

Unstructured Data Analysis

Lecture 11: Wrap-up predictive model evaluation, classical classifiers; intro to neural nets & deep learning


George Chen

Today

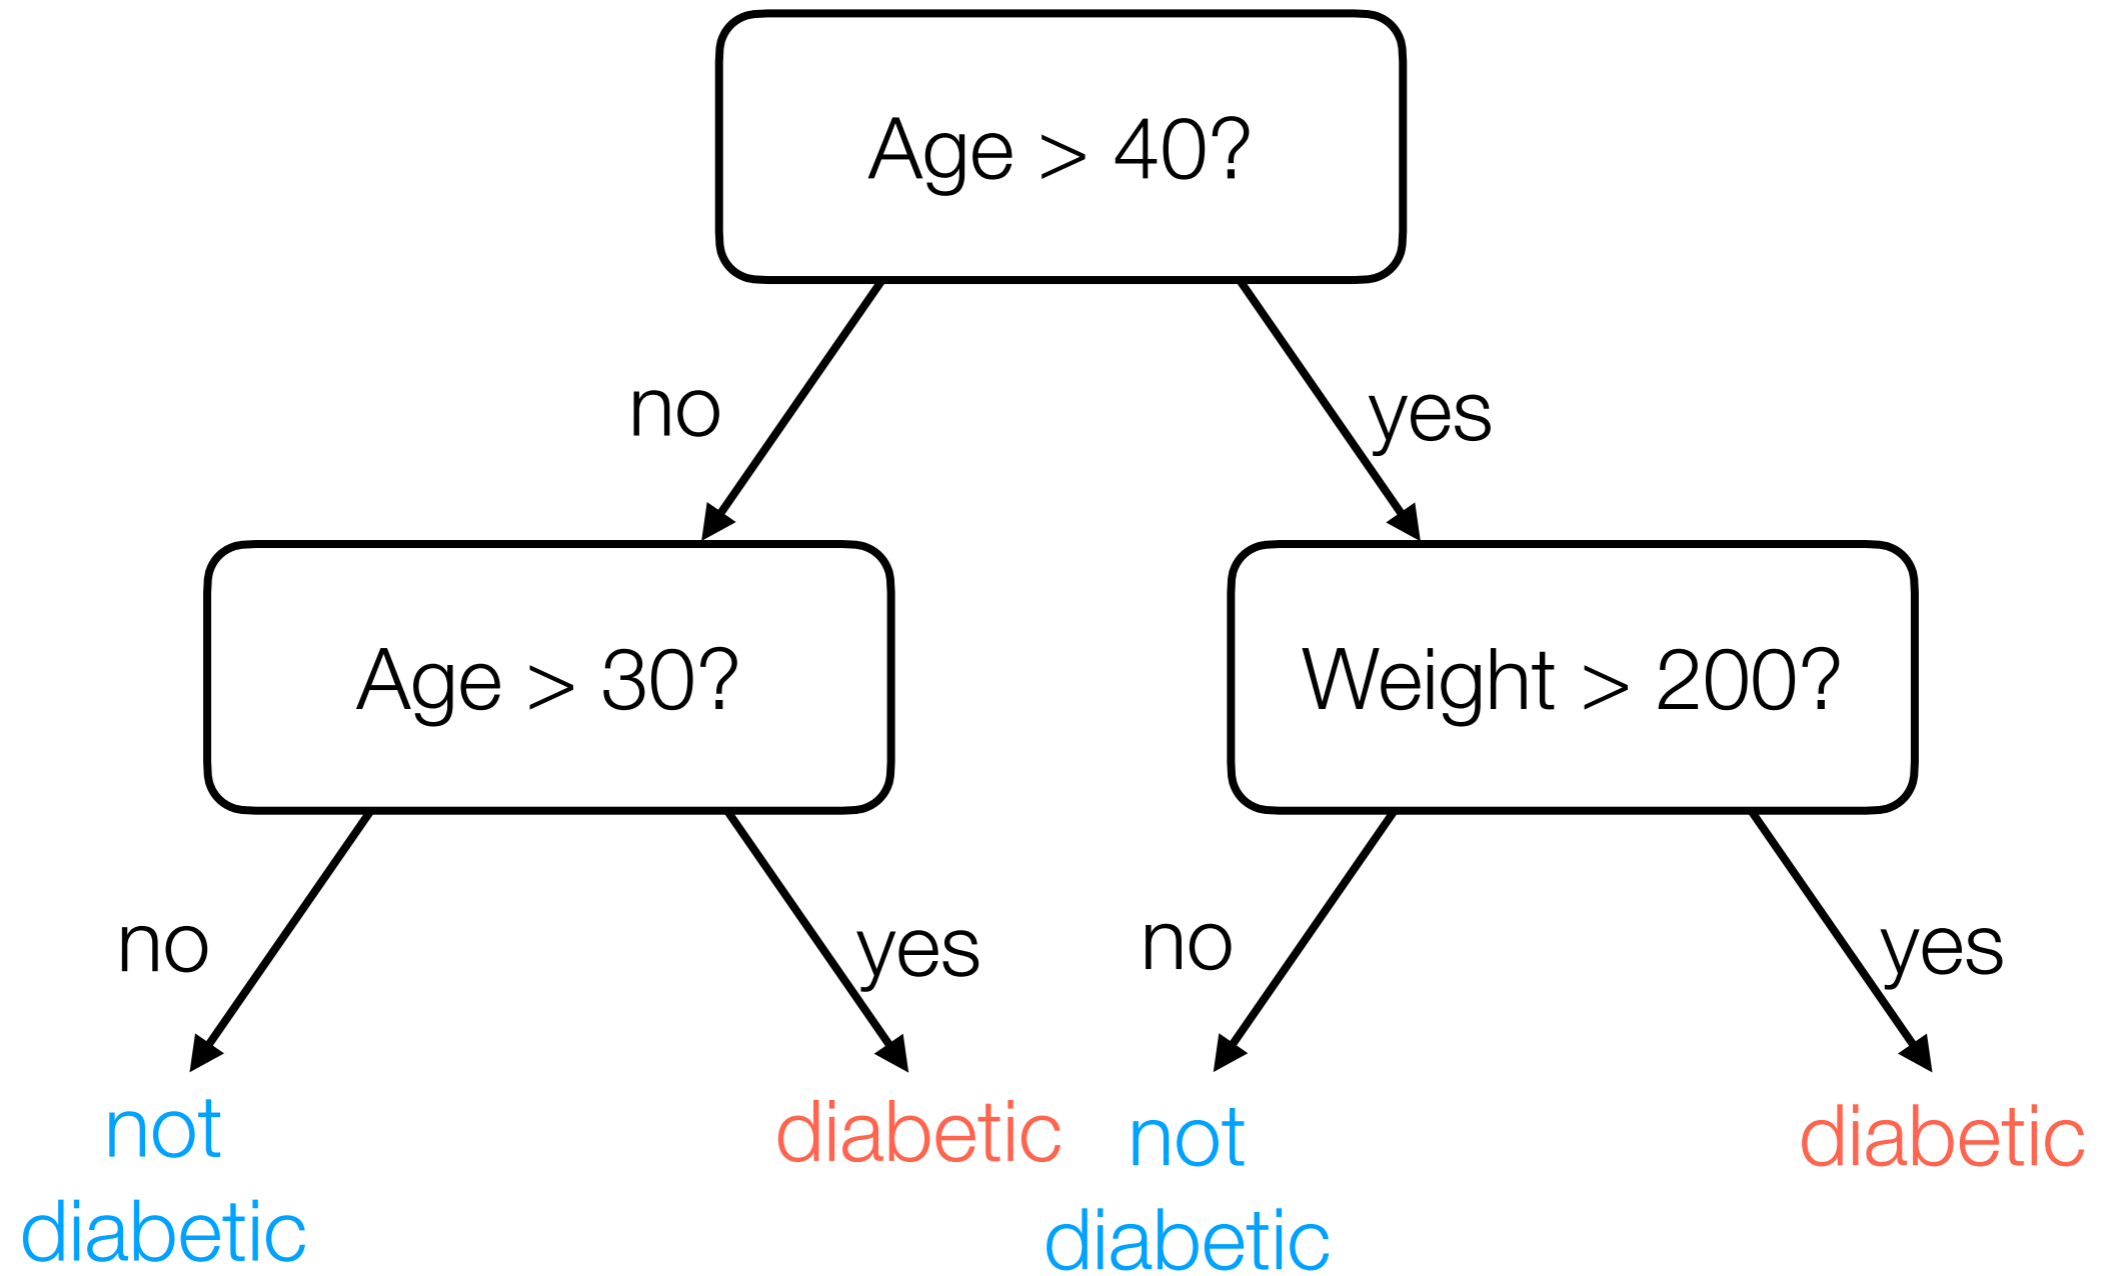
- Wrap-up coverage of how to evaluate whether a predictive model is good & classical classifiers
- In many datasets (especially small, structured ones), neural nets & deep learning could work poorly... in such cases, often decision-tree-based methods can work well
random forests, gradient boosting (e.g., XGBoost)
- Start coverage of neural nets & deep learning

Decision Trees & Forests

Example Made-Up Data



Example Decision Tree




Learning a Decision Tree

- Many ways: general approach actually looks a lot like divisive clustering *but accounts for label information*
- I'll show one way (that nobody actually uses in practice) but it's easy to explain

Learning a Decision Tree


1. Pick a random feature
(either age or weight)



2. Find threshold for which red and blue are as “separate as possible” (on one side, mostly red; on other side, mostly blue)

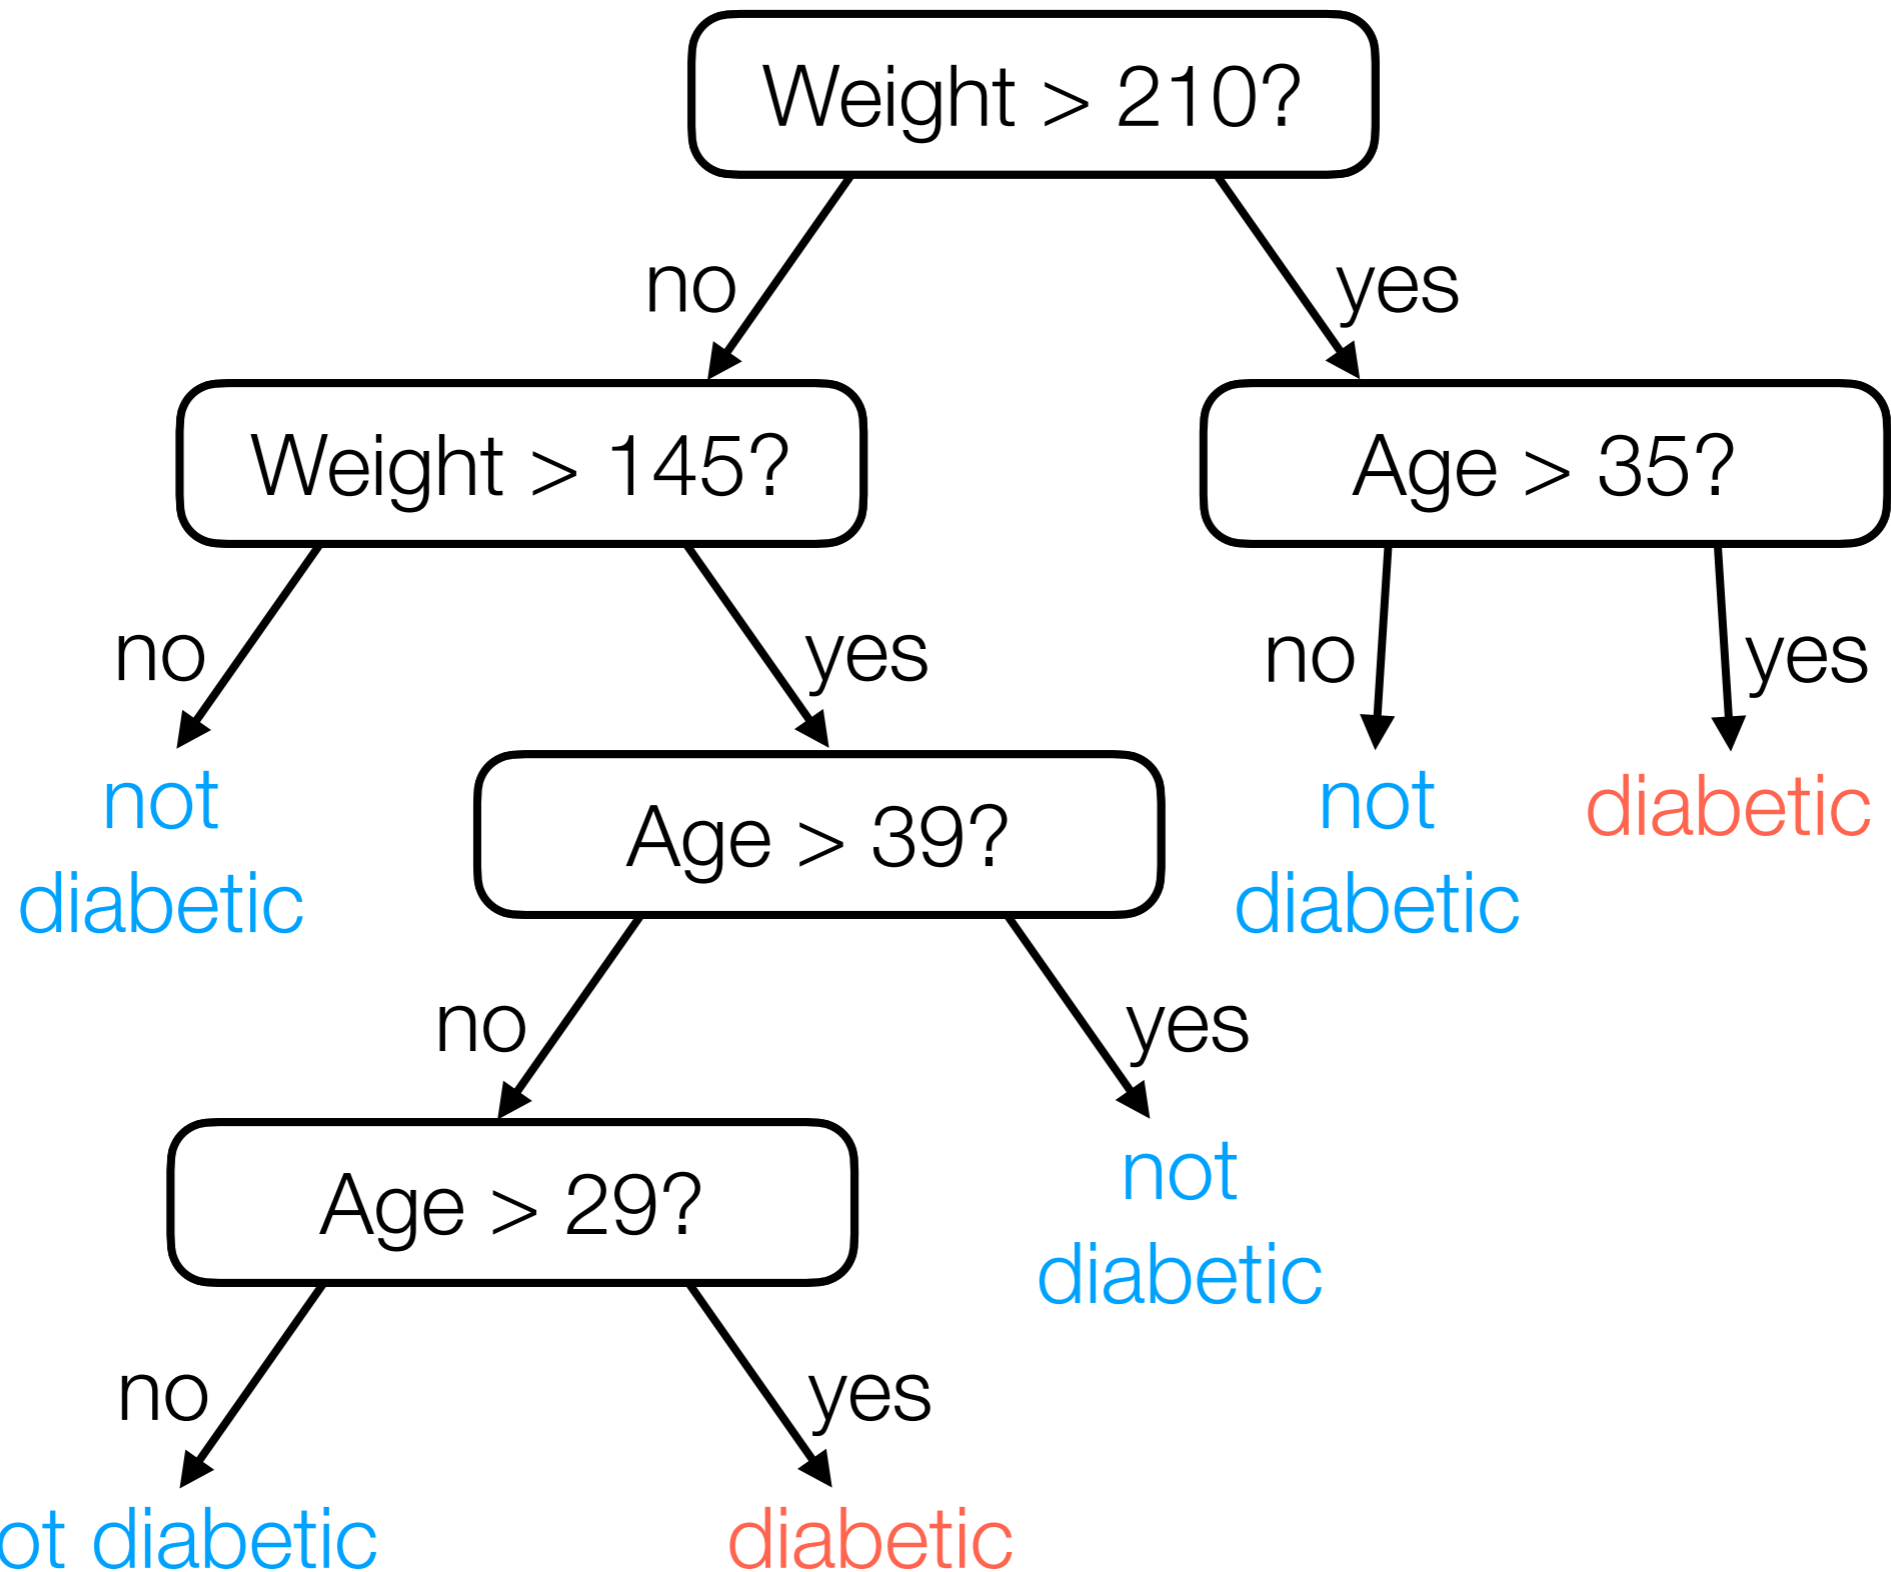
Learning a Decision Tree

Within each side, recurse until a termination criterion is reached!



Example termination criteria: $\geq 90\%$ points within region has same label, number of points within region is < 5

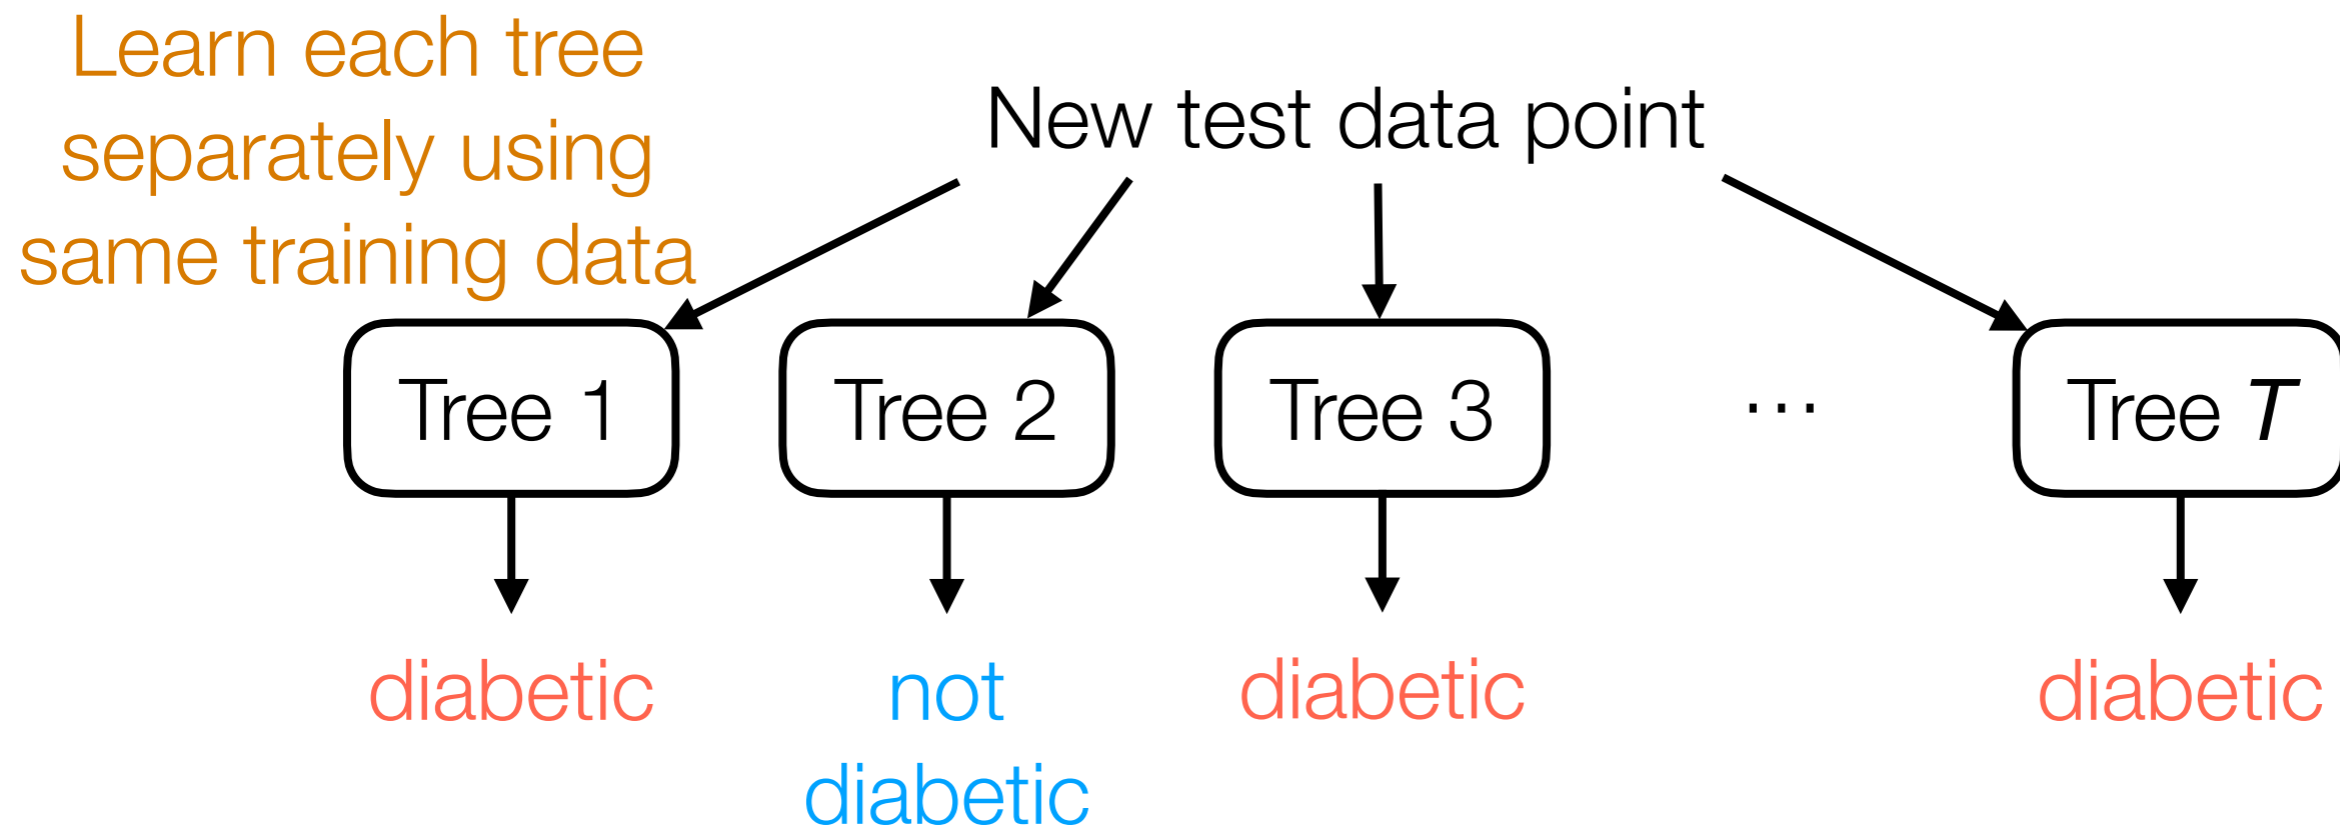
Decision Tree Learned



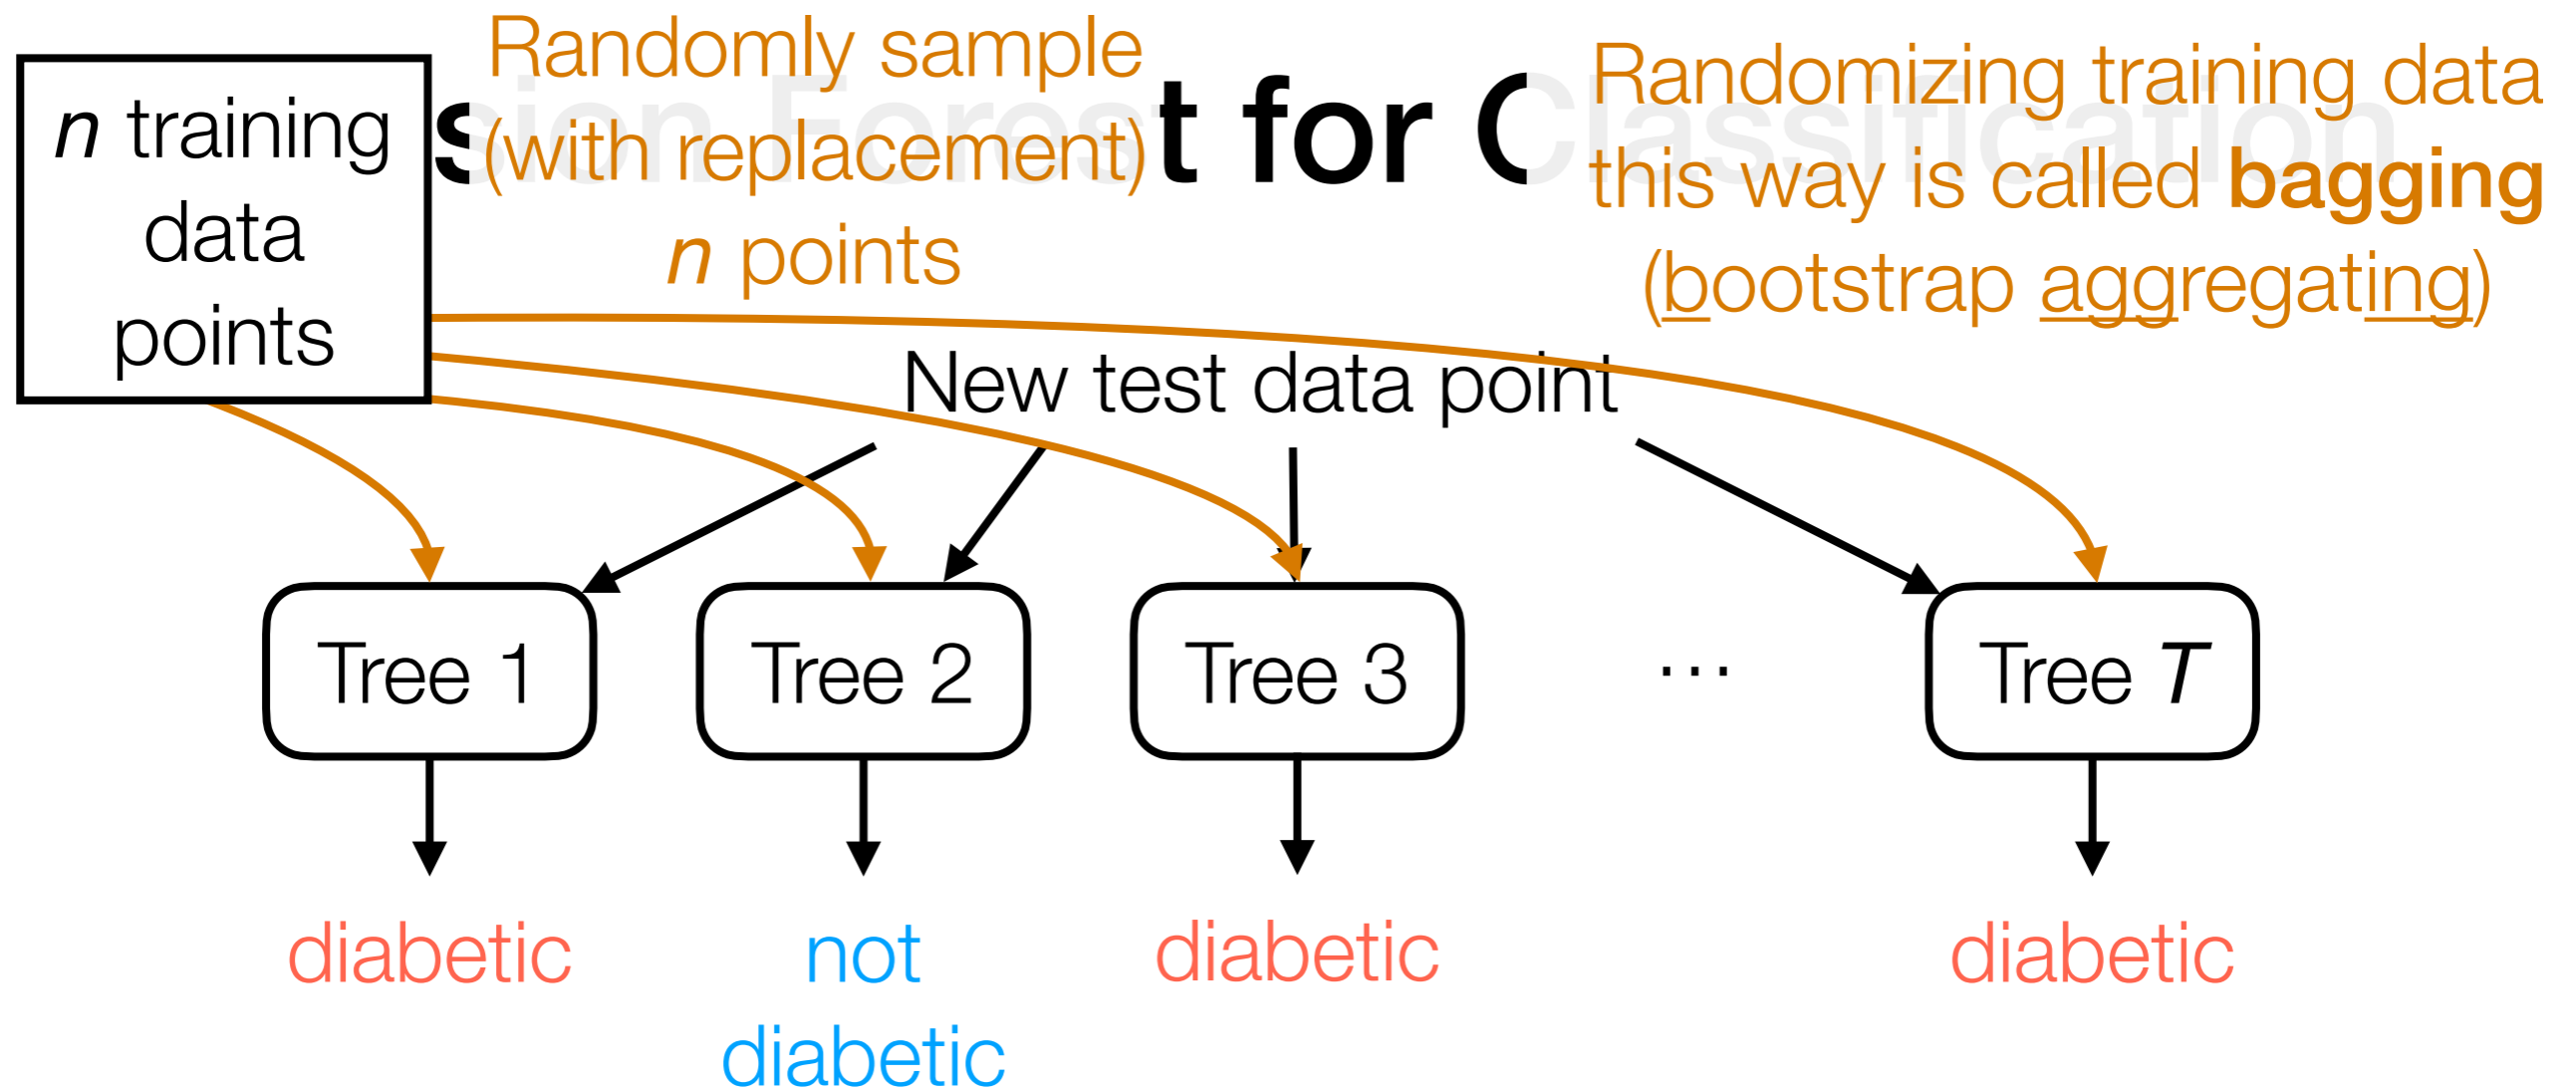
For a new person with feature vector (age, weight), easy to predict!

Decision Forest for Classification

- Typically, a decision tree is learned with randomness (e.g., we randomly chose which feature to threshold)
 - by re-running the same learning procedure, we can get different decision trees that make different predictions!
- For a more stable prediction, use many decision trees



Final prediction: majority vote of the different trees' predictions



Question: What happens if all the trees are the same?

Adding randomness can make trees more different!

- **Random Forest:** randomize training data used for each tree, randomly choose a few features to try to split on (and among these features, choose the best one to split on)

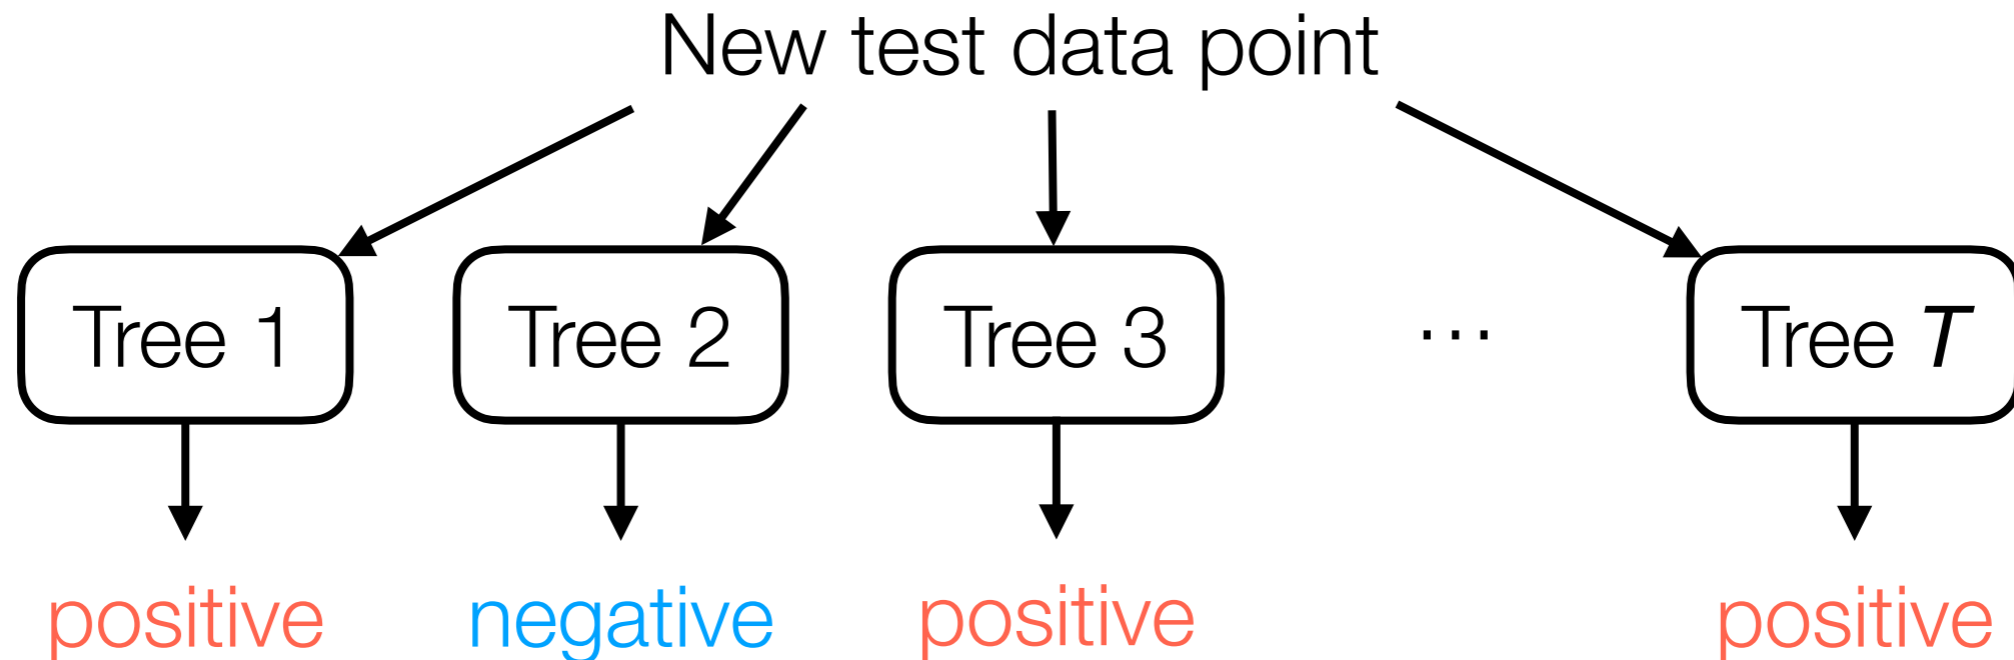
Back to the demo

Another Way to Benchmark

- In the demo: we just saw that we can compare test set prediction accuracy across different algorithms and also look at confusion matrices
- For binary classification, we can do a more detailed analysis

Binary Classification: ROC Curves

For simplicity, think of the random forest for now

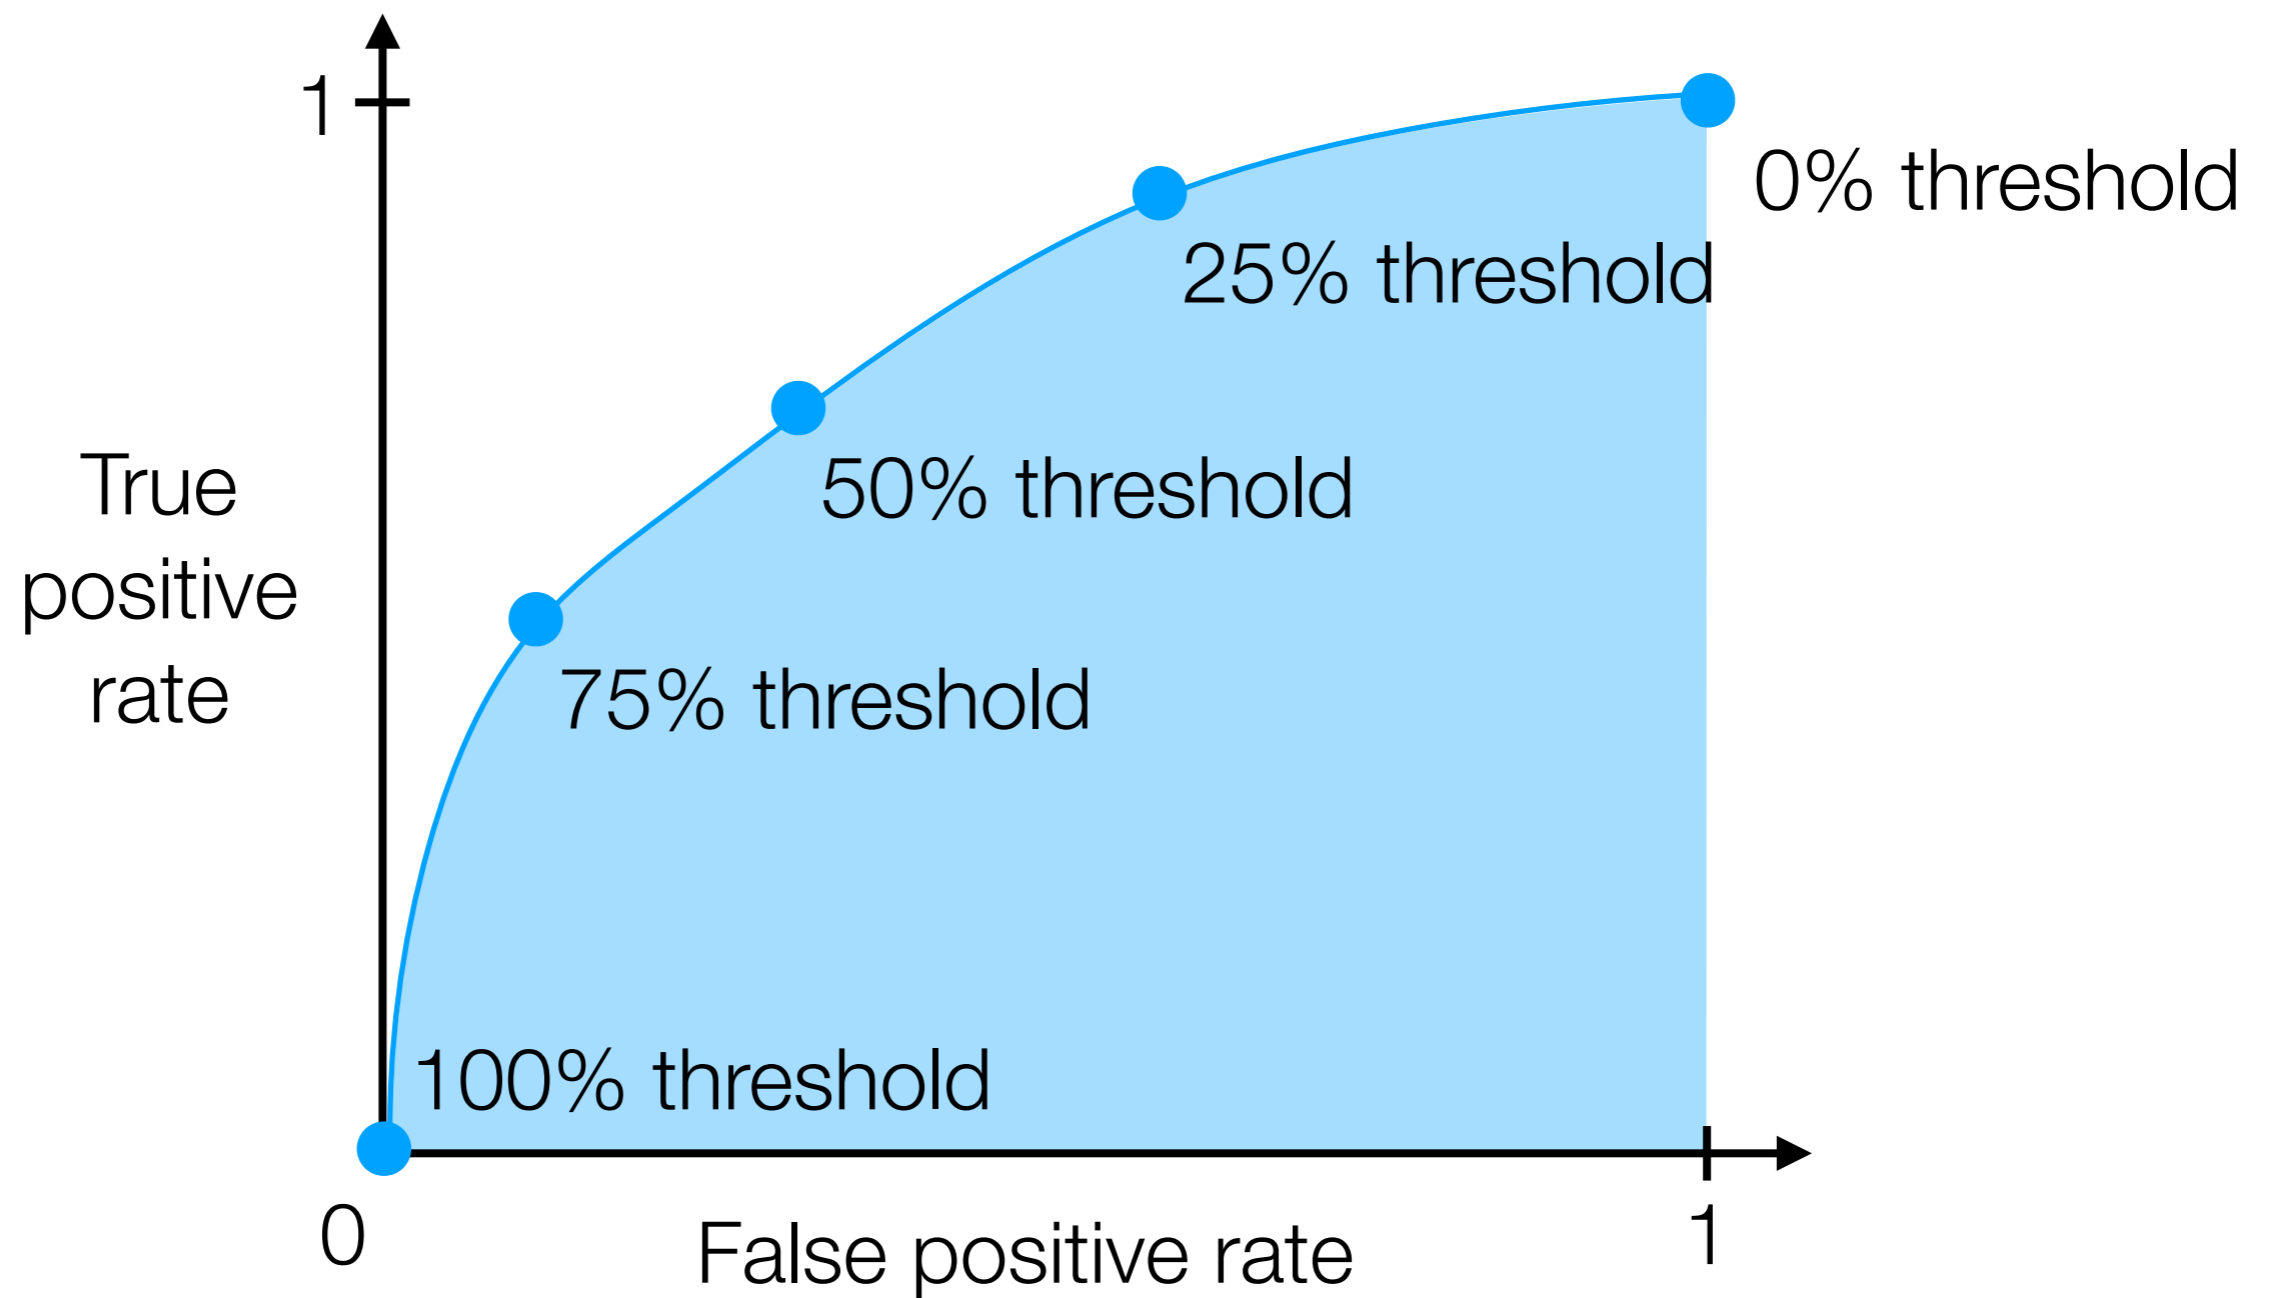


Final prediction: majority vote of the different trees' predictions

$\geq 50\%$ of trees need to say **positive** for final prediction to be **positive**

We can vary this 50% threshold!


Binary Classification: ROC Curves



Error rates are computed on test data


Binary Classification: ROC Curves

A classifier with the green curve is better than the one with the blue curve

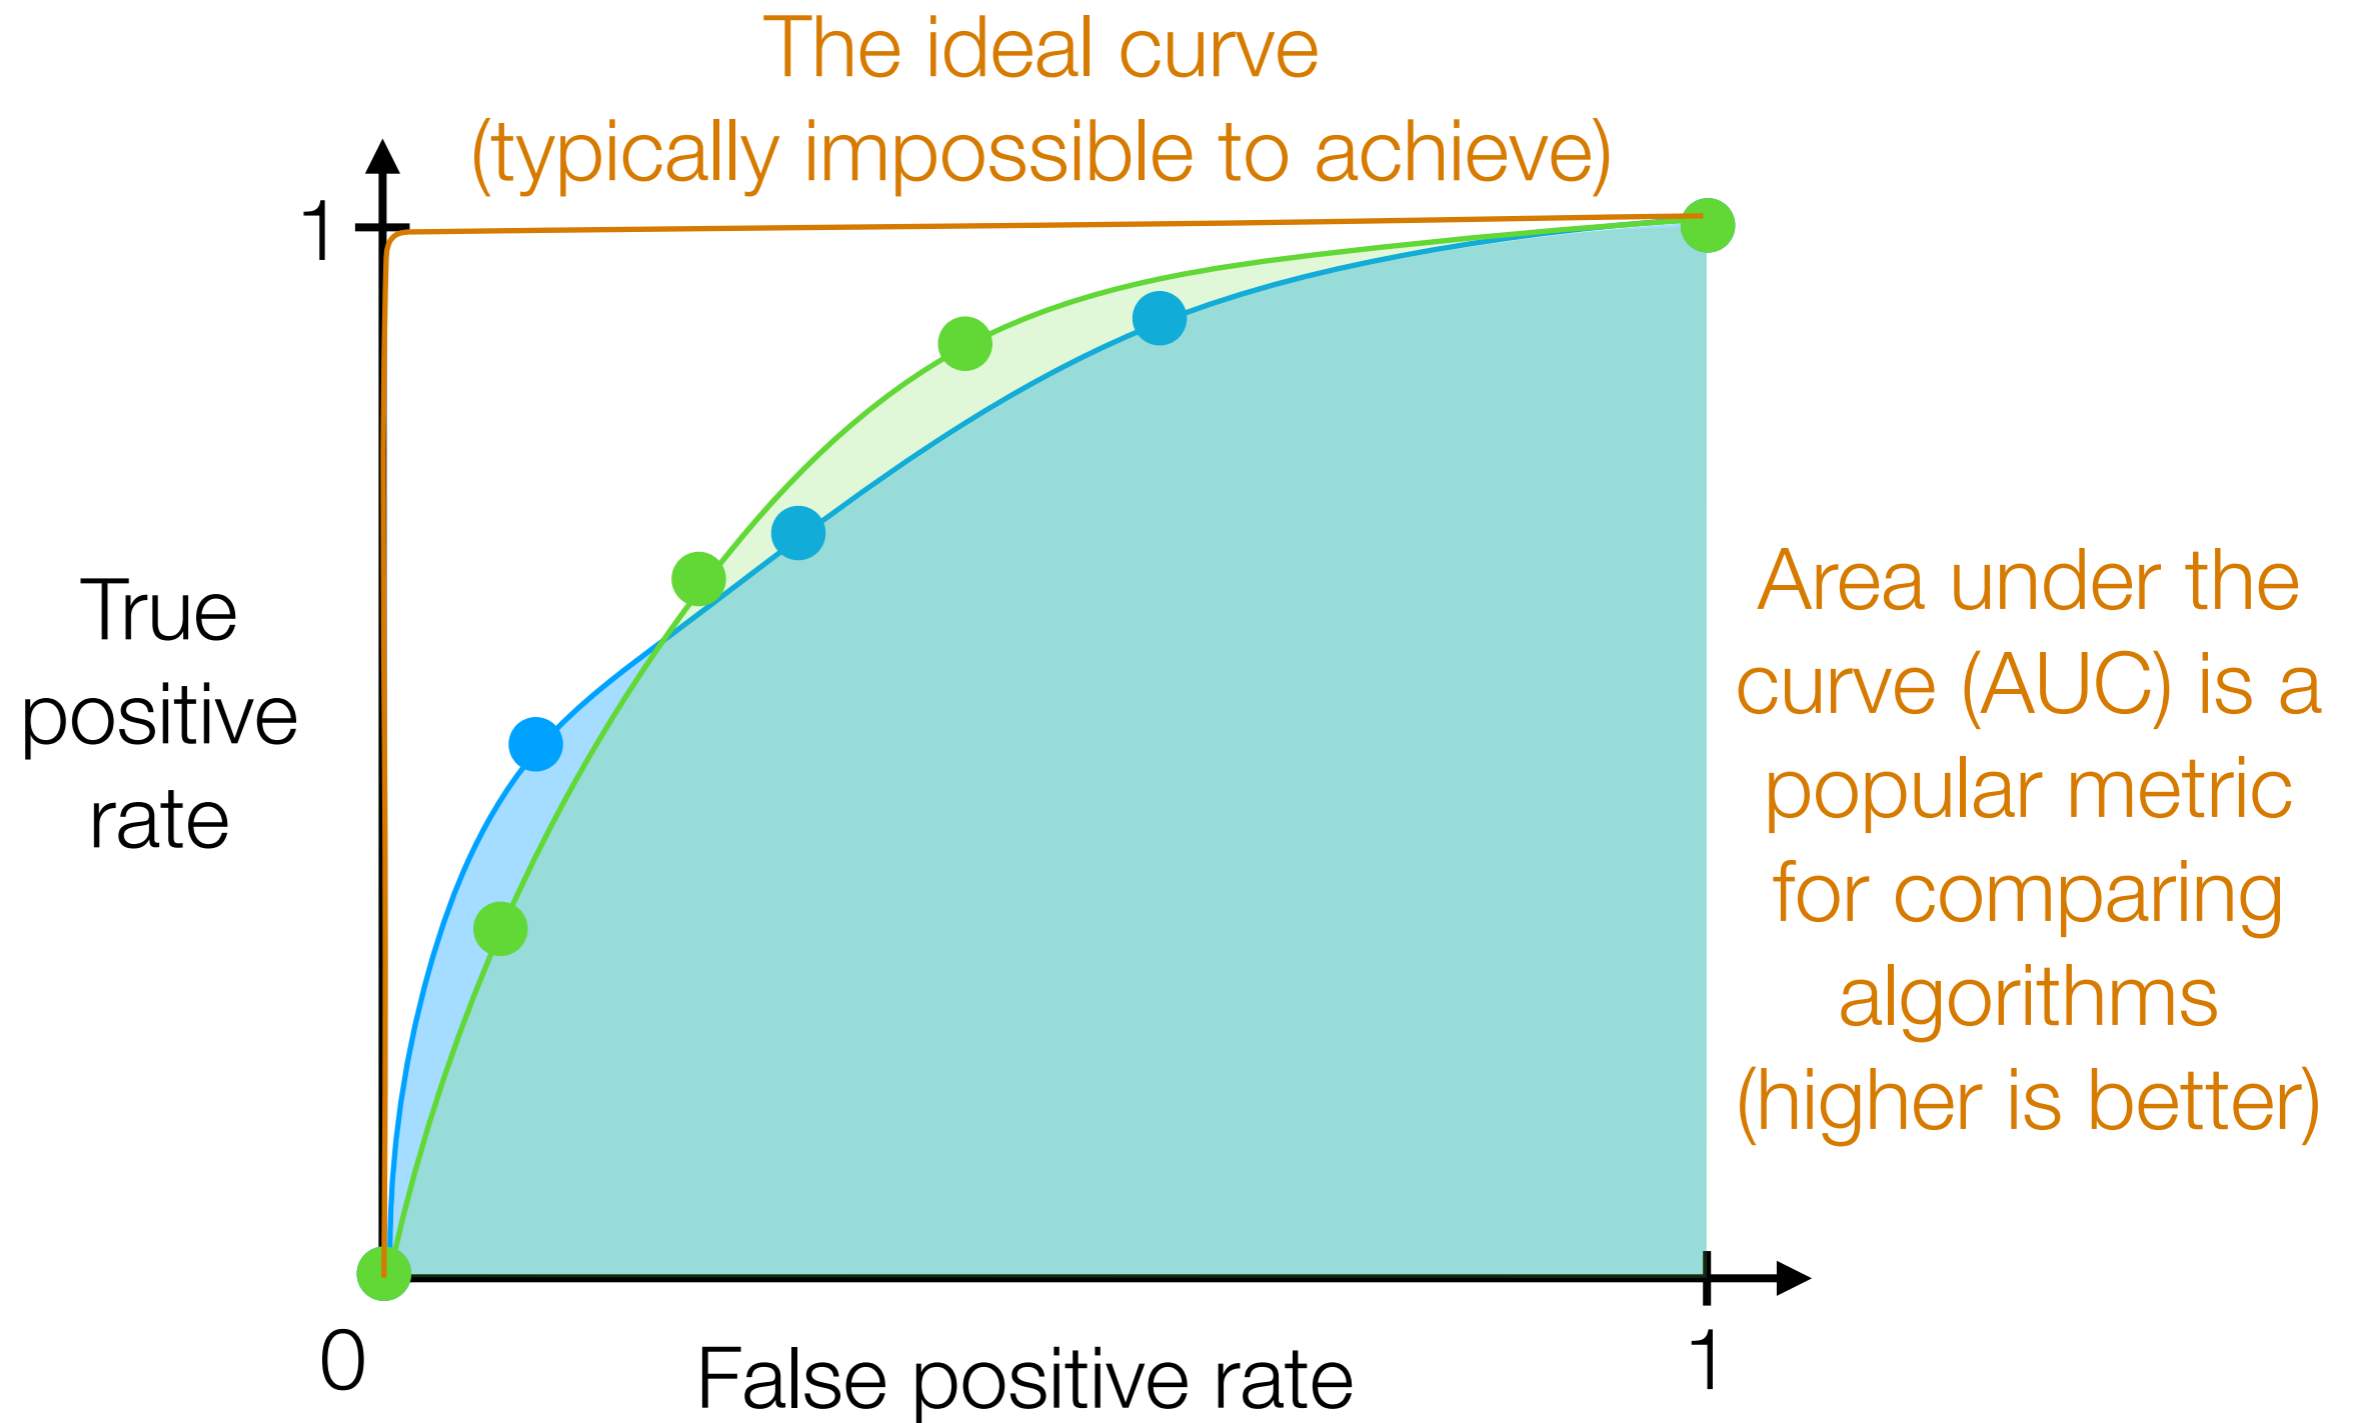


Binary Classification: ROC Curves

It's possible that algorithms are better in different regimes



Binary Classification: ROC Curves



Binary Classification: ROC Curves


What we just saw:

- For a classifier that we can set the threshold probability to different values, we can plot an ROC curve
- True positive rate (TPR) and false positive rate (FPR) are evaluated on test data

Other variants are possible:

- Plot precision vs recall instead of TPR vs FPR
- Can actually plot ROC/precision-recall curves sweeping over hyperparameters aside from threshold probability!
- For ROC/precision-recall, rather than evaluating on test data, can evaluate on validation data during training *to help choose hyperparameters*

Binary Classification: ROC Curves



Can also be computed on validation data instead of test data!

Intro to Neural Nets & Deep Learning

IMAGENET

Over 10 million images, 1000 object classes



2011: Traditional computer vision achieves accuracy ~74%

2012: Initial deep neural network approach accuracy ~84%

2015 onwards: Deep learning achieves accuracy 96%+

Russakovsky et al. ImageNet Large Scale Visual Recognition Challenge. IJCV 2015.

Deep Learning

Extremely useful in practice:

- Near human level image classification (including handwritten digit recognition)
- Near human level speech recognition
- Improvements in machine translation, text-to-speech
- Self-driving cars
- *Better* than humans at playing Go



Google DeepMind's AlphaGo vs Lee Sedol, 2016

GAMING

TECH

ARTIFICIAL INTELLIGENCE

DeepMind's StarCraft 2 AI is now better than 99.8 percent of all human players

16 

AlphaStar is now grandmaster level in the real-time strategy game

By [Nick Statt](#) | [@nickstatt](#) | Oct 30, 2019, 2:00pm EDT



SHARE



Turing Award Won by 3 Pioneers in Artificial Intelligence



From left, Yann LeCun, Geoffrey Hinton and Yoshua Bengio. The researchers worked on key developments for neural networks, which are reshaping how computer systems are built. From left, Facebook, via Associated Press; Aaron Vincent Elkaim for The New York Times; Chad Buchanan/Getty Images

By **Cade Metz**

March 27, 2019




Is it all hype?



panda
~58% confidence

+ .007 ×




adversarial
noise

=



gibbon
~99% confidence

Source: Goodfellow, Shlens, and Szegedy. Explaining and Harnessing Adversarial Examples. ICLR 2015.

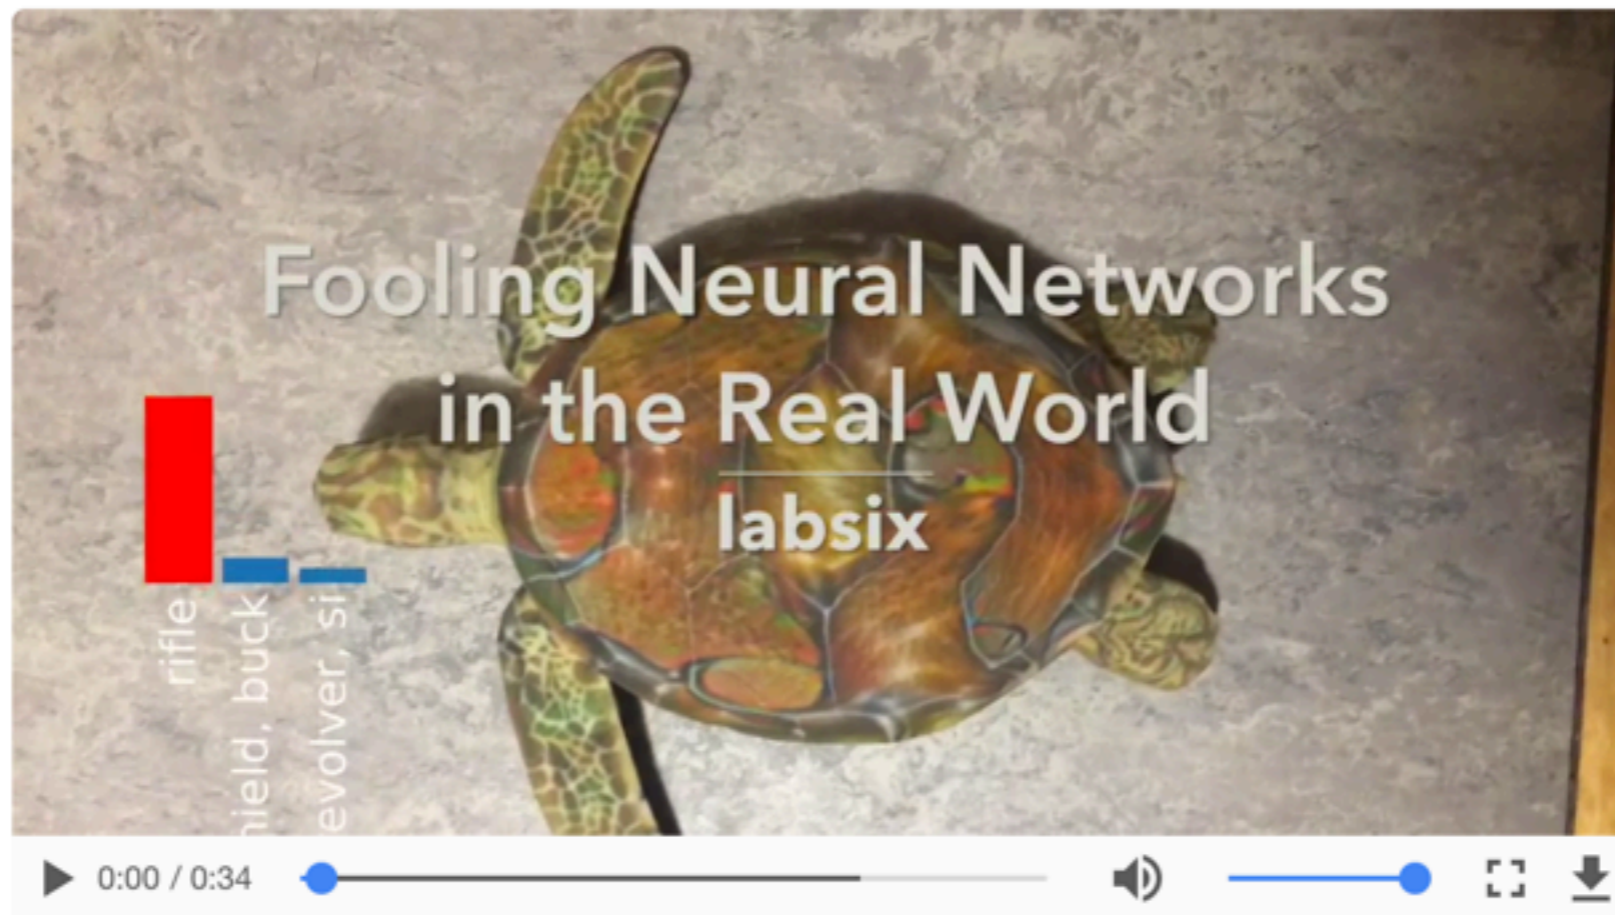


Source: Papernot et al. Practical Black-Box Attacks against Machine Learning. Asia Conference on Computer and Communications Security 2017.

Fooling Neural Networks in the Physical World with 3D Adversarial Objects

31 Oct 2017 · 3 min read — shared on [Hacker News](#), [Lobsters](#), [Reddit](#), [Twitter](#)

We've developed an approach to generate *3D adversarial objects* that reliably fool neural networks in the real world, no matter how the objects are looked at.



Neural network based classifiers reach near-human performance in many tasks, and they're used in high risk, real world systems. Yet, these same neural networks are particularly vulnerable to *adversarial examples*, carefully perturbed inputs that cause

Source: labsix




Source: <https://www.cc.gatech.edu/news/611783/erasing-stop-signs-shapeshifter-shows-self-driving-cars-can-still-be-manipulated>



Source: Gizmodo article "This Neural Network's hilariously bad image descriptions are still advanced AI". September 16, 2015. (They're using the NeuralTalk image-to-caption software.)


GENERAL MORE MODELS ▾



The image shows a black and white cow lying in a lush green field. In the background, there are rolling green hills under a cloudy sky. Other cows are visible grazing in the distance.

General	VIEW DOCS
cow	0.992
cattle	0.983
mammal	0.979
grass	0.978
livestock	0.966
farm	0.964
landscape	0.963
pasture	0.954
grassland	0.949
agriculture	0.948
no person	0.945

Source: Pietro Perona




General [VIEW DOCS](#)

no person	0.991
beach	0.990
water	0.985
sand	0.981
sea	0.980
travel	0.978
seashore	0.972
summer	0.954
sky	0.946
outdoors	0.944
ocean	0.936

cow is not among top objects found!

Source: Pietro Perona

GENERAL FACE NSFW COLOR MORE MODELS



PREDICTED CONCEPT	PROBABILITY
group	0.979
adult	0.977
people	0.976
furniture	0.960
room	0.957
business	0.903
indoors	0.901
man	0.896
seat	0.895

VIEW DOCS

elephant is not among top objects found!

Source: David Lopez-Paz

Another AI Winter?

~1970's: First AI winter over symbolic AI

~1980's: Second AI winter over "expert systems"

Every time: Lots of hype, explosion in funding, then bubble bursts



Michael Jordan [Follow](#)

Michael I. Jordan is a Professor in the Department of Electrical Engineering and Computer Sciences and the Department of Statistics at UC Berkeley.

Apr 18 · 16 min read




Photo credit: Peg Skorpinski

Artificial Intelligence—The Revolution Hasn't Happened Yet

Artificial Intelligence (AI) is the mantra of the current era. The phrase is intoned by technologists, academicians, journalists and venture capitalists

<https://medium.com/@mijordan3/artificial-intelligence-the-revolution-hasnt-happened-yet-5e1d5812e1e7>

What is deep learning?