

Exercise 8: Visualization

Name:

UTID:.....

Let's talk about politics. Specifically: What makes a democrat a democrat? What makes a republican a republican? Their face obviously.

The test this theory you go and download all 2200 images of current and past congresswoman and congressman together with their party affiliation from www.congress.gov. Then you train a network on those images. Students in my graduate class did this and trained a model with a 72% accuracy. This is quite remarkable from just a face picture.

a) Can you beat this model? For each image below guess the party affiliation of the person.



a)



b)



c)



d)



e)

b) By now you probably realized that this is kind of a hard task. Now let's look at what the network uses to solve it. This might tell us something about our data and how it is structured. Briefly describe one method from class that could help you visualize what the network looks at when making its decision.

c) Can you think of a modification, or extension of one of the visualization techniques that would give you even more insights into the problem of predicting party affiliation from a portrait? (Be creative, there is not a single right answer).

d) (US citizens only) Go early voting.