

Deakin Research Online

This is the published version:

Moncrieff, Simon, Venkatesh, Svetha and West, Geoff 2007, Dynamic privacy in a smart house environment, in *ICME 2007 : Proceedings of the 2007 IEEE International Conference on Multimedia and Expo*, IEEE, [Beijing, China], pp. 2034-2037.

Available from Deakin Research Online:

<http://hdl.handle.net/10536/DRO/DU:30044593>

Reproduced with the kind permissions of the copyright owner.

Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Copyright : 2007, IEEE

Dynamic Privacy in a Smart House Environment

Simon Moncrieff, Svetha Venkatesh and Geoff West
Department of Computing Curtin University of Technology
GPO Box U1987, Perth, 6845, W. Australia

Abstract—A smart house can be regarded as a surveillance environment in which the person being observed carries out activities that range from intimate to more public. What can be observed depends on the activity, the person observing (e.g. a carer) and policy. In assisted living smart house environments, a single privacy policy, applied throughout, would be either too invasive for an occupant, or too restrictive for an observer, due to the conflicting goals of surveillance and private environments. Hence, we propose a *dynamic* method for altering the level of privacy in the environment based on the *context*, the situation within the environment, encompassing factors relevant to ensuring the occupant's safety and privacy. The context is mapped to an appropriate level of privacy, which is implemented by controlling access to data sources (e.g. video) using data hiding techniques. The aim of this work is to decrease the invasiveness of the technology, while retaining the purpose of the system.

I. INTRODUCTION

Smart houses seek to enhance a person's environment and way of life. Assisted living, a form of smart house, seeks to monitor a home to ensure the occupant's safety, enabling the aged and invalid population to remain in their homes for longer. However, surveillance applications in such environments require privacy measures if the technologies are to be accepted by the occupants. This is due to the private nature of the home, and the invasive nature of surveillance. In this paper, we introduce a method for dynamically applying privacy measures to the monitoring of a smart home based on the situation within the home. For example, what video data should a carer be allowed to access at a given moment.

There is a large corpus of work examining design strategies for privacy in ubiquitous and pervasive computing [1], [2]. Rather than attempting the difficult task of defining privacy, from examining previous work we identify four key properties required for the *design* of privacy sensitive applications. A dynamic aspect (dependent on situation), and flexibility (able to accommodate different preferences or perceptions) are required for the implementation of the privacy management within the system [1]. Feedback and control mechanisms are required for communication between the privacy system and the occupant [2], and have been identified as integral to the user's acceptance of such systems. While we have developed a privacy sensitive smart house that addresses all four properties, in this paper we focus on the dynamic aspect.

A dynamic approach is required to address the conflict between the needs, or wants, of the occupant, and the ability of the system to achieve its purpose. For the occupant, intrusions into their life should be minimal. However, to perform the function of the system, an observer, e.g. a carer, should be

given sufficient information to carry out their duties, e.g. ensure the occupant's safety. This conflict between privacy and purpose can be highlighted as follows; a situation in which a carer has access to all data would enable them to perform their duty, but, given that a private space is under observation, would likely be unacceptable to the occupant. Consider the example of the bathroom, a hazardous environment that requires detailed monitoring, but also requires a high degree of privacy. Thus, access to data cannot be statically determined across situations. An approach that can dynamically control access to data is required, with access in the worst case scenario differing to access when the situation is considered *normal*.

We present a framework for a privacy sensitive smart house that dynamically determines what data an observer can access given the situation (**context**). The dynamic aspect is primarily achieved using two concepts. Firstly, the *context space*, which encapsulates the context of the environment, encompassing aspects of social interaction, spatial location, and the activities of the occupant. We quantify the context using the multi-modal analysis of audio, and binary sensor data. Secondly, the *privacy space* is a n dimensional representation of the privacy levels for each information source (dimension in the privacy space). In this paper we consider a 3D privacy space, consisting of video, audio, and binary sensor data (e.g. reed switches). The privacy levels are implemented using data hiding techniques that obscure access to the data at different resolutions. For example, no information is shown at the highest level of privacy, and all data is accessible at the lowest privacy level. We term points in the privacy space the *data access level*, which has an inverse relationship to the privacy level, i.e. a higher data access level has a lower level of privacy. The multi-resolution and multi-modal properties of the privacy space enable the determination of dynamic privacy policies.

The novelty of this work lies in the use of the environmental context to dynamically alter the applied level of privacy within a smart house. Privacy in smart house environments is becoming an increasingly salient issue, as ubiquitous technologies and analysis methods evolve along with the growing need for assisted living. Consequently, the significance of this work lies in actively addressing privacy issues in such environments.

II. PRIVACY IN UBIQUITOUS COMPUTING ENVIRONMENTS

Previous approaches to addressing privacy in assisted living applications have limited monitoring to sensor agents [3], or limited to sensor type, e.g. by using anonymous binary sensors [4]. Consequently, to examine applications that actively address privacy issues, we focus on previous work

in the wider field of ubiquitous computing and surveillance. From examining the elements of such applications, we have developed a generic framework (Figure 1) for implementing privacy sensitive environments. The framework consists of three components; while the majority of previous approaches do not contain all components, they can be considered to be a subset of the framework. The first component consists of determining the context associated with a property of the environment, e.g. using a combination of motion sensors (detect the presence of a person), and an RFID tag (identify the person), to determine access privileges [5], or the presence of multiple people [2]. The context component is used to determine if privacy measures are required or not. The second component, the rule set, uses a rule based approach to determine the privacy level. The rule set interprets the context data in a binary fashion, determining if privacy measures are required, e.g. privacy measures are required if a person is authorised to be in a particular area [5]. The preferences are then used to determine the *level* of privacy according to pre-determined criteria, e.g. an observer's authorisation level [6], a user's predefined preferences [7], or a user's requested privacy level [5]. The privacy level is passed to the data filter component, which implements privacy using data hiding techniques to filter the data at either two levels, e.g. camera on or off [2], or at multiple resolutions, e.g. obscuring video data to mask the occupants identity at different levels, using methods such as bounding boxes, and shadows [5].




Fig. 1. A generic framework for developing privacy sensitive applications.

This generic framework has a number of advantages; incorporating context and multi-resolution data filtering introduce a dynamic aspect. However, the dynamic nature is limited to determining if privacy is implemented or not, as determined by the context, while predefined, static preferences are used to determine the privacy filtering level. This limits the scalability of the context with respect to multiple situations.

III. SYSTEM DESIGN

We propose to dynamically adjust the *privacy level* according to the context. To achieve this, the generic framework was used to develop a specific framework for a privacy sensitive smart house environment, shown in Figure 2. In addition, this framework introduces two key ideas for expanding previous approaches, the context space, and the privacy space.




Fig. 2. The framework for developing privacy sensitive applications for the smart house environment (compare to Figure 1).

A. Context Space

In comparison with previous approaches, we adjust the data access level according to context. Thus the aim of the context is to influence the *level* of privacy filtering. Contextual data is data that is semantically linked to situation. The context space incorporates multiple aspects of the environment (contexts) within a single multi-dimensional construct. Each context type can be discrete or continuous, and, in combination, account for numerous situation types. We identify four types of context that are relevant to ensuring the safety of the occupant and to reducing the invasion of the occupant's privacy.

1) *Spatial Context*: The spatial context (*where*) indicates the location of the occupant within the smart house, which is divided into public areas (e.g. kitchen, lounge), and private areas (e.g. bedroom, bathroom). The spatial context is determined as private (0), public (1), or not present (-1) depending upon the location of the occupant. The purpose of the social context is, firstly, to distinguish between private and public areas, as privacy issues arise depending upon the activities associated with different rooms, e.g. due to the privacy sensitive activities that occur in the bathroom (private), a higher level of privacy is required. Secondly, the spatial context and the purpose of the locations, combined with other contexts, can be used to identify potential risks.

2) *Social Context*: The social context (*who*) determines the presence of social interaction through the detection of conversation¹. Conversation is determined by initially detecting speech audio clips (0.25s) using a decision tree trained to distinguish between speech, noise, and silence [8]. Then, for each 1s window of audio (ω), the number of speech clips in a large window (120s, centred on ω) is determined. If the proportion of clips exceeds the threshold σ (0.3), ω is labelled conversation. The purpose of the social context is: (1) to allow/deny access to the conversation depending on the observer, e.g. a carer should not be given details of the content of the conversation and (2) to influence the level of monitoring as (a) multiple people suggests the occupant is already being monitored, reducing the need for a high level of monitoring, and (b) there is a potential risk associated with the presence of other people (e.g. strangers), so a degree of monitoring remains necessary.

3) *Hazard Context*: The hazard context (*how* and *what*) is used to increase awareness in situations associated with the use of hazardous devices, a device that has to be attended to while active, or on, e.g. a stove or fridge. The hazard context consists of two components. The first indicates whether or not a hazardous device is active, determined by the binary sensor associated with the device. This component is discrete, with a value of 0 or 1. The second, termed *anxiety* [9], is a probabilistic framework that learns the temporal patterns of interaction with a hazardous device, and associated devices in the environment, to detect abnormal interactions when the hazardous device is active. The anxiety is continuous, and

¹Not associated with a telephone, determined using a sensor to detect the activation of the telephone

has a range of 0 to 1. If the occupant interacts with the hazardous device, or responds positively to a query regarding their status, the anxiety is set to 0. Otherwise, the anxiety increases as a greater deviation from a normal interaction is observed. An anxiety of 1, an abnormal interaction, indicates a potential danger to the occupant; either an accident has occurred (preventing an interaction with the hazardous device), or the occupant has forgotten the hazardous device is active. The latter represents a danger as such devices need to be attended, e.g. if a stove is left on, it becomes a fire hazard.

4) *Activity Context*: The activity context (*what*) is used to raise awareness if the occupant becomes unusually inactive. The activity context is determined using the *anxiety* to learn typical patterns of interaction with the environment, signalled by binary sensors and the audio sensor (audible interactions with the environment indicating general activity [9]). We treat the audio sensor as a hazardous device, with an absence of audio activity resulting in an increase in the activity context value. A lower value indicates a lesser need for monitoring as the occupant is generally active. However, a high value indicates an unusual lack of activity (a high expectation that the occupant should be active). An observer needs to be given sufficient data access to distinguish between inactivity due to a passive activity (sleeping or reading), or due to the occupant being incapacitated due to injury.

B. Rule Set

The rule component is used to map a point in the context space to a point in the privacy space. The mapping represents the flexible component of the framework due to the incorporation of preferences (purpose). The same context can be mapped to different points in the privacy space depending upon the observer's purpose, e.g. doctors, family, and carers will all have different levels of data access depending upon the context. The mapping between the context and privacy spaces equates to a mapping between semantic concepts. Currently, a decision tree is used to generate the rules used to perform the mapping, trained using the context calculated for data captured within a smart house and labelled with an appropriate privacy policy. A detailed discussion of the risk analysis used to determine the labelling is beyond the scope of this paper. However, an example mapping, for a carer observer and a single occupant, will be used to examine the dynamic properties of the framework.

C. Data Filter and Privacy Space

The rule set defines the data access level for each information source, the data filter component then implements the appropriate data hiding technique to represent a point in the privacy space. The privacy space consists of an axis for each monitoring modality present in the environment, in this case three axes are present, one for each modality; audio, video, and binary sensors. The privacy space is discrete, with the co-ordinates of each axis representing a data hiding technique. Hence, a point in the privacy space defines the data access

level for each data type. This enables a multi-modal, multi-resolution data filter to be described using a single point. While the filtering method used for corresponding data access levels differs between data sources, the level of detail is equivalent. Five levels of filtering are used, increasing in detail from 0 to 4. Access to past data is introduced at levels 3, filtered at the appropriate level, providing information on events leading to the current state of the environment. The access levels are:

- 0– All data is blocked.
- 1– *Audio*: Background and foreground (used to indicating activity) sounds are displayed as labels [10]. *Video*: The (x, y) co-ordinates of the occupant are shown on a plan of the smart house. *Sensor*: The room in which the occupant is interacting with sensors, and rooms with active hazardous devices.
- 2– *Audio*: The background audio signal, with labels representing environmental foreground audio and speech. *Video*: A bounding box is imposed on the video, replacing the image of the occupant. *Sensor*: The device last interacted with, and the active hazardous devices.
- 3– *Audio*: The background and foreground environmental audio, with a label representing speech. *Video*: The image of the occupant is replaced by a shadow image [5], revealing information on the posture of the occupant, and interactions with the environment. *Sensor*: The device interacted with, and the form of the interaction, and the active hazardous devices.
- 4– All data is presented to the observer.

IV. EXPERIMENTATION

The proposed privacy system was investigated by recording a number of everyday scenarios within a smart house environment simulated within a lab, consisting of two rooms, a kitchen/lounge area (*public*) and a bedroom (*private*). Devices (e.g. stove, fridge) were placed in the environment to simulate the real world functionality of the room; binary sensors, including pressure mats, were used to detect device interaction. The environment is further augmented with video cameras (10), and a microphone in each room. Audio was captured at $44.1kHz$, $16bit$, in wav format. The sensor logs, and video and audio streams were captured, stored, and processed offline.

A number of scenarios were captured in three recording sessions, which were split into two sets, training data (two sessions, $95.2mins$ and $104mins$), and testing data ($95.3mins$). Each session consisted of numerous scenarios. These included the extremes of each context, e.g. anxiety values ranging from 0 to 1. Combinations of scenarios were similarly present, e.g. changing spatial location while increasing anxiety. The continuous capture of data across scenarios resulted in periods of transitional context, which combined with the scenarios to produce a large range of context within the data.

Each element of the context was then calculated (Section III-A), at a resolution of $1s$, for each data sequence, to generate a feature vector of the form (*Hazard2 (Anxiety)*, *Hazard1*, *Social*, *Activity*, *Spatial*). A ground truth data access level (point in the privacy space) was then assigned to sections within each session according to the context type of the

section (according to the mapping described in Section III-B). A privacy space point was then assigned to each second of data according to the label associated with the section encompassing the data. This process resulted in a context point, and corresponding privacy point for each second of data.

The data was then used to train and test a decision tree classifier (*J48* [11]) to explore the efficacy of dynamic privacy. The attribute data used for training consisted of the context data. The corresponding class data was determined by considering each point of the privacy space to be a class type, with 14 privacy space classes identified. A number of experiments were performed to test the extent to which ambiguous context data is present (1), determine the efficacy of the rules for mapping between the context and privacy space for unseen data (2), and determine general performance (3).

The classification results for the experiments (Table I) show a strong match between the ground truth privacy label and the privacy label determined by the context. Figure 3 shows context and corresponding data access levels for audio and video for a segment of training data set 1 encompassing a hazardous event (cooking, from point A to B), and a period of inactivity. The spatial context was 1 (public space), and no social interaction was present. The figure shows the dynamic nature of the context and the corresponding dynamic characteristics of the privacy measures (data access levels) with respect to the changing context, and the variation present between the data access levels for different data sources. The *classed* plots display the classified data access levels superimposed on the *ground truth* plot, thus, visible sections of the *classed* plots represents deviation from the ground truth. Such deviations are generally short pulses or delays at boundaries between changing access levels due to smoothing.

TABLE I

CONTEXT TO PRIVACY MAPPING RESULTS.

Experiment	Training Data	Testing Data	Result
1	<i>Train 1 and 2</i>	<i>Train 1 and 2</i>	95.36%
		<i>Train 1</i>	95.27%
		<i>Train 2</i>	95.44%
2	<i>Train 1 and 2</i>	<i>Test 1</i>	91.90%
3	<i>Train 1 and 2</i>	10 fold split	94.60%

V. CONCLUSION

As both the need for assisted living applications and the sophistication of sensor technology increases, privacy concerns need to be addressed for the active monitoring of smart homes. Privacy management is important if the monitoring is to be accepted by the occupant. In this paper we have detailed a framework for dynamically altering the privacy level in a smart house, based on the environmental context. The dynamic aspect of the framework is required for the surveillance of private/home environments to decrease the intrusive nature of the technology, while maintaining the functionality of the framework. The context comprises several quantifiable aspects used to identify risks to the occupant's safety, and privacy sensitive activities. The privacy measures are implemented using data hiding techniques for each modality present, to

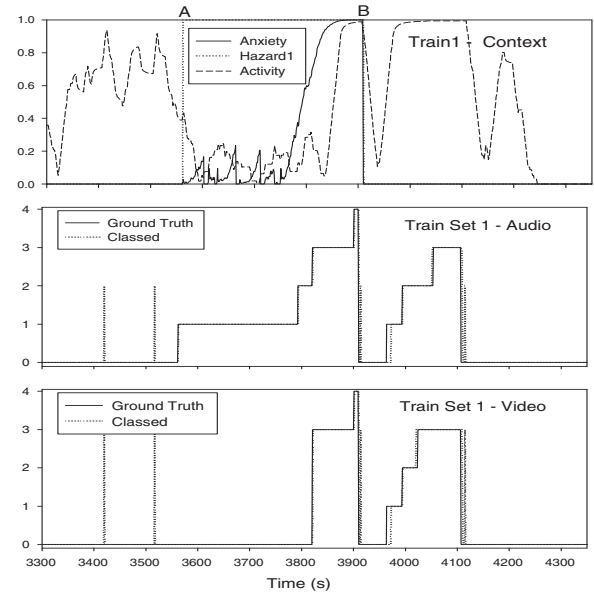


Fig. 3. A segment of the Train 1 data set depicting a hazardous event, displaying corresponding context space elements, and privacy filtering levels for audio and video.

obscure data at multiple resolutions. In using the context, the method is capable of dynamically altering the data access level for a given observer.

ACKNOWLEDGMENT

ARC Discovery Grant: Homes that Sense and Support (DP 0449437)

REFERENCES

- [1] L. Palen and P. Dourish, "Unpacking "privacy" for a networked world," in *SIGCHI Conference on Human Factors in Computing Systems (CHI03)*, Ft. Lauderdale, Florida, USA, April 2003, pp. 129–136.
- [2] C. Neustaedter and S. Greenberg, "The design of a context-aware home media space for balancing privacy and awareness," in *5th International Conference on Ubiquitous Computing*, 2003, pp. 297–314.
- [3] S. Das and D. J. Cook, "Health monitoring in an agent-based smart home," in *Proceedings of the International Conference on Smart Homes and Health Telematics (ICOST)*, Singapore, September 2004.
- [4] D. H. Wilson, "Assistive intelligent environments for automatic health monitoring," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, September 2005.
- [5] J. Wickramasuriya, M. Alhazzazi, M. Datt, S. Mehrotra, and N. Venkatasubramanian, "Privacy-protecting data collection in media spaces," in *ACM Multimedia 2004*, New York, NY, October 2004.
- [6] A. Senior, S. Pankanti, A. Hampapur, L. Brown, Y.-L. Tian, and A. Ekin, "Enabling video privacy through computer vision," *IEEE Security and Privacy*, vol. 3, no. 3, pp. 50–57, 2005.
- [7] D. A. Fidaleo, H. Nguyen, and M. Trivedi, "The networked sensor tapestry (nest): a privacy enhanced software architecture for interactive analysis of data in video-sensor networks," in *ACM 2nd international Workshop on Video Surveillance and Sensor Networks, VSSN '04*. ACM Press, New York, NY, October 2004, pp. 46–53.
- [8] S. Moncrieff and S. Venkatesh, "Narrative structure detection through audio pace," in *IEEE Multimedia Modeling 2006*, Beijing, China, 4–6 Jan 2006, pp. 20–27.
- [9] S. Moncrieff, S. Venkatesh, G. West, and S. Greenhill, "Incorporating contextual audio for an actively anxious smart house," in *Int'l Conf. Intelligent Sensors, Sensor Networks and Information Processing*, Melbourne, Australia, December 2005, pp. 373–378.
- [10] S. Moncrieff, S. Venkatesh and G. West, "Persistent audio modelling for background determination," in *IEEE International Conference on Multimedia and Expo (ICME 2005)*, Amsterdam, Netherlands, 2005.
- [11] I. H. Witten and E. Frank, *Data Mining: Practical machine learning tools with Java implementations*. Morgan Kaufmann, 2000.