Winter 2018/19 **COSC 4P80: Assignment 1**

Due date: Friday, February 8th @5:00 pm

Task: Using multi-layer feed forward neural networks to detect the presence of humans in infrared readings

The objective of this assignment is to learn about supervised learning in feed-forward neural architectures. To achieve this, you'll be analyzing data from a thermopile array to detect the presence of humans. More practically, you'll be developing a network that can correlate the provided input patterns to the expected classifications.
You'll be separating the data into training and testing sets, to gauge how well your network generalizes, and writing up a report on the total process.

Thermopile Arrays

Thermopile arrays are passive sensors that register infrared from the environment. The thermocouples are arranged in such a way that incident radiation can be detected across a grid, in a fashion similar to an image sensor. That is to say a thermopile array can be used as a makeshift form of infrared camera.
Because passive infrared radiation is highly correlated to temperature, that means they act as a form of heat detector. There are myriad uses of such technologies, but most of the more interesting uses focus on the detection of living beings (or determining the precise locations thereof). For example, unlike traditional infrared motion sensors, thermopile arrays can be used as security sensors; detecting the presence of humans and/or animals, even without motion. Additionally, with sufficient finesse, they can be used to isolate particular sizes or ranges of heat (and thus filtering out smaller targets like birds and such).

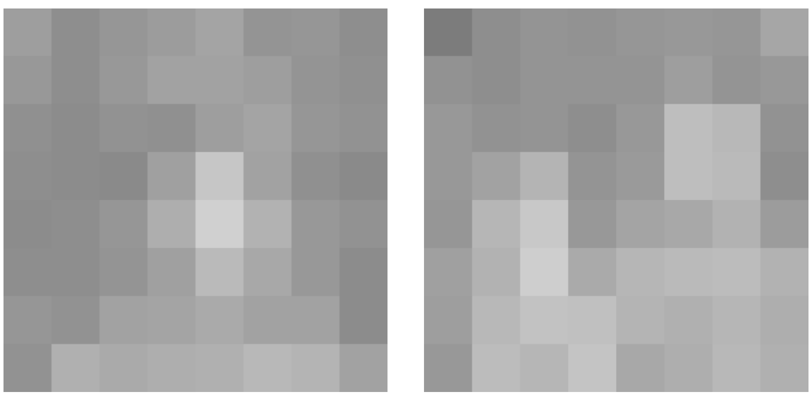

Pictured on the left is a single human face, approximately-centred. On the right is a pair of faces. Note that a human's surface temperature is normally lower than his or her core body temperature. As such, the distinction might not be quite as drastic as one might expect. However, this should be more than adequate for basic testing purposes.

Detecting Humans

A Grid-EYE sensor is an 8×8 thermopile array, which allows for 64 temperature readings per sample. The normal commercial use is for things like digital signage and kiosks (to either turn on or change their displays upon the advance of a person). Ostensibly, it should be easy to detect the presence or absence of humans within a frame. For example, since humans are warm-blooded, they tend to have a surface temperature greater than the ambient temperature of the room (for typical room temperatures).
In general, because infrared (being radiation) tends to spread out and disperse, it should be a reasonably simple task to determine the distance of a human (identified as a subject by being higher than background, but with proximity defined by the extent to which it's higher). Furthermore, one could likely determine whether a larger infrared source was likely to be from a single close source, or multiple distant sources (based on things like the actual width of the field, compared to the intensity).
Of course, heat isn't the *only* source of infrared. For example, taking a reading outdoors during the daytime will yield strong reflected infrared radiation from some surfaces, and a substantial signal from scattered radiation from the sky.

Since infrared flooding from the sky is still far outside the range of what one would expect from a typical human, your experiments will confirm whether or not this is a practical test outdoors (or if infrared reflecting off faces throws it off).

For this assignment, you've been given a collection of data files, including both indoor and outdoor samples.
For outdoor samples, there are three categories: no subject, single subject from 1' away, and single subject from ~3' away.
For indoor samples, there are samples with no subject; with a single subject from 1', 3', and 6'; two subjects from 3'; and three subjects from 3'.

You'll be devising experiments to answer questions such as the presence of subjects, number of subjects, distances, or even environment (outdoors vs indoors).
Some of these questions will likely be easier to answer than others.

Implementation

You are required to ***use*** feed forward neural networks that are able to classify the thermopile data. The data is currently arranged in a fashion that was simple to *capture*, but may not be the most appropriate for immediate *use*. You are free (and encouraged) to combine, scale, or filter the data in whatever fashion you deem appropriate, with the following caveats:

- In your writeup, document how you handled/manipulated the data
- If you applied any prethresholding (e.g. to filter the intense infrared from the outdoor shots), it must be done in a generalized deterministic fashion (e.g. `all values over 4000')
- Basically, you must still leave the primary classification to the neural network itself

The datafile format (which, again, you're free to change) is as follows:

- A single-line description
- A line of '=' characters
- Samples with one-line-per-complete frame, consisting of:
  - A single number indicating the *thermistor* reading (this is used for determining the actual general ambient temperature for calibration/correction); you're free to ignore this value if you like
  - 64 tab-delimited integer (short) values; one-per-pixel

As a suggestion, you may or may not wish to work out some simple program to convert lines from the data files into visualizations (similar to the included screenshots). It *might* help give insight into which types of samples are causing difficulty and why. For most papers, a couple visualizations would be appropriate for the experimental writeup anyway.

Requirements

- A feed-forward network, based on backpropagation
  - Highly advisable to include momentum
  - Try different variations on learning parameters (e.g. learning rate, momentum, etc)
  - Consider using a different activation function if you prefer
  - Try to get this working as well as you can before you continue to the next part
- Include another variation for learning
  - e.g. quickProp, rProp, etc.
  - You *can* use an entirely different training technique (e.g. GAs), but I *strongly* advise against that
- Construct tests that attempt to classify. Suggested categories include:
  - Indoors:
    - Presence or absence of subject (I'd suggest the 3' dataset for this)
    - Distance of single subject (it's up to you whether you include no-subject in this)
    - Number of subjects (1, 2, 3, or none)
  - Outdoors:
    - Decide for yourself what test you'd like (e.g. presence or absence; distance; both)
  - For all tests, remember to perform *multiple runs* for each scenario to ensure reliability of results
  - You should split the data into separate training and testing sets
- Notice that none of the above says you need to be the one to implement the neural network yourself
  - This is, first and foremost, an assignment about methodologies and experimentation

- Writeup – formal report using LaTeX
    - Your report will be treating this as legitimate research; *not as an assignment*
    - Abstract, introduction, and problem definition
        - Ensure that a person who was unfamiliar with the problem would fully understand the problem you're trying to solve
        - Don't forget to explain things that are already explained in this assignment specification
    - Background on neural networks
        - Ensure that I have a rough understanding of what a neural network is, and whatever learning rules you used to train them
        - (But don't go overboard here)
    - Methods of Analysis
        - Give a brief explanation of what statistical mechanisms you're using
            - e.g. comparison of means, ANOVA, contingency table analysis, etc.
        - Try to include a brief explanation of how you're actually evaluating/ranking the performance directly (e.g. **false positives** vs **false negatives**)
    - Experimental Results/Discussion
        - What sorts of parameters did you use? How did the results turn out?
            - Try to include at least a few comparisons here
        - What sorts of problems were easier or harder?
        - Can you reasonably hypothesize why certain tasks were harder? Can you think of a way you could test that theory in the future?
        - This will be your largest section, and should also include a few charts/graphs if possible
    - Conclusion or synopsis
        - Does what it says on the box
    - Tip: You *can* change some of these headings (and can especially *add* headings) if you don't feel they quite fit the approach you prefer. Ensure, however, that the paper remains formal, clear, and complete
    - The expected length is *probably* around 10–14 pages
    - Final Thoughts: I'm not concerned whether you use IEEE style, AMS, etc. I only care that it's well-written and reasonably formal (and consistent). Remember to provide citations for any materials you use (this also applies to the implementation itself)

Evaluation will depend largely on how much you put into this. It's important to point out that the actual *success* of your tests is secondary at best. Poor results that are well-analyzed, documented, and discussed are vastly superior to great results that lack any accompanying proof or discussion.
The more you include (within reason), the better you'll do. Remember to proofread, though, because proper communication is necessary for fully understanding analysis.

Note that you are free to use whatever language/tools you like, so long as the TA can test your code on the lab hardware and/or sandcastle (or you make alternate arrangements).

Submission

Submission is **electronic**:

- Your code
- Your raw experimental data
    - If you have many pages of this, just include a smaller sample (e.g. from a few experiments)
- Your report
- Simple instructions on how to use your program (unless it's *incredibly* self-explanatory)

As a reminder, the submission script is `submit4p80`

Don't forget: cheating is bad. Plagiarism will lead to charges of academic misconduct, and angry scowls.

Credit for structure: Dr. Beatrice Ombuki-Berman