Collaborative Explanation and Response in Assisted Living Environments Enhanced with Humanoid Robots

Antonis Bikakis[1], Patrice Caire[2], Keith Clark[3], Gary Cornelius[2], Jiefei Ma[3], Rob Miller[1], Alessandra Russo[3] and Holger Voos[2,*]

[1] *Department of Information Studies, University College London, London WC1E 6BT, U.K.*

[2] *Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg, L-1359 Luxembourg, Luxembourg*

[3] *Department of Computing, Imperial College London, London SW7 2RH, U.K.*

{a.bikakis, r.s.miller}@ucl.ac.uk, {patrice.caire, holger.voos}@uni.lu, garycorn90@gmail.com, {k.clark, j.ma, a.russo}@imperial.ac.uk

Keywords: Assisted Living, Distributed Reasoning, Evidence Gathering, Explanation Generation, Service Robots.

Abstract: An ageing population with increased social care needs has provided recent impetus for research into assisted living technologies, as the need for different approaches to providing supportive environments for senior citizens becomes paramount. Ambient intelligence (AmI) systems are already contributing to this endeavour. A key feature of future AmI systems will be the ability to identify causes and explanations for changes to the environment, in order to react appropriately. We identify some of the challenges that arise in this respect, and argue that an iterative and distributed approach to explanation generation is required, interleaved with directed data gathering. We further argue that this can be realised by developing and combining state-of-the art techniques in automated distributed reasoning, activity recognition, robotics, and knowledge-based control.

1 INTRODUCTION

Electronic health services give opportunities for providing better care, particularly for the elderly. But their development gives rise to a number of research challenges, for example in terms of privacy, user-friendliness, conviviality and security. A prominent area of research and development has been in Ambient Intelligence (AmI) systems, in particular for assisted living environments. The long term vision of the AmI research community is to provide systems that intelligently and unobtrusively assist human inhabitants with tasks in their everyday environment. This environment is dynamic and complex, and in order to operate effectively, an AmI system must have some ability to identify and explain the events occurring within it, sometimes in the face of incomplete, uncertain or seemingly inconsistent information.

Humans have a number of key abilities for coping with events in their environment. One is the ability to quickly filter out events of interest or that need a response from the normal environmental background. Another is the ability to identify possible causes or explanations for such events, and then act appropriately to eliminate or confirm them. This often involves taking action to obtain further information. In this paper we propose an approach to AmI systems that mirrors these abilities, using state-of-the art techniques in automated distributed reasoning, activity recognition, robotics, and knowledge-based control. Our particular focus is on directed data gathering, triggered by the generation of tentative explanations. Robots are a particularly useful tool in this respect, because of their multi-functional and mobile capabilities.

The research questions we wish to address are: (1) how to identify events of interest, e.g. ones that could lead to emergency situations, (2) how to determine the causes of these events, and (3) how to then decide what to do. Although traditionally these problems are handled separately and sequentially, we propose an iterative process which interleaves these activities, using a knowledge-based approach underpinned by a distributed, abductive inference engine.

The remainder of the paper is structured as follows. Section 2 illustrates the challenges in addressing the above questions with an example scenario. Section 3 summarises the current state of the art, Section 4 presents our proposed methodology, and Section 5 concludes and outlines our next steps.

[*] Authors are listed alphabetically.

2 OUR VISION

2.1 Running Example

Eighty-year-old Wally lives alone, with health declining. An intelligent, agent-based home care system has been installed to help him, including an alarm connected to a local call centre that Wally can trigger himself while at home, and two humanoid Robotic Care Assistants (RoCAs). These small, gentle-looking mobile humanoid robots rest at charging stations unless tasked by the system. One day at 11am the automated Home Care Agent (HCA) receives an input indicating that Wally has pressed his alarm button. The HCA queries the help centre to ask if the operators there have managed to communicate with Wally via his home or mobile phones. They have not, and so are waiting the pre-agreed 10 minutes for the system to generate explanations on which to act.

The HCA therefore tries to generate explanations for the alarm. It asks the two RoCAs (one upstairs, one downstairs) to try to locate both Wally and the alarm button. The latter is painted with a distinctive pattern that the robots can easily recognise. At the same time the HCA tries to locate Wally's mobile phone, which is integrated into the system as an agent, using GPS. It locates this 0.5km away, moving up Fairtree Road at 30km/hr. The mobile's log does not show any recent outgoing calls, and confirms that an unanswered call was made from the call center. It also indicates that the mobile has been set to "silent".

The HCA's profile of Wally contains the following information: (i) Wally habitually switches his mobile to silent and does not answer calls while on public transport, (ii) a list of places that Wally often visits, including the house of his friend Ernie. Wally's profile states that he habitually takes the no. 27 bus along Fairtree Road when visiting Ernie, so the system generates and ranks the following explanation as highly likely: *The alarm was triggered erroneously or is faulty, because Wally is on the bus to visit a friend.*

Meanwhile, neither robots nor house cameras have been able to locate Wally. A house camera reports a location for the alarm button, but a nearby robot checks then refutes this. So the HCA sends a message to the help centre with Ernie's contact details and a suggestion that they call him to ask if Wally is visiting. Ernie confirms that a visit is planned, and shortly after phones back to confirm Wally's arrival. An expensive ambulance call-out is thus avoided.

Finally the HCA double-checks via a text exchange with Wally that he did not press his alarm, then sends a text to the alarm system maintenance service asking for an appointment to check the system.

2.2 Research Challenges

In cases like the running example, answering the three questions we posed in Section 1 brings about several research challenges. Although in this example the input indicating that Wally has pressed the alarm button can easily be identified as an event that requires further investigation, identifying events of interest is not always straightforward. Furthermore, detecting the causes of such events, and deciding how to react to them in such complex environments requires handling challenges such as:

- *How to decide which from the possibly large amount of available information is relevant to the event*. In our example: how can the system determine that the fact that Wally's mobile phone is silent is relevant to explaining the triggered alarm.
- *How to collect the relevant information*. In the running example locating the alarm button inside the house requires the two robots to move and search in places that may not be reached by static sensors, and to recognise the distinctive pattern painted on the alarm button.
- *How to combine the available information to generate possible explanations*, especially when this information comes from diverse sources and in different formats. In our example, the conclusion that Wally took bus 27 to visit Ernie is based on a combination of different types of information such as Wally's location, recent calls and habits, coming from different sources, i.e. Wally's mobile phone and the HCA.
- *How to deal with imperfections in the available information*, e.g. inaccuracy, incompleteness or conflicts. In the running example, the house camera and a robot report conflicting information about the location of the alarm button.
- *How to make the system robust to failures*. In our example, the system should be able to generate sensible explanations for the triggered alarm even if one of the RoCAs or Wally's mobile phone fails.
- *How to protect the privacy of all involved people*. In our scenario, Ernie must have given his consent to Wally's HCA to share his contact details with the help centre in case of emergency.
- *How to react to events taking into account the individual requirements of all involved parties*. In our example, apart from the overall goal of taking care of Wally, Wally may wish to keep his personal information private, while the system managers may be required to run the system efficiently.

3 STATE OF THE ART

In the literature, one can find several proposals for reasoning with imperfect information in AmI environments (Bettini et al., 2010). Many of them use Machine Learning (ML) techniques such as Bayesian Networks (Petzold et al., 2005), Time Series predictions (Das et al., 2002), Markov Models (Gellert and Vintan, 2006) and Neural Networks (Pansiot et al., 2007) for specific reasoning tasks such as identifying a particular user activity. Most recently, the use of Reinforcement Learning has been proposed to provide an "implicit feedback loop" between context prediction and actuation decisions (Boytsov and Zaslavsky, 2010). Although ML solutions are acceptably accurate, due to the lack of an explicit knowledge representation, they cannot provide high-level explanations for automatically chosen actions, nor can they be easily extended and modified in a modular fashion by the users or managers of an AmI system. Moreover, they cannot cope with run-time dynamicity, as any enhancement to the system or changes to the environment necessitate re-training of the system.

Rule-based approaches overcome some of the above limitations. A variety of decidable and tractable formalisms can be used to create knowledge-based domain models, thus facilitating efficient, formal reasoning about context (Broda et al., 2009). Their formality, expressiveness, modularity and extensibility allow them to better satisfy the needs for interoperability among heterogeneous components, adaptability to changes in the environment, and maintainability of large knowledge bases (Bikakis and Antoniou, 2010b).

Most current rule-based solutions adopt a centralised reasoning approach, but the need for distributed reasoning has also been acknowledged mainly for better scalability and robustness. Most distributed approaches, though, are either not fully decentralised, e.g. knowledge is distributed, but reasoning is local and agents do not exchange context information (Román et al., 2002); or are limited in their reasoning capabilities. For example, reasoning in (Viterbo and Endler, 2012) is distributed in that different computational nodes cooperate to infer a global context state; the reasoning process is however limited to ontology-based inference, and is not resilient to inaccurate or conflicting information. Finally, all existing knowledge-based reasoning approaches assume a single direction information flow: from sensor input, through to data analysis and decision-making, and then to reaction and control, irrespective of the extent to which each of these phases are decentralised (Snchez-Garzn et al., 2012; Valero et al., 2013).

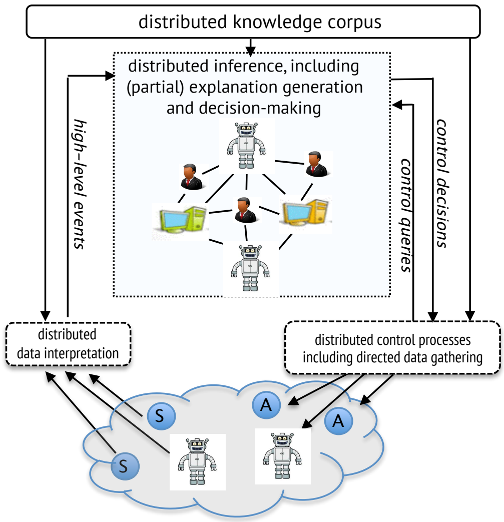

Figure 1: Top Level Architecture.

4 PROPOSED APPROACH

The proposed approach would integrate (1) computational logic techniques for distributed reasoning and explanation generation in the presence of partial knowledge, inconsistency and noisy data, (2) control procedures capable of performing proactive and robust responses of the system in an interleaved manner with the event-driven knowledge-based inference and (3) enhanced cognitive-driven robotic capabilities such as visual scene understanding and privacy and security awareness. Figure 1 illustrates our top level architecture, where 'S's and 'A's signify various sensors and actuators in the environment.

4.1 Distributed Defeasible Abduction

The first component of our proposed system would extend the DAREC engine, proposed in (Ma et al., 2010; Ma et al., 2011), with contextual defeasible inference (Bikakis and Antoniou, 2010a; Bikakis et al., 2011). DAREC is a general-purpose system that permits collaborative reasoning between agents over decentralised incomplete knowledge. It is particularly suited to compute collaboratively explanations of given observations. For instance, in cognitive robotics it can be used to collectively *abduce* explanations, in terms of descriptions of the world, from sensor data. Agents can recruit other agents on-the-fly, based on their local knowledge and reasoning capabilities, and recover from other agents' failures during a distributed computation task. The DAREC distributed algorithm is flexible and efficient, able to perform constraint satisfaction and customisable with application-dependent coordination strategies to better balance collaborative inference and inter-agent communication according to the particular domain and computational infrastructure. However, one of the main underlying assumptions is *global consistency*. An explanation has to be guaranteed to be consistent with respect to observed information, constraints and local knowledge of the agents involved in the collaborative inference. In practice, conflicting information may occur when integrating information from different sources.

Contextual Defeasible Logic (CDL), on the other hand has been used, also in a decentralised manner, to tackle the issue of distribute deductive inference in the presence of inconsistent information using non-monotonic inference and priorities among the conflicting chains of reasoning (arguments) (Bikakis and Antoniou, 2010a; Bikakis et al., 2011). Such priorities may represent different levels of confidence in the content of the arguments or different levels of trust in the information sources. These approaches, however, cannot compute collaborative explanations from distributed observations. We propose overcoming these two limitations by reformulating the DAREC distributed abductive inference mechanism in CDL. The main idea is to be able to collectively compute explanations of distributed observations but also determining the contextual information for which that explanation is plausible. This would for instance enable the generation of an explanation from the sensor data, that Wally is not in the house despite the camera reporting an indoor location for the alarm button.

4.2 Knowledge Representation

To communicate effectively, the agents of our proposed system would need to share a common vocabulary and ontology for referring to objects in their environment and the relationships between them. Developing such an ontology, and expressing sufficient "commonsense" knowledge about the environment with it, is itself a major research challenge. The ontological framework for expressing such knowledge should facilitate basic spatial, temporal and causal reasoning. For example, the system might at some point need to utilise the knowledge that spilling boiling water onto an electrical device (e.g. Wally's alarm button) typically damages the device, and causes it to be hot and wet, and to remain wet for some time. The Event Calculus (EC) (Kowalski and Sergot, 1986; Miller and Shanahan, 2002) is a prime candidate for this type of knowledge representation and reasoning. The EC is a logical mechanism for inferring the cumulative effects of a sequence of events recorded along a time line, or (used abductively (Shanahan, 1989)) inferring possible causes of a temporal sequence of observations (e.g. sensor readings). The EC has been extended with several features that are particulaly relevant in the present context. One is the ability to infer compound or "high level" events (e.g. 'Wally has gone upstairs') from a sequence of "smaller" events (e.g. the triggering of movement sensors in sequence from the bottom to top of the staircase) (Artikis and Paliouras, 2009; Alrajeh et al., 2013). Another is the ability to reason *epistemically* about the agent's own future knowledge should it perform sensing or data-gathering actions (e.g. phoning Ernie will result in knowing whether Wally is visiting him) (Ma et al., 2013). For these reasons, our proposed system would embed EC-style ontology inside the agent's knowledge.

4.3 Knowledge based Control

To address the first three challenges in Section 2.2 our proposal would allow for proactive multi-agent explanation generation and evidence gathering. As described in Section 4.1, distributed defeasible abductive reasoning would enable multi-agent computation of context-dependent explanations to distributed observations in the presence of uncertain and conflicting information. This process would be most effective when done proactively and in an interleaved manner with multi-agent evidence gathering. Our underlying infrastructure of sensors/actuators network and mobile robots would be capable of performing *proactive* data gathering: responding to requests for specific data gathering, deemed by the knowledge-based inference to be relevant for computation of more likely explanations. It could also trigger requests for inference tasks for knowledge-driven control.

Knowledge-driven control can be achieved using a multi-thread agent architecture that embeds Nilsson's Teleo-Reactive (TR) procedures (Nilsson, 1994) for robot control. A TR procedure is an ordered sequence of condition-action rules, in which the conditions can access the current agent's store of observed and inferred (i.e. deduced or abduced) dynamic beliefs: sensor percepts, told and remembered beliefs about the environment and its inhabitants. The rule actions are either device control actions, calls to TR procedures including recursive calls, or forking of information gathering or distrubuted reasoning threads. In each called procedure the first rule with an inferable guard is fired eventually resulting in device actions or thread forking. Device actions typically continue until new device actions are determined. The belief store of an agent is continuously and asynchronously updated as the (reactive) control procedures execute. On each update the last rule firings are reconsidered starting with the firing of the initially called procedure of each task thread. This unique operational semantics means that TR control is robust and opportunistic. If helped a TR controlled robot will automatically skip actions, if hindered it will redo actions.

Reasoning about the expected effects of control gives the system a means for intelligently monitoring its own performance. For instance, integration of sensing/actuation and knowledge-base inference may enhance the visual perceptions of mobile robots. Computer vision tasks (e.g. 3D reconstruction, object recognition and tracking), can be complemented with knowledge inference to provide the service robot with the ability to draw conclusions, generate plausible explanations from what it sees and gain real-time visual scene understanding. The robot would be able to reason about the objects in the environment and its inhabitants, explain current observations, or even reconstruct a narrative of the scene by deriving and using information that it cannot see, thus making more intelligent decisions about its actions.

Communication between multi-agent components could make use of an agent acquaintance model acquired using publish/subscribe (Robinson and Clark, 2010) to recruit agents likely to be able to contribute to an inference task. As new agents are added to the system, they subscribe for event notifications of interest to them, updating subscriptions to reflect focus of interest. These agents then exert control over the monitored system by posting action request notifications to be routed to other agents and devices. No component needs to know the identities of other components, or even what other components there are. All that has to be decided is the ontology for notifications and subscriptions.

4.4 Robotics Considerations

In applications such as ours, hardware considerations have to be taken into account. To satisfy the system requirements our service robots would include different kinds of capabilities from efficiency and cost to conviviality (Caire et al., 2011). So we need to examine a range of potential service robots which could fulfil the requested tasks. For example, to provide the basic functionality of looking for Wally in case of emergency, the simplest robot could consist of a Kinect camera attached to a mobile robot platform. For a more sophisticated robot, a tablet attached to a pole could be added as a feature. This would enhance its interactions with Wally, and provide a basis to be used as a healthcare assistant. A higher end robot could consist of a humanoid robot with more subtle and varied interaction capabilities, possibly equipped with 3D camera. Such a service robot could be potentially be used as a personal companion at home.

4.5 Privacy and Security

Of course, and particularly in AAL and scenarios such as ours, privacy and security issues are always a concern. Our approach in this domain is to first, ensure that communication on the health network is encrypted. Any secure method, such as Transport Layer Security (TLS) and Secure Sockets Layer (SSL) would be appropriate. The TLS protocol allows distributed applications to communicate across a network in such a way as to prevent eavesdropping and tampering, while the SSL ensures confidentiality, integrity, and authenticity of individual packets.

Second, strict access policy can be put in place to ensure that, for example, only doctors have direct access to patients. Provisions could then be set to, for example, relax the strict policy and allow doctors to, under certain conditions, share access, or part of it, with others, e.g. nurses. Approaches such as Role-Based Access Control or Attribute-Based Access Control could be used, e.g. (Kateb et al., 2014). The latter method offers the interesting perspective whereby a subject's requests to perform operations on objects are granted or denied based on the attributes of the subject, the object and the environment.

In our scenario, the robots only navigate within the house. Privacy with respect to video and sound is therefore not an issue. Furthermore, they transmit video and sound solely when there is an emergency. Areas such as the bathroom, could by default, i.e. when no emergency has been triggered, be declared *no-go* areas. In all other cases the video and sound data would only be used locally by the robot to interact with its environment.

As for the recording devices themselves, blurring filters may be used on cameras to allow only a general view of the scene, i.e. no details. This could be used in some cases, such as alarms, to check whether Wally is in the bathroom. As for microphones, switches have to be implemented with an *off* default setting. Indeed, it is preferable to keep microphones switched off at all times except for emergencies, e.g. where the robots are looking for Wally.

5 CONCLUSION

In this paper we argue that the ability of an AmI system to explain the causes of perceived events in its environment is key to its success, and is best achieved by an iterative process of tentative explanation generation interleaved with directed evidence gathering. We further argued that a number of recently developed knowledge-based technologies and methods could be extended and combined into a next generation of AmI systems in the area of assisted living, in particular utilising the latest generation of mobile robots acting as agents in a distributed computational setting. We outlined some of the associated research challenges and described a composite approach to their solution using the latest methods in distributed abductive and defeasible reasoning, teleoreactive control, and event-based knowledge representation. Such a system would capitalise on the advantages of a distributed, knowledge-based approach, such as transparency of computation, easy adaptability and extendability, no single point-of-failure, and closeness of fit with human-level reasoning.

REFERENCES

Alrajeh, D., Miller, R., Russo, A., and Uchitel, S. (2013). Reasoning about triggered scenarios in logic programming. In *Proc. ICLP'13*.

Artikis, A. and Paliouras, G. (2009). Behaviour Recognition Using the Event Calculus. *Artificial Intelligence Applications and Innovations III*, pages 469–478.

Bettini, C. et al. (2010). A survey of context modelling and reasoning techniques. *Pervasive and Mobile Computing*, 6(2):161 – 180.

Bikakis, A. and Antoniou, G. (2010a). Defeasible Contextual Reasoning with Arguments in Ambient Intelligence. *IEEE Transactions on Knowledge and Data Engineering*, 22(11):1492–1506.

Bikakis, A. and Antoniou, G. (2010b). Rule-based contextual reasoning in ambient intelligence. In *Proc. RuleML 2010*, volume 6403 of *LNCS*, pages 74–88. Springer.

Bikakis, A., Antoniou, G., and Hassapis, P. (2011). Strategies for contextual reasoning with conflicts in Ambient Intelligence. *Knowledge and Information Systems*, 27(1):45–84.

Boytsov, A. and Zaslavsky, A. B. (2010). Extending context spaces theory by proactive adaptation. In *Proc. NEW2AN 2010*, volume 6294 of *Lecture Notes in Computer Science*, pages 1–12. Springer.

Broda, K., Clark, K., Miller, R., and Russo, A. (2009). Sage: A logical agent-based environment monitoring and control system. In *Proc. AmI 2009*, volume 5859 of *LNCS*, pages 112–117. Springer.

Caire, P., Alcade, B., van der Torre, L., and Sombattheera, C. (2011). Conviviality measures. In *Proc. AAMAS 2011*.

Das, S. K., Cook, D. J., Bhattacharya, A., III, E. O. H., and Lin, T.-Y. (2002). The role of prediction algorithms in the mavhome smart home architecture. *IEEE Wireless Communications*, 9(6):77–84.

Gellert, A. and Vintan, L. (2006). Person movement prediction using hidden markov models. *Studies in Informatics and Control*, 15(1):17–30.

Kateb, D. E., Zannone, N., Caire, P., Moawad, A., Nain, G., Mouelhi, T., and Traon, Y. L. (2014). Conviviality-driven access control policy. *Requirements Engineering Journal*, pages 1–20.

Kowalski, R. A. and Sergot, M. J. (1986). A Logic-Based Calculus of Events. *New Generation Computing*, 4:67–95.

Ma, J., Miller, R., Morgenstern, L., and Patkos, T. (2013). An epistemic event calculus for asp-based reasoning about knowledge of the past, present and future. In *Proc. LPAR-19*.

Ma, J., Russo, A., Broda, K., and Lupu, E. (2010). Distributed abductive reasoning with constraints. In *Proc. DALT 2010*.

Ma, J., Russo, A., Broda, K., and Lupu, E. (2011). Multi-agent confidential abductive reasoning. In *Proc. ICLP 2011*.

Miller, R. and Shanahan, M. (2002). Some Alternative Formulations of the Event Calculus. In *Computational Logic: Logic Programming and Beyond, Essays in Honour of R. A. Kowalski, Part II*, pages 452–490.

Nilsson, N. (1994). Teleo-reactive programs for agent control. *Journal of Artificial Intelligence Research*, 1:139158.

Pansiot, J., Stoyanov, D., McIlwraith, D. G., Lo, B., and Yang, G. (2007). Ambient and wearable sensor fusion for activity recognition in healthcare monitoring systems. In *Proc. BSN 2007*, pages 208–212. Springer.

Petzold, J., Pietzowski, A., Bagci, F., Trumler, W., and Ungerer, T. (2005). Prediction of indoor movements using bayesian networks. In *Proc. LoCA 2005*, volume 3479 of *LNCS*, pages 211–222. Springer.

Robinson, P. and Clark, K. (2010). Pedro: a Publish/Subscribe Server using Prolog Technology. *Software: Practice and Experience*, 40(4):313329.

Román, M., Hess, C., Cerqueira, R., Ranganathan, A., Campbell, R. H., and Nahrstedt, K. (2002). A middleware infrastructure for active spaces. *IEEE Pervasive Computing*, 1(4):74–83.

Shanahan, M. (1989). Prediction is deduction but explanation is abduction. In *Proc. IJCAI'89*.

Snchez-Garzn, I., Milla-Milln, G., and Fernndez-Olivares, J. (2012). Context-aware generation and adaptive execution of daily living care pathways. In *Proc. IWAAL 2012*, volume 7657 of *LNCS*, pages 362–370. Springer.

Valero, M. n., Bravo, J., Garca, J. M., de Ipia, D. L., and Gmez, A. (2013). A knowledge based framework to support active aging at home based environments. In *Proc. IWAAL 2013*, volume 8277 of *LNCS*, pages 1–8. Springer.

Viterbo, J. and Endler, M. (2012). *Decentralized Reasoning in Ambient Intelligence.* Springer Briefs in Computer Science. Springer.