40

My result for modified Billy (Spring 22)

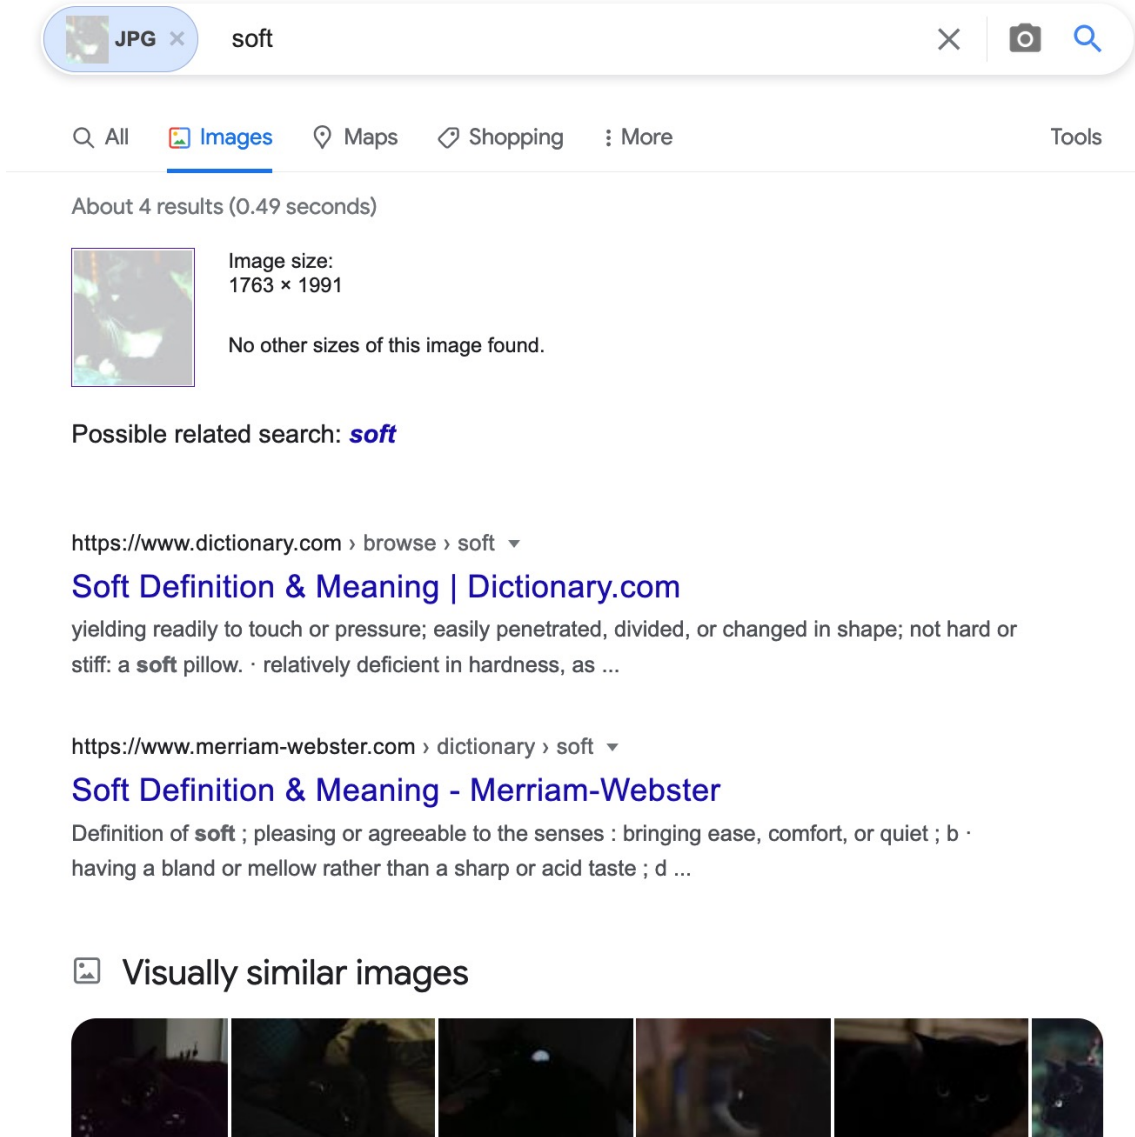

My result for modified Billy (Spring 23)

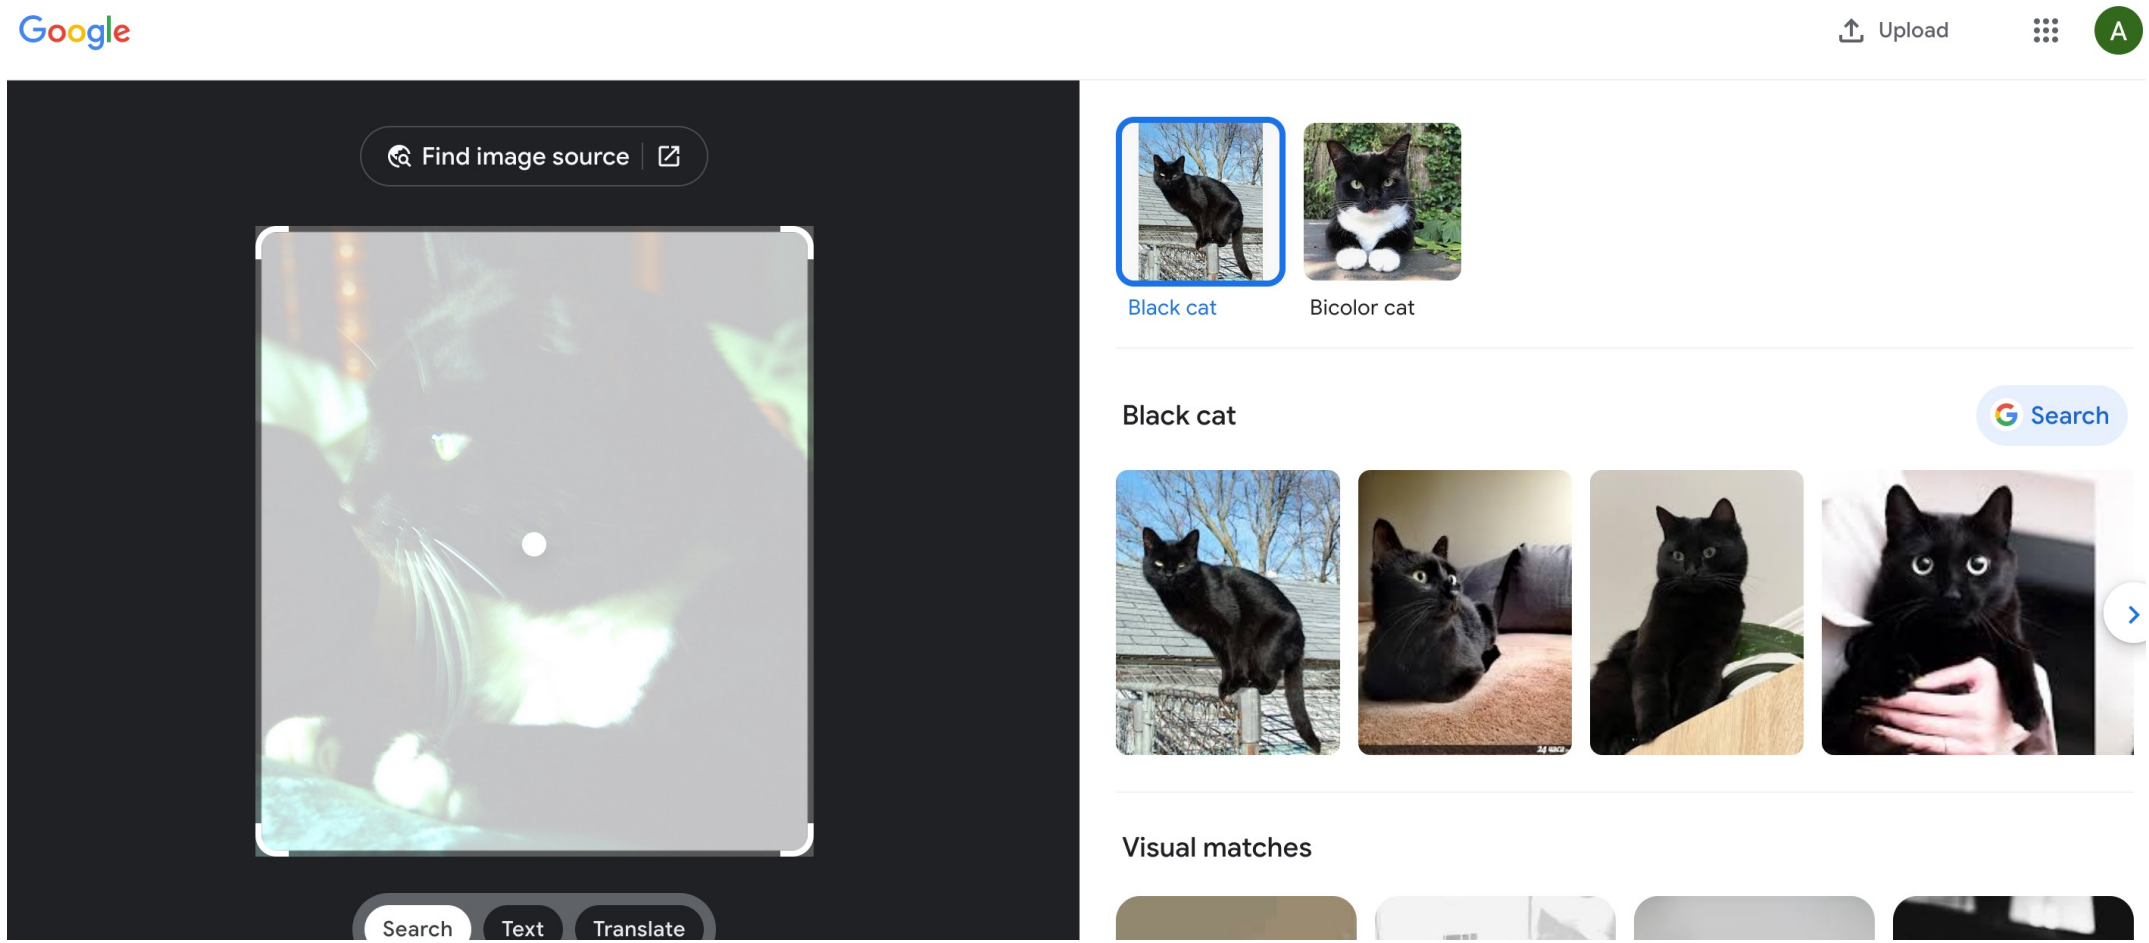

How does Google Images work?

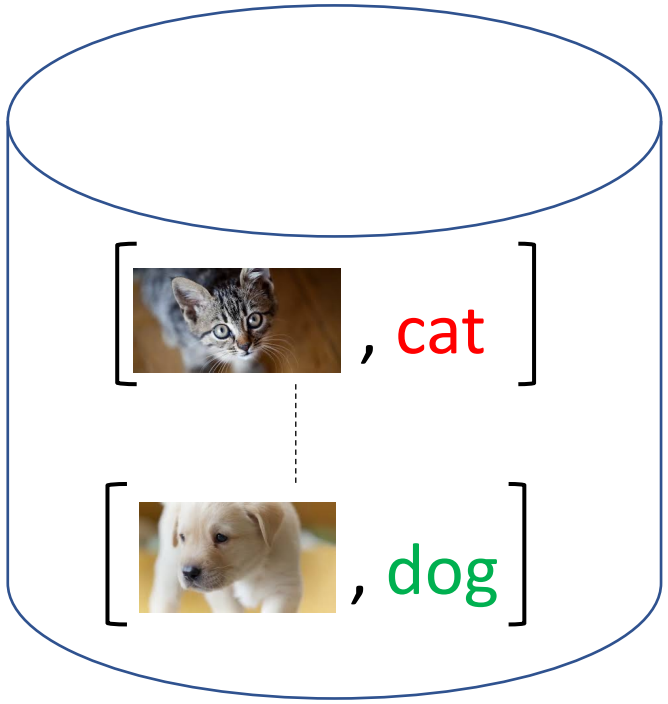

Training

When a new image comes in

Cc

Go

gs

When an algorithm isn't...

Suresh Venkat Follow
Oct 2, 2015 · 5 min read

The popular press is full of articles about "algorithms" and "algorithmic fairness" and "algorithms that discriminate, (or don't)". As a computer scientist (and one who studies algorithms to boot), I find all this attention to my field rather gratifying, and not a bit terrifying.

What's even more pleasing is that the popular explanation of an algorithm follows along the lines of the definition we've been using since, well, forever

> *An algorithm is a set of steps (the instructions) each of which is simple and well defined, and that stops after a finite number of these steps.*

If we wanted a less intimidating definition of an algorithm, we turn to the kitchen:

AOC is right!

Correct for modern algorithms

“Algorithms can be biased”

Modern Algorithm = Machine Learning (ML)

Start with a Question

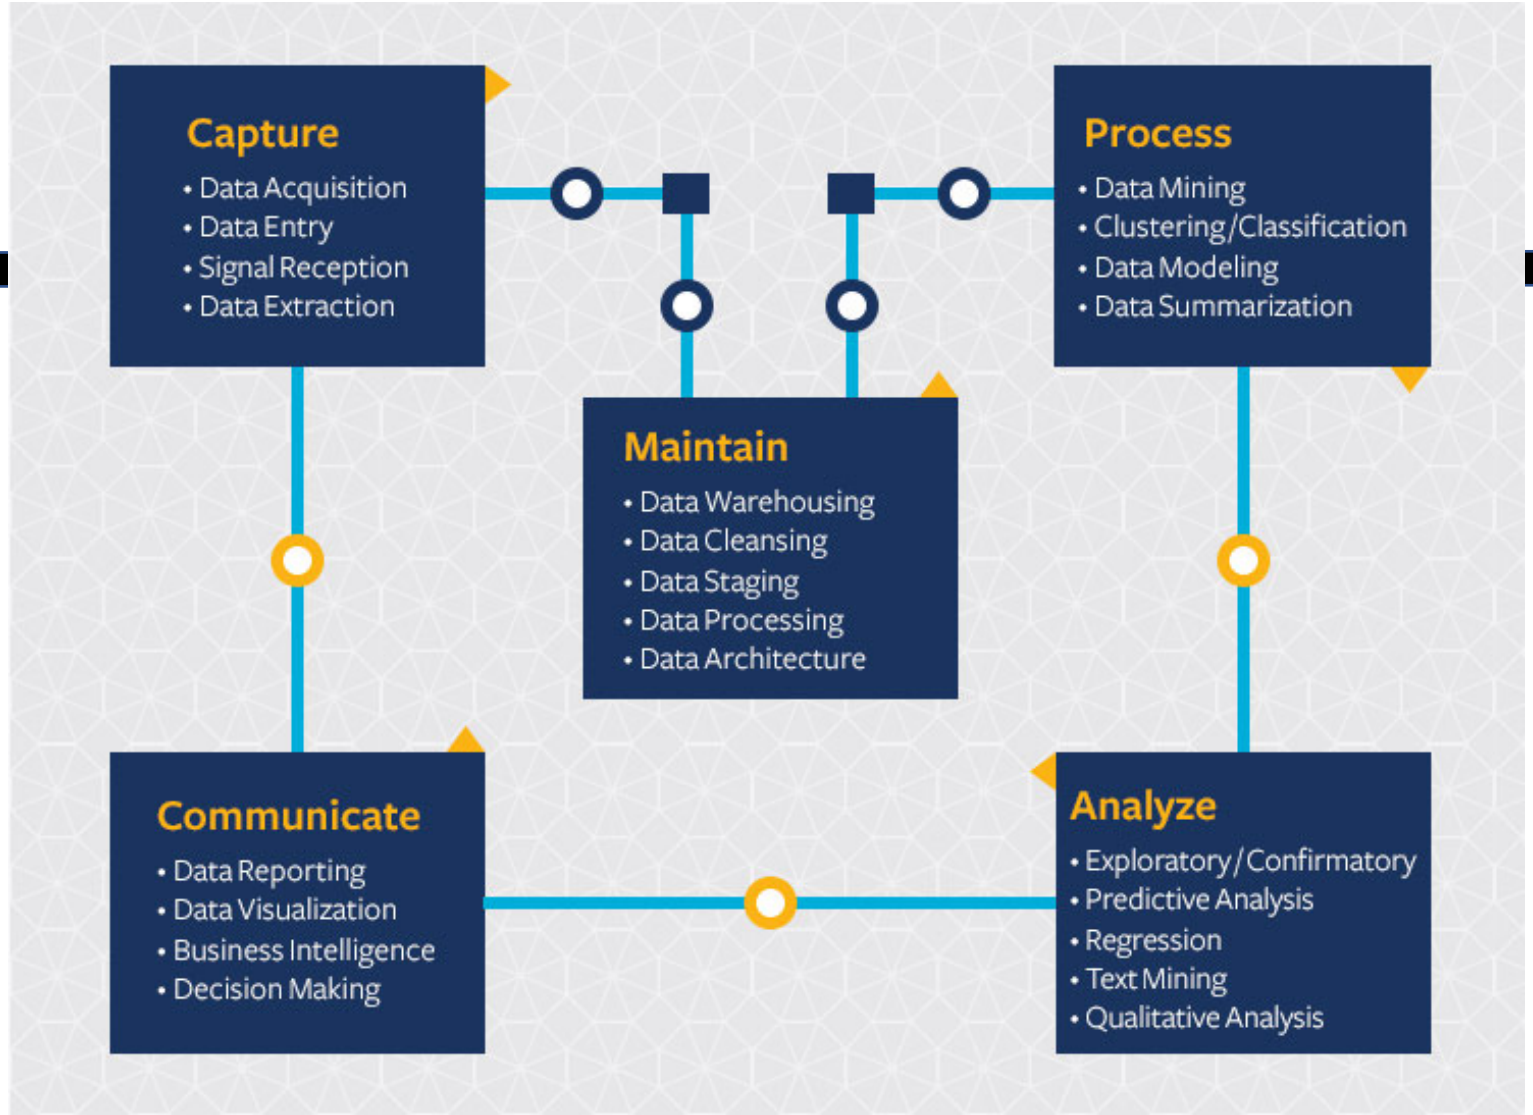

Deploy!

The Machine Learning Pipeline

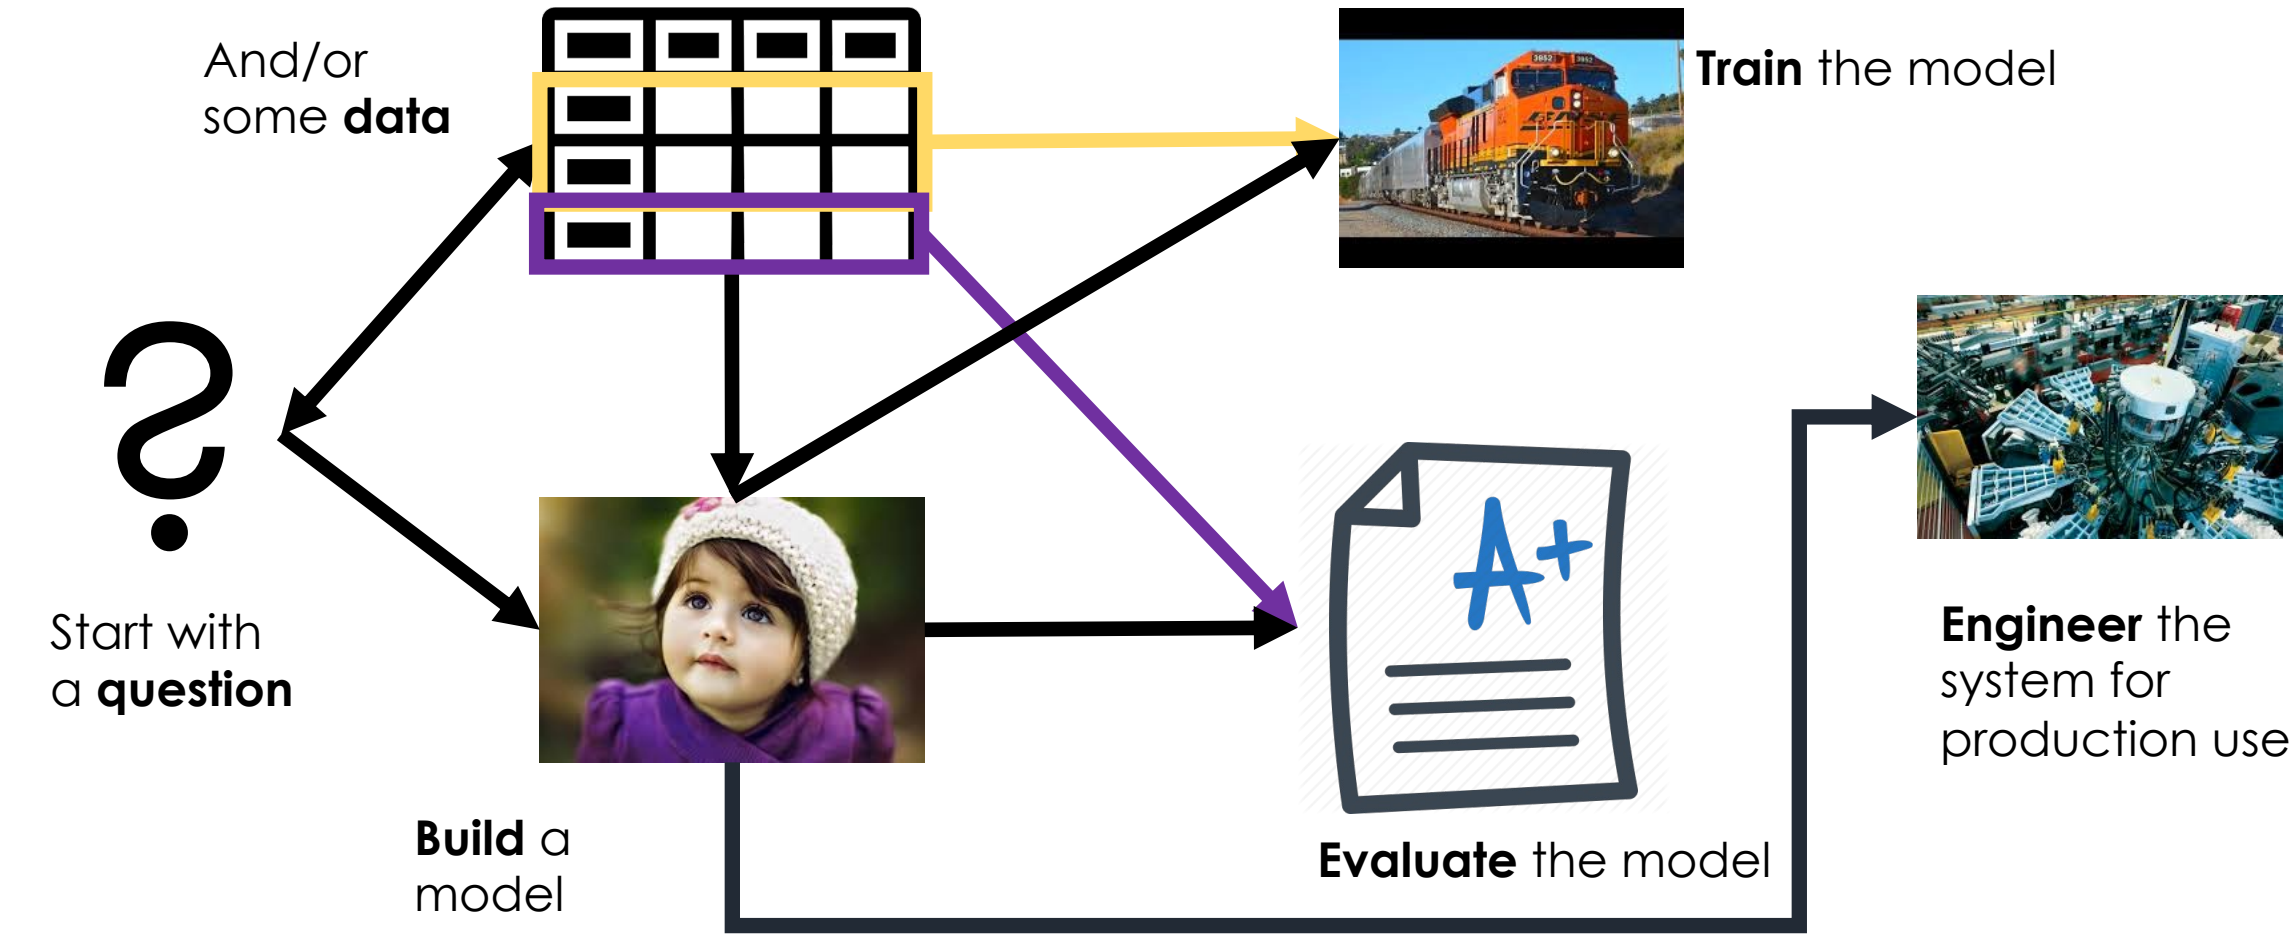

Slide by Kenny Joseph

Back to cats vs. dogs

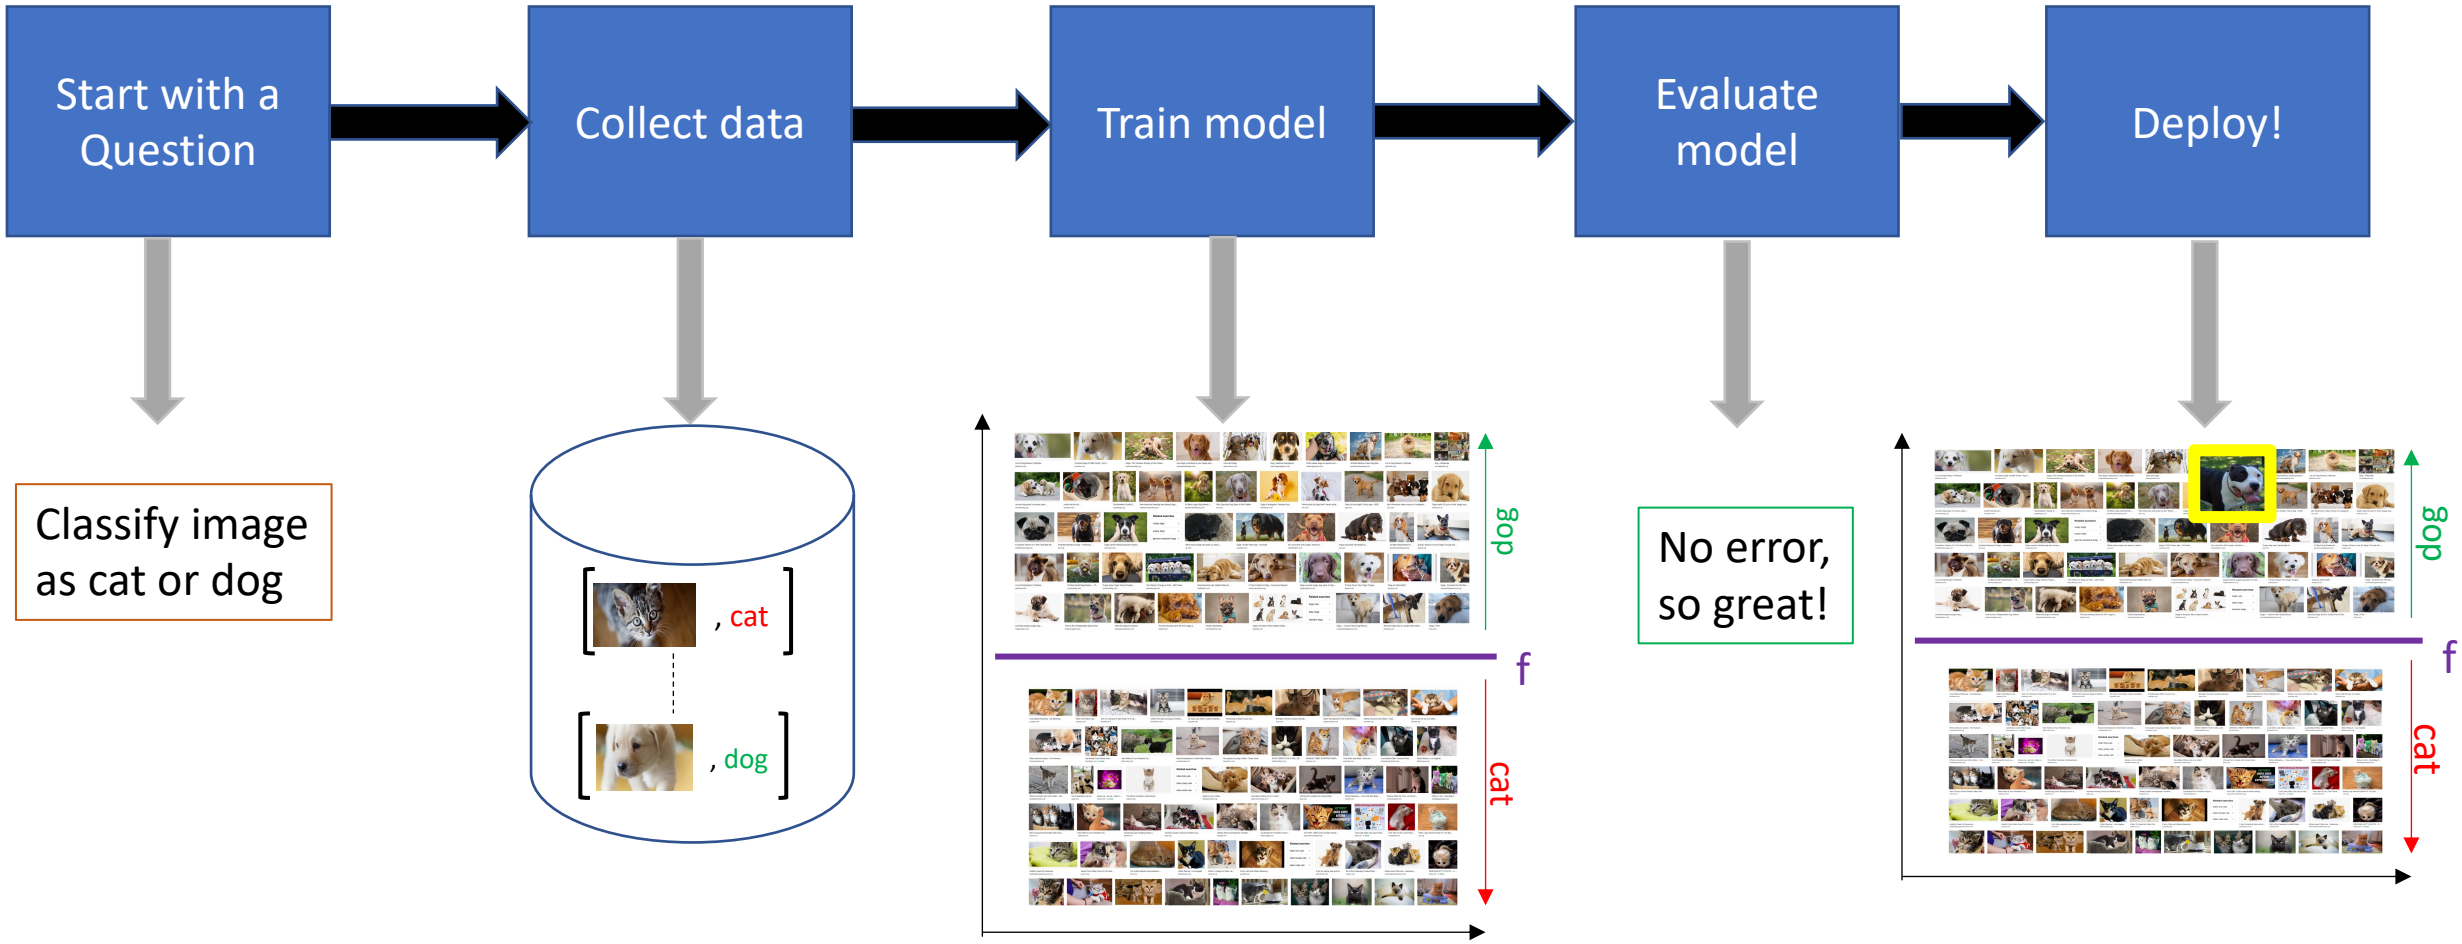

What is missing from this picture?

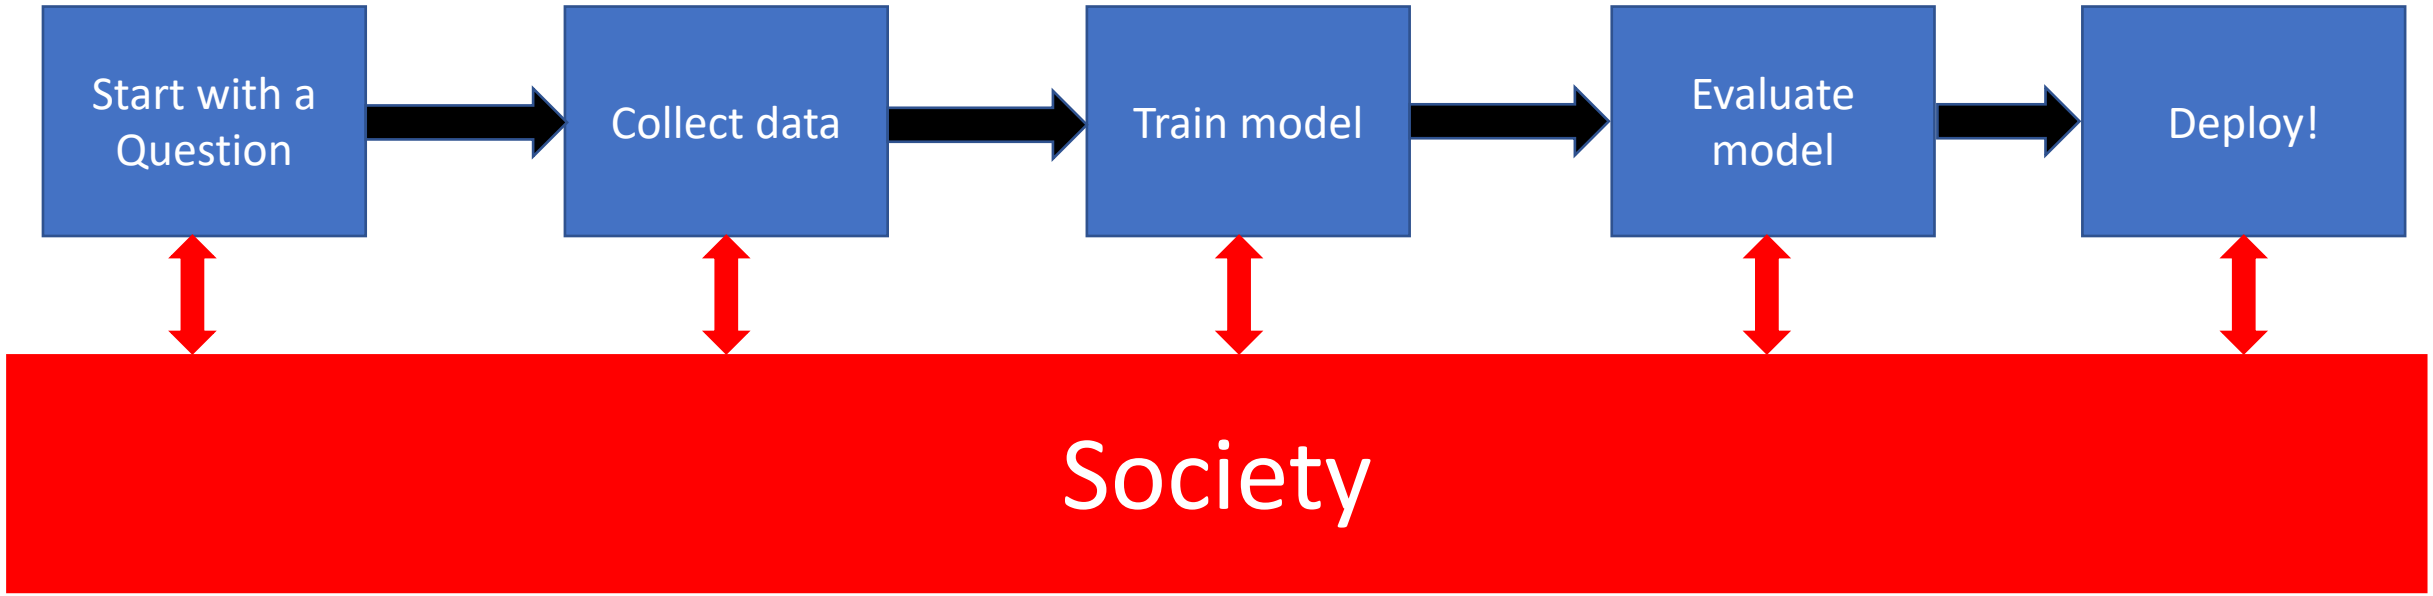

A real view of the ML Pipeline

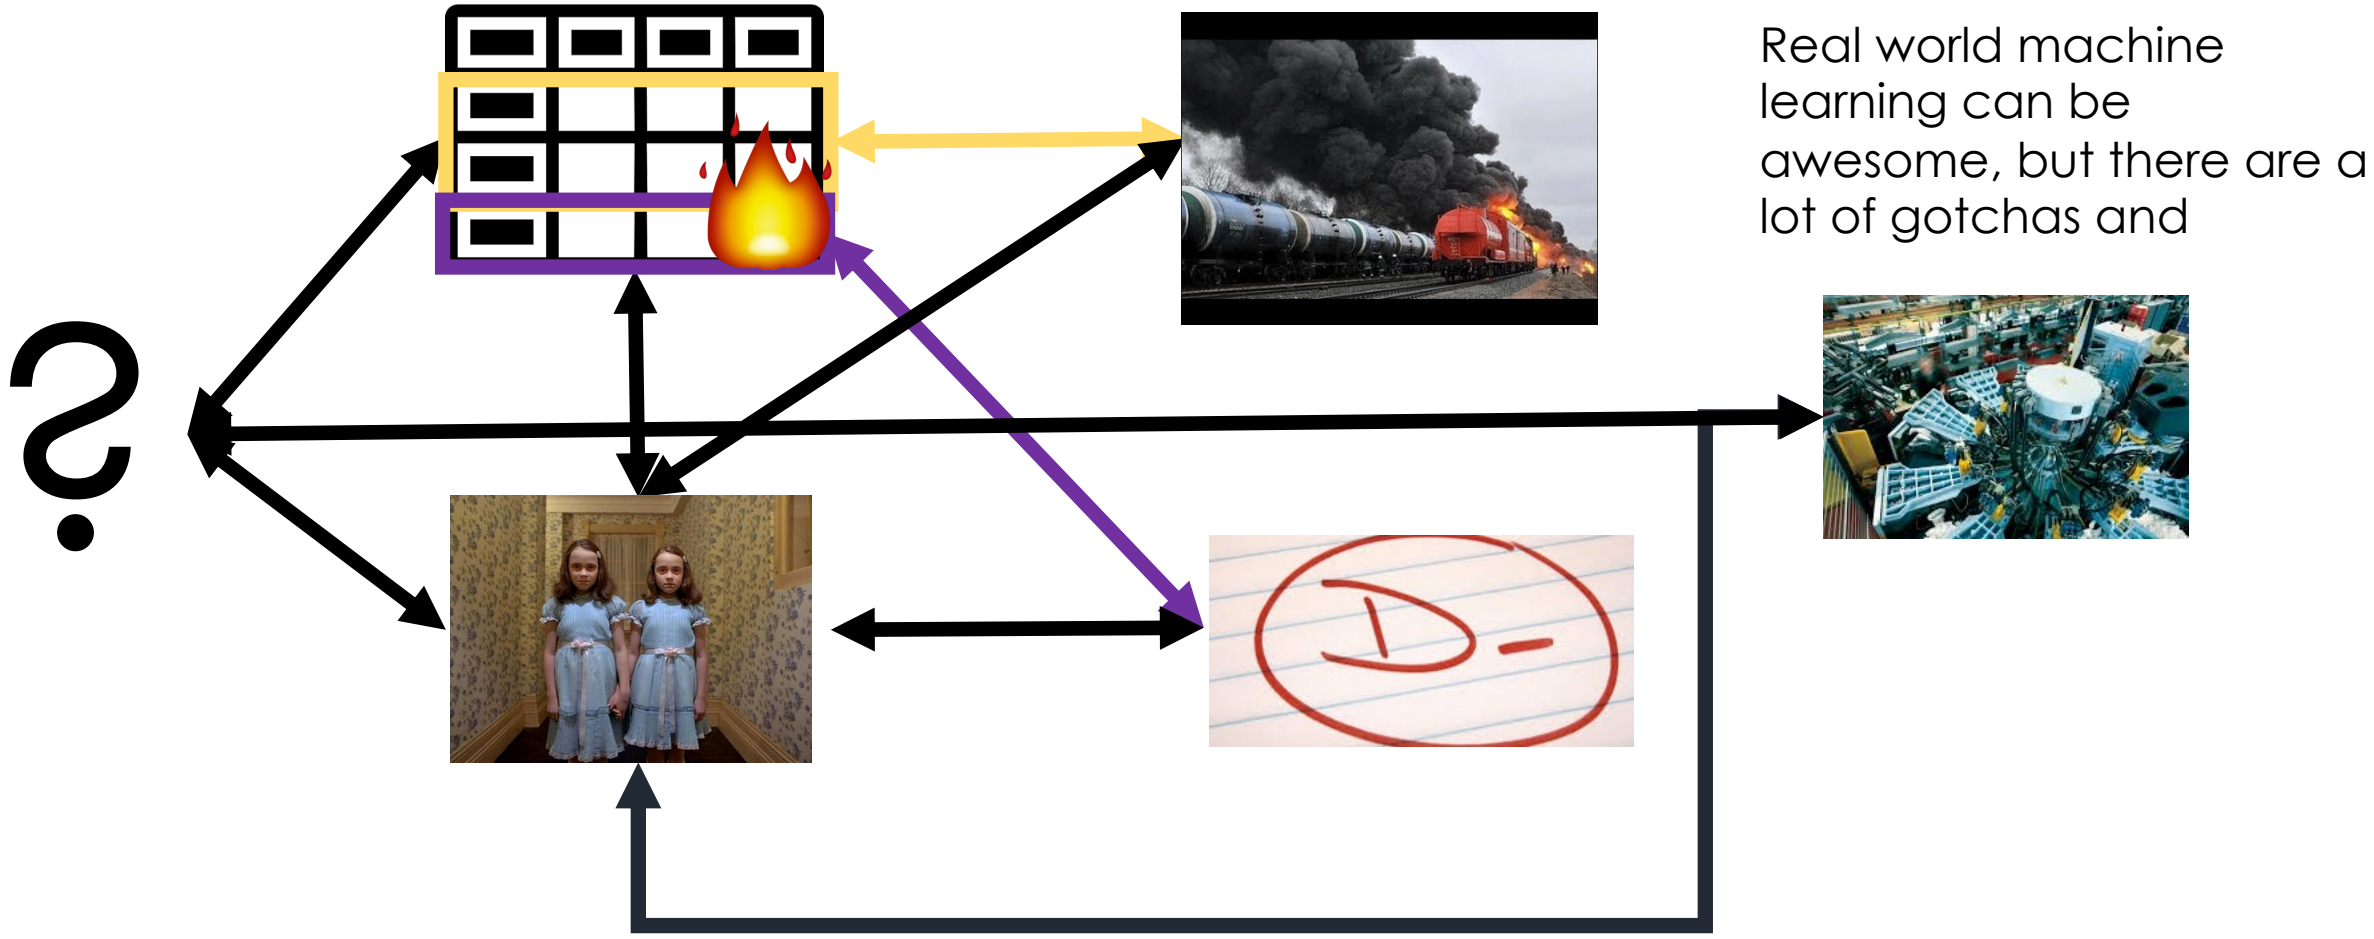

Real world machine learning can be awesome, but there are a lot of gotchas and

Slide by Kenny Joseph

Start with a Question → Collect data → Train model → Evaluate model → Deploy!

Real world goal

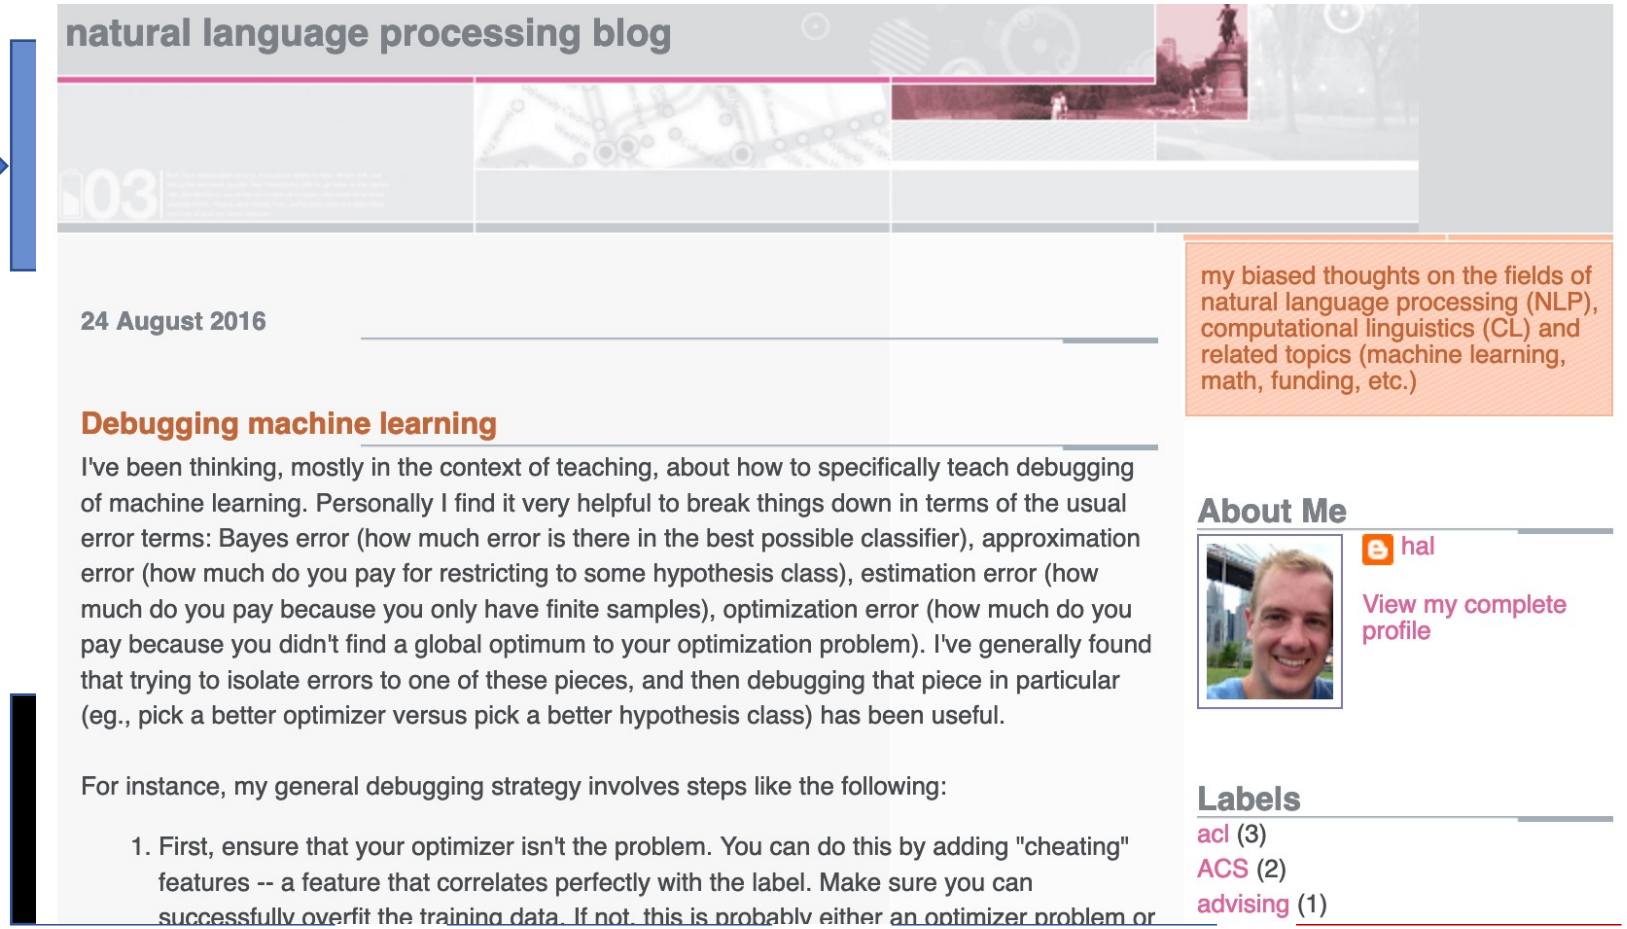

Data representation → Target class/model → Training data set

What is missing?

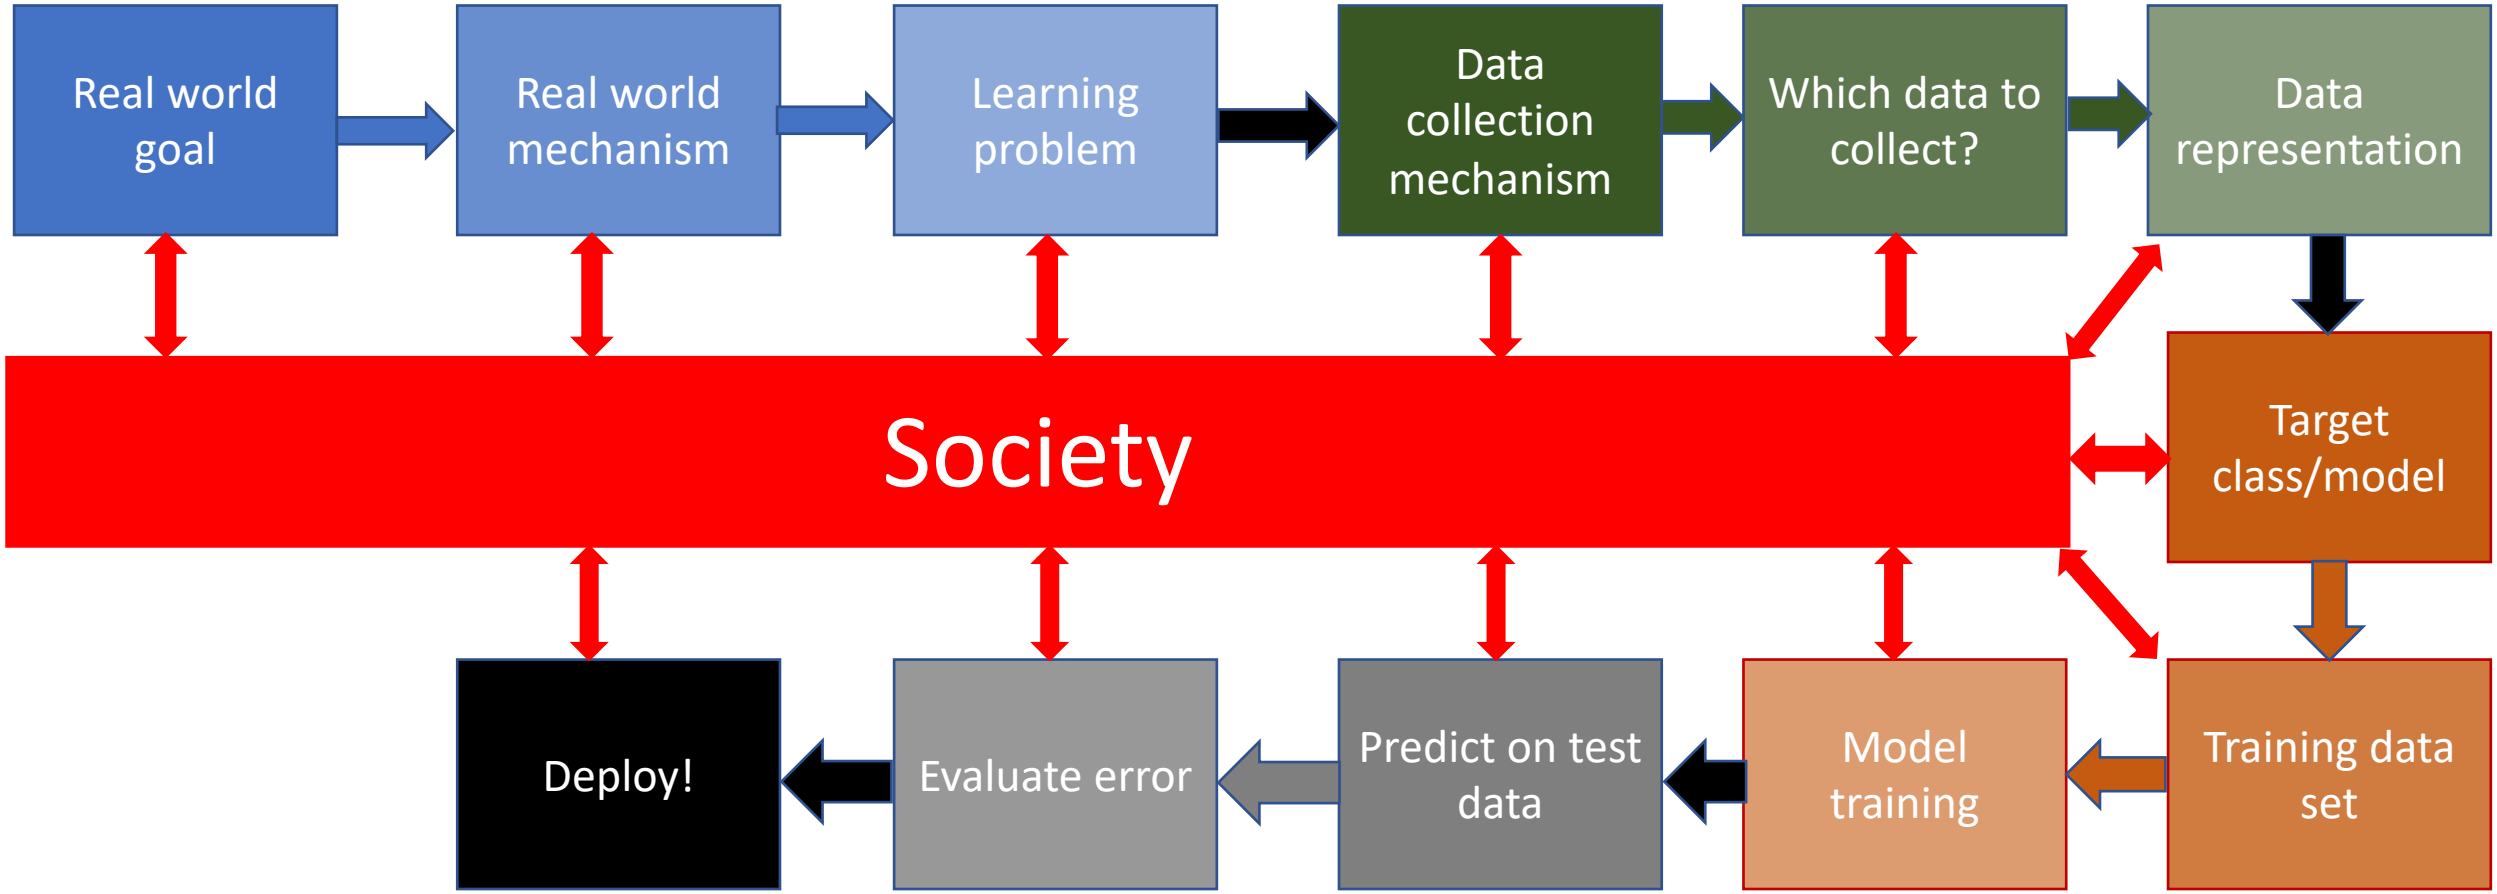

Not the only ML+society pipeline in town

INSIDE AI

Black-Boxed Politics:

Opacity is a Choice in AI Systems

Katarzyna Szymielewicz Follow
Jan 17 · 23 min read

Written by: Agata Foryciarz, Daniel Leufer, Katarzyna Szymielewicz

Illustrations by: Olek Modzelewski

Artificial intelligence captures our imagination like almost no other technology: from fears about killer robots to dreams of a fully-automated, frictionless future. As numerous authors have documented, the idea of creating artificial, intelligent machines has entranced and scandalized people for millennia. Indeed, part of what makes the history of 'artificial intelligence' so fascinating is the mix of genuine scientific achievement with myth-making and outright deception

Not the only ML+society pipeline in town

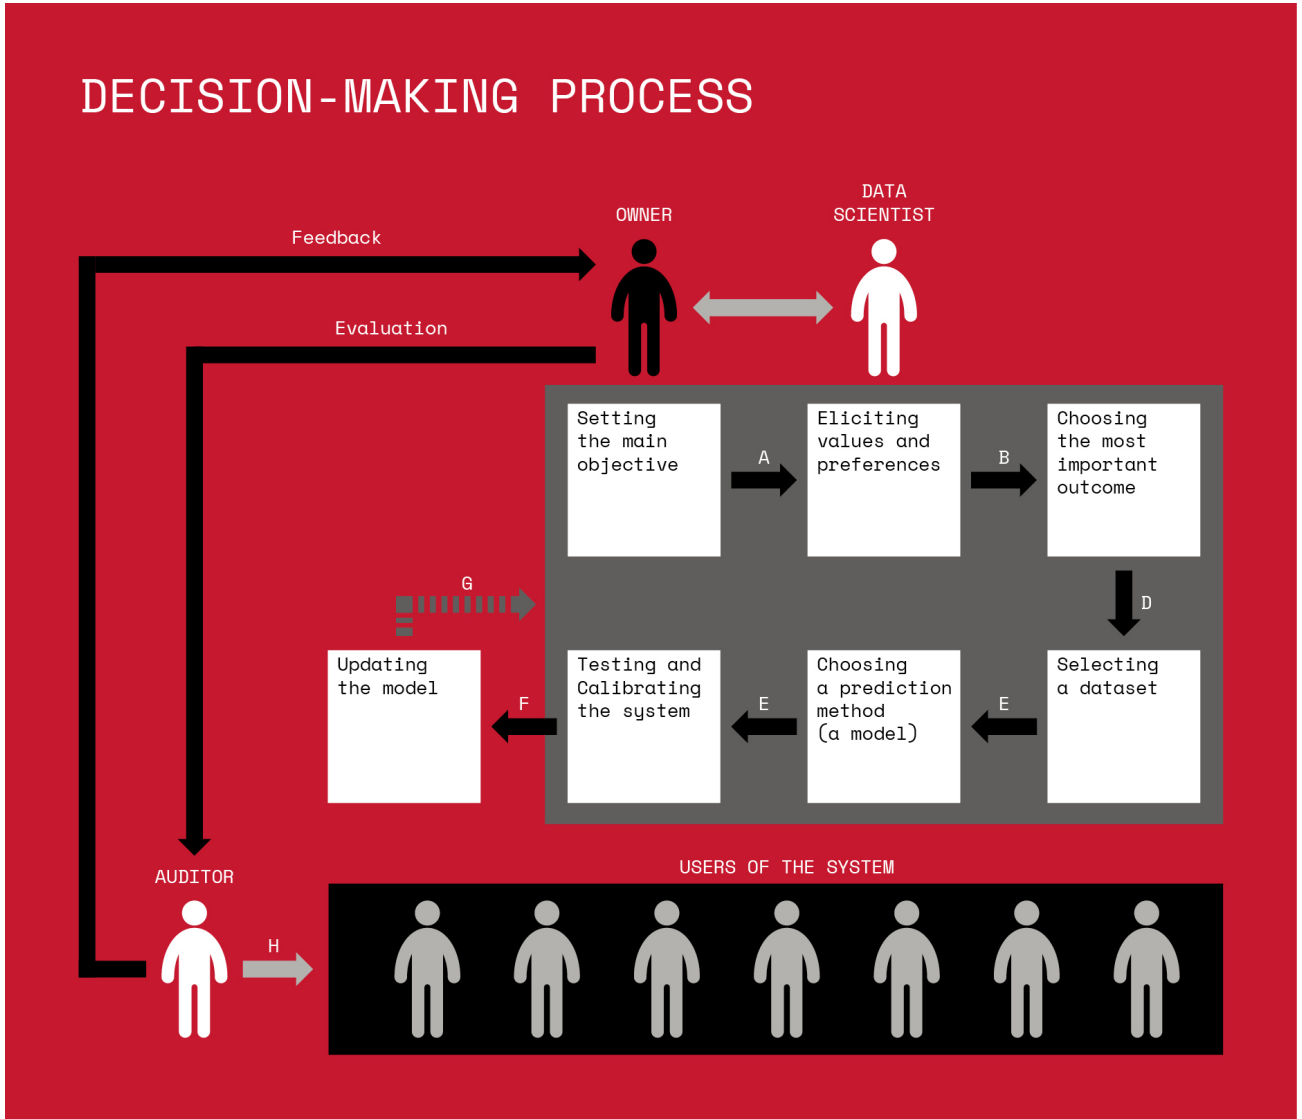

A walkthrough

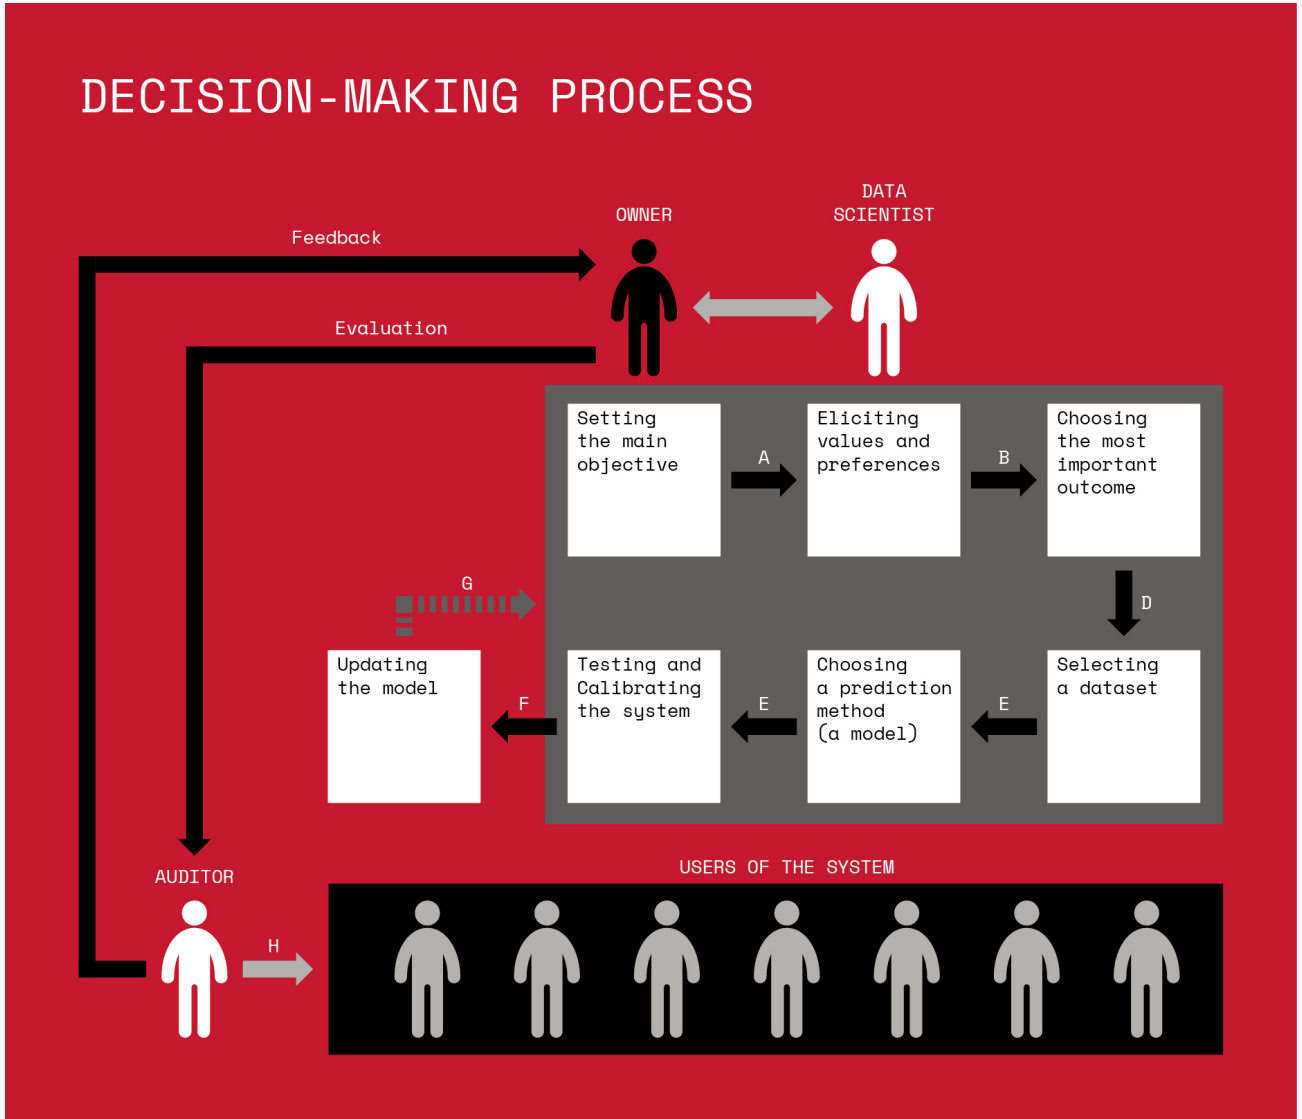

A walkthrough

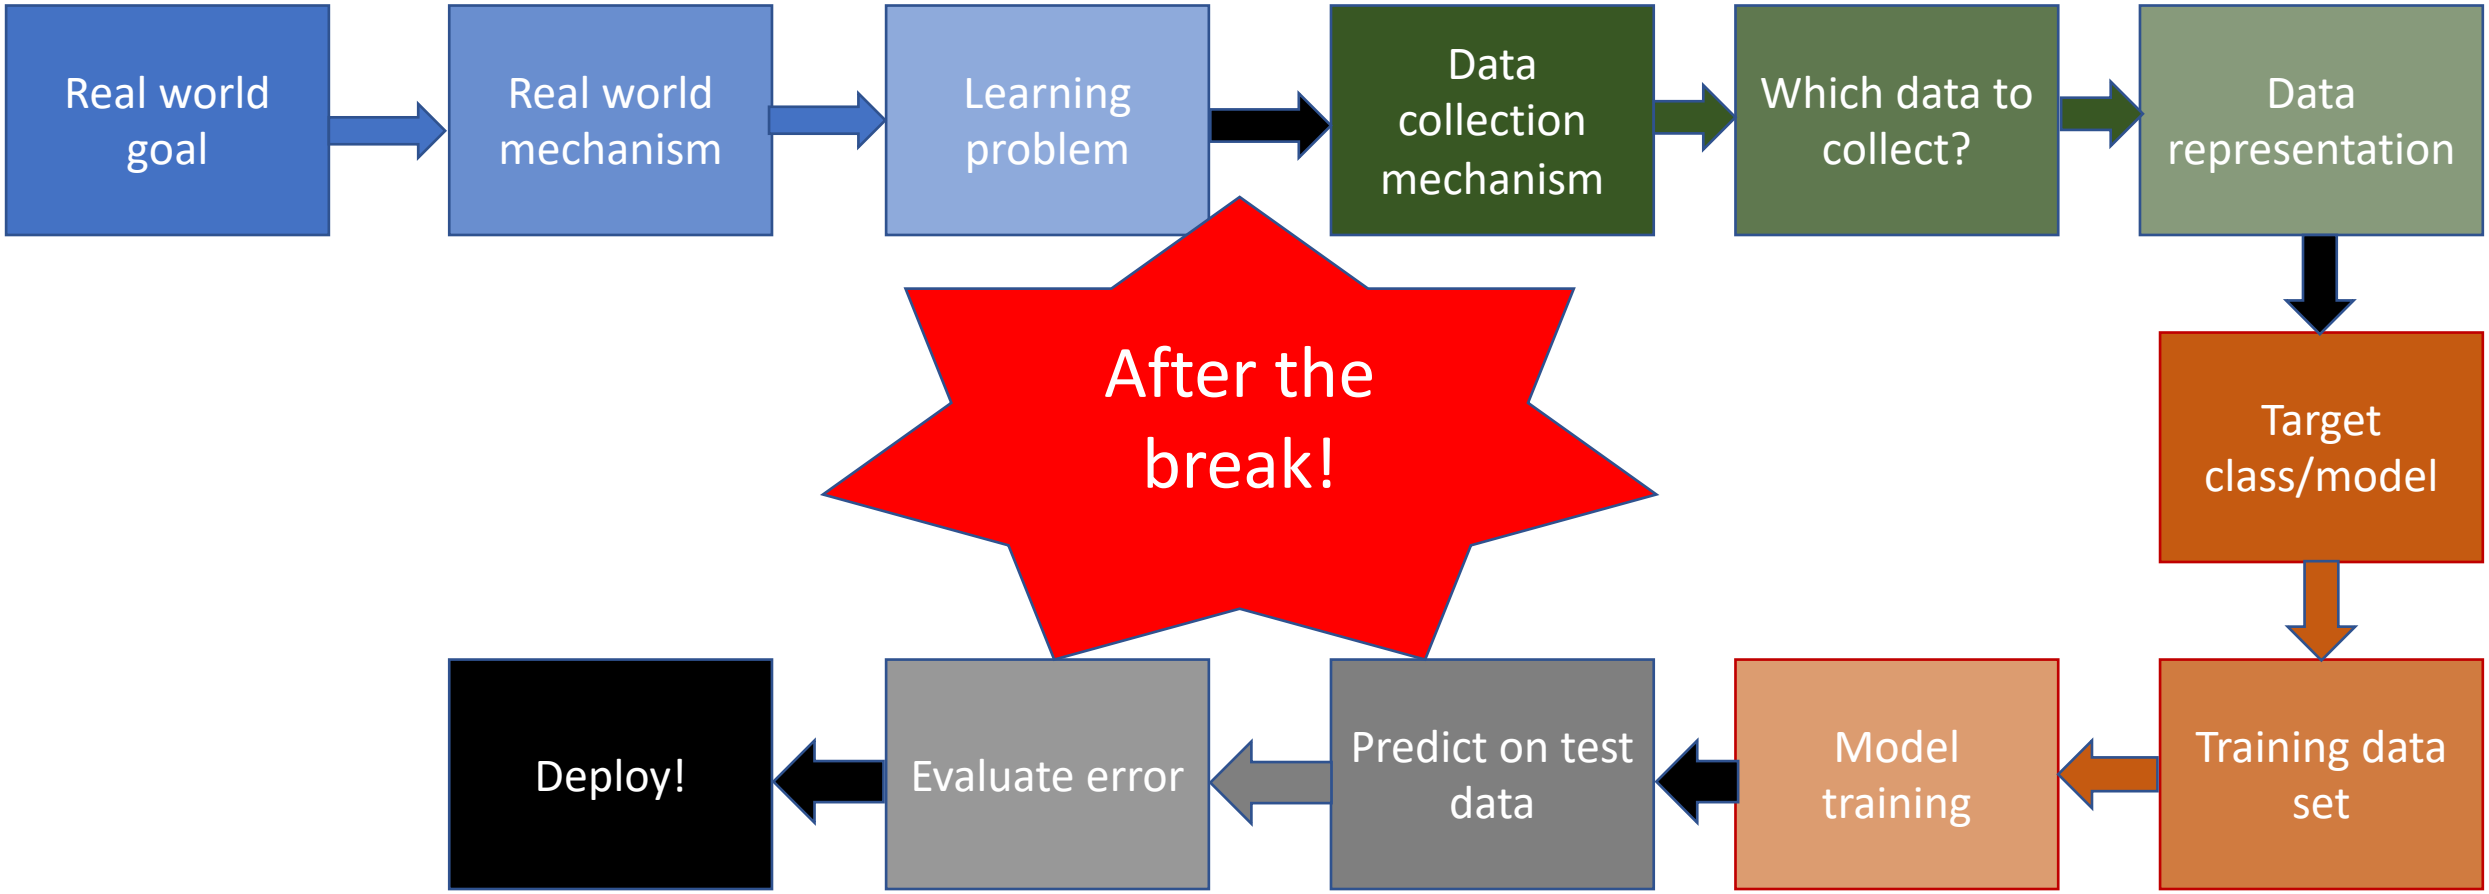

Have you heard of COMPAS?

COMPAS (software)

From Wikipedia, the free encyclopedia

COMPAS, an acronym for Correctional Offender Management Profiling for Alternative Sanctions, is a case manag
Equivant) used by U.S. courts to assess the likelihood of a defendant becoming a recidivist.[1][2]

COMPAS has been used by the U.S. states of New York, Wisconsin, California, Florida's Broward County, and oth

Contents [hide]

Risk Assessment [edit]

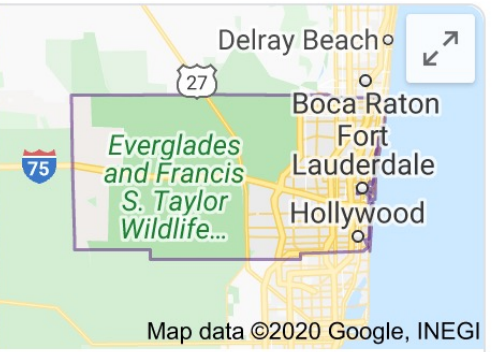

Broward County

County in Florida

Broward County is a county in southeastern Florida, US. According to a 2018 census report, the county had a population of 1,951,260, making it the second-most populous county in the state of Florida and the 17th-most populous county in the United States. The county seat is Fort Lauderdale. Wikipedia

Incorporated cities: 24

Population: 1.936 million (2017)

Mayor: Mark D. Bogen

PRO PUBLICA

Donate

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

A sample of their result

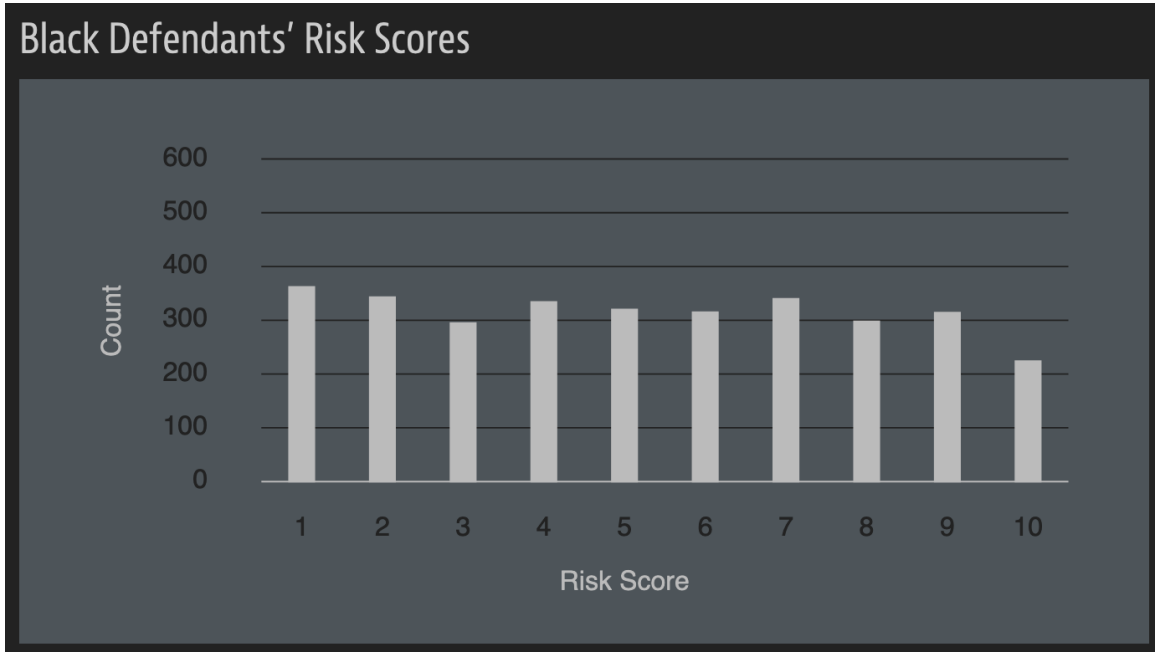

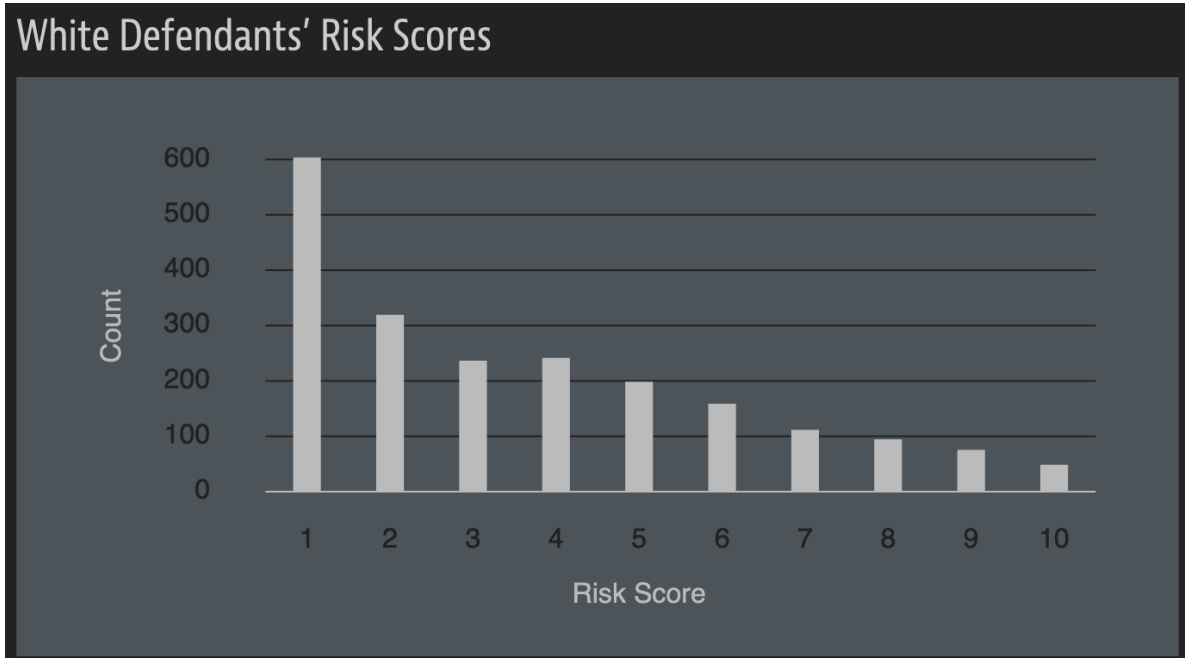

False Positives, False Negatives, and False Analyses: A Rejoinder to "Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks."

Anthony W. Flores
California State University, Bakersfield
Kristin Bechtel
Crime and Justice Institute at CRJ
Christopher T. Lowenkamp
Administrative Office of the United States Courts
Probation and Pretrial Services Office

A walkthrough

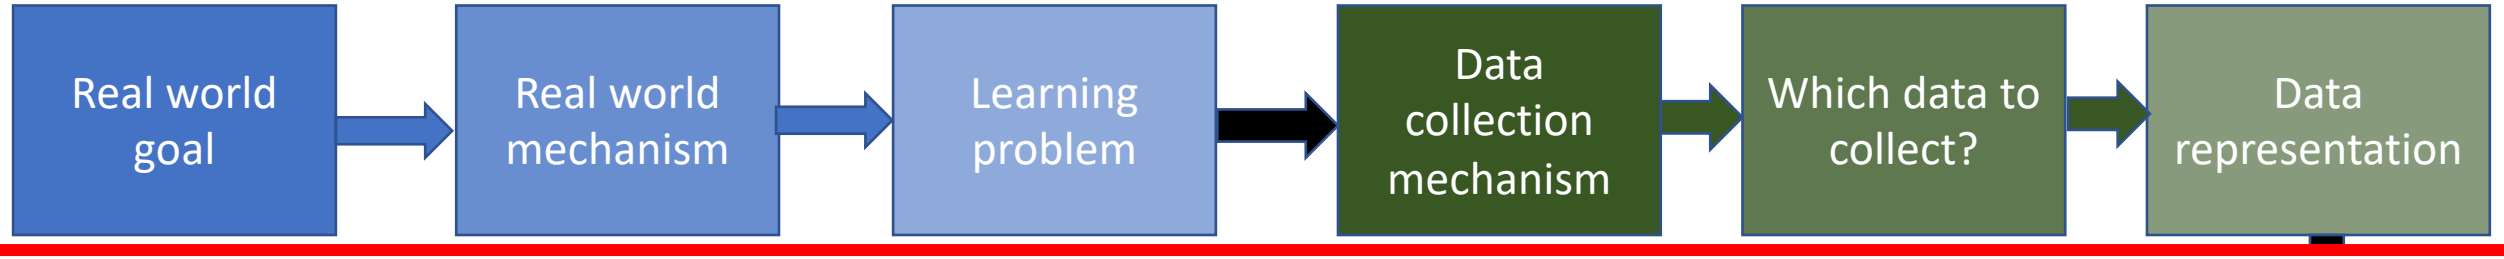

The Problem

Imagine a situation where the creator of COMPAS had access to the COMPAS dataset. In particular, you are in the team that wants to predict recidivism based on the COMPAS dataset. How would you go about doing it?

Well, let's just walk through the ML pipeline to see how you would go about doing this.

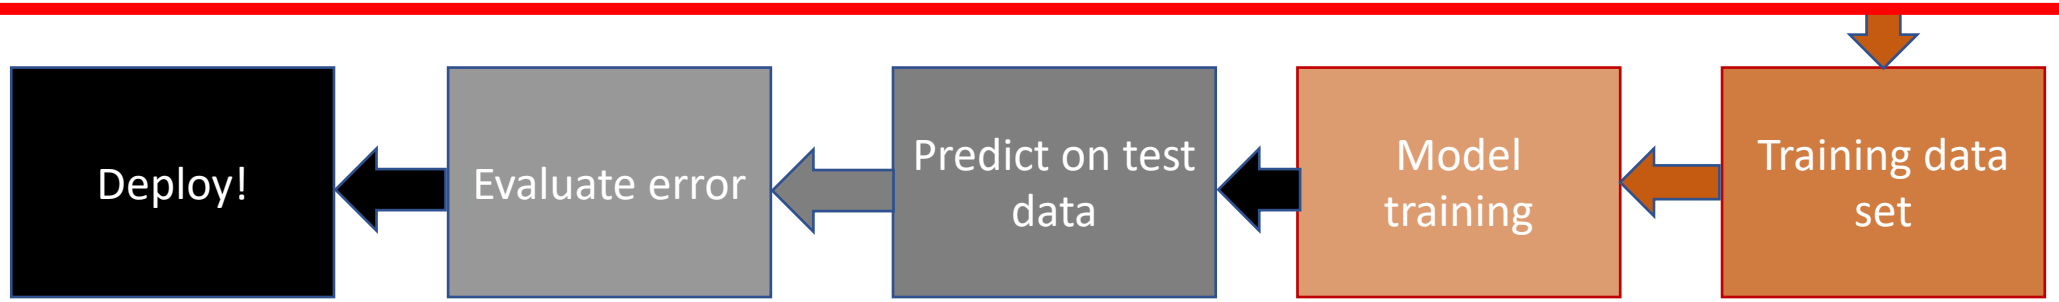

Real world goal

Real world goal

Real world goal

Reduce crime in society.

The Problem

Imagine a situation where the creator of COMPAS had access to the COMPAS dataset. In particular, you are in the team that wants to predict recidivism based on the COMPAS dataset. How would you go about doing it?

Well, let's just walk through the ML pipeline to see how you would go about doing this.

Real world mechanism

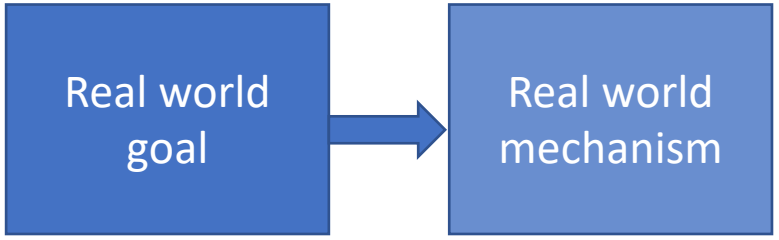

Real world mechanism

Based on some studies (or not!), your superiors decided that repeat offenders contribute most to crime. This in turn they decided would mean that if one could identify who would commit a crime again in the future, then one could use this information when making judgment on the current crime. Thus, they decided they wanted a system that can identify folks who will re-offend in the future and then promptly handed off the problem to your group to solve it.

The Problem

Imagine a situation where the creator of COMPAS had access to the COMPAS dataset. In particular, you are in the team that wants to predict recidivism based on the COMPAS dataset. How would you go about doing it?

Well, let's just walk through the ML pipeline to see how you would go about doing this.

Learning problem

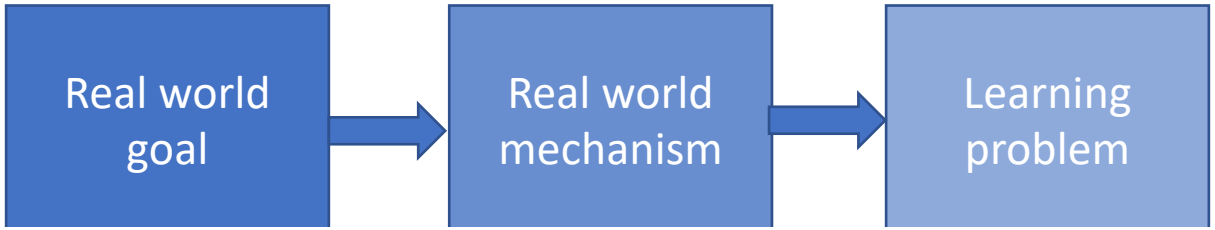

Learning problem

Your group decides on the simplest learning problem: given a defendant *predict* if they will re-offend or not (in other words you are doing *binary classification* (binary because you are "labeling" defendants as either going too re-offend or not going to re-offend and you are doing classification because you are putting people into the two bins-- i.e. giving them a binary label and hence assigning them a "class."

There is another related option (which is what COMPAS: instead of assigning defendants to two scores: they assign a score from 1 (being least likely to re-offend) to 10 (most likely to defends). This range of score (rather than a binary classification) could potentially be more useful to the end user of your system.

However, for our discussion (and indeed for most of the rest of the course), we will focus on binary classification.

The Problem

Imagine a situation where the creator of COMPAS had access to the COMPAS dataset. In particular, you are in the team that wants to predict recidivism based on the COMPAS dataset. How would you go about doing it?

Well, let's just walk through the ML pipeline to see how you would go about doing this.