**Massachusetts Institute of Technology**

# 6.038 Artificial Intelligence in Practice

Problem Set 1A: Tuesday, February 3, 1998
DUE: Thursday, February 5, 1998

## 1 Background

The objective of this assignment is to get you to review some basic machine-learning methods as well as the methodology of attacking machine-learning problems in general.

Much of this problem set (both Part A and B) is based on one developed by Tom Mitchell and Avrim Blum (of CMU) as part of their course in Machine Learning. The problem is also featured in the textbook by Tom Mitchell "Machine Learning" published by McGraw Hill, which is a good introduction to the field of machine learning.

## 2 Face Recognition

You can find a set of images of 20 people on the course Web page `http://www.ai.mit.edu/courses/6.038/`. Each of the images is an array (128 columns $\times$ 120 rows) of brightnesses. Each entry (pixel) in the array has a number in the range 0 to 255. Note that the images of each person vary in the head position, expression, etc. There are at least 28 images of each person in the full data set.

Our objective is to build a system that can recognize (**new images of**) the people in these images or possibly to recognize certain characteristics, e.g. whether the person in the image is wearing sunglasses or looking to the left. At this point, we want you to think about some of the issues in achieving this. Part B of this assignment will ask you to actually do some of this.

We will be considering three basic machine learning methods you learned about in 6.034:

- Nearest Neighbors (NearN),
- Identification Trees (IDT), and
- Neural Nets (NN).

We will also consider several alternative problems:

- Recognizing whether a given image belongs to a specific person, say Joe, in any orientation, with glasses or not, etc.
- Recognizing all of the 20 people in the training set (in any orientation, etc).
- Recognizing that an image is of someone wearing sunglasses, or someone looking to the left, etc.

Please hand in a short writeup of your answers to the following questions. Please be as specific as possible in your answers, but be brief; you should not need more than a couple of pages total.

1. Which of the three methods listed above would you use to tackle the three problems. Justify your choice.

2. How would you construct feature vectors from the given images to train your chosen method. Justify your choice.

3. For each of the problems given above, specify parameters that would be needed by your chosen method, e.g. distance metric for NearN, number and size of layers for NN, etc. If you don't know how to pick the parameter a-priori, indicate how you would go about finding out good values when given the data.

4. For your chosen method, suggest how you go about (a) training and (b) estimating the performance of the method on new data. In particular, be specific about how you partition the data, etc. Indicate what issues you need to be concerned about.