

Chapter 17

Faces Project: The Context and Consequences of Feature Recognition, Detection, and Classification

17.1 Overview

This is a project that asks you to extend the work you have already done on facial recognition and feature detection by analyzing the performance and considering the consequences of an existing algorithm within a real(ish) context. This is a fairly new formulation of this project and we are giving you the freedom (and responsibility) to choose an interesting path and follow it judiciously.

The LinAlgCo owns the rights to all uses of linear algebra. They've recently become aware of the use of linear algebra in *feature recognition, detection, and classification* ("FeaRDeClass") algorithms. The company is concerned about the ethical implications of the widespread use of these algorithms. They don't want to tarnish the good name of linear algebra. You have been hired as a consultant to address these concerns.

Specifically, they've asked you to do the following:

1. Identify a specific context in which linear algebra is being used to do *FeaRDeClass*.
(Examples: Smile detection using linear regression, identifying missing persons from photos in social media, unlocking your phone with your face using Eigenfaces, classifying movie preferences...)
2. Pose a question about the context in which *FeaRDeClass* is being performed and the possible ethical implications. You do not have to be able to answer this question, but it should guide your investigations.
(Examples: Privacy, bias, misuse, reinforcing negative structures...)
3. Answer some part of your question by analyzing the results of a *FeaRDeClass* algorithm. You can use any *FeaRDeClass* algorithm as long as you can explain how it works. You need to do some quantitative analysis, but the specifics are up to you.
(Examples: Does the facial recognition have higher accuracy with group X than group Y? Are STEM documentaries more likely to be recommended to men than women?)

Your consulting team is expected to produce a formal report, due to LinAlgCo by Monday March 2nd at 9:00am.

17.2 What we expect you to do

1. Start with some background research on contexts for *FeaRDeClass* and their associated ethical issues. This research will help you to choose what to focus on, and you should also reference this research when discussing the context of your project in the introduction of your report.

2. Choose an important question related the context and ethics of your chosen *FeaRDeClass*. This should be rooted in a real context, and you do not have to be able to answer it. Break off a small sub-question that you think you can answer in one and a half weeks through an analytical approach that utilizes eigenfaces, another facial recognition algorithm, or linear regression. Consider the ethical consequences of your chosen topic. Under what circumstances could the technology be harmful? Whom might it harm, and how?
3. Plan, execute, and document some analysis (which could include modifying/creating an algorithm) to answer your sub-question.
4. Explain the mathematical algorithm you are using in detail, explaining the various steps and what the purpose of each step of the process is. Use equations!
5. Explain how the results of your analysis inform the question you are trying to answer. Tie the results of your sub-question back into your larger question and chosen context. What can you conclude from the analysis you did? Recommend areas for future investigation.
6. You should understand the metrics against which your programs should be measured. How do you characterize the accuracy of your approach? Against what should you compare this accuracy? How do you quantify the consequences of your approach?
7. Communicate the context, analytical approach, and findings via a formal technical report to the LinAlgCo.

17.3 Resources

1. Your existing eigenfaces algorithm or the example solution posted. Let us know if you need help getting eigenfaces working.
2. The smile detection algorithm, which uses linear regression. Your version or the walkthrough from class can be modified to do something similar.
3. Training and test images for your class and past QEA classes.
4. The 10k faces database. This includes >2,000 images that have been classified in terms of demographics and other info (like whether people are facing the camera) and a software tool to narrow the database by classifiers (e.g., to only smiling men). The downside of this database is there is only one photo of each person.
5. The internet. In addition to doing context/background research, you can go find a different algorithm or face database if you prefer, but be aware that this will take extra time!
6. Your teaching team. Remember that we are here to support your learning! Bounce ideas off of us in office hours. Don't let MATLAB get you down; ask for help early and often.

17.4 Deliverable

You need to produce a written report, but we've broken it into a few sub-assignments to keep you on track and create opportunities for feedback.

1. Due Monday 2/24: An informal document outlining your chosen context, big question, sub-question, and the algorithm you will use to answer your sub-question, plus a plan for what analysis you will do and what kind of results you will get. (You should also get started on the analysis before Monday, but you don't have to turn any results in yet.) This document should serve as an outline for your final report.
2. Due Thursday 4/27: A draft of your written report. The teaching team will give you feedback. The more complete the draft, the more helpful your feedback will be! You should bring a printed version to class.
3. Due Monday 3/2: A final version of your written report.

17.5 Project report

You will generate a professional-looking and edited report to send to LinAlgCo that summarizes and justifies the decisions you have made. The executives at LinAlgCo are familiar with linear algebra and mathematics notation, but you should not assume they know anything about *FeaRDeClass* algorithms in general or the particular one you are studying.

The goal of the report should be to help LinAlgCo understand the context and consequences of the specific *FeaRDeClass* algorithm you have questioned and analyzed. You are NOT writing a story about what you did in the project. Aim for content and clarity, not length.

Structure

The report should have the following six sections (and you might want to break them into subsections):

1. Summary

What will I find in this report?

Open with a one paragraph summary that orients LinAlgCo to what they will find in the report. It should make clear why the report was written, what each section will accomplish, and what the key insights and results are. You should also summarize your recommendations.

2. Introduction

What is this project?

Your introduction should: (1) Provide the background and context for the *FeaRDeClass* that you have chosen to analyze. When and where is it used? By whom? What are the general technological or social issues associated with its use? (2) Explain your algorithm technically. How does it work? Bear in mind your audience. (3) Lay out the general ethical implications of the algorithm that you are investigating. Under what circumstances could the technology be helpful or harmful? Whom might it help or harm, and how? (4) Clearly state, within the broader ethical context, what question or issue you are exploring and what sub-question you are quantitatively investigating.

3. *Methods*

How does your approach work, and what did you do with it?

Having introduced the reader to terminology and ideas, this section should lay out the approaches you are using, both in terms of the chosen algorithm and the analysis you are doing with it. Use equations and define all variables.

4. *Detailed Findings*

What are the main results and consequences of your work?

This section should contain your main results and consequences of your work which you have quantified. This section should contain some clear, informative, labeled, and captioned plots and images that demonstrate your findings. Quantitative results should be clearly connected to the context of the investigation. Why are your findings meaningful? Reflect on the downsides of the technology and the people it could hurt, and suggest some strategies for improvement.

5. *Recommendations*

What are the key takeaways? Summarize the key findings of the report, situate them in the greater context, and identify areas for future investigation. This section should be concise—they details go in the previous section.

6. *References*

Provide full citations for sources referenced in the paper. Format doesn't matter here as long as you provide sufficient information about each of your sources.

17.6 *Grading rubric*

- 1 pt. Summary presents a clear, high-level overview of paper.
- 3 pts. Question being investigated is clearly rooted in a real context, as discussed in the introduction and justified with references. Discussion of potential for harm is thorough.
- 2 pts. Algorithm is clearly explained using equations and words.
- 2 pts. Analytical approach is clearly explained.
- 2 pts. Findings are justified with appropriate figures and discussion.
- 2 pts. A clear connection is drawn between the findings and the original sub-question, greater question/issue, and greater context.
- 2 pts. Paper is logically organized and writing, figures, and equations are polished.