

Exercise 11: Face recognition

Name:

UTID:.....

You just had a great business idea. You are going to create a portal where you and all your friends can hang out, share photos, and other current news about your lives. As a true longhorn you're calling your page facehook(em)¹. The true genius of your idea is the business model. You're going to use facehook to track your users and profile their interests, in order to bombard them with a merciless stream of mind-numbing targeted ads. However, you don't just want to track your users, but also all their friends in all the pictures they upload to your portal. To do that you need face recognition. You download an off the shelf face detector with a nearly 100% detection accuracy. All that remains to be done is identifying which of your friends are tagged in which picture. How could you tackle this problem using a deep network?




Figure 1: Example output of your face detector. Each white box is an input to your system.

a) Could you use Binary Classification to solve this problem? If yes, how? If no, why not?

b) Could you use (multi-class) Classification to solve this problem? If yes, how? If no, why not?

c) Could you use Regression to solve this problem? If yes, how? If no, why not?

¹sorry for the lame name

Since facebook is still quite small, you also want to target any (future) users not currently signed up for facebook. You decide to create ghost profiles² of users currently not signed up. You obviously only store this information in case they ever sign up (and to make heaps of money serving them ads until they do).

d) Do any of the above techniques still work? If yes, how? If no, why not?

Today in class we'll look at techniques that deal with large and dynamically changing label sets. Most of which you could use in facebook.

²also known as shadow profiles in the real world.