

I know what you did in the lounge: invasion of privacy by the Free Food Cam

Marwan Abi-Antoun

mabianto@cs.cmu.edu

Abstract

The goal of this study is to determine how the Free Food Cam which provides continuous pictures of the graduate student lounge in the School of Compute Science at Carnegie Mellon University Carnegie Mellon University can be used to invade the privacy of unsuspecting faculty, staff and graduate students caught in its crosshairs. In particular, we found several cases where the webcam can determine how long a given person was in the lounge, on a coffee break, or reading the paper: if the person is a staff, this information is publicly available, including to their supervisor.

Introduction

A publicly accessible webcam, the Free Food Cam [1], offers a 24-hour view of the graduate student lounge of the Compute Science department at Carnegie Mellon University. The webcam offers a good photo quality at a relatively close range, making it possible to identify many faces on the camera (See Figure 1). However, the camera has a narrow range of vision offering only a partial view of the lounge. Anecdotally, the webcam was originally installed to prevent the theft of an espresso machine when it was originally installed. Over time, the webcam derived its name from its use to monitor any food leftovers that are brought to the lounge. In this report, we are mainly interested in how one can use this webcam to monitor the unsuspecting and invade their privacy by reporting information on them without their express consent.

Methods

The lounge is a relatively popular area where many occupants in Wean Hall tend to socialize. Due to the specific observed location, the occupants are assumed to be mostly faculty, graduate students and staff within the School of Computer Science at Carnegie Mellon University, and associated departments. Except for the occasional visitors, the population is mostly closed, which opens the door to surveillance by “intimate stalkers” since an observer is likely to recognize most of the people who appear on camera.

A Java program was designed to query the webcam URL every 10 minutes, over a week, and store the resulting image locally on disk. Each file was named after the time it was captured (including day, hour, minute and second¹). The full dataset is available at [2]. The captured data was manually analyzed to determine the following information:

- Count the number of persons visible in each still picture; if a person is only partially visible, they still get counted;
- Indicate whether food was present at the time;
- Determine if one or more group of people (*clique*) are seen together in multiple frames;
- Describe type of activity (eating, reading, socializing, etc.);
- Note any “persons of interest” (e.g., various high-influence individuals).



Figure 1: Sample image from the Free Food Cam (image Tue_Mar_28_2006__11_36_15.jpg).

¹ An example filename is Fri_Mar_31_2006__12_06_22.jpg As an afterthought, the filename should have been more descriptive to include whether it's AM or PM. However, the timestamp on the file (“DateModified”) captures that information. The images were sorted by their DateModified attribute before they were analyzed.

Due to the subjective nature of many of these observations, there was little room for automating their measurement. In particular, we determined that a crude face recognition software will likely confuse a dark bowl containing food with the face of a person: more generally, an automated approach is not guaranteed to work, since it will depend on the segmentation, crop, and will need to normalize for variations in pose, illumination and facial expression [3]. In addition, a human looking at the picture can determine if the same person appears in multiple frames, even if they have their back to the camera, based on the color and type of clothing. The number of persons visible versus time of day was plotted in Figure 3.

Results

An analysis of the above data shows the following trends:

- There are hardly any people in the lounge between the hours of 10:30PM and 8:30AM;
- The peak of the activity happens between 12:00PM and 4:00PM;
- The peaks of activity do not always correspond to the availability of free food;
- Despite the name of the camera, free food is not often available, and when it is provided, it lasts between 10-30 minutes. For instance, after analyzing the data for a week, we found evidence of visually appealing food in only a few instances²;
- The student lounge seems to get particularly busy at times, with large groups (more than 4 persons) having some kind of discussion or meeting;
- With the limited dataset, Friday tends to have more sustained activity throughout the day; in fact, visual inspection of the picture data showed a chess game being played in the evening;
- Groups tend to linger longer in the lounge than individuals.

Discussion

There are potentially several privacy-invasive uses of the above data. For instance, in some cases, the same person is visible in multiple consecutive snapshots. If the person is a staff member, this data can be used by that person's supervisor as an indicator that the staff member is taking unusually long coffee breaks. For instance, in one case, a staff person is seen taking a coffee break of at least 30 minutes³ (See **Error! Reference source not found.**).

Similarly, a graduate student is seen reading the newspaper for at least 30minutes⁴. In the latter case, it would not be hard to conceive the future of "Black Friday", the end-of-semester long graduate student evaluation, to include mention of excessive idle time spent in the student lounge, that could be better spent on research. For instance, if a student told her advisor that she cannot meet due to a class commitment, and is seen by her advisor hanging out in the student lounge for extended periods of time, she might have some explanation to do. Unlike running into her said advisor over time, the student may not be aware that someone may be recording data from the lounge camera, and reviewing it, looking for specific clues.



Figure 2: The same staff person observed in the lounge on an extended coffee break lasting at least 30 minutes: the face of the person was blocked off to protect their identity.

² One morning instance was Tue_Mar_28_2006__11_36_15.jpg. One evening instance was Wed_Mar_29_2006__07_16_19.jpg

³ See files Mon_Mar_27_2006__11_16_12.jpg, Mon_Mar_27_2006__11_26_12.jpg, Mon_Mar_27_2006__11_36_12.jpg

⁴ See files Tue_Mar_28_2006__05_26_17.jpg, Tue_Mar_28_2006__05_36_17.jpg, Tue_Mar_28_2006__05_46_17.jpg.

Limitations.

There were several limitations to our operational definitions:

- As mentioned earlier, the accounting of “cliques” is not very precise; at the outer edges of the field of vision, it is hard to determine the number of people; in some cases, a person may only be partially visible or turning their back to the camera;
- If a person temporarily walks out the camera’s sight and comes back in the following frame, we currently consider as if the person did not leave the lounge;
- When counting groups of people, the groups tend to change: in many cases, there were more than one group visible in a given frame; furthermore, the group membership tend to change, with some people leaving the group, and other people joining in;
- The interval size was constant at 10 minutes: it might be more instructive to have a finer grained interval during busier hours (e.g., 8:00AM to 8:00PM), and larger intervals during the wee morning hours (12:00AM – 6:00 AM). After analyzing the data over an entire week, we think it would be more appropriate to dynamically change the interval to capture more frequent data for the times when the lounge is more likely to be occupied. More specifically, the coarse timing interval is not precise enough to determine if the same person is present 21 minutes or 29 minutes?
- In many cases, we suspect that a person of interest may be outside of field of vision. Some people may be aware of the presence of the webcam (in fact, they may have monitored it to determine to make the trip to the lounge to sample the free food). As a result, they may intentionally avoid standing directly in the camera’s field of vision, or intentionally turn their back to camera;
- Data was only captured for a week. It might be more instructive to capture the data over longer periods of time, in order to draw statistically significant figures;
- The camera we used is focused on only one area of the lounge, and therefore, cannot give a good indication of the overall level of activity in the lounge. However due to the relatively small size of the lounge, we think that one can safely extrapolate the general occupation patterns of the rest of the lounge.
- There were several corrupt images. In some cases, one can still discern the number of people, in other cases, the image is practically useless. In particular, the camera seems to be mounted on a movable surface that is occasionally moved, severely distorting the visual field⁵.

Further Analysis. The captured data can be further analyzed along various axes. We can think of the following dimensions:

- Breakdown of occupants in terms of student, staff and faculty. Unfortunately, this requires recognizing the classification of over a thousand faculty, staff and student that comprise the school. Furthermore, there will be a small number of indeterminate values, to account for outside visitors.

Conclusion

In conclusion, the Free Food Cam can be used to monitor and record the behavior of faculty, staff and students in the graduate student lounge, with varying degrees of invasion of privacy. We think the invasion of privacy is particularly acute in the case of staff members, if they are caught off guard taking extended coffee breaks (over 30 minutes) in the middle of the day. Similarly, graduate students are not immune from the prying eyes of their advisor.

References

- [1] <http://zimbs.srv.cs.cmu.edu/coke/ffc.html>
- [2] CamSpy Data. Available at <http://www.cs.cmu.edu/~mabianto/CamSpy.zip> (24 MB)
- [3] CMU Pose, Illumination, and Expression (PIE) database http://www.ri.cmu.edu/projects/project_418.html

⁵ See Fri_Mar_31_2006__07_26_23.jpg Fri_Mar_31_2006__07_16_23.jpg




Figure 3: Plot of the number of persons visible in the Free Food Cam's field of vision over a week.