

An Address-Event Fall Detector for Assisted Living Applications

Zhengming Fu, *Student Member, IEEE*, Tobi Delbruck, *Senior Member, IEEE*, Patrick Lichtsteiner, *Member, IEEE*, and Eugenio Culurciello, *Member, IEEE*

Abstract—In this paper, we describe an address-event vision system designed to detect accidental falls in elderly home care applications. The system raises an alarm when a fall hazard is detected. We use an asynchronous temporal contrast vision sensor which features sub-millisecond temporal resolution. The sensor reports a fall at ten times higher temporal resolution than a frame-based camera and shows 84% higher bandwidth efficiency as it transmits fall events. A lightweight algorithm computes an instantaneous motion vector and reports fall events. We are able to distinguish fall events from normal human behavior, such as walking, crouching down, and sitting down. Our system is robust to the monitored person's spatial position in a room and presence of pets.

Index Terms—Address-event, AER, assisted living, CMOS image sensor, elderly home care, fall detection, motion detection, temporal-difference, vision sensor.

I. INTRODUCTION

HUMAN society is experiencing tremendous demographic changes in aging since the turn of the 20th century. The current life expectancy in the US is 77.85 years, and is extending as medical care is improved. According to a report of U.S. Census Bureau, there will be a 210% increase in the population with age of 65 and over within the next 50 years [1]. The substantial increase in the ageing population will cause society to face two challenges: increase of ageing people will require more investment in elderly care services; the decrease of working population will cause shortage of skilled caregivers of elders. In the future, this imbalance between the number of elderly people and that of the caregivers will be exacerbated when life expectancies increase. Intelligent elderly care systems deliver one solution to reduce the workload of elderly caregivers without compromising the quality of services.

In the past, various solutions were proposed based on emerging technologies. *Video monitoring* is a commonly-used solution in nursing institutions. But considerable human resource is required in order to monitor activities. Patients' privacy is also compromised when they are monitored. Another common solution is to have patients raise alarms when they

are in trouble by pushing a button on a wearable or pendant device [2]. This solution depends on the patient's capability and willingness to raise alarm. For example, a fall may result in unconsciousness. A dementia patient may not be able or willing to push the button when necessary [3]. Both scenarios would limit this "push-the-button" solution in applications. Other solutions include *wearable devices*, such as motion detector, accelerometers, etc [4]–[9]. They are with patients all the time, continuously collecting and streaming out physical parameters. An alarm is raised when predefined conditions of these signatures are satisfied. The effectiveness of wearable sensors, however, is also restricted by the willingness of patients to wear them.

Fall is a major health hazard for the elders when they live independently [10]. Approximately 30% of 65-year-old people fall each year. This number becomes higher in medical service institutions. Although less than one fall in ten results in an injury, a fifth of fall incidents require medical attention. Another recent publication indicates that 50% of patients in nursery institutions fall each year, while 40% of them fall more than once [11]. How to effectively assess, respond, and assist elderly patients in trouble becomes an important research topic in medical elderly care services [12].

Elderly care systems aim to effectively evaluate and respond to the behavior of elderly people when they live alone. These systems have the following requirements.

- 1) The sensor systems should be *non-intrusive* to patient life. The impact of elderly care systems on patients' lives is expected to be reduced to the minimum. From the system's perspective, elderly care systems are expected to be small enough to be placed in appropriate locations. An ideal elderly care system operates with zero maintenance.
- 2) The sensor systems should preserve patient privacy. Most people under care expect that their privacy is respected. No private information should be released until an emergency is detected. Many elders are against using commercial-off-the-shelf (COTS) cameras or microphones in their home, because they feel they are monitored and their privacy is compromised. In elderly care sensor nodes, most information analysis and decision-making should occur within the detection nodes. This eliminates the necessity to transmit information outside the detector and protects patient privacy.

Fig. 1 illustrates the fall detector setup. The detectors take multiple side-views of the scene in order to detect accidental activities and raise alarms. The vision systems are mounted on the wall at a height of 0.8 m, which is approximately the same height of a light switch. Our approach is innovative for two

Manuscript received October 1, 2007; revised . First published July 25, 2008; current version published September 10, 2008. This paper was recommended by Associate R. Etienne-Cummings.

Z. Fu and E. Culurciello are with the Department of Electrical Engineering, Yale University, New Haven, CT 06511 USA (e-mail: zhengming.fu@yale.edu; eugenio.culurciello@yale.edu).

T. Delbruck and P. Lichtsteiner are with the Institute for Neuroinformatics (INI), Zurich CH-8057, Switzerland.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TBCAS.2008.924448




Fig. 1. Address-event fall detectors are used for assisted living applications. The detectors are mounted on the wall at a height of 0.8 m, which is approximately the same height of a light switch.

reasons: First, an asynchronous temporal contrast vision sensor reports pixel changes with a latency on the order of milliseconds. Second, a lightweight computation algorithm plus a fast readout allow us to compute an instantaneous motion vector and report fall events. This cannot be done with a frame-based temporal-difference image sensor because the frame rate is constant, and redundant information in images saturate the transmission bandwidth. Notice that in this paper we will refer to a motion detection system performed with temporal differences image sensors only.

This paper is divided into seven sections. In Section II, we describe the design overview for elderly home-care systems. Section III describes the temporal contrast (motion-detection) vision sensor and the test platform used in the fall detection work. In Section IV, we evaluate the asynchronous temporal contrast vision sensor in tracking fast movement. In Section V and Section VI, a lightweight moving-average algorithm to compute centroid events is presented. This algorithm is then evaluated as a fall detector. Section VII describes the design concerns in a fall detector system. Section VIII concludes the paper.

II. AN ATC VISION SENSOR

The core technology used in our research is an asynchronous temporal contrast (ATC) vision sensor. A temporal contrast vision sensor extracts changing pixels (motion events) from the background [13] and reports temporal contrast, which is equivalent to image reflectance change when lighting is constant. A temporal contrast vision sensor can extract motion information because, in normal lighting conditions, the intensity of a significant number of pixels changes as a subject moves in the scene [14]–[16]. In the ATC vision sensor used here, every pixel reports a change in illumination above a certain threshold with an asynchronous event, i.e., pixels are not scanned with a regular frame rate but every pixel is self-timed. In case of an event, the corresponding pixel address is transmitted. After the event is acknowledged by an external receiver, the pixel resets itself.

A key feature of ATC is the temporal contrast response which means that the sensor reports scene reflectance changes (caused e.g. by moving objects), discarding local absolute illumination. A major advantage of this ATC image sensor is that it pushes information to the receiver once a predefined condition is satisfied. This feature is important in high-speed vision systems




Fig. 2. The 64×64 address-event temporal contrast vision sensor used in the fall detector system [17], [18].




Fig. 3. Temporal contrast image from the (b) ATC image sensor and (a) one intensity frame. The subject is swaying left to right. The ATC imaging system is placed in front of the subject with a distance of 3 meters and a height of 0.8 m.

because a pixel sends information of interest immediately, instead of waiting for its polling sequence. A pixel generates a higher rate of events when it experiences larger changes in light intensity.

Fig. 2 shows the ATC image sensor system [17]–[20] used in the fall detection experiment. The temporal contrast vision sensor contains a 64×64 array of pixels and responds to relative changes in light intensity. The imaging system streams a series of time-stamped address-events from the vision chip, and sends them to a PC via a USB interface. The data is reported in the address-event format with 12 bits (6 for X address, 6 for Y address in a 64×64 image sensor.) The silhouette of a moving subject can be reconstructed on a PC (The address-event vision reconstruction software is available from <http://www.jaer.wiki.sourceforge.net/>). The vision system uses a Rainbow S8 mm 1:1.3 lens, and the lens format is $2/3''$. Fig. 3 shows an image from the ATC image sensor and its targeted scene. The imaging system is placed in front of the subject with a distance of 3 m and a height of 0.8 m. The image sensor features a high dynamic range of 120 dB. The sensor consumes 30 mW of power at 3.3 V, which is comparable to most low-power COTS image chips on the market [21]–[23]. The power consumption is approximately 120 mW for the USB device in Fig. 2. Notice that the camera is used as a line-powered fixed device in home and laboratory installations. We suppose that the system will be professionally installed by caregivers.

III. COMPARISON BETWEEN ATC AND FRAME-BASED TEMPORAL DIFFERENCE VISION SENSORS

In this section we compare an ATC image sensor with a frame-based system (a COTS web-camera) and characterize them by tracking an object in free fall. We demonstrate that the ATC image sensor performs better in high-speed tracking for two reasons: Firstly, the ATC image sensor has higher temporal resolution and delivers timely information on motion events. Secondly, the ATC image sensor ranks data based on importance and selectively sends informative data on motion only. This reduces information size and communication bandwidth. A COTS camera samples images at low rates (30 fps), resulting in at most a few temporal difference frames of information for each fall event. This little data is not enough to compute accurate velocity and acceleration measurement to distinguish a fall and is a major impediment to the use of COTS camera for this application. The image data is also blurred by the camera speed.

In order to use a COTS camera to detect motion, some image data manipulation is necessary. For comparative purpose, we wrote a real-time temporal difference image emulator using an COTS camera [24]–[26]. The software can be downloaded from <http://www.eng.yale.edu/elab/FallDetect.html>. Image frames from the COTS camera are down-sampled to 64×64 pixels, and pairwise subtracted to mimic a temporal difference imager. Using the same frame twice in two subsequent differences is not necessary, since the event resolution is not increased, but the overall number of events increases, at the expense of more computation after readout. For this manipulation 8,192-bit subtractions and thresholdings are performed by a PC ($64 \times 64 = 4096$ byte subtractions and an identical number of thresholding operations). The threshold of the COTS was set to match the one from the ATC (10 for an 8 bit pixel output). The temporal difference frames are then converted into an address-event stream, in order to compare them to the ATC output. This is performed by reporting only the address of the pixels that have changed by a threshold. This comparison is fair because address-event is the most efficient way to report sparse matrices of events.

Fig. 4 shows the measured responses as the ATC vision sensor tracks a box in free fall. The sensor communicates motion events at 1330 event/s when it is monitoring the object's fall. The event rate reduces to 221 event/s in the quiet period when no motion is present in the scene. These noise events are due to source/drain junction leakage in pixel transistors, and are sparse, uncorrelated in space and time. The noise events are represented by circles in Fig. 4. The noise events can be filtered out by the fact that they are spatio-temporally uncorrelated [27] but we chose not to do so to keep the computational model closely matched to cheap embedded architectures. In this experiment 1590 events have been collected during the 1.1s fall, 94% of which describe the fall, the rest are noise.

Fig. 5 shows the measured address-event outputs from the frame-based image emulator as it tracks a box's fall. The event rate is 150 event/s in average and ten times less than the ATC vision sensor. Every frame contains a lot of redundant information due to the unchanged background. Fig. 5 reports only 232




Fig. 4. (a) Measured responses while the ATC vision sensor tracks an object thrown in the air and then falling. The object is 3 m away and the camera is installed at 0.8 m. (b) Noise events (represented by cycles) and fall events (represented by dot) distribution when the ATC vision sensor tracks the object free-falls.

events during the 1s fall, with no added noise. Notice also the spread of events in the Y axis for each frame: it is up to 15 pixels out of a total of 64 resulting in a 23% spread for COTS. On the other hand, it is only 3 pixels in the reconstructed ATC frame for a spread of 4.6% (see Fig. 4(b)); more precisely the fall events between 3 and 3.2 s). This data shows that the ATC system can perform at least 5 times more precise vertical velocity calculations than a COTS sensor. Notice that we can generate an ATC frame for comparison purposes by collecting events for 30 ms and then generating an histogram frame.

ATC vision sensors have two main advantages when compared to frame-based image sensors: first, the ATC vision sensor has a higher temporal resolution in high-speed tracking applications. In the experiment the ATC vision sensor shows a 10 times higher event rate as it tracks the free fall. The uniform frame rate of the COTS camera imposes an upper limit on the temporal difference sampling rate. Second, the ATC vision sensor has a higher bandwidth efficiency because it selectively sends information. Given this experimental setting, with an image resolution of 64×64 , the ATC vision sensor saves over 84% bandwidth for transmission of the image data. ($[12 - \text{bit address} \times 1590 \text{ events in } 1 \text{ s}] = 19080 \text{ bits in the ATC vision sensor versus}$




Fig. 5. Address-event responses from a frame-based temporal difference image emulator as it is tracking a free-falling object. The camera has a resolution of 64×64 and runs at 30 frame/s. The imaged object is 3 m away and the camera is installed at 0.8 m. The uniform frame rate of the camera imposes a limit on the temporal difference sampling rate, which is 150 event/s on average.

$[1 \text{ bit/pixel} \times 64 \times 64 \times 30 \text{ frame/s}] = 122,080$ bits in the frame-based image emulator). In order to output a temporal difference image, extra bit-operations are performed by an external processor.

IV. DETECTING FALLS USING ADDRESS-EVENT VISION SENSOR

In this section, we use a temporal average of the motion events from the ATC vision sensor, here referred as *centroid event*, to track fall hazards and evaluate its dynamics. We describe a light-weight averaging algorithm and demonstrate that the centroid event can instantaneously capture signatures of fall hazards and differentiate them from other human behavior.

A. Centroid Event Computation

In machine-vision research, many successful algorithms have been proposed to profile human behavior. However, these algorithms are difficult to be implemented on sensor nodes with limited computing power. Centroids are an effective way to estimate object motion in space. Centroids can be computed as temporal averages of a series of events. A single centroid event address, (x_c, y_c) , during a fixed period can be calculated using (1). Where N is the number of events in a given window, and (x_i, y_i) are the event addresses. The spatial average is performed over a time window of 30 ms or a minimum $N = 10$.

$$x_c = \left\lfloor \frac{\sum_{i=1}^N x_i}{N} \right\rfloor, \quad y_c = \left\lfloor \frac{\sum_{i=1}^N y_i}{N} \right\rfloor \quad (1)$$

Centroids are computed as moving averages of a series of events [28], [29]. This can be done with high temporal resolution and low hardware cost. One possible hardware implementation is using a FIFO buffer, which stores the events that occur in a fixed time period. As a new event comes in, a computation cycle starts with removing the expired events and appending the incoming event in the buffer. All events in the buffer are averaged using (1) to get a centroid event at this time.




Fig. 6. Vertical address (Y) of events when the camera is monitoring a fall. The fall occurs in 0.9 s, and 5043 events are transmitted.

B. Evaluation of a Fall

Fall detection is of particularly importance when a person lives alone. This is one of the most critical scenarios in assisted-living environment and requires immediate attention. When more than one person are in the same room, one of them can call medical assistance if a fall hazard occurs.

The data rate of the ATC vision sensor, i.e. *event rate*, is correlated with motion speed, size and light contrast in the scene. When the scene (lighting condition) is set, the event rate is useful for characterizing the motion in the scene. For example, a faster motion causes more events to be generated during a time period. Due to the different event rates, the number of events to be averaged varies depending on the motion in the scene. Fig. 6 shows event responses directly measured from the ATC imager when it tracks a person's fall. Before the fall, the subject swings side-to-side in a normal walking manner. This side-to-side motion can be compared to the data collected during the fall. Each side-to-side movement generates approximately 460 event/s, while the fall causes 5600 event/s. The burst of events is due to the fast motion in the scene. For comparison, when the same fall is monitored by a frame-based temporal difference emulator, 1500 events are communicated in one second. This data is computed using a COTS camera placed at a distance of 3 m and running at 30 frame/s. This comparison shows that the ATC image sensor reports 3.7 times higher temporal resolution than a frame-based imager when tracking the fall.

Fig. 7 shows centroid event responses when the imager monitors a person's fall and crouch-down. A crouch-down scenario is considered because it is a motion similar to a fall but it happens on a longer time scale. In the experiment, a short walking motion with side-to-side oscillations happens before the fall and the crouch-down. In Fig. 7(a), a fall causes the vertical address to decrease from 30 to 5 in 0.9 s. An instantaneous event rate of 5600 event/s indicates significant changes in the scene. Fig. 7(b) shows changes in the centroid when the imager observes a person crouching down, and then getting up. In this case, the centroid vertical coordinates reports a slower vertical




Fig. 7. Vertical address of the centroid events (Y_c) (a) when a person falls, and (b) when a person crouches down. In (a), when the person falls, the centroid vertical address decreases by 25 pixels in 0.9 s. The event rate is approximately 5600 event/s. In (b), when the person crouches down, the centroid vertical address decreases by 20 pixels in 2 s. The event rate is approximately 310 event/s. The fall causes faster decrease in the vertical address than the crouch-down.

velocity and a slower event rate. The Y address decreases from 30 to 15 in 2 s at an event rate of 310 event/s.

In order to numerically evaluate the dynamics in the ATC imager's event outputs, a centroid vertical velocity is computed in (2), where $\Delta t = t_i - t_j < T$ and T is a fixed time period.

$$V_y = \frac{\Delta y_c}{\Delta t} = \frac{(y_{c,i} - y_{c,j})}{t_i - t_j} \quad (2)$$

It is important to provide a measurement of the centroid that is invariant to the distance between camera and the imaged subject. In order to be invariant to distance, the vertical velocity (V_y) is divided by the height of the subject in pixels, y_d , as shown in (3). The height is the difference between the maximum and minimum standard deviations of the vertical address during a fixed 30 ms time period. The unit of a normalized vertical velocity is second^{-1} .

$$V_{y,norm} = \frac{V_y}{y_d} = \frac{\frac{\Delta y}{\Delta t}}{y_d} = \frac{(y_i - y_j)/y_d}{t_i - t_j} \quad (3)$$

Fig. 8 is the most important experimental results of this research. It shows centroid vertical velocities in four scenarios: 1) a person crouches down; 2) a box free-falls; 3) a person falls forward; and 4) a person falls backward. The velocities are normalized for distance using (3). The centroids in Fig. 8 show both positive and negative velocities because of the difference between the moving-average and the physical centroids when tracking a fast moving subject. The ATC vision sensor stochastically fires events [30]. The good estimation of physical centroid requires time to collect events. We set the averaging windows small in order to keep a high temporal resolution in tracking. This causes the events' average to oscillate around the physical centroid when the vision sensor monitors a fast motion. For example, in a human fall case, the first average of events, which

mostly describe the lower part of the person, is followed by the second average of events, which mainly describe the upper body. Even though the vertical address of the physical centroid decreases, the estimated centroid, i.e. average, could still show a positive velocity.

Notice that, in Fig. 8, a fall shows a peak velocity of -3 s^{-1} in the vertical address decrease. The peak velocity is close to that of the free-falling box. When the vision sensor monitors a person crouching down, the centroid vertical velocity reaches a peak of -1 s^{-1} , a value that is three times smaller than the fall's velocity. By estimating the peak vertical velocity, a fall is differentiated from other human behavior. More importantly, the fall shows more than three times peak-to-peak vertical velocity of the side-to-side swing before the fall, which is comparable to the response to normal walking.

V. RESULTS AND DISCUSSION

We evaluated various scenarios using laboratory trials. The ATC vision sensor monitored a group of three people individually at a distance of 3 m. The view was from the side and the vision sensor was mounted at a height of 0.8 m, which is approximately the same height of a light switch on the wall. Each actor performed a predefined task list, including fall and other normal human behavior. The fall scenarios we tested in this work included a variety of fall types, such as fall forward, fall backward, and fall sideways. We also monitored other scenarios besides the falls, which frequently happen in home-assisted living applications. These tasks include a person walking, crouching down, sitting down and a cat walking. Pets are considered in this work because they are common company to the elders.

Experimental results are illustrated in Fig. 9. They show the comparison of the initial centroid vertical address and peak centroid vertical velocity. The space in Fig. 9 is divided into four areas. The centroids in Area II are those of human behavior,




Fig. 8. Centroid vertical velocities ($V_{y,norm}$) in address-event streams demonstrate the difference among four scenarios: (a) a person crouches down; (b) a box free-falls; (c) a person falls backward; and (d) a person falls forward. The velocities are normalized for distance between the imaged subject and the image sensor. The red lines in the figures outline the boundary of the peak vertical velocities.




Fig. 9. Centroid address-event evaluation matrix of five common home-assisted living scenarios. The centroid vertical velocities are normalized for distance.

including crouching-down, walking, and getting-up. They have higher coordinates than the pet centroids in Area I, which are closer to ground. The centroids in Area III are reported as fall hazards. They have higher peak vertical velocities than other human behavior.

The centroids of people are around the middle of the camera view when people move. The centroids fluctuate between 30 and 40. This coincides with the expected height of human body. The pet’s centroid moves at a lower coordinate, generally closer to ground. Both the fall and crouch-down demonstrate negative vertical velocities, which are due to the decrease in vertical address. The fall’s centroid decreases at a much higher speed of -3 s^{-1} . The fast decrease in the centroid vertical address distinguishes a fall from other normal human behavior.

Table I shows the statistics of the event rate when the vision system observes human behavior. When the system monitors a fall, the burst event rate is twice that of a person walking. This is because that the light intensity in more pixels changes in a unit time period when the system monitors a faster motion. When there is no motion in the scene, the event rate is as slow as 300 event/s with a larger variance of 30.5 event/s. This is due to the noise events associated with source/drain junction leakage in pixels’ reset transistors.

In order to design high-accuracy fall detectors, there are a number of issues that need to be considered related to system design and implementation. When monitored subjects are too close to a fall detector, they can block the camera’s view, and

TABLE I
STATISTICS OF THE EVENT RATE IN DIFFERENT EXPERIMENTAL SCENARIOS

Behavior	Average Event Rate (event/s)	Variance (event/s)
Walk	2100	11.3
Crouch Down	3500	15.2
Fall Down	5120	10.2
Sit Down	3150	15.2
No Motion	300	30.5




Fig. 10. Relation between (1) the event rate and the distance from the camera and (2) the percentage of active pixels and the distance. When the monitored person is 1 meter away from the detector, 80% of the active pixels in the sensor send events. In this situation, the detector cannot make accurate estimation of the motion centroid.

make it impossible for the detector to make accurate motion estimations. In order to resolve this problem, two fall detector systems can be installed on the opposite corners of a monitored room. When the subject is too close to one of the detectors, the other detector can accurately perform the fall detection.

Fig. 10 shows the relation between the distance from the camera and the event rate. It also reports the percentage of active pixels as a function of subject distance. Both the event rate and the percentage increase as the monitored person gets closer to the detector. When the person is 1 m away from the detector, 80% of the pixels in the sensor send events. In this situation, the detector cannot make accurate estimations of the motion centroid because it cannot compute the position of the centroid and the dimension of the moving target.

Multiple fall detectors can be deployed with reasonable budget. The detector we used in this work is a self-contained system. The vision system includes a customized vision sensor, a microcontroller with regulators. Due to its low computation complexity, a low-power and low-cost 16-bit microcontroller [31] is commercially available for the centroid computation and thresholding. The power budget of the detector is approximately 31 mW, including 30 mW for the image sensor [19] and 1 mW for the 16-bit microcontroller. This is an attractive solution when compared to complex computer-vision techniques, which require significant computation resource to run these algorithms.

Pets are common company to the elders when they live alone. In most cases, pets are not helpful to the elders when they are in




Fig. 11. Flow of a fall detection and reporting operations.

trouble. It is one of our concerns to distinguish between pet and human in motion. A pet's motion centroid is usually at a lower coordinate than a human. In the experiment, when both subjects are 2 m away from the camera, the vertical address of human centroid fluctuates around 30 to 40, while a pet is approximately at 10.

Fig. 11 illustrates a fall detecting flowchart. A negative threshold of vertical velocity is predefined. As long as the event centroid change passes the threshold, a potential fall alarm is raised. The detector causes a "potential fall detected" message and starts a timer, for example, for 30 s. The detector will continue to track motion in the scene until the timer expires. In many cases, the fall event is not traumatic to the elder, and he can get up and walk away as normal. In this case, the centroid moves up with a positive vertical speed. This is similar to the getting-up scenario shown in Fig. 8(a). This is similar to other scenarios, like an object falling in the field of view or a pet jumping around. As the person walks around normally, the motion centroid moves around the middle of the camera view. The person can also cancel the potential alarm. This also addresses sudden changes in lighting condition, as reflective objects or light switched off. There are two other possible scenarios: the person is striving to get up but fails; a pet moves around the injured elder. Both cases cause the motion centroid to move close to ground. The detector system continue to hold the potential alarm message until the timer expires. When this happens, the potential alarm is confirmed and communicated to assistance providers.

VI. CONCLUSION

In this paper, we evaluate an address-event temporal contrast fall detector. We propose using this detector to detect falls, a major health hazard in elderly home-care applications. This

approach innovates in two ways: First, an asynchronous temporal contrast vision sensor features high temporal resolution and reports pixel changes with a latency on the order of milliseconds. Second, a lightweight moving-average algorithm allows the detector to compute instantaneous motion vectors and report fall hazards immediately with low computational effort. A Matlab code for the fall detector system can be downloaded at <http://www.eng.yale.edu/elab/FallDetect.html>.

The fall detector presented in this paper can detect falls in a home assisted living environment, and possesses the following advantages:

- 1) The motion detector is small in size (like a web-camera) and it is thus non-intrusive. The installation of these nodes in elderly people's apartment will cause minimal change in their living patterns. Deployment is easy because the device can be affixed to a wall near a power socket. Wireless deployment is not necessary, since the necessity to change batteries would be a burden on the patient or caregiver.
- 2) The motion detector protects the patient's privacy. An address-event imager takes no image snapshot, and filters out detailed visual appearance of the patients. All images are processed locally. No data is communicated until an emergency is detected.

REFERENCES

- [1] U.S. Census Bur. Population Data Public Web Site [Online]. Available: <http://www.census.gov/population/www/>
- [2] T. Hori, Y. Nishida, and H. Aizawa, "Sensor network for supporting elderly care home," *IEEE Proceeding of Sensors*, vol. 2, pp. 575–578, Apr.–Jun. 2004.
- [3] S. Allin, Bharucha, J. Zimmerman, D. Wilson, and M. Roberson, "Toward the automatic assessment of behavioral disturbances dementia," in *2nd Int. Workshop on Ubiquitous Computing for Pervasive Healthcare Appl. (UbiHealth 2003)*, October 2003, pp. 25–29.
- [4] A. Tabar, A. Keshavarz, and H. Aghajan, "Smart home care network using sensor fusion and distributed vision-based reasoning," in *4th ACM Int. Workshop on Video Surveillance and Sensor Networks (VSSN 2006)*, October 2006, pp. 145–154.
- [5] G. Yang, B. Lo, J. Wang, M. Rans, S. Thiemjarus, and J. Ng, "From sensor networks to behaviour profiling: A homecare perspective of intelligent buildings," in *IEE Seminar for Intell. Buildings*, Nov. 2004, pp. 10–15.
- [6] B. Lo, J. Wang, and G. Yang, "From sensor networks to behaviour profiling: Ubiquitous sensing for managed homecare of the elderly," in *IEEE 3rd Int. Conf. on Pervasive Computing (PERVASIVE)*, May 2005, pp. 1–7.
- [7] M. Alwan, P. Rajendra, S. Kell, and D. Mack, "A smart and passive floor-vibration based fall detector," in *Inf. and Commun. Technol.*, Apr. 2006, pp. 1003–1007.
- [8] T. Tamura, "Wearable accelerometer in clinical use," in *27th Annu. Conf. on Medicine and Biology*, Sep. 2005, pp. 7165–7166.
- [9] T. Degen, H. Jaeckel, M. Rufer, and S. Wyss, "Speedy: A fall detector in a wrist watch," in *Proc. IEEE 7th Int. Symp. on Wearable Computers*, Sep. 2003, pp. 184–187.
- [10] E. Duthie, "Falls," in *Medical Clinics of North America*, Apr.–Jun. 1989, pp. 1321–1335.
- [11] R. Tideiksaar, *Falling in Old Age: Prevention and Management*. New York: Springer-Verlag, 1998.
- [12] A. Sixsmith and N. Johnson, "A smart sensor to detect the falls of the elderly," *IEEE Pervasive Comput.*, vol. 2, pp. 42–47, April–June 2004.
- [13] V. Gruet and R. Etienne-Cummings, "A pipelined temporal difference imager," in *IEEE Int. Symp. on Circuits Syst. ISCAS*, May 2002, vol. 3, pp. 683–686.
- [14] C. Higgins and V. Pant, "A biomimetic VLSI sensor for visual tracking of small moving targets," *IEEE Trans. Circuits Syst. I: Reg. Papers*, vol. 51, pp. 2384–2394, Dec. 2004.
- [15] F. Okamoto, Y. Fujimoto, T. Nagata, and M. Furumiya, "A CMOS imager with new focal-plane motion detectors," in *Proc. 1999 IEEE Int. Symp. on VLSI Circuits*, Jun. 1999, pp. 139–140.

- [16] G. Barrows, K. Miller, and B. Krantz, "Fusing neuromorphic motion detector outputs for robust optic flow measurement," in *Int. 1999 Joint Conf. on Neural Netw. (IJCNN '99)*, Jun. 1999, vol. 4, pp. 2296–2301.
- [17] P. Lichtsteiner and T. Delbruck, "A 64×64 AER logarithmic temporal derivative silicon retina," in *IEEE Ph.D. Research In Micro-Electronics & Electronics (PRIME)*, 2005, pp. 202–205.
- [18] Lichtsteiner and T. Delbruck, " 64×64 event-driven logarithmic temporal derivative silicon retina," in *2005 IEEE Workshop on Charge-Coupled Devices and Adv. Image Sensors*, 2005, pp. 157–160.
- [19] P. Lichtsteiner, C. Posch, and T. Delbruck, "A 128×128 120 db 30 mw asynchronous vision sensor that responds to relative intensity change," in *2006 Int. Solid State Circuits Conf. (ISSCC 2006)*, 2006, pp. 508–509.
- [20] P. Lichtsteiner and T. Delbruck, "A 100 db dynamic range high-speed dual-line optical transient sensor with asynchronous readout," in *2006 IEEE Int. Symp. (ISCAS 2006)*, 2006, pp. 1659–1662.
- [21] Micron Semiconductors Image Product Internet Guideline [Online]. Available: <http://micron.com/products/imaging/>
- [22] Agilent Camera Modules and Image Sensors for Sensor Networks [Online]. Available: <http://www.home.agilent.com/>
- [23] Omnivision Technology Semiconductors Image Product Guideline [Online]. Available: <http://www.ovt.com/products/>
- [24] T. Teixeira, A. Andreou, and E. Culurciello, "An address-event image sensor network," in *2006 IEEE Int. Symp. on Circuits and Syst.*, May 2006, pp. 955–958.
- [25] T. Teixeira, D. Lymberopoulos, E. Culurciello, Y. Aloimonos, and A. Savvides, "A lightweight camera sensor network operating on symbolic information," in *Proc. First Workshop on Distributed Smart Cameras 2006*, April 2006, pp. 55–58.
- [26] T. Teixeira, E. Culurciello, E. Park, D. Lymberopoulos, A. B. Sweeney, and A. Savvides, "Address-event imagers for sensor networks: Evaluation and modeling," in *Proc. Inf. Process. in Sensor Netw. (IPSN)*, April 2006, pp. 458–466.
- [27] Java AER Open Source Project [Online]. Available: <http://jaer.wiki.sourceforge.net/>
- [28] R. Mejiya, *Time Series Analysis: Theory and Practice*.
- [29] T. Yamamoto, *Time Series Analysis in Economics*.
- [30] A. B. Apsel and A. G. Andreou, "Analysis of data reconstruction efficiency using stochastic encoding and an integrating receiver," *IEEE Trans. Circuits Syst.: Analog Digit. Signal Process.*, vol. 48, pp. 890–897, October 2001.
- [31] 16-Bit Ultra-Low-Power Microcontroller 60 kB Flash HW Multiplier [Online]. Available: <http://focus.ti.com/docs/prod/folders/print/msp430f149.html>

Zhengming Fu (S'01) received the B.E. degree from Shanghai University, Shanghai, China, the M.S. and M.Ph. degree in electrical engineering from Yale University, New Haven, CT, where he is currently working toward the Ph.D. degree.

He is a research assistant affiliated with Yale E-lab. His research interests are analog/mixed-signal and low-power circuit design, CMOS intelligent image sensor for sensor network applications.



Tobi Delbruck (M'99–SM'06) received the in undergraduate degree from University of California in physics and applied mathematics and the Ph.D. degree from Caltech, Pasadena, in 1993 in computation and neural systems, with Carver Mead.

He is a group leader at the Institute of Neuroinformatics (INI), part of ETH Zurich and the University of Zurich, Switzerland, and visiting scientist at Caltech, Pasadena. His main interest is in developing neuromorphic electronics, particularly vision sensor chips. He co-invented the standard neuromorphic adaptive photoreceptor circuit and bump circuit. Subsequently he worked for several years for Arithmos, Synaptics, National Semiconductor, and Foveon, where he was one of the founding employees. In 1998 he moved to Switzerland to join INI. In 2002 he was lead developer of the tactile luminous floor used in INI's exhibit Ada: Playful Intelligent Space.

Dr. Delbruck has been awarded 8 patents, and has over 30 refereed papers in journals and conferences, 4 book chapters, and 1 book. He has been awarded 6 IEEE awards, including the 2006 ISSCC Jan Van Vessel Outstanding European Paper Award. He was named IEEE Circuits and Systems Society Distinguished Lecturer 2007.



Patrick Lichtsteiner (M'06) received the Diploma (equivalent to M.S.) degree in physics and the Ph.D. degree from the Swiss Federal Institute of Technology, Zurich, Switzerland, in 2002 and 2006, respectively.

He is currently a Postdoctoral fellow at the Institute of Neuroinformatics, University of Zurich and Swiss Federal Institute of Technology, Zurich, Switzerland. His research interests include CMOS imaging, neuromorphic vision sensors and high speed vision.

Dr. Lichtsteiner and his colleagues have been awarded 4 prizes for IEEE conference papers, including the 2006 ISSCC Jan Van Vessel Outstanding Paper Award.



Eugenio Culurciello (S'97–M'99) received the Ph.D. degree in electrical and computer engineering in 2004 from The Johns Hopkins University, Baltimore, MD.

In July 2004 he joined the Department of Electrical Engineering at Yale University, New Haven, CT, where he is currently an assistant professor. He founded and instrumented the E-Lab laboratory in 2004. His research interests are in analog and mixed-mode integrated circuits for biomedical applications, sensors and networks, biological sensors,

silicon-on-insulator design, and bio-inspired systems.