ML and Society

Mar 8, 2022

Discussion Summaries due tomorrow

By 5pm on Autolab

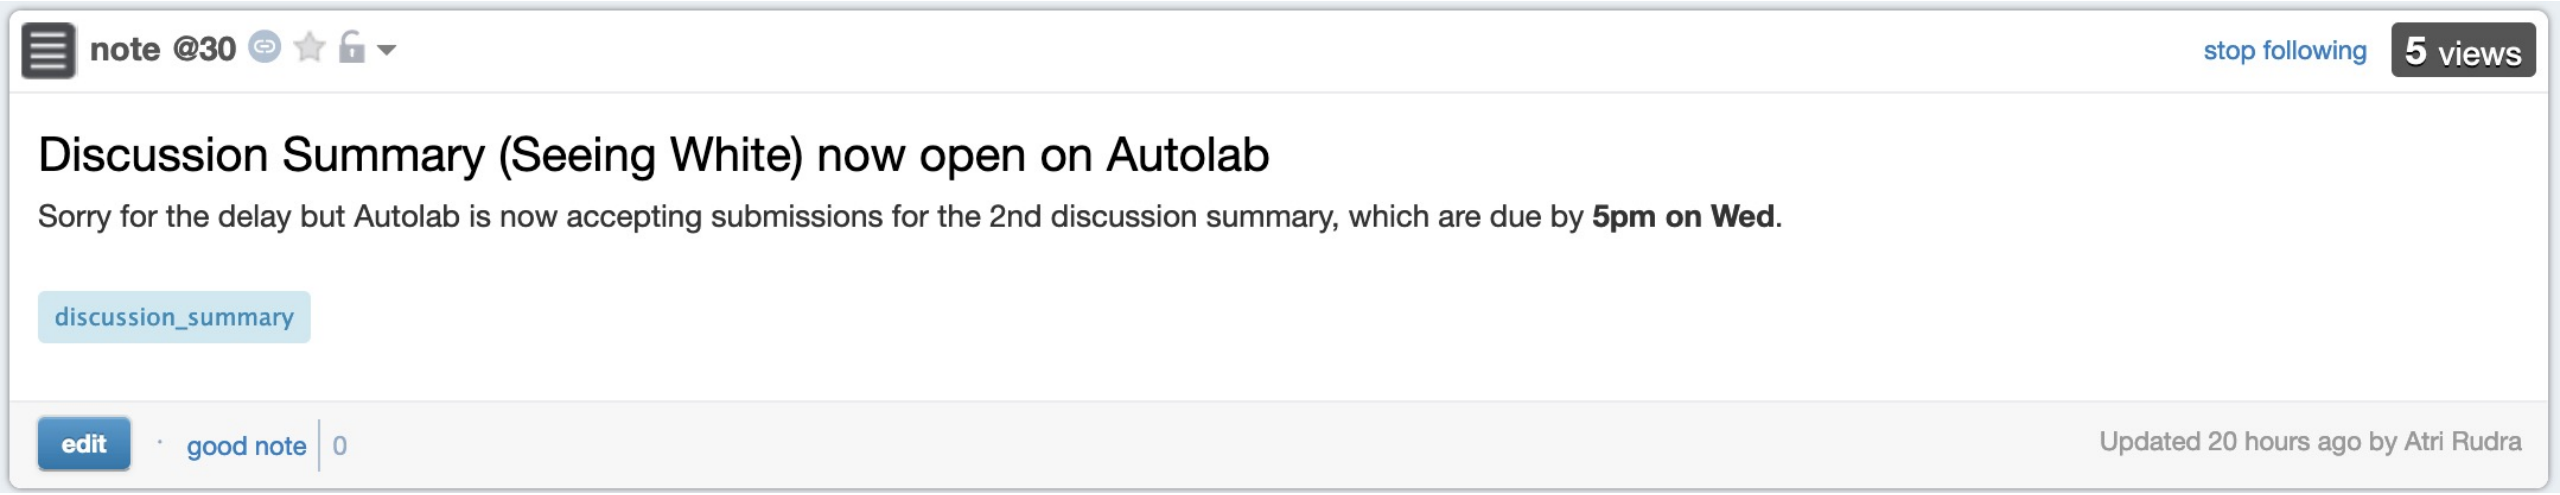

Delay in Impossible Project deets

Where we left off last Tue…

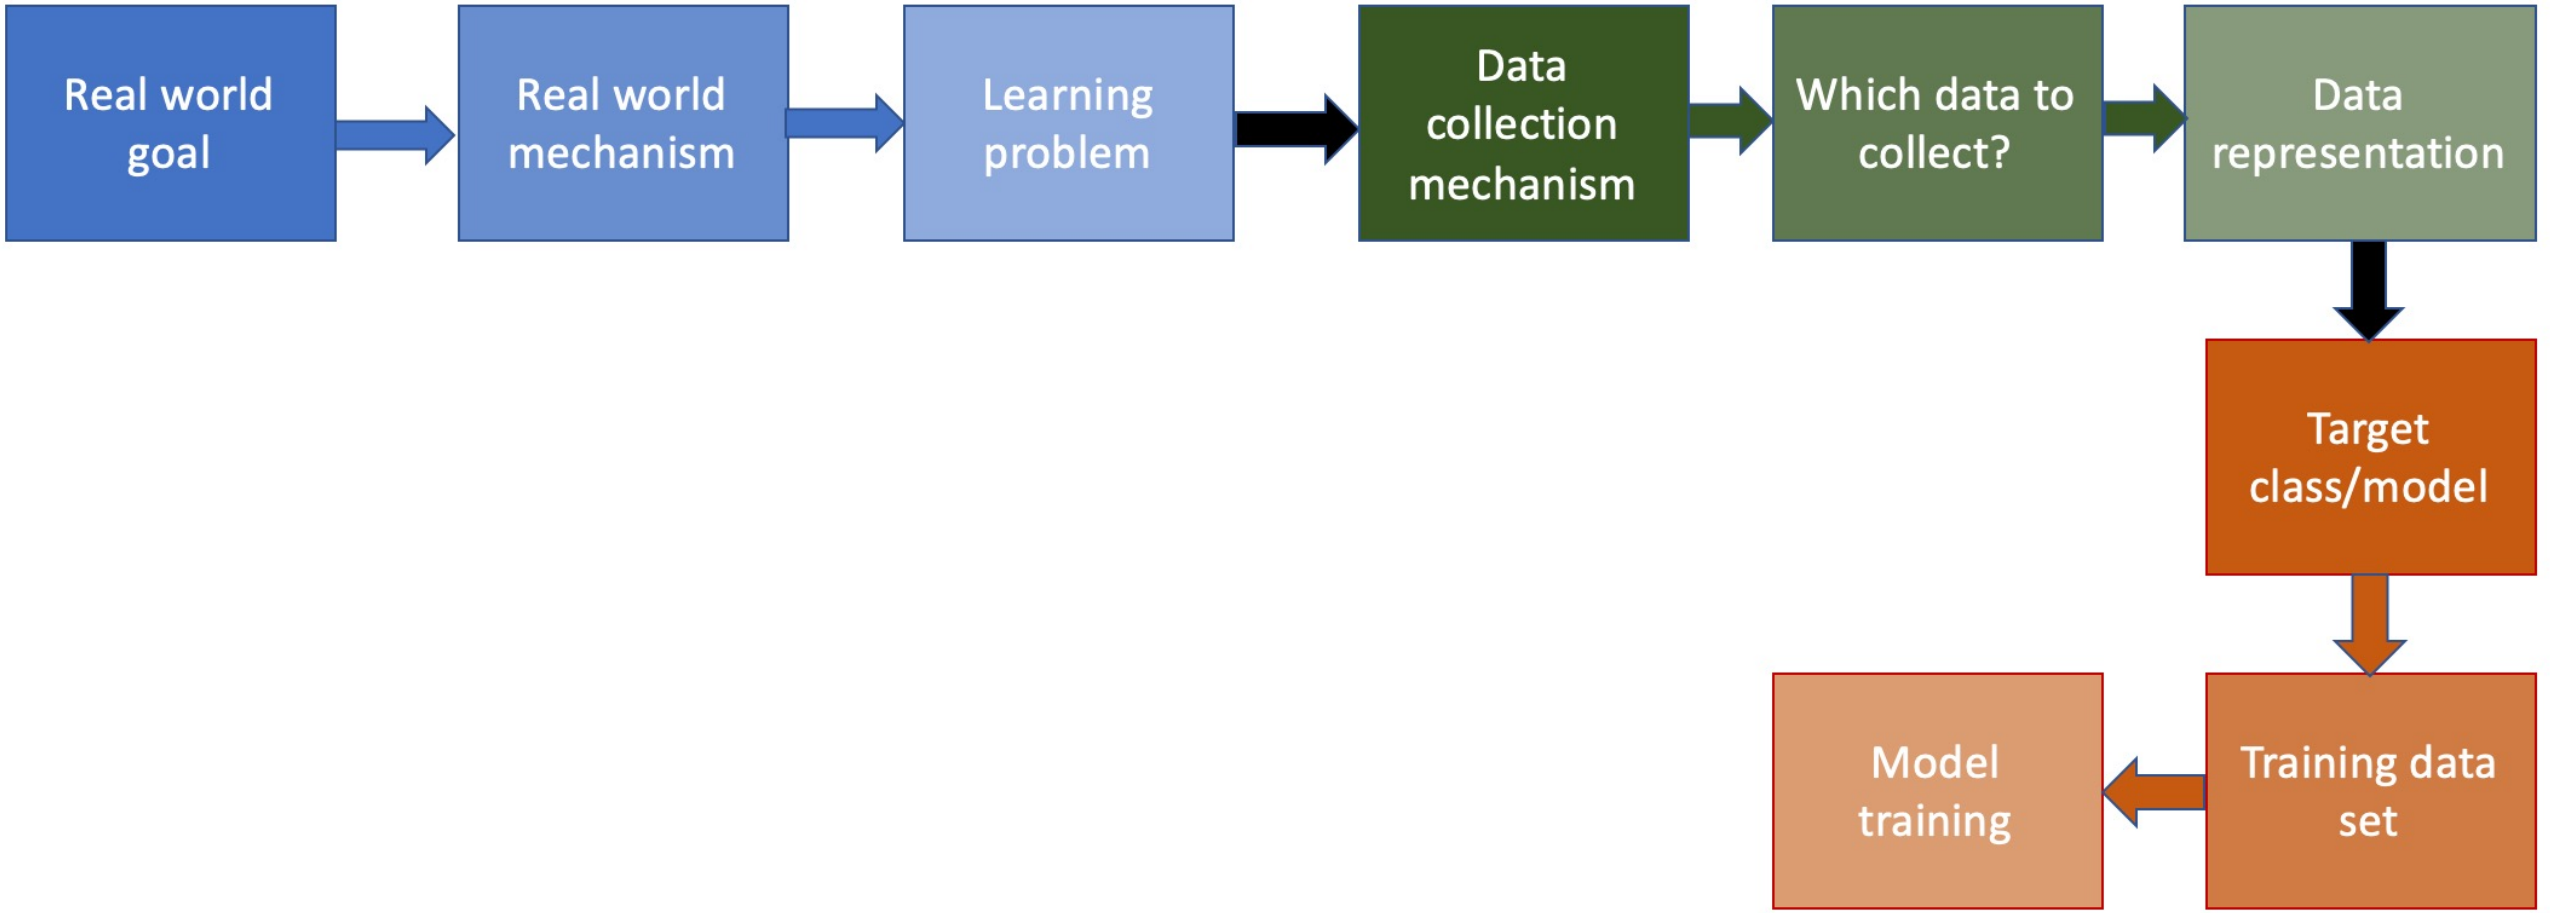

Next step in the ML pipeline

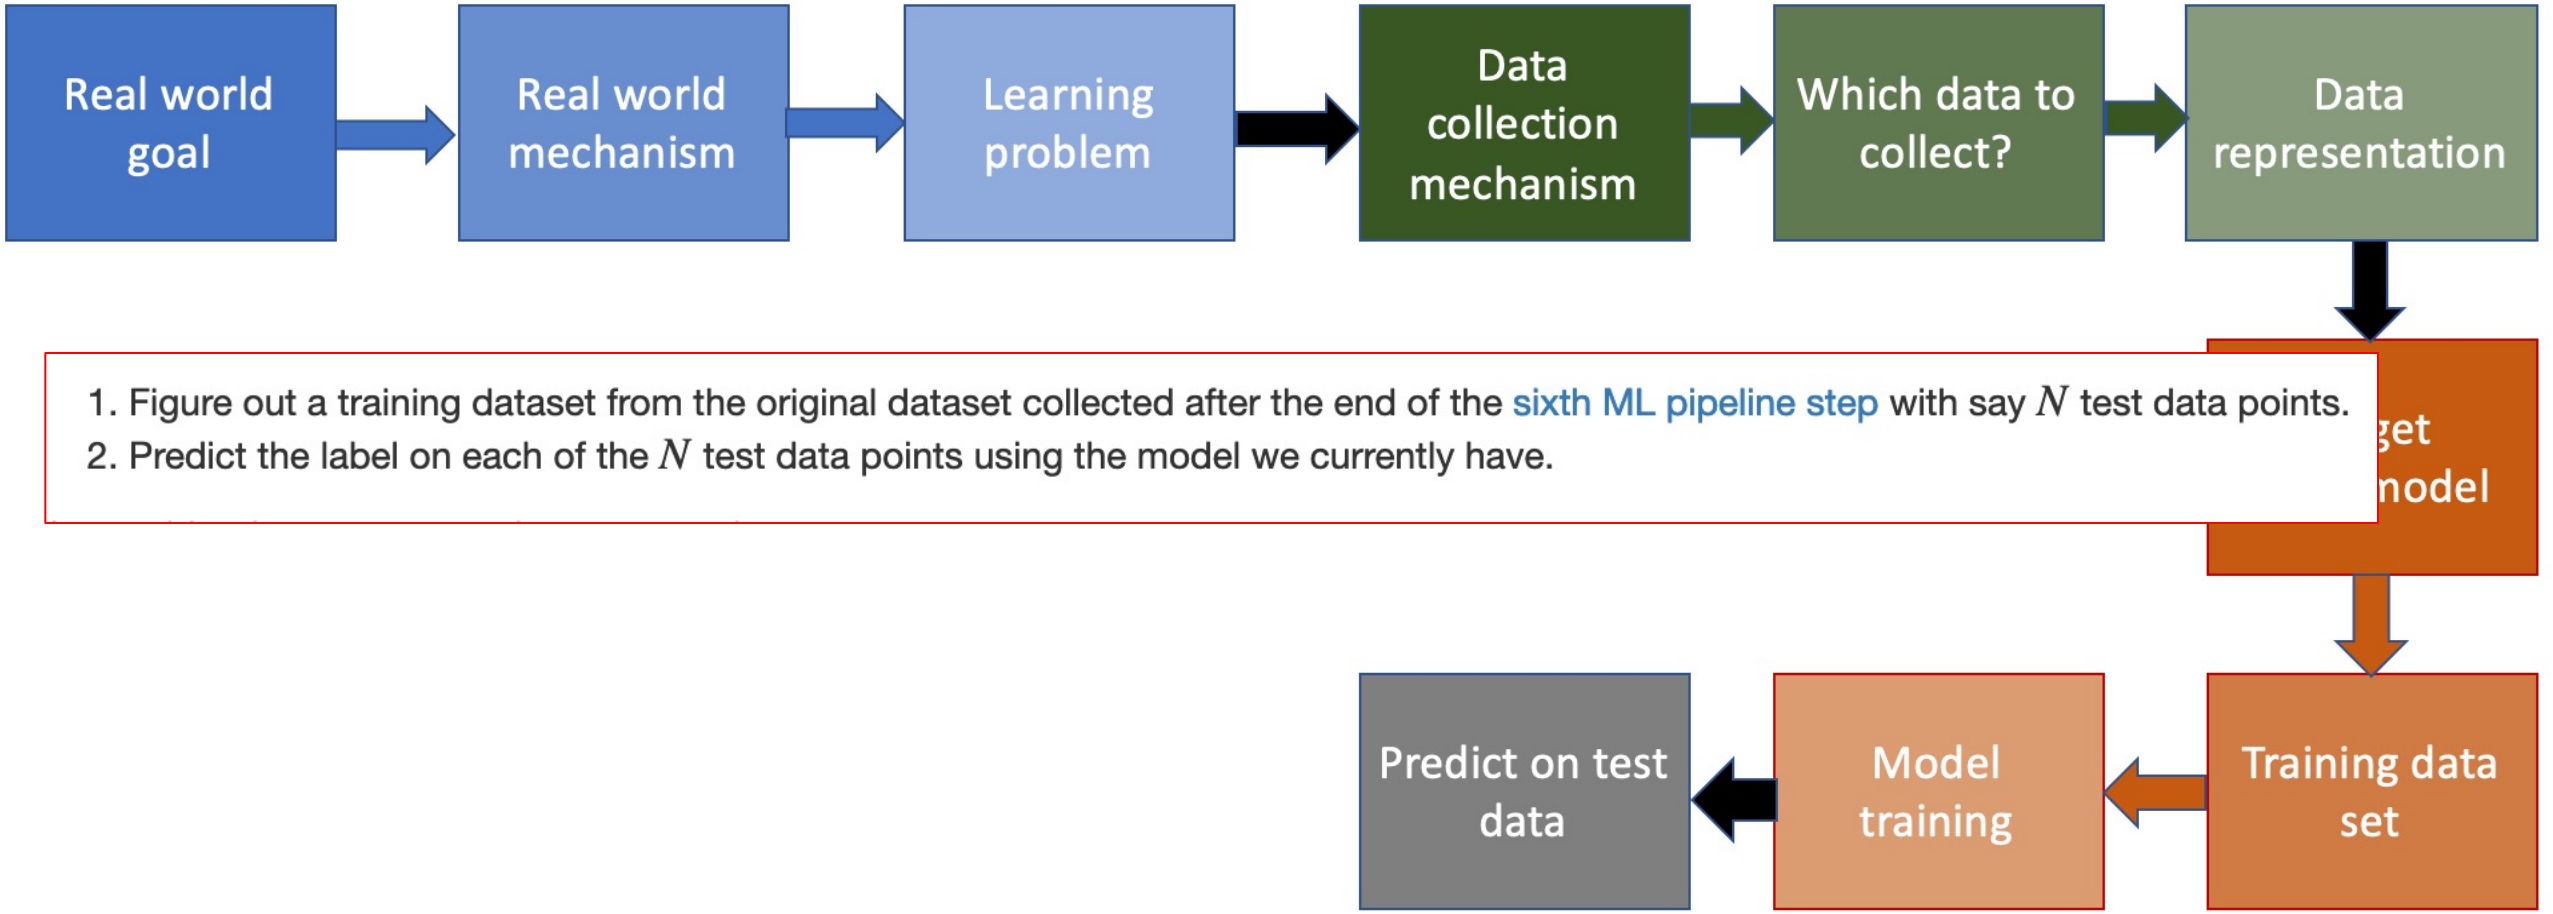

Predicting on dataset

Predicting on training dataset

When considering model classes, we have already considered the problem of prediction (e.g. see our earlier discussion on efficient prediction for linear models, decision trees and neural networks. Then doing prediction on the N test data points is simple:

- Run the prediction algorithm for the given model on each of the N test data points and record the predicted label.

Thus, the overall time taken by this step is N times the time taken for one prediction, which overall is reasonable if one prediction can be done efficiently.

Picking the training dataset

Actually, this paper is easy: since we have already picked a subset of the original dataset as the training dataset, **pick the rest** as the testing dataset.

If you do something bit more fancy than picking a random subset of the dataset as the training dataset (such as cross validation), then picking the training dataset also becomes bit more interesting, but we will not explore this option in this course.

Last step in the ML pipeline!

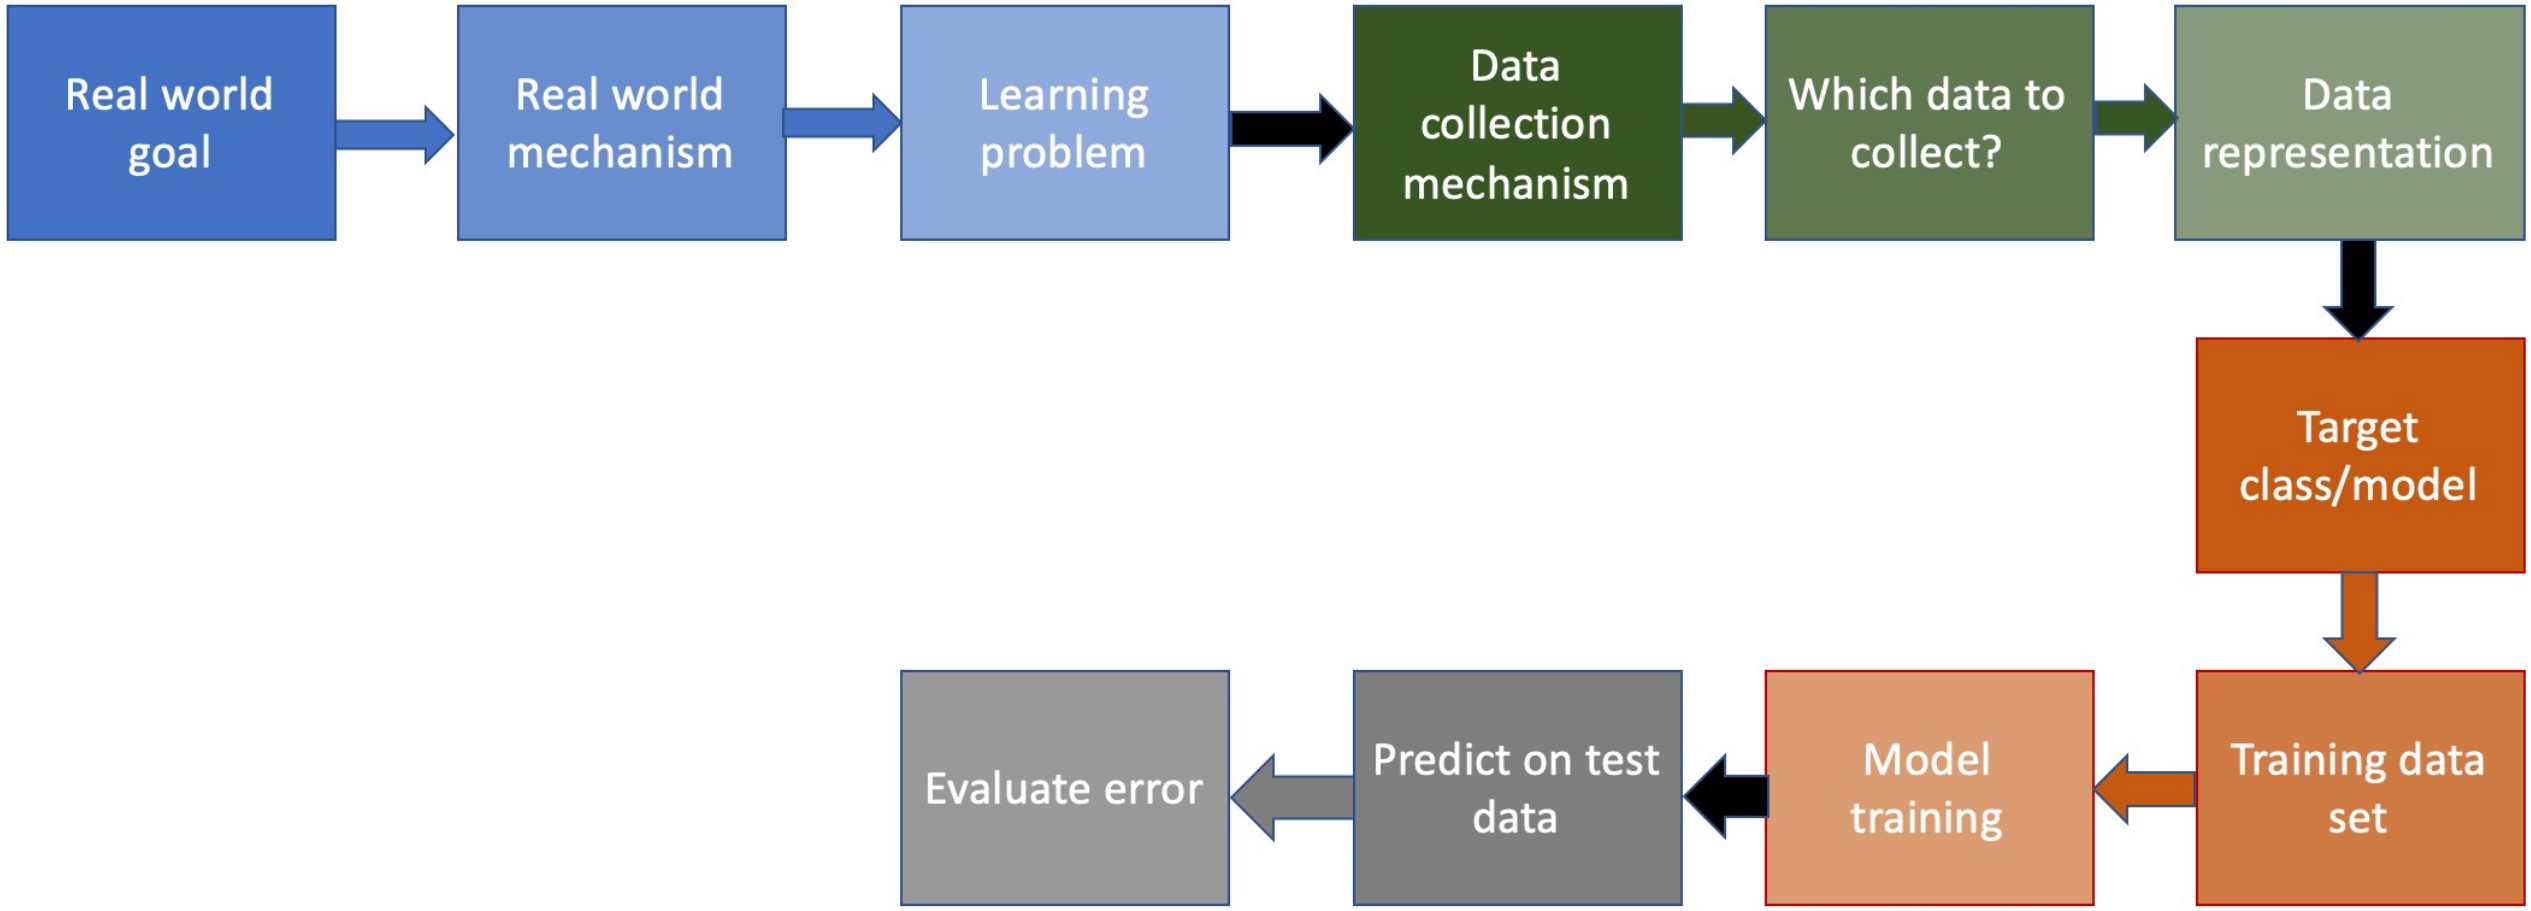

Using accuracy (# misclassified points)

Measuring accuracy

For each of the N training dataset points, check if the model classifies it correctly according to its original label in the dataset. In other words, we count the number of misclassification error.

We should note that by no means this is the only notion of accuracy (and indeed in the next part of the course we will explore more such options) but this is definitely the most commonly used one so for now, we will just mention this option.

Why did we use a convex loss function in the model training phase?

One question that might be troubling you is why we picked the number of misclassified points in the testing dataset as our notion of final model error but did not use the same notion of number of misclassified point as the loss function in the model training phase?

We have already talked about this but it bears repeating-- the difference is due to efficiency considerations. Figuring out the optimal model using the number of misclassified points does not have any efficient algorithms, which is why we used a convex loss function for which we can find an optimal model using gradient descent. On the other hand, for evaluating the error of the chosen model on the testing dataset, we only need to run N predictions, which is efficient.

Feb 2020 tweet from Geoffrey Hinton

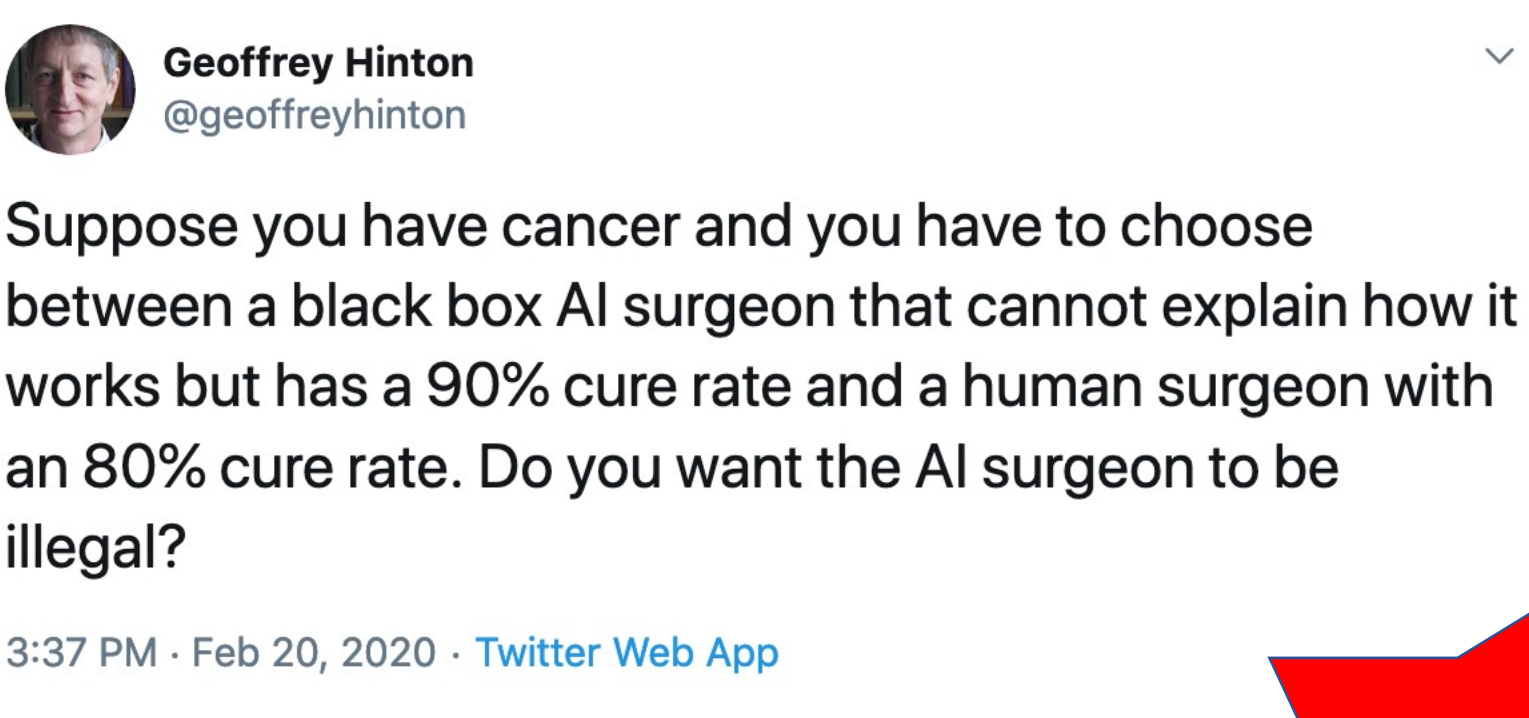

Discuss!

Some twitter responses

Eric Topol @EricTopol · Feb 20

Replying to @geoffreyhinton and @jeremyphoward

Geoffrey,
That's a lot like the question I posed to @pmddomingos in #DeepMedicine. You can guess his answer, which would be mine too:

> We already accept black boxes in medicine. For example, electroconvulsive therapy is highly effective for severe depression, but we have no idea how it works. Likewise, there are many drugs that seem to work even though no one can explain how. As patients we willingly accept this human type of black box, so long as we feel better or have good outcomes. Should we do the same for AI algorithms? Pedro Domingos would, telling me that he'd prefer one "that's 99 percent accurate but is a black box" over "one that gives me explana-

Nick Cammarata @nicklovescode · Feb 20

Replying to @geoffreyhinton

I think a more visceral question is "which would you have do your surgery"

I'd take the 90% only if I knew the distribution it was trained on is very similar to me, which it often isn't in clinical trials. The 80% would give me more confidence in its ability to transfer to me

19 40 1.2K

Moritz Hardt @mrtz · Feb 20

The role of AI in medicine aside, ethics rarely reduces to trolley problems like these.

> **Geoffrey Hinton** @geoffreyhinton · Feb 20
>
> Suppose you have cancer and you have to choose between a black box AI surgeon that cannot explain how it works but has a 90% cure rate and a human surgeon with an 80% cure rate. Do you want the AI surgeon to be illegal?

Suresh Venkatasubramanian Retweeted

Alexandra Olteanu @o_saja · Feb 20

Suppose you have limited time and you have to choose between using your expertise to solve real problems and focusing on useless/unrealistic hypotheticals in domains you are clueless about. What would you want to do?

> **Geoffrey Hinton** @geoffreyhinton · Feb 20
>
> Suppose you have cancer and you have to choose between a black box AI surgeon that cannot explain how it works but has a 90% cure rate and a human surgeon with an 80% cure rate. Do you want the AI surgeon to be illegal?

3 11 112

Passphrase for today: Kathy Pham

Kathy Pham

From Wikipedia, the free encyclopedia

Kathy Pham is a Vietnamese American computer scientist and product management executive. She has held roles in leadership, engineering, product management, and data science at Google, IBM, the Georgia Tech Research Institute, Harris Healthcare,[1][2][3] and served as a founding product and engineering member of the United States Digital Service (USDS) in the Executive Office of the President of the United States at The White House. In 2021, Pham was named the Deputy Chief Technology Officer of the Federal Trade Commission.[4]

Pham is a Fellow and Faculty member at the Harvard Kennedy School where she created and teaches Product Management and Society. She has held positions as Fellow at Mozilla, Fellow at the Rita Allen Foundation, and Fellow at the Harvard Berkman Klein Center where she co-founded the Ethical Tech Group and was part of the Ethics and Governance of Artificial Intelligence Fellows in partnership with the MIT Media Lab.[5][6][7][8] At Mozilla, Pham co-leads the Responsible Computer Science Challenge and co-founded the Mozilla Fix the Internet Incubator.[9] Pham founded Product and Society, which focuses on product management, ethics, and the public interest.[10] Pham has been part of a championship StarCraft II team,[11] and placed 1st the Imagine Cup competition, representing the United States with a sentiment analysis (EmotionAI) engine.[12]

Contents [hide]

Kathy Pham

Pham in 2015

Education	Georgia Institute of Technology, Supélec MIT Media Lab Harvard University
Board member of	Anita Borg Institute for Women and Technology, Make the Breast Pump Not Suck, Civic Signals

Early life and education [edit]

Pham's parents were Vietnamese boat people, who spent several years in refugee camps before immigrating to the United States. Her brother, United States Marine Corps Captain David Pham, was presented the Purple Heart medal during combat operations in Afghanistan.[13][14]

Pham attended Windsor Forest High School in Savannah, Georgia where she was a member of the volleyball team and graduated as Salutatorian of

Preview of Seeing White…