User-centric, embedded vision-based human monitoring: A concept and a healthcare use case

Tahir Nawaz
Computational Vision Group
School of Systems Engineering
University of Reading
Reading, United Kingdom
t.h.nawaz@reading.ac.uk

Bernhard Rinner
Institute of Networked and Embedded Systems
Alpen-Adria Universitaet
Klagenfurt, Austria
bernhard.rinner@aau.at

James Ferryman
Computational Vision Group
School of Systems Engineering
University of Reading
Reading, United Kingdom
j.m.ferryman@reading.ac.uk

ABSTRACT

In an Internet of Things (IoT) camera-based monitoring application the transmission of images away from the video sensors for processing poses security and privacy risks. Hence, there is a need for an advanced trusted user-centric monitoring system that pushes the application of security and privacy protection closer to the sensor itself and which enables an enhanced control on data privacy. To this end, this white paper proposes a new approach that involves sensor edge computing to enable sensor-level security and privacy protection and allows observed individuals to interact and control their data without impacting on the quality of the data for further processing. Overall, an IoT vision system is presented that employs a network of fixed embedded cameras in a highly trusted manner, possessing both privacy-protecting and data security features. As a potential application, we discuss an Ambient Assisted Living (AAL) healthcare use case demanding privacy and security for outpatients.

Keywords

Edge computing; Privacy and security protection; White paper.

1. INTRODUCTION

The use of Internet of Things (IoT) camera-based human monitoring systems is expected to become increasingly ubiquitous [26, 33] in a wide variety of applications [1, 7, 33, 34] considering the rapid advancements in ICT. This inevitably leads to an enhanced need of the incorporation of a notion of 'trust' (from a user's point of view) in such systems. In an IoT camera-based monitoring scenario, user classes could primarily be of two types: (1) observed individuals, which are people being monitored remotely in an environment; (2) remote observers, which are the authorized personnel enjoined to remotely monitor individuals' activities. Other type of user classes could be unauthorized remote observers and non users. For the acceptance of IoT camera-based monitoring by users, trust is of high importance. The term, trust, is here defined to possess two key features: security (i.e. ensuring the integrity and confidentiality of the recorded image data); and privacy protection (i.e. altering the recorded imagery to ensure that the observed individuals are not/less recognizable). Present-day camera-based monitoring approaches [3, 21, 34] generally aim to achieve such trust by providing security and privacy protection at a stage away from the sensing platform. Indeed, such solutions could result in a heightened risk of leakage of recorded image data and hence a compromise in the privacy of observed individuals. This risk also limits such monitoring approaches from becoming entirely pervasive by hindering their use in scenarios requiring high levels of privacy [31]. Moreover, existing approaches are generally more remote-observer-centric than observed-individual-centric: they tend to focus more on technological improvements to address mainly the needs of remote observers than on the acceptability of technology by addressing also the needs of the observed individuals.

Ambient Assistive Living (AAL) is a key application domain where the notions of trust and user centricity are highly desirable. A fundamental AAL task involves monitoring activities of outpatients (the observed individuals) either autonomously by algorithms or manually by the healthcare professionals (the remote observers). Driven by European and international initiatives, there exists a large body of research on human activity monitoring for AAL. In the majority of these works (e.g., [6, 9]), security and privacy protection is—if at all—addressed from a technology perspective.

This paper proposes a new concept that attempts to make a radical impact in the form of a strongly user-centric, highly-trusted and pervasive monitoring by (1) adopting a user-centric approach that essentially balances the conflicting needs of both observed individuals as well as remote observers by empowering the former with an increased level of control over their recorded image data, (2) equipping sensors with 'edge computing' by pushing the security and privacy protection features closer to the sensor thus significantly minimizing the aforementioned risk, (3) making the privacy protection approach adaptive to enable selection of varying protection levels as deemed necessary in different scenarios. We also discuss the use of the proposed concept in an AAL healthcare use case.

2. RELATED ADVANCEMENTS

This section reviews existing related advancements and approaches (with a particular focus in the AAL domain) from the perspective of users (Sec. 2.1) and technological innovations (Sec. 2.2).

2.1 Users' perspective

While a substantial research has been carried out on sensor-based monitoring in different domains [1, 8, 33, 34], healthcare has un-

ICDSC '16, September 12-15, 2016, Paris, France

DOI: http://dx.doi.org/10.1145/2967413.2967422

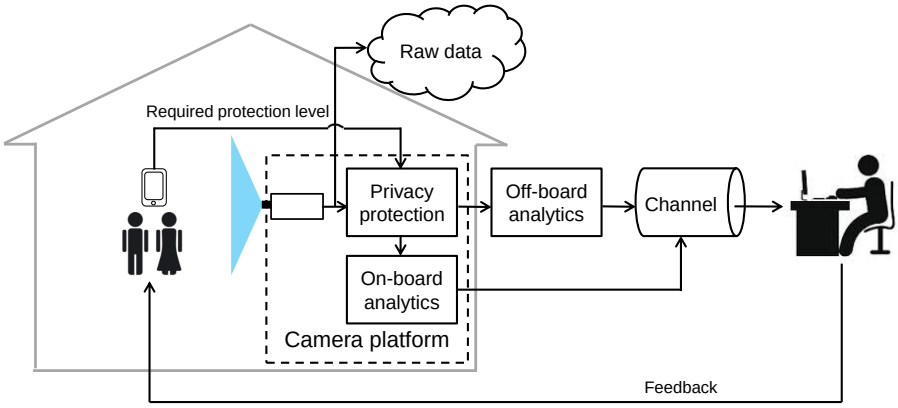

Figure 1: Pipeline for a user-centric trusted human monitoring system.

derstandably also drawn considerable attention with regards to Ambient Assisted Living (AAL). For example, the Activities of Daily Living (ADLs) of patients (with varying health issues) are monitored at their homes using multiple sensing modalities [34]. Prognosis and diagnosis of dementia is performed by monitoring eye movements of patients using cameras installed at their homes [2]. Wearable camera devices are used to monitor dementia patients [20]. An important initiative was made with a focus on applications in different domains including healthcare monitoring in purpose-built smart homes [11]. Other works also exist [4] with applications in patient monitoring with Alzheimer's disease [10], Parkinson's disease [15], sleep problems [18] and mental issues [22].

Existing research focuses primarily on the development and advancement of technology to account for the (monitoring) needs of the remote observers without placing enough focus on the acceptability of technology by the observed individuals by effectively addressing their privacy concerns too. There appears to be a need to adopt an enhanced user-centric approach that could balance the needs of both user types and empower observed individuals with a greater control over the monitored data without limiting the monitoring requirements of remote observers. Moreover, from the remote observers' point of view, the need also remains to provide enabling conditions for an increased pervasive automated monitoring.

2.2 Technological advancements

IoT-based monitoring solutions with different sensing modalities: The use of IoT-based solutions is steadily increasing for different monitoring scenarios particularly in AAL where different sensing modalities are employed such as wearable sensors [10, 15, 16, 20, 21, 34], ambient sensors [11, 18, 21, 34], 3D (Kinect) sensors [30, 34], and cameras [6, 9, 21, 34].

Indeed the use of solely camera-based solutions within an IoT framework is still at a nascent stage [34] and has advantages over other sensing modalities in such applications [8]. An enhancement of IoT-based solutions for healthcare monitoring using smart cameras with advanced onboard protection and analytic capabilities would therefore be desirable.

Privacy protection in videos: Over the last decade, various methods have been developed to protect privacy in visual data. Most techniques operate only on the visual data (i.e., images and videos) to protect sensitive regions and rely on image processing algorithms such as scrambling by JPEG-masking, in-painting, pixelation, blanking, replacement with silhouettes, blurring, warping or morphing (e.g., [25, 33]). The privacy protection capability of these methods strongly depends on the detection performance of sensitive regions and their level of modification. Several approaches have been recently proposed to evaluate the protection techniques and its effect on the utility of the visual data (e.g., [5, 13, 19, 24]. An embedded camera platform was developed [31] with cartooning [14] as an example for on-board privacy protection. It would however still be preferable to push privacy protection closer to the image sensor and initiate protection already in the sensory edge.

Security of image/video data: Security in video surveillance has been discussed for several decades (e.g., [28, 29]). Integrating security functionality onboard 'smart' cameras has been proposed to leverage the security across entire camera networks. More recently, camera platforms with hardware-based security functionality based on trusted platform modules (TPM) [32] or physically unclonable functions (PUF) [12] have been developed to provide a stronger protection than pure software approaches. Moreover, camera platforms have been developed, which provide hardware-based security functionalities such as trusted boot, authenticity, confidentiality and integrity of video data [27, 32]. It would still be desirable to integrate hardware-based security capabilities in embedded and resource-constrained camera platforms. Moreover, it would be interesting to investigate resource availability for providing the actual security functionality.

3. USER-CENTRIC TRUSTED SENSING ENVIRONMENT

An effective IoT solution employing vision sensors (cameras) is not well explored [34] and could greatly aid remote human monitoring. A key challenge in an IoT camera-based monitoring system, however, would be to guarantee and develop an acceptable level of 'trust' of users. Moreover, existing solutions [8, 33] focus mainly in addressing the requirements of remote observers while not placing as much emphasis on the needs of observed individuals. Indeed, an effective consideration of conflicting needs of the two user classes is important for a successful solution.

Fig. 1 presents the pipeline of the proposed concept for a user-centric trusted sensing environment. The proposed concept is aimed to overcome the limitations of current IoT-based monitoring systems by establishing a user-centric trusted sensing environment that integrates the following three aspects: (1) a trusted camera platform that pushes security and privacy protection to the sensory edge by performing the two functionalities onboard of the camera sensors;

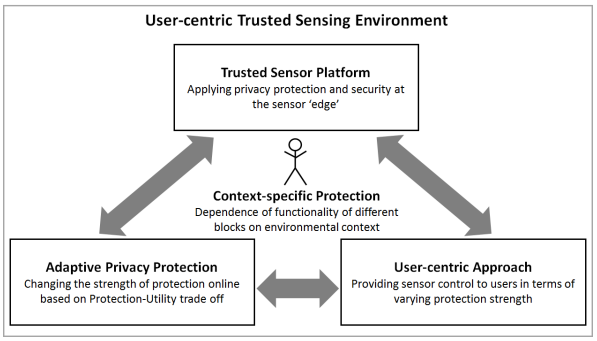

Figure 2: Key components of the proposed user-centric trusted sensing environment.

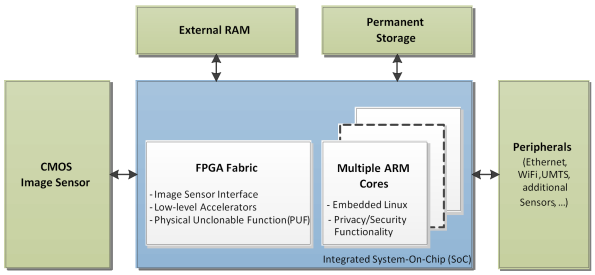

Figure 3: Architecture of the trusted sensor platform.

(2) adaptive privacy protection methods on the trusted platform in order to adjust the level of protection to current needs of users by dynamically exploring the utility-protection tradeoff; (3) the trusted sensing environment delivering trusted feedback to users about the captured data and performed analysis and enables some control on the deployed adaptive protection mechanism. Fig. 2 depicts the proposed approach that comprises of three key components as described next.

3.1 Trusted sensor platform

A fundamental hypothesis of this work is that user-centric trust in a resource-limited sensing environment can be established by making security and privacy protection inherent properties of the sensor platform. The key idea is to protect access to the sensor and encapsulate dedicated security and privacy functionality in a secure sensing unit embedded on the camera platform. The secure sensing unit has exclusive access to the image sensor's raw data. This approach enhances prior work [31] by integrating the camera platform on a hybrid ARM/FPGA System-on-Chip (SoC) and by exploiting dedicated hardware properties of the SOC-platform in the form of physically unclonable functions (PUFs). The key advantage of this approach is to provide hardware-supported security functionality without requiring a dedicated hardware component.

Fig. 3 depicts an overview of the sensor platform architecture. The central component is the hybrid ARM/FPGA SoC (e.g., Xilinx Zynq and UltraScale or Altera SoC) containing multiple general-purpose ARM cores which can execute the camera computer vision applications, management tasks and network communication. The image sensor interface, computer vision accelerators and PUFs form a part of the FPGA fabric. External components include volatile and non-volatile external storage as well as communication interfaces such as Ethernet, WiFi or UMTS. PUFs provide a unique fingerprint of an integrated circuit and serve as basis for security functionality on our platform. Ring-oscillator PUFs are realized on the FPGA fabric for secure key generation which serves as a root for implementing security functionality on the ARM core [17]. In particular the following functions are realized: (i) digital signing, (ii) encryption, (iii) time stamping and secure system boot, and which are exploited to achieve a trusted sensor platform by checking its hardware/software state and securing all data transfer from the platform. Privacy protection functionality is realized on the ARM core as well (cf. the following section).

3.2 Adaptive privacy protection techniques

The two key aspects to consider for evaluating a privacy protection technique are protection and utility [24, 33]. Protection refers to a quantification of the extent of identity information (that would make an object recognizable) hidden in the image data by a privacy protection method. Completely hiding identity information may however not be desirable as there may be a need to preserve structural information for performing some sort of behavioral analysis. Utility, at an elementary level, refers to a quantification of the preservation of the structural information in the image data by a privacy protection technique. However, a more application-specific definition of utility may involve evaluating how well the desired low-level and/or high-level tasks could be performed on the protected data. An ideal privacy protection technique may aim to maximize both protection and utility. However, in practice an appropriate trade off between utility and protection is investigated for a particular scenario. Depending upon the requirements from users, different scenarios could need different levels of protection and utility in an application. Such an approach is expected to enable a more pervasive monitoring even in scenarios with higher protection needs.

3.3 User-centric approach

We propose to adopt a strong user-centric approach that aims to effectively address and balance the needs of both remote observers and observed individuals' users. From the point of view of remote observers, the approach enables them to remotely monitor individuals by means of the provided privacy-protected image data when needed. Additionally, remote observers could provide feedback to the observed individuals via mobile devices. From the point of view of observed individuals, the approach enables them to control sensors in terms of choosing their desired level of protection (deemed necessary for a particular scenario) from a set of allowed protection levels while ensuring a 'reasonable' utility-protection trade off that satisfies the minimum needs of remote observers. The observed individuals can request a desired protection level by means of a simple interface or mechanism. The request can then be served on board using the protection functionality of sensor platform. It is evident that a fundamental aspect of the approach is to meet the needs of different user types.

As depicted in Fig. 1, a complete system pipeline containing the above three components of course includes also other stages spanning the whole lifetime of data transmission. Moreover, as a part of an effective monitoring, some video analytics could also be performed ranging from low-level analysis (person detection and tracking) to high-level analysis (activity recognition) in order to aid remote observers with the identification of different activities of interest. Taking into account the resource constraints, the analytics could be mostly performed on the 'protected' data off board and a part of them could be performed on board thus ensuring an enhanced level of trust from the point of view of patients that are being observed.

Table 1: Summary of the video sequences. Key. NF: number of frames; ADLs: activities of daily living.

Sequence	Area type	NF	Frame size	Activities
S1	Kitchen	1729	2160×3840	ADLs
S2	Kitchen	715	2160×3840	ADLs, person falling
S3	Hallway	579	2160×3840	ADLs
S4	Hallway	325	2160×3840	ADLs, person falling
S5	Sitting	1359	2160×3840	ADLs
S6	Sitting	1308	2160×3840	ADLs, person falling

4. HEALTHCARE USE CASE

The growing focus on AAL research is inevitable due to a fast-ageing population in the world. It is also notable that the associated costs involved in the provision of the on-the-spot assistance and monitoring of the Activities of Daily Living (ADLs) are considered to be unaffordable and unsustainable; hence a need for a remote monitoring using an effective IoT-based solution employing a network of sensors (cameras). The proposed approach in Sec. 3 could therefore provide an effective solution for a AAL-based healthcare use case. In such a case the two user types are as follows: observed individuals are *outpatients* requiring remote care, whereas remote observers are *healthcare professionals*.

4.1 Privacy requirements

From the perspective of healthcare professionals, a remote sensing solution is desired to allow highly pervasive visual monitoring (that may be required for outpatients requiring a constant remote care) to aid with the identification of different outpatients' activities e.g. walking, sitting, sleeping, falling etc. However, such pervasive monitoring may compromise the privacy requirements of outpatients. Indeed, depending on the area in a residence, outpatients could require a different level of protection strength: a much higher level of privacy is expected to be desired in a bedroom than in a kitchen or a living room. Hence, an effective remote monitoring solution in such a use case should address the needs of both user types as ensured in the proposed approach. It essentially boils down to achieve an appropriate balance between protection and utility levels for a scenario under consideration. This is further demonstrated next in the form of a utility-protection analysis on video recordings in real representative residential settings.

4.2 Utility-protection analysis in real residential settings

We used a set of video sequences recorded in different areas (kitchen, hallway, sitting) of a real residential environment. Specifically, there are are six sequences, two in each of the three areas, with people performing normal activities of daily living (ADLs) as well as some unusual ones that could require care/attention of healthcare professionals (in this case such an activity is a fall of a person). For such residential scenarios, it would be desirable that technology employs low-cost affordable cameras. The recordings used a GoPro HERO4 camera (wide-angled lens) with full resolution (2160×3840) from a fixed viewoint. Table 1 provides a summary of the video sequences. Below we present preliminary results of applying different privacy protection techniques and in turn analyzing and evaluating the utility-protection trade off to achieve a varying levels of privacy protection for different scenarios. This could indeed pave a way towards enabling observed individuals (outpatients) with a control to choose their desired protection strength.

We study the effect of applying three privacy protection methods including cartooning, blurring and pixelating on all video sequences. Cartooning involves applying an initial blurring on an input image data followed by mean-shift filtering and edge recovery using the already generated gradient mask with sobel edge detector [14]. The kernel size at the initial blurring stage (A) and the spatial radius (sp) and color radius (sr) at the mean-shift filtering stage are given as follows [14]: $A_i = [i \cdot A_{orig}/50]$; $sp_i = [i \cdot sp_{orig}/50]$; $sr_i = [i \cdot sr_{orig}/50]$; where i is the filter intensity: $i \in [1,50]$ and the parameters $A_{orig} = 7$, $sp_{orig} = 20$ and $sr_{orig} = 40$ [14]. Additionally, as done in [14], for establishing some correspondence and a fair comparison among different techniques the kernel size used in the case of blurring and pixelating for a particular filter intensity, i, is equal to sp_i as defined above for cartooning. We apply each privacy protection technique globally on full frames in every sequence for a full variation of filter intensity, i. To account for on-board resource, computational and bandwidth limitations, image data could desirably be processed in a downsampled form [23]. We therefore apply privacy protection techniques on the frames in a downsampled form (10% of the original resolution) to analyze different techniques under an extreme resolution setting. As for a quantitative comparison of privacy protection techniques, we employ a recent method [24] that evaluates the aspects of protection and utility to enable an analysis of utility-protection trade off. Given a sequence, at each frame a protection score is computed as an appearance dissimilarity between original and privacy-protected frames. Protection scores are then averaged across the sequence to provide sequence-level protection score; the lower the score the lower the protection. Likewise, a utility score is computed as a structural similarity between original and privacy-protected frames. Utility scores are then averaged across the sequence to provide sequence-level utility score; the lower the score the lower the utility. Note that for an initial analysis we use here a traditional way of computing protection and utility that involves encapsulating fidelity of appearance and structure. A more application-specific way could involve evaluating algorithmic performance on the protected data as a means of assessing protection and utility. Fig. 4 plots utility (U) *vs.* protection (P) scores of the three privacy protection techniques for a variation of i on all sequences. Fig. 5 shows some sample qualitative results for the privacy protection techniques.

The utility-protection plots (Fig. 4) could help to fulfill the needs of both the outpatients and healthcare professionals. From an outpatient's point of view, the protection axis (x-axis) of plots is of more importance, which allows switching among varying levels of protection and depending on the requirement of an outpatient an appropriate protection strength could be selected. From a point of view of healthcare professionals, the utility axis (y-axis) holds more importance that enables choosing among varying levels of utility. Indeed, depending on a scenario, the system could give priority to either the level of utility sought by healthcare professionsals or the level of protection desired by outpatients. For example, in a situation where an a constant care of an outpatient is necessary, a higher utility (e.g. $U = 0.85$) might be desirable and cartooning would therefore be a preferred option as it maximizes protection from outpatient's perspective (Fig. 4). Alternatively, in a situation that demands a high privacy requirement from outpatient (e.g. $P = 0.70$) and a compromise in utility is permissible, pixelating would be the desired option. Therefore, such control over the utility and protection from both user types instills a strong notion of user centricity into the proposed solution.

5. DISCUSSION AND FUTURE WORK

We put forth a concept for a highly trusted, user-centric and pervasive human monitoring approach that pushes the privacy protection and security closer to the sensing platform for an enhanced

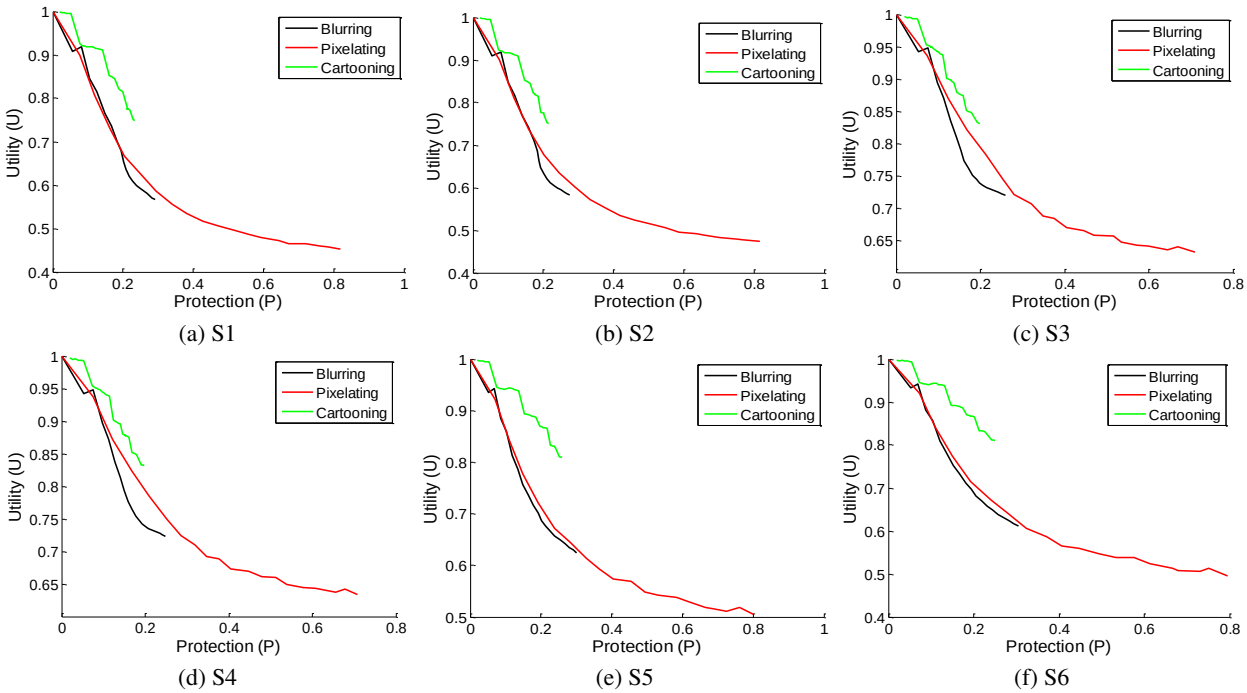

Figure 4: Utility score (U) plotted *vs.* protection score (P) obtained by different privacy protection techniques for a variation of filter intensity on all sequences.

level of trust; accommodates monitoring requirements of remote observers and privacy needs of observed individuals; instills user-controlled adaptability at the sensor platform to allow varying levels of privacy protection using the Utility-Protection trade off. As a potential application of the proposed concept we discussed its effectiveness in an AAL healthcare use case. We used video sequences in a real residential environment to perform an initial utility-protection analysis that forms a basis towards incorporating a notion of user centricity into the proposed concept. The sequences cover activities of daily living (ADLs) as well as some unusual activities (person falls) in key residential areas: kitchen, hallway, sitting. We performed an analysis of varying levels of utility and protection provided by different privacy protection techniques on all sequences that could facilitate in balancing the conflicting needs of two user types: outpatients and healthcare professionals.

We believe this white paper could lay a foundation for several future directions in order to fully realize the proposed concept of a user-centric trusted environment. Firstly, there is a room for conducting a thorough study to investigate effective methods of applying privacy protection and security of image data towards on-board camera processing. Secondly, the needs remains to effectively instil user-controlled adaptability at the sensor platform that allows varying levels of privacy protection using the 'Utility-Protection trade off' analysis. Thirdly, it would be interesting to investigate trade-off between on-board and off-board data analytics. Finally, in order to ensure the acceptability and application of the proposed concept, it would be important to investigate relevant societal, ethical and legal aspects.

6. ACKNOWLEDGMENTS

This research has received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement no. 312784, and from the Austrian Research Promotion Agency (FFG) under grant number 842432.

7. REFERENCES

[1] Fastpass project. https://www.fastpass-project.eu/. Accessed March 2016.

[2] Modem project. http://gow.epsrc.ac.uk/ngboviewgrant.aspx?grantref=ep/m006255/1. Accessed March 2016.

[3] P5 project. http://www.foi.se/en/customer–partners/projects/p5/p51/. Accessed March 2016.

[4] G. Acampora, D. J. Cook, P. Rashidi, and A. V. Vasilkos. A survey on ambient intelligence in healthcare. *Proceedings of the IEEE*, 101(12), 2013.

[5] A. Badii, A. Al-Obaidi, M. Einig, and A. Ducournau. Holistic privacy impact assessment framework for video privacy filtering technologies. *Sig. & Ima. Proc.*, 4(6):13–32, 2013.

[6] N. B. Bo, F. Deboeverie, M. Eldib, J. Guan, X. Xie, J. Nino, D. V. Haerenborgh, M. Slembrouck, S. V. de Velde, H. Steendam, P. Veelaert, R. Kleihorst, H. Aghajan, and W. Philips. Human Mobility Monitoring in Very Low Resolution Visual Sensor Network. *Sensors*, 14:20800–20824, 2015.

[7] A. Cavoukian. Surveillance, then and now: Securing privacy in public spaces. Technical report, 2013.

[8] A. A. Chaaraoui, P. Climent-Perez, and F. Florez-Revuelta. A review on vision techniques applied to human behaviour analysis for ambient assisted living. *Exp. Sys. Appl.*, 39(12):10873–10888, 2012.

[9] A. A. Chaaraoui, J. R. Padilla-Lopez, F. J. Ferrandez-Pastor, M. Nieto-Hidalgo, and F. Florez-Revuelta. A vision-based

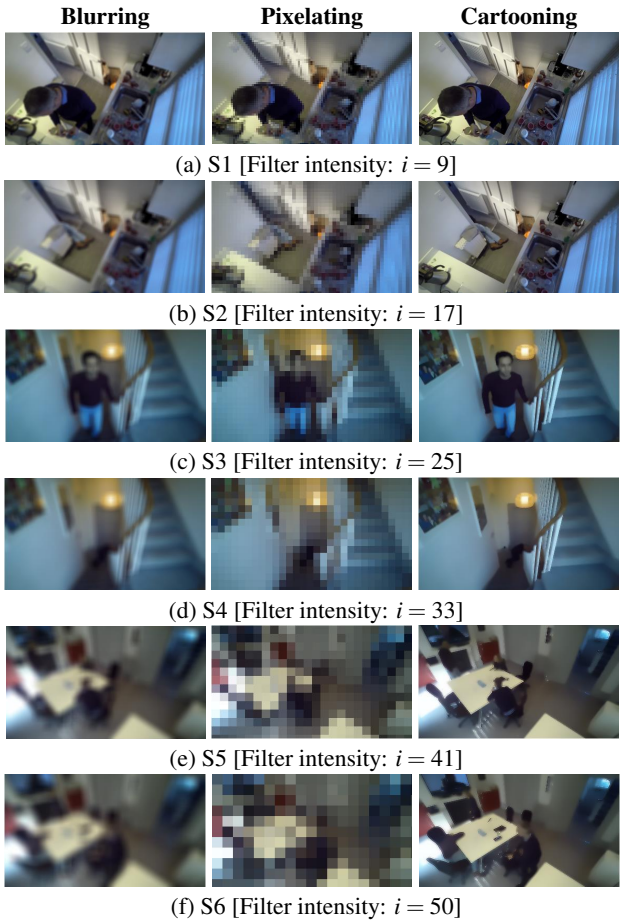

(a) S1 [Filter intensity: $i = 9$]

(b) S2 [Filter intensity: $i = 17$]

(c) S3 [Filter intensity: $i = 25$]

(d) S4 [Filter intensity: $i = 33$]

(e) S5 [Filter intensity: $i = 41$]

(f) S6 [Filter intensity: $i = 50$]

Figure 5: Sample qualitative results for different privacy protection techniques on all sequences with an increasing filter intensity ($i = 9, 17, 25, 33, 41, 50$). Column 1: blurring; column 2: pixelating; column 3: cartooning.

system for intelligent monitoring: Human behaviour analysis and privacy by context. *Sensors*, 14:8895–8925, 2015.

[10] Y. Charlon, N. Fourty, W. Bourennane, and E. Campo. Design and evaluation of a device worn for fall detection and localization: Application for the continuous monitoring of risks incurred by dependents in an alzheimer's care unit. *Exp. Sys. Appl.*, 40(18):7316–7330, 2013.

[11] D. Cook and M. Schmitter-Edgecombe. Assessing the quality of activities in a smart environment. *MIM*, 48(5):480–485, 2009.

[12] J. Delvaux, R. Peeters, D. Gu, and I. Verbauwhede. A survey on lightweight entity authentication with strong pufs. *ACM Comp. Surv.*, 48(2), 2015.

[13] F. Dufaux and T. Ebrahimi. A framework for the validation of privacy protection solutions in video surveillance. In *Proc. of ICME*, 2010.

[14] A. Erdelyi, T. Barat, P. Valet, T. Winkler, and B. Rinner. Adaptive cartooning for privacy protection in camera networks. In *Proc. of AVSS*, Seoul, August 2014.

[15] D. Giansanti, V. Macellari, and G. Maccioni. Telemonitoring and telerehabilitation of patients with parkinson's disease: health technology assessment of a novel wearable step counter. *Telemed. e-health*, 14(1):76–83, 2008.

[16] P. Gupta and T. Dallas. Feature selection and activity recognition system using a single triaxial accelerometer. *IEEE TBME*, 61(6):1780–1786, 2014.

[17] M. Hoeberl, I. Haider, and B. Rinner. Towards a secure key generation and storage framework on resource-constrained sensor nodes. In *Proc. of EWSN Work.*, Graz, 2016.

[18] A. Kealy, K. McDaid, J. Loane, L. Walsh, and J. Doyle. Derivation of night time behaviour metrics using ambient sensors. In *Proc. of PervasiveHealth*, Venice, 2013.

[19] P. Korshunov, C. Araimo, F. D. Simone, C. Velardo, J.-L. Dugelay, and T. Ebrahimi. Subjective study of privacy filters in video surveillance. In *Proc. of IEEE Work. MMSP*, 2012.

[20] R. Megret, V. Dovgalecs, H. Wannous, S. Karaman, J. Benois-Pineau, E. E. Khoury, J. Pinquier, P. Joly, R. Andre-Obrecht, Y. Gaestel, and J.-F. Dartigues. The immed project: wearable video monitoring of people with age dementia. In *Proc. of ACM MM*, Firenze, 2010.

[21] M. Mubashir, L. Shao, and L. Seed. A survey on fall detection: Principles and approaches. *Neurocomp.*, 100(2013):144–152, 2013.

[22] M. Nambu, K. Nakajima, M. Noshiro, and T. Tamura. An algorithm for the automatic detection of health conditions. *IEEE EMBM*, 24(4):38–42, 2005.

[23] T. Nawaz and A. Cavallaro. A protocol for evaluating video trackers under real-world conditions. *IEEE TIP*, 22(4):1354–1361, 2013.

[24] T. Nawaz and J. Ferryman. An annotation-free method for evaluating privacy protection techniques in videos. In *Proc. of AVSS*, Karlsruhe, 2015.

[25] J. R. Padilla-Lopez, A. A. Chaaraoui, and F. Florez-Revuelta. Visual privacy protection methods: A survey. *Exp. Sys. Appl.*, 42(9):4177–4195, 2015.

[26] F. Pittaluga and S. J. Koppal. Privacy preserving optics for miniature vision sensors. In *Proc. of CVPR*, Boston, MA, June 2015.

[27] B. Rinner and T. Winkler. Privacy-protecting smart cameras. In *Proc. of ICDSC*, Venice, Italy, 2014.

[28] A. Senior, S. Pankanti, A. Hampapur, L. Brown, Y.-L. Tian, A. Ekin, J. Connell, C. F. Shu, and M. Lu. Enabling video privacy through computer vision. *IEEE Sec. Priv.*, 3(3):50–57, 2005.

[29] D. N. Serpanos and A. Papalambrou. Security and privacy in distributed smart cameras. *Proceedings of the IEEE*, 96(10):1678–1687, 2008.

[30] R. Wang, G. Medioni, C. J. Winstein, and C. Blanco. Home monitoring musculo-skeletal disorders with a single 3d sensor. In *Proc. of CVPR Work.*, 2013.

[31] T. Winkler, A. Erdelyi, and B. Rinner. Trusteye.m4: Protecting the sensor - not the camera. In *Proc. of AVSS*, Seoul, August 2014.

[32] T. Winkler and B. Rinner. Securing embedded smart cameras with trusted computing. In *EURASIP JWCN*, 2011.

[33] T. Winkler and B. Rinner. Security and privacy protection in visual sensor networks: A survey. *ACM Comp. Surv.*, 47(1), 2014.

[34] N. Zhu, T. Diethe, M. Camplani, L. Tao, A. Burrows, N. Twomey, D. Kaleshi, M. Mirmehdi, P. Flach, and I. Craddock. Bridging e-health and the internet of things: The sphere project. *IEEE Intel. Sys.*, 30(4):39–46, 2015.