Nels Beckman

Data Privacy

Assignment Six: Use of Publicly Available Web Cameras

Abstract:

The advent of the web camera seems to bring with it a whole host of privacy concerns, many of which we have only just begun to address. It seems that as more and more publicly accessible web cameras are erected in a larger variety of situations, these privacy questions are no longer abstract thought exercises, but rather real and pressing issues. In this paper I will show that it is possible to use a publicly available web camera, whose ostensible purpose is mostly innocuous, to monitor the employees of a specific company and judge their performance without their knowledge in an automated manner.

Introduction:

One could say that the publicly accessible web camera or webcam is the ultimate voyeur's tool. You can observe the habits and behaviors of people all over the world from the privacy of your own home. Anecdotally, we might also suggest that as bandwidth increases and as digital camera technology matures and drops in price, more high quality images are available to those who might want to abuse it.

In this paper I will concentrate on one webcam in particular. New Horizons Computer Learning Center is an educational center that focuses on teaching computer skills to working adults. Their Gainesville, Florida branch, accessible at http://www.nhgainesville.com/, links to a variety of publicly accessible web cameras that are located right inside their educational center and put their employees and students right in the viewfinder. Furthermore, their camera feed is of high quality (640x480 images) and updates frequently (on the order of once per second). The purpose of these webcams is likely innocuous, but as we will see it can be abused.

Background:

Some forms of employee monitoring are already becoming common. For instance, at some workplaces, employee Internet usage is monitored to ensure that employees are not using the Internet for personal reasons. While this would be something that was hard to monitor manually because of the large number of data involved, because it can be automated it is a practice that has been considered by many businesses. Cameras in the work place have existed for a long time, primarily to monitor for suspicious activity. It is in general too much work to review enough workplace video to make a judgment about an employee's performance. However, if that video could be monitored automatically, the situation might be much different.

Methods:

For the purposes of this experiment, however, we will say that the receptionist is failing to do her work during the period of time where there is no one at the desk to greet incoming customers. Based on this simple idea, I will show that a remote evaluation of employee performance is possible.

In order to set up this experiment, we must first develop a set of operational definitions. We will define whether or not the receptionist is 'doing her job' in the following manner; If there is anyone seated in one of the two receptionist's chairs the receptionist is 'doing her job,' otherwise not. We have define this in such a way because it is possible and indeed likely that the receptionist might delegate the duty of greeting customers for some amount of time in order to do some other task or go on an official break. However, when there is no one seated at the receptionist's desk, we say that the receptionist has failed in her responsibility because she is not able to greet and direct incoming customers, and because she has not taken it upon herself to assign someone else to do the job temporarily. These concepts are illustrated in the three figures shown. In the first figure, we see the primary receptionist. In the second, we see a second person to whom we say the responsibility has been delegated. In both we say that the receptionist was correctly performing her duties. In the final figure, however, we say that she is not. Based on these definitions we rate the employee's performance by the percentage of time where there was no one manning the front desk.

We have developed a tool in the Java programming language that makes automated employee monitoring possible, and based on a large data set of images, we have run some experiments. The program itself can be described as follows:

We start by collecting a number of images from the webcam itself. This was done by using an automated web browser that at the command-line could download an image from a specific web address. This tool was then combined with the Unix automation tool *Cron*. Cron allowed us to automatically download the same webcam image every minute so that over the course of a day we could collect many images. From this collection of images, we separated out about 165 images where there was no receptionist present. From this collection of images, we used the tool provided to us by the professor to create an average image. This tool goes through all of the images and takes the pixel-wise average and then outputs a new image based on the

average pixels. This helps me to eliminate variations that occur because of lighting and other small discrepancies. Next, we run another tool provided to us by the professor that computers the pixel-wise standard deviation for all of the images.

Finally, we run our tool on the images themselves. The purpose of our tool is to automatically determine if the receptionist is present or not. There is a multi-step process that occurs when the tool runs. We have selected a region of the image that is behind the desk where we believe the receptionist will most likely appear if he is performing his work correctly. This region is pictured in the adjacent figure, and we use the pixel coordinates of this box in our analysis.

1. The first step is to run through all of the pixels in this box and note the pixels that are *X* times the standard deviation above the average image. The intuition here is that if the receptionist is present, the pixels that make up his image will be of above-average brightness when compare with the empty picture.
2. We run a median filter on the image. This step was suggested by Ralph Gross, our teaching assistant, and its purpose is to eliminate some of the noise from the picture that might otherwise get in the way. (A median filter assigns the median value of the 9 pixels surrounding and including a given pixel to that pixel.)
3. Then we perform a connected component analysis of those pixels that were *X* times the standard deviation above average. This step was also suggested by Ralph. The idea here is that we run through the pixels in the image, and group them into blobs. If any of the neighbors of a pixel that we have already checked are in a component, then we add that pixel to the same component, otherwise is forms a new component.
4. Finally, we throw away any components that have fewer than *Y* pixels in them. The idea here is to eliminate any very small components that are just noise.
5. At this point, if any components remain, we say that this must be the image of the receptionist, and we flag the receptionist as being present.

Results:

First we should note that the high quality of the webcam images would allow me (if I actually knew any of the employees of New Horizons) to correctly identify all employees pictured. Of course if the company had only one receptionist, it would not matter how poor the images were (up to a point) since you would still know who the receptionist was. On the other hand, the frequency of the webcam updates and the timestamp in the upper portion of the screen allowed me to pinpoint events to the second and would give me almost indisputable evidence

against an employee. This is the kind of data that is hard to refute, at least along the simple dimension of receptionist presence.

In this, the final experiment of this project, we ran our analysis on 452 images that had already been hand-classified. Furthermore, we ran this analysis for several values of *X* and *Y*. We ran the experiement with *X* being both two and three (standard deviations above average). Similarly, we varied the minimum number of pixels required for a component to be saved (*Y*) from 250 to 850. The accuracy can be summarized in the following graph:

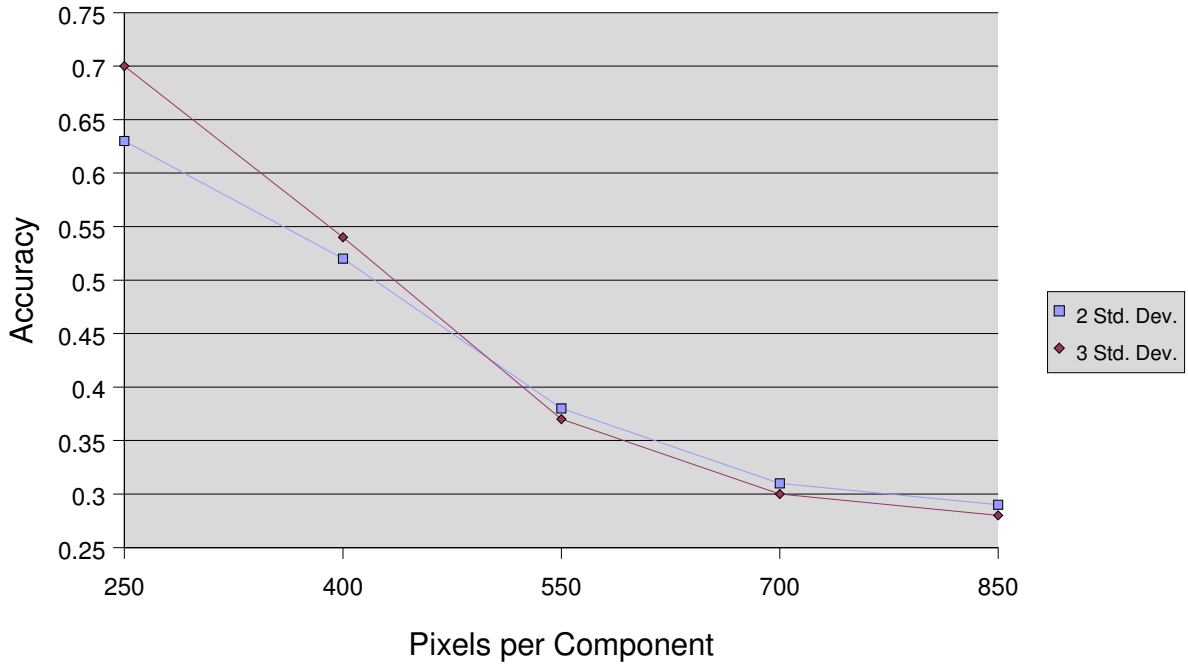

We calculate the overall performance using a few measures. We will say that a true positive is when our program correctly determines that the receptionist is present. Likewise, a true negative is when the program correctly determines that there is no receptionist present. Based on these definitions, we can see that the accuracy of this technique was as high as **70%**. (For the case where the accuracy was highest, we measured the precision as being **69%** and the recall as being **97%**.)

Discussion:

We should see the results of this experiment as being a proof of concept. We have shown that it is in fact possible to automatically perform a performance evaluation of one specific aspect of the receptionist's job. We were able to successfully use a webcam, that in all likelihood was installed for innocent reasons, to determine how well an employee does his job without his knowledge. This was by no means a perfect experiment. The fact is, basing employee reviews on a tool that is only accurate seventy percent of the time is not particularly prudent. This is the type

of thing that could lead to major lawsuits if the employees were to find out. However, this program was written in a short amount of time, and it is quite possible that refinements could be made to the same basic technique, and that these refinements would yield a much better automated analysis. Already, we have experimented with certain parameters, such as minimum component size and number of standard deviations above average. Maybe tweaking other parameters, for example the size of our rectangular region, would leave to even better performance. Certain conditions make this an excellent choice of webcam for this experiment. We have a still camera and relatively stable lighting conditions. However, the lighting conditions were not always stable enough. From printing out sets of 'change images' that showed only the pixels in the final components, I would occasionally see large components that were made up entirely of lighting differences. Similarly, sometimes a customer's head sticking into the recatangular region was interpreted by my algorithm as being the receptionist. This is illustrated in the following two images. On the left we have a partial image of an acting receptionist. On the right, changes in lighting have fooled our analysis. It is not entirely clear how we would change the current algorithm to account for these particular issues.

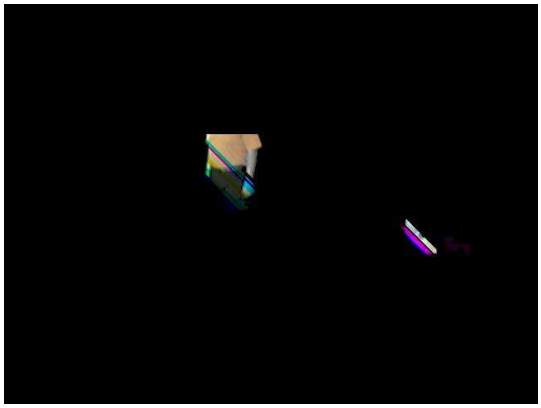

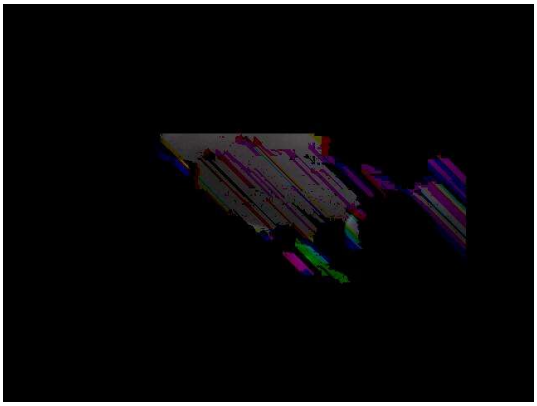

Conclusion:

Webcams should be suspect on privacy issues precisely because they are publicly available. If, for instance, the employers of New Horizons were to tell their employees that they wanted to install a private camera, viewable only by them, the employees most likely would be up in arms. What other conceivable purpose could there be for this type of closed-circuit camera other than monitoring employees? On this other hand, with a publicly accessible camera, employees could always offer some plausible sounding motivation. I do not believe that most employers are so unscrupulous as to engage in this particular form of deceit. However, what I do worry about is cameras that are installed with an innocent intent, which later on are discovered to have another potential usage. It is exactly these types of privacy issues that should be opened up for public debate. While the automated analysis we ran was not perfect, its performance was good enough that we should begin to discuss these sorts of issues, because it is likely that in the near future others will be able to provide improved versions of similar techniques.