STOVE TOP THERMAL MONITORING FOR ASSISTED LIVING AT HOME

Ming Ye Yuan, James. R. Green, and Rafik Goubran

Systems and Computer Engineering, Carleton University

INTRODUCTION

The elderly segment of the population is growing worldwide. In Canada, seniors (i.e. persons over 65) comprise the fastest growing segment of the population [1] and are expected to form 18% of the Canadian population by the year 2021 [2]. This trend will place an increasing strain on the health care system as people tend to require more health care as they age [3]. Large cost savings can be achieved if elderly persons are able to continue to live independently, a situation which is often far preferred by patients to the alternatives of institutionalization in hospitals or long term care facilities. Diverse technologies are being developed to facilitate this so-called 'aging in place' [4-8]. However, most monitoring system initiatives thus far have focused on the bedroom, bathroom, and exit doors, and there remains a significant need for safety monitoring in the kitchen. Specifically, the stove is the number one cause of fire accidents in the residence [9]. As people age, some are faced with declining mental acuity due to age-related illnesses. This in turn can lead to dangerous behaviour around the stove top where the stove is forgotten and left unattended, burners may be left on, and pots may boil dry. A stove top monitoring system would go a long way towards increasing kitchen safety by detecting and correcting dangerous situations before a fire occurs.

At Carleton University, we are creating a stove top monitoring system based on thermal imaging to track stove top status and associated human activity in order to generate alerts as appropriate. Ultimately, this system will continuously monitor the status of all burners (i.e. on/off, occupied/unoccupied, warming/cooling, etc) and detect human activity associated with the stove top. These generated events will be analyzed by a separate system and alerts will be generated ranging from gentle audible reminders to emergency action. This paper focuses on our efforts to detect human activity on a stove top using an infrared camera. Monitoring of other events (pot presence, pot boiling dry, etc) and the subsequent alert generation remain as future work, and will be presented in a separate paper.

SYSTEM AND METHODS

The system components include an infrared camera (FLIR model A40, http://www.flir.com/CA/) mounted on a tripod above a stove top, connected to a laptop computer via a FireWire connection. All processing is done in MATLAB at this point. The system operates in two modes: calibration and monitoring. During calibration, the system automatically determines the number, perimeter, centroid, and size of each burner. The system does not assume that all burners are circular, which accommodates irregular burner shapes and skewed images due to off-center camera placement. To calibrate the system, all burners are turned on at full heat for 60 seconds. The captured image is converted to a binary mask by applying a temperature threshold of 80 degrees Celsius. Resulting blobs are dilated to connect individual hot spots within a burner, and any remaining blobs of less than 7 pixels are removed as these small isolated blobs are likely to represent thermal noise. The remaining blobs represent the burners and the perimeter, size and centroid are calculated for each. Please see the results section below for sample results using both a 2-burner and a 4-burner stove top.

When in monitoring mode, the camera continuously captures and analyzes images. Various frame rates have been investigated. A minimum frame rate of 1 frame per 5 seconds was chosen such that typical human activity is likely to be captured by at least one frame. To detect human activity, each frame is thresholded with a temperature of 30 degrees Celsius. Burner regions are excluded, as defined by the masks generated during calibration, including a margin of twice the nominal burner radius to account for possible spreading of heat or the presence of an oversized pot. A number of regions representing thermal noise are often included in the resulting image. Two strategies are employed to differentiate between human activity and thermal noise. The first strategy leverages the fact that human activity originates from outside of the field of view of the camera (i.e. the stove top). The field of view is set such that a small margin surrounding all four sides of the stove top is included.

This ensures that thermal noise on the cooking surface does not extend to the edge of the image, since the air gap surrounding most cooking surfaces prevents significant heat conduction beyond the cooking surface itself. Even in the absence of an air gap, such as a surface-mounted cooking surface embedded in a counter, heat conduction to the surrounding surfaces is expected to be minimal due to the change in material between the cooking surface and the surrounding counter. Therefore, all potential regions of human activity are filtered such that only blobs connected to the image borders are kept. This significantly reduces the rate of false positives, however occasionally thermal noise reaches the image border (e.g. through dense steam escaping an uncovered pot). The second strategy to reduce false positives takes advantage of the fact that human activity tends to be transient while thermal noise is not. Therefore when a new candidate blob appears in the image, it is tracked for 30 seconds. If the blob remains within the field of view for more than 30 seconds, it is rejected as thermal noise. These two strategies (i.e. border filtering and time filtering) result in excellent sensitivity and specificity as detailed below.

RESULTS

Calibration Mode

Several experiments were conducted using two different stove tops: a 2-burner electric coil stove top and a 4-burner electric ceramic stove top. Calibration was successful on both stoves, with the number of burners identified automatically. Calibration results are illustrated in Figure 1 below.

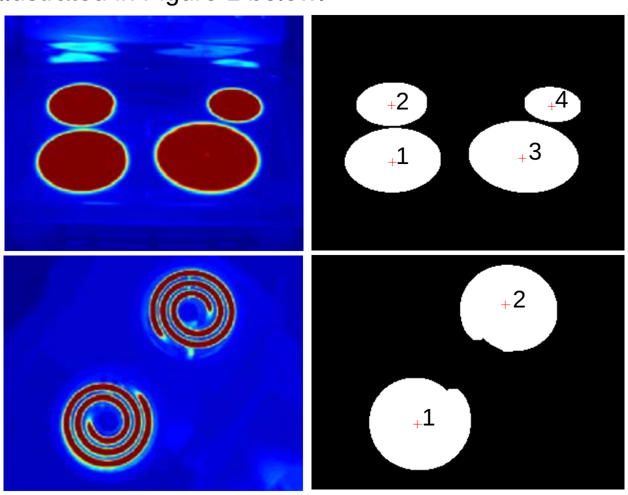

Figure 1: Calibration results for 4-burner (top) and 2-burner (bottom) stove tops showing correct burner identification.

Monitoring Mode

To evaluate the performance of the monitoring mode, four experiments were conducted: 2A and 2B were conducted on the 2-burner stove top (see Figure 3), while 4A and 4B used the 4-burner ceramic stove top (see Figure 2). In all experiments, the calibration protocol described above was first executed to determine the burner locations. In experiments 2A and 2B, only one burner was active and was covered with a pot of water. In experiments 4A and 4B, two burners were active, both covered by pots of water. The four experiments are described in Table 1 below. For the purpose of evaluating the human activity detection algorithms, a frame was considered to truly contain human activity if either a portion of a human (e.g. hand) was visible within the frame, or if an object was manipulated by an unseen hand (e.g. a pot was moved). An 'event' is a contiguous series of frames displaying human activity.

Table 1: Experiment Description

	Experiment Description		**Human Activity**	
	Frame Period	***Total Frames***	***Events***	***Frames***
2A	10s	83	3	19
2B	5s	37	2	8
4A	1s	449	15	136
4B	1s	399	13	79

The human activity monitoring algorithm was applied to the data from each experiment and frames were noted where human activity was detected. The sensitivity and specificity for each experiment are provided in Table 2 below on a per-frame basis. As can be seen, the specificity was excellent (mean=99%) indicating a very low false positive rate. Likewise, the sensitivity of the algorithm is very good on all experiments (mean=90%). The only experiment that saw a significant number of missed positive frames was 4A. These were caused by the surface temperature of the user's arm falling below the threshold of 30 degrees due to thick clothing. As discussed below, this had no impact on per-event accuracy since the onset and termination of human activity was correctly detected as the bare hand entered and left the image frame.

Table 2: Human Activity Monitoring Performance

	Per Frame	
	Sensitivity	***Specificity***
2A	0.89	1.0
2B	1.0	0.96
4A	0.75	1.0
4B	0.94	0.99

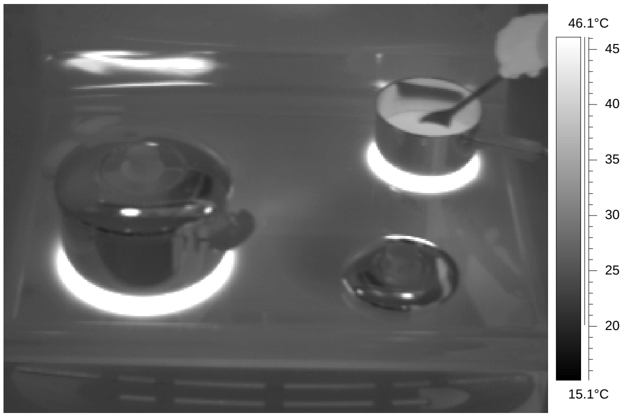

Figure 2: Thermal image of 4-burner stove top showing human activity (top-right) and thermal noise (top-left).

Since this system is meant to detect human activity events, it makes more sense to evaluate the performance of the system in terms of events rather than individual frames. A single event may span several frames. For example, when a user reaches into the field of view of the camera, removes a pot lid, stirs the contents, replaces the lid, and leaves the scene, this single event may span 20 seconds. In our experimental results, at least one frame from each true human event was correctly detected by the algorithms. Hence, the system correctly identifies ALL true events leading to a sensitivity of 1.0. When considering the specificity of the system in terms of events, only two events were falsely predicted (i.e. 2 frames were predicted to contain human activity that did not overlap a true human activity event). Both of these false positive events were due to individual frames (28 and 81) where dense steam from an uncovered pot was observed at the frame border. In both cases, the billowing effect of the steam caused its centroid to differ sufficiently over a period of 30 seconds that it also passed the temporal filter. Therefore this thermal noise was incorrectly classified as human activity. All other instances of thermal noise were rejected by either the border or temporal filter stages as illustrated in a sample frame (Exp 2A, frame 53) in Figure 3. One false positive frame was observed in experiment 2B, but it only served to extend a real event by 1 frame.

Of the two fundamental types of error (I vs. II), overprediction is the most dangerous since an unattended stove may fail to generate alerts if human activity is falsely detected. The underprediction error has must lesser consequences where a user may simply receive unnecessary reminders to attend to their stove. In total, two false positive events occurred while all 33 true human activity events were correctly identified leading to highly acceptable rates of each error type.

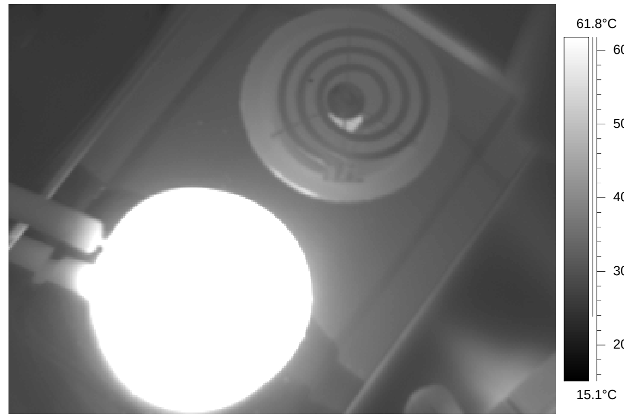

Figure 3: Thermal noise (from a laptop vent) is visible in the bottom right corner of the image. Although is this noise is connected to the image border, thereby bypassing the border filtering step, it was correctly rejected by the temporal filter since its centroid remains constant over 30 seconds.

DISCUSSION

It is important when considering which technologies may be deployed in a home monitoring scenario to consider the perceived privacy of the user as this has a large impact on the acceptance of the technology by the user. For example, many users will feel uneasy with the installation of a pan-tilt-zoom camera in their home, even if they have assurances that the data is only processed locally and never leaves their home. With this in mind, we have limited the field of view of the camera to cover only the stove top with a small margin around it. Furthermore, the camera orientation is also fixed, thereby increasing the user's confidence that their privacy is not being unduly invaded.

Positioning the camera immediately over the cooking surface, by attaching it to the hood vent for example, has advantages for image processing in that the burners will not be skewed by camera angle. However, such a position may lead to an accumulation of cooking grease on the camera lens, thereby occluding the view over time. Instead, the camera will ultimately be positioned slightly off-center and the image processing algorithms have been designed to handle ellipsoid burners rather than assuming a specific geometry such as perfect circles.

This system is designed to be retrofit on an existing stove top. In this way, this system differs from other potential solutions which would monitor stove status (i.e. burner state and temperature) through direct measurement of current flow to the burners and via thermocouples on the stove surface. This is critical in order to keep the cost of the monitoring system at a minimum rather than requiring the user to purchase a new appliance with embedded sensors. This is an especially important consideration for elderly users who are often on fixed incomes. Of course, any commercial system would make use of commercial off-the-shelf components rather than a high-end infrared camera and such a setup will be evaluated in future work.

While the initial target audience for this device are elderly persons with declining mental acuity, other populations may benefit from such a stove monitoring system. For example, persons who are developmentally delayed or who have cognitive disabilities may also benefit from such a system.

CONCLUSIONS

A stove top monitoring system is being developed to assist the elderly to live independently while ensuring that stove use follows best practices for safety. Early results show that the detection of human activity can be achieved with high sensitivity and specificity, with no missed positive events and only a single false positive event observed over four experiments. The final system will combine knowledge of human activity with stove state (burners on/off, warming/cooling, occupied/unoccupied, etc) in order to generate appropriate alerts to greatly increase safety in the kitchen. Such a system is expected to increase the quality of life for elderly users by making it possible to safely live independently in their own home.

ACKNOWLEDGEMENTS

This work was supported by a grant from the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

[1] M. W. Rosenberg and E. G. Moore, “The health of Canada’s elderly population: current status and future implications,” *Journal of the Canadian Medical Association*, vol157, pp. 1025-32, 1997.

[2] Statistics Canada. 2007. Census trends for Canada (table). Census Trends. *2006 Census*. Statistics Canada Catalogue no. 92-597-XWE. Ottawa. Released December 4, 2007.

[3] Public Health Agency of Canada. http://www.phac-aspc.gc.ca/seniors-aines/pubs/fed_paper/fedreport1_01_e.htm, web accessed February, 2008

[4] “E-health services in home and community settings,” *Workshop of Health Canada*, Office of Health and the Information Highway, Aylmer, QC, Canada, March 17-18, 2002.

[5] S. Hussain, S. Dar, “Architecture for smart sensors system for tele-health,” *IEEE International Workshop on Health Pervasive Systems*, Lyon, France, June 29, 2006.

[6] N. Noury, “AILISA: experimental platforms to evaluate remote care and assistive technologies in gerontology,” *IEEE Proceedings of the 7th International Workshop on Enterprise Networking and Computing in Healthcare Industry*, Busan, South Korea, June 23-25, 2005, pp. 67-72.

[7] P. Harmo, T. Taipalus, J. Knuuttila, J. Vallet, A. Halme, “Needs and solutions - home automation and service robots for the elderly and disabled,” *IEEE International Conference on Intelligent Robots and Systems*, AB, Canada, August 2-6, 2005, pp.3201-3206.

[8] A. Arcelus, M. Jones, R. Goubran, and F. Knoefel, “Integration of Smart Home Technologies in a Health Monitoring System for the Elderly,” *First International Workshop on Smart Homes for Tele-Health,,* 21-23 May 2007, Niagara Falls, Ontario.

[9] V. M. S. Lee, W. Wong, and C. Lau, "Home accidents in elderly patients presenting to an emergency department," *Accident and Emergency Nursing*, vol. 7, p. 7, 1999.