

Coaching Through Smart Objects

Chris Baber
School of Engineering, University
of Birmingham
Birmingham, UK
c.baber@bham.ac.uk

Ahmad Khattab
School of Engineering,
University of Birmingham
Birmingham, UK
AXK208@student.bham.ac.uk

Joachim Hermsdörfer
Department of Sport and Health
Sciences,
Technische Universität München,
Germany
joachim.hermsdoerfer@tum.de

Alan Wing
School of Psychology, University
of Birmingham
Birmingham, UK
a.m.wing@bham.ac.uk

Martin Russell
School of Engineering,
University of Birmingham
Birmingham, UK
m.j.russell@bham.ac.uk

ABSTRACT

We explore the ways in which smart objects can be used to cue actions as part of coaching for Activities of Daily Living (ADL) following brain damage or injury, such as might arise following a stroke. In this case, appropriate actions are cued for a given context. The context is defined by the intention of the users, the state of the objects and the tasks for which these objects can be used. This requires objects to be instrumented so that they can recognize the actions that users perform. In order to provide appropriate cues, the objects also need to be able to display information to users, e.g., by changing their physical appearance or by providing auditory output. We discuss the ways in which information can be displayed to cue user action.

Author Keywords

Tangible user interface; Activity recognition; Multimodal cueing.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI).

INTRODUCTION

In this paper we regard ‘coaching’ as a way of encouraging people to act. Broadly, coaching involves a set of processes which are aimed at helping an individual (or group of individuals) improve, develop or learn skills. Often coaching will involve a dialogue between the individual and their coach. Replacing a (human) coach with a digital counterpart, therefore, raises some interesting questions concerning the ways in which to determine the

improvement, development or learning by the individual. Ensuring that the coaching is tailored to the abilities of the individual is an essential for digital coaching [1]. From this, a basic specification for a digital coach would include the ability to determine which action is performed by the individual, to evaluate the actions (against some quality criterion), and to provide advice, guidance or cueing that could lead to improvement (or alteration) in the performance of the actions.

Patients with neurological deficit, as the result of degeneration, damage or injury, can often struggle with Activities of Daily Living (ADL). These problems can range from confusion over the sequence in which tasks should be performed, to forgetting which action can be performed using a given object, to failing to manipulate a given object. These errors could involve putting breakfast cereal into a coffee mug (when making coffee) or attempting to pour from a milk container before opening it [2]. Around 1/3 of stroke survivors have difficulty with ADL but the errors that they make in planning and execution of ADLs are predictable [3]. If it is possible to predict errors (rather than these being random or chance occurrences) then it is possible to predict *when* an error could be likely to arise and to provide some cueing to prevent the erroneous action and to encourage a preferred, i.e., non-erroneous, action.

Returning to the basic specification outlined above, for a system to provide coaching to help patients with neurological deficit to perform ADL, we need to recognize which actions people are performing (which typically means having some means of sensing that an action is being performed), to determine the quality of the performance (which requires not only activity recognition but also some evaluation of how this performance meets some criterion), to evaluate performance in terms of an anticipated outcome or intention, and to provide some means of guiding, cueing or otherwise providing advice and feedback in response to the action and in anticipation of subsequent action.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
PervasiveHealth '17, May 23–26, 2017, Barcelona, Spain

© 2017 Association for Computing Machinery.
ACM ISBN 978-1-4503-6363-1/17/05...\$15.00
<https://doi.org/10.1145/3154862.3154938>

In the EU-funded project *CogWatch* (figure 1) we developed technology that supported ADL through recognition of activity and cueing to reduce errors [4, 5, 6].

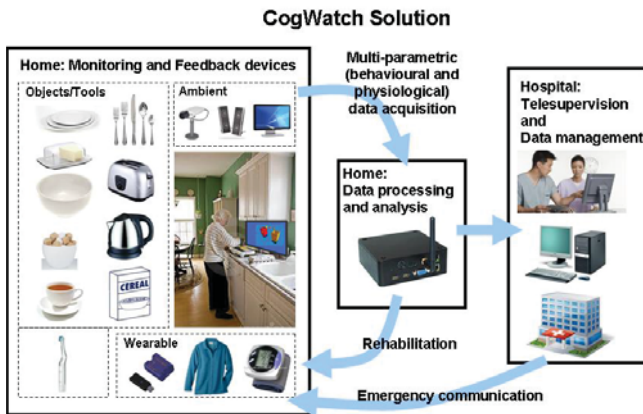


Figure 1: Schematic of the CogWatch concept

Figure 2 shows the CogWatch system being used in the laboratory. Sensors (accelerometers and force sensitive resistors) on the objects placed on the table detect the actions that a person makes with them. Data from the sensors, together with hand tracking data from Microsoft Kinect, are used to create Hidden Markov Models for activity recognition [7]. In order to determine when to provide a cue to the user, the activity recognition output is compared with the prediction of the actions which would be appropriate for a goal. The prediction is based on Partially-Observed Markov Decision Process (POMDP) models of task sequence [8, 9].



Figure 2: Interacting with the CogWatch system

Thus, when a person performs a sequence of actions, such as making a cup of tea, each action they perform is recognized and compared with a set of plausible actions. If an action is not part of this plausible sequence it could be defined as an error, e.g., because an action is repeated or because it was not appropriate at that point in the sequence. If this occurs then the user receives a prompt on the visual.

The results from evaluation sessions showed that patients struggle to complete ADL, such as tea-making, without support, e.g., a majority of patients failed to successfully prepare the required tea. Even when patients

were able to consult printed, visual instructions on the step-by-step sequence of actions, they still failed to complete the tasks. We believe that this shows that this printed form of support is ineffective for this particular task and population. In contrast, almost all of the trials with CogWatch support resulted in patients successfully completing the tea making tasks [4]. However, even in these CogWatch trials patients made errors (most of which they were able to correct) and took significant time to complete the activity. While the *CogWatch* project demonstrated that patients were able to respond effectively to the cues presented to them, the system relied on the use of a visual display to provide these cues as shown in figure 2. One explanation of the time and errors that were noted in these trials is that patients might have found it difficult to divide their attention between the physical actions involved in performing the tasks using the objects, and the more abstract task of reading instructions and relating these to their actions. Consequently, we explored whether the cues could be provided by the objects themselves.

CUEING ACTION FROM OBJECT BEHAVIOUR

We assume that objects could be designed to provide visual, tactile or auditory cues to the user. This builds on prior work on smart objects [10] and Tangible User Interfaces [11]. A smart object typically has awareness (defined as the ability to sense where it is, how it is being used etc.), representation (defined as the ability to make sense of its awareness), and interaction (defined as the ability to respond to the user or other objects). In *CogWatch*, as discussed previously, awareness was achieved through the integration of sensors on objects and representation was through the developed on HMM and POMDP. In order to support interaction, we extend the design of these objects to present information to users.

The development of Tangible User Interfaces (TUI) over the past two decades has been dependent on the availability of miniature sensors and processors. Much of this work has focused on the development of objects as input devices or objects as forms of ambient display. Contemporary work, particularly at MIT [12], has been exploring ways in which objects can be physically transformed. In terms of application in the healthcare domain, an ambient display has been developed to alert teenagers with Attention Deficit / Hyperactivity Disorder (ADHD) to support everyday planning of activity [13]. Similarly, ambient displays can be used to provide reminders to patients concerning the time to take medication [14, 15]. The concept of a Rehabilitation Internet of Things (RioT) [16] uses COTS wearable sensors to provide data about the physical activities of individuals wearing these devices.

Before describing the ways in which we have implemented tangible, ambient and other cues in the design of our objects, the next section considers preliminary experimental analysis of the impact of cues, using different modalities, on the activity of patients.

Visual cue of object's state and required tasks



Figure 3: LED on Drawer Handle (from [17])

In order to indicate the state of an object, one can use visual cues to show its temperature or whether it is turned on or off, or open or closed. Figure 3 shows the use of Light-Emitting Diode (LED) in the handle of a drawer to alert a person to this drawer. The light could indicate that this particular drawer contains the saucepan that the person needs for making a sauce.

We explored the relationship between LEDs and the state of objects in a simple problem-solving exercise [18]. The objective is to ensure that four boxes had satisfied their goals (figure 4). Each box had a set of rules to specific its 'goal', defined by the position of the box and its proximity to other boxes. Each box has 3 LEDs representing its state: one to indicate if the goal has been satisfied, one to indicate 'communication status' (in terms of connection with the table), and one to indicate 'proximity'.

We were interested in whether people would try to learn the 'rules' that the boxes were using or whether they would find it easier to learn the pattern, or arrangement of the boxes, and whether the rules or patterns could generalize to new configurations. The argument for this comparison was that the patterns could be considered in terms of affordance. Not only were the patterns easier to understand (which suggests that the visual cues provides useful semantic information) but also participants found it easier to generalize patterns than the rules. We had assumed that the patterns would be learned as 'local' descriptions whereas the rule would be 'global' descriptions, and therefore easier to generalize. The finding that people could extrapolate generalizable understanding (even if they could not always articulate these in the form of rules) suggested that learning of the patterns also provided global information (contrary to our expectations).




Figure 4: Visual feedback on networked objects [18]

Biological sound cue

Asking patients with Parkinson's Disease to walk in time to the (prerecorded) sound of footsteps on gravel can lead to better support with gait problems than walking in time to a metronome [19]. This suggests that there is some element of the 'natural' sounds which, in addition to the marking of time, can improve performance.

In terms of ADL, Bienkiewicz et al. showed that the noise made during the performance of a task (such as the sound of a saw, the sound of pouring water, or of stirring with a spoon) can support apraxic patients in recalling a motor program which is otherwise not accessible [20]. We presume that the effect of the cue is even much stronger if the objects itself omits the biological sound.

Functional grasps cues

We grasp tools according to the intended use [21]. Thus, while we may grasp a hammer in different ways when we want to transport it, we will grasp it at the handle with the thumb towards the heavy part when we want to use it immediately to drive a nail into a wall. While in apraxic patients a functional grasp does not guarantee the correct use of the tool [22], such a grasp serves as a strong attractor that increases the likelihood of executing the correct gesture [23]. Even for neurotypical participants, people are faster at performing a manual response to an object when they use the hand that is aligned with the handle of a manipulable object compared to its functional end [24]. This suggests that having some means of indicating which part of an object to grasp could be useful.

The action could be cued by simple modifications to the handle, e.g., by having LEDs on the handle to distinguish this from other parts and turning on the LEDs to draw the user's attention to the handle. Furthermore, if the object could change orientation or its shape, this might further encourage a grasp required for a specific task.

Haptic and Tactile Cues

As Poupyrev et al. [25] note, there is a tendency for TUIs to respond to users primarily through visual or auditory displays and there has been less work on displays which can

change their physical appearance. The development of small, easy-to-use actuators makes it possible for shape-changing objects to be created.

In addition to cueing when to perform an action, it is possible to influence the ongoing performance of an action in order to correct or compensate the manner in which the action is performed. Figure 5 shows a commercial product which is designed to compensate for tremor, such as might arise from Parkinson's Disease.



Figure 5: Stabilizing Spoon [<https://www.liftware.com/>]

Having the object changing its physical behaviour, e.g., through vibration, could be used to cue the user to which object to pick up (by making the object wobble on the table) as well as compensating for the movements performed by the user. As well as having the entire object move it is possible to focus on sections. For example, we experimented, for example, with an arrow on the lid of a jug which would point to the direction in which the person should move the jug (the arrow was connected to a servomotor driven by a magnetometer which responded to magnets placed on the table or in other objects).

CREATING AFFORDING SITUATIONS

In Psychology and Human-Computer Interaction, the concept of 'affordance' relates to the ways in which a person's action is performed in response to their perception of the environment [26, 27]. Gibson [28] introduced the term affordance into psychology, suggesting that we perceive the world in terms of opportunities for action. What an object affords is determined by the physical properties of the objects (e.g., shape, orientation, size), and by the action capabilities of the agent and by the intention that the use of this object will support. Thus, an affordance is the relationship between an individual's ability to act and the opportunities provided to that person in the given situation in pursuit of a given intention [29]. That is, a cup of

particular dimensions can be grasped by a person of particular abilities in the context of performing a task with a particular intended goal: a person with hemiparesis might struggle to use a pinch grasp on a small teacup handle; a person with tremor might find it difficult to raise a full cup to their mouth (without spillage). A person laying the table will pick up the cup differently than a person who intends to drink from it.

From this perspective, one can consider coaching in terms of the form of encouragement of, and support for, actions which are appropriate to a given context (defined by objects, person's abilities and intentions). We aim to create 'affording situations' in which the appropriate action is subtly cued by the objects that the person needs to use. In this way, cueing is embedded naturally into the familiar world of the person's home, and coaching is a matter of responding to these cues.

The manner in which a person interacts with an object can be used to infer the physical and cognitive difficulties they might be experiencing in performing the task. As the person interacts with the object, data (from sensors on the object) will be used to define the nature of the atomic tasks being performed and, from these data, the intended outcome can be inferred. Having inferred a possible outcome, objects can modify their state to invite or cue subsequent actions, or provide feedback on the task as it is performed.

DESIGNING INTELLIGENT OBJECTS

For the CogWatch project, objects used in ADL were kept as normal as possible in appearance and function, to avoid causing further confusion to the patients. This meant that the sensors had to be small and discrete. For several of the objects used in the archetypical ADL of making a cup of tea, we developed an instrumented coaster. This design allowed us to package the sensors and circuitry into a device that is fitted to the underside of the object, where it has very little visual impact and does not obstruct the use of the object (figure 6). This is inspired by the well-known MediaCup concept [30].

The Coaster is fitted with Force Sensitive Resistors (FSRs) which are used to not only determine when the object is on the table or lifted, but can also be used to estimate how much liquid is being poured into a container. In addition to FSRs, a triaxial accelerometer is used to record movement. The set of sensors is controlled by a Microchip dsPIC30F3012 microcontroller, which has an integrated 12 bit analogue digital converter (ADC) that is used for digitizing the sensor readings. The microcontroller is programmed to digitize, compress and prepare the sensor data and manage the transmission of the data via Bluetooth. The data are buffered on the microcontroller so that they can be re-transmitted (avoiding data loss) if the wireless connection is interrupted for a short period. An ARF7044 Bluetooth module is used to transmit the sensor data to a host computer via a Bluetooth wireless connection (figure 7).




Figure 6: CogWatch coaster fitted under a mug (a) next to another dismantled coaster (b) and underside of coaster showing force sensitive resistors (c).

The use of sensors on objects to recognize the actions that are being performed has been explored in a range of projects in addition to CogWatch. In an extension to this project, we are developing objects which combine action recognition with action cueing in order to support coaching. In this case, the aim is to modify the objects in ways which can indicate to the user that they need to perform a specific action with a specific object. Clearly this work is still in its infancy, and in this paper we are discussing some of the conceptual developments and initial prototypes.

Context-aware predictions

The overall aim of our work is to create affording situations in which the behavior of the person can be cued by the behavior of the object which, in turn, responds to the (inferred) intentions of the user. In this case, coaching will involve building a (computer) model of the intentions of the user (based on the previous activity of the person etc.) and a set of criteria for how these intentions ought to be met. The criteria could be learned by the computer or, more likely, could be defined by occupational therapists who would be working with the patient. In this case, the idea would be that a given ADL could be decomposed into specific actions, and each action could be defined in terms of quality of performance.

In the CogWatch project, tasks were recognized using Hidden Markov Models, HMM (figure 7). In this implementation, multiple HMM decoders ran in parallel to ‘listen’ for output from sensors attached to objects. All sensor data were collated (through the Bluetooth receiver) into a single vector. Specific ‘tasks’ were identified in terms

of Probability Density Functions (PDFs), learned from training the recognizer on instances of each ‘task’. Thus, figure 7 shows the model for ‘tilt the kettle until the cup is full’, and uses data from the Force Sensitive Resistor (FSR) on the kettle (to indicate that the handle has been grasped and the kettle lifted), from the accelerometer on the kettle (to indicate tipping), and from the FSR on the cup to indicate that the cup is full. If the kettle continues to be tilted, then the cup could overflow and so a warning message could be provided to the user.

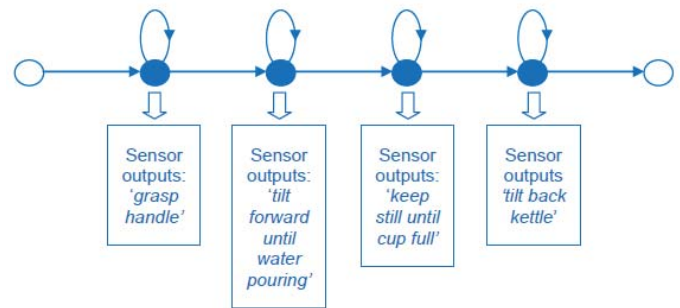


Figure 7: HMM for ‘tilt kettle until cup is full’

In addition to various ‘tasks’, to CogWatch system works with an Task Model (built using Partially-Observable Markov Decision Processes) which identify states of the system, including both tasks and ‘errors’ (i.e., ‘tasks’ which could be performed but which are undesired for a given goal). This requires the POMDP to identify the likely goal being pursued and to identify deviations from the path towards that goal.

Subsequent development will incorporate data from more sensors into the vector for both activity and task models. We see this as following a simple taxonomy in which activity can be classed as Global (e.g., entering or leaving a specific room at a specific time of day), Local (e.g., activity which can be performed in that room at that time), and Object-Specific (e.g., activity which can be performed using a particular object). Thus, one could identify which room a person is in, and then infer a set of plausible tasks that could be performed, and then infer which objects would be required to perform those tasks and how these objects could be used. In this way, the task model is seeded with a set of plausible tasks, and this could enhance recognition and prediction of error. In addition, our work to date has only considered whether an action has been performed, and makes no judgement as to the quality of that action. Consequently, it would be useful to know if the action has been performed well or whether there have been problems in performance, e.g., [31].

Adapting: patients and objects

A challenge for this work is to consider how patients might respond to the visual (or auditory) appearance of the objects. We note that prior work (discussed above)

suggests that patients can respond positively to appropriate auditory cues, and that these can encourage task-relevant actions. One would assume that providing cues at appropriate decision points could also be beneficial. However, there remains the possibility that patients could become confused or disoriented by the behavior of animate objects. None of this work has been exposed to patients and thus we do not know how they might respond and react to lights, sounds or moving parts on everyday objects. This remains an open question and one to be explored in subsequent work. We hope that patients could either recall learned behaviour in response to the objects, or could learn new behaviours that the objects cue.

In addition to the patients learning behavior, we would expect the objects themselves to become capable of learning the behavior of individual patients. In this way, it would be possible for the level of coaching that is required to be modified. If the patient shows improvement in performance, say responding quickly and accurately to cues or even anticipating a cue, then there is no need for that cue to be provided. Alternatively, if the patient seems to struggle with a task, then additional cues might be necessary. From the previous discussion on context-aware predictions, this level of adaptation could be a simple matter of creating sub-task models of activity (in order to evaluate the quality of performance) and some simple decision logic defining when to play a given cue. This would represent a modest alteration to the current design, in which cues would be selected on the basis of defined thresholds. A more interesting development would be to have activity recognition (at sub-task level) being performed longitudinally so as to track the manner in which patient behavior changes (either improving or deteriorating).

Privacy Preservation

Given the capability to record, interpret and share data relating to the actions of patients, it is important to be able to provide some safeguards on the ways in which such data. While this has not been, as yet, a focus of our work, it is important to recognize the challenges that privacy and data security pose. Bluetooth, as we used in the CogWatch devices, advertises the MAC address of devices, which could (for instance) mean that an address could be paired with a given person (which has prompted concerns over ‘tracking’ of individuals). Use of random, private MAC addresses (resolved using Identity Resolution Key encryption) can reduce this problem. However, this does not guarantee that data sharing could not be compromised. In our more recent development work, we use the Blynk server to connect devices. In effect, this creates a local, private cloud (which allows us to connect many devices with a Secure Sockets Layer).

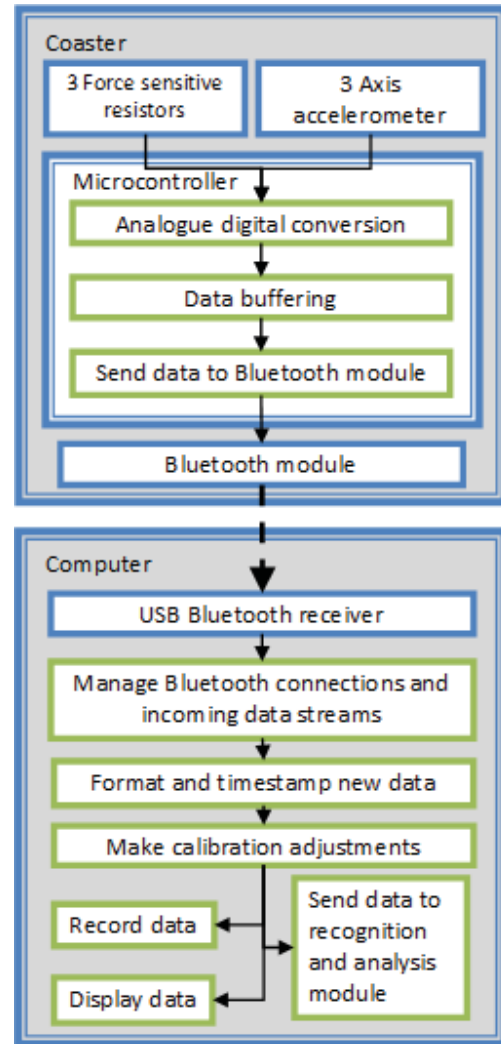


Figure 7: CogWatch Coaster system design.

A Collection of Animate objects

We have extended the capability of the objects designed for our work by incorporating actuators and displays which can provide the cues outlined in the previous sections (figure 8).

For example, assume that you will make a hot drink by boiling a kettle and then pouring this into a mug into which you have already added coffee granules. This can be decomposed into a sequence of steps; each step has a set of successive sub-steps which are more likely to lead to the goal. Simple performance measures, such as the time taken to complete each step, whether there appears to be hesitation during this performance, whether there are physical effects (such as tremor or jerky motion), or whether the action has been performed correctly. When a person appears confused, e.g., when they fail to act, or when they make an error, i.e., when they perform a step which is not one of the recommended ones, then they might require a prompt. In some cases, the prompt could form part of a training programme in which repeated exposure could help the person remember a sequence. This would be potentially

useful for recovering stroke patients. In other cases, it might not be possible for the person to learn, and so the prompts would be presented whenever they performed this sequence. This would be potentially useful for Alzheimer's patients. When the action has an error, then the challenge is how to interrupt the sequence of actions that the person is performing in order to either correct the sequence, or prevent further errors.

Figure 8 shows a set of objects used in our initial experiments. The first, and most obvious, point to note is that these are all first-order prototypes and have been designed to support experimental evaluation with neurotypical participants, rather than patients. The second point is that many of these feature coloured LEDs as visual feedback. Note that some of the objects have green lights and one mug has red lights. We recognize that green and red may be problematic for colour-blind people and this could be reconfigured in later designs. However, in our initial trials we do not tell participants what the lights mean but rather wait to see what interpretation the participants provide. What is not apparent from this image is that the objects are also fitted with motors or speakers, and that they can be controlled remotely (from an app running on an

Android phone). Thus, for example a simple vibration-motor in the spoon (which is in the opened drawer) vibrates to make the spoon wobble, and a motor (taken from the drive of a CD drawer) raises and lowers the jug handle. The intention is that the physical movement of the objects can both attract the attention of the user and also provide a cue as to the action to perform. For example, in one trial we place the jug in front of the participant and raise the handle: depending on which side the handle is facing, participants tend to use the hand on the same side as the handle (even when this is their non-preferred hand). If the participant does not respond to the handle raising, then an auditory cue (of the sound of pouring water) is played to draw their attention to the jug and prompt a lift and pour action, and if this fails to elicit a response a verbal prompt (with a voice recording of the phrase 'pick me up') is played. Trials are repeated for the various actions and sequences of actions and we are exploring how well people learn to respond to the objects, and how the different forms of cue influence activity. From this, we are examining how coaching (in terms of learning novel or unusual sequences, as well as reinforcing well-known sequences) can be supported by the use of object properties.



Figure 8: A collection of animate objects

DISCUSSION AND FUTURE WORK

In this paper we have considered ways in which familiar objects can be adapted to support affording situations. The overall concept is that the action that a person performs will be influenced by the context in which they perform that action. Context will, in turn, be influenced by the person's goals, the previous tasks that they have performed, the objects available to them and the state of these objects. By recognizing the actions that the person has performed we seek to infer their goal (or intention) and from this information we modify the appearance of the objects to cue credible actions in a sequence. Future work will take the notion of cueing further so that, in addition to examining tasks in sequence we will also consider wider contextual definitions of these tasks as the basis for cueing. For example, knowing the time and that the person is walking into the kitchen, we can cue them to open a cupboard (by illuminating the cupboard handle) so that they can take out a mug (which will, once the cupboard door opens, also light up). Once the cup has been retrieved, then we can cue a sequence of steps in the preparation of a drink.

The overall goal of this work is to create, as far as practicable, subtle cues through which actions can be suggested (when this is necessary for a given user). In the case of the recovering stroke patient, the use of these cues might reduce in line with their recovery such that after a few months of using this system, there might be few or no cues presented because the patient has recovered full ability for these tasks. In the case of patients with Alzheimer's or other forms of dementia, the cueing might need to be continued.

The concept of affording situations is intended to highlight the importance of context in understanding affordance; it is not simply a matter of saying that a 'jug affords lifting or pouring'. While these are actions that are, of course, possible with the jug, in order to know that this particular jug can be lifted or poured by this particular person, one also needs to know the capabilities of the person. Furthermore, in order to know whether either lifting or pouring is an appropriate action to make, one also needs to know the intention of the person (or the intention that could be plausible in that situation). Our aim is to develop technology that can recognize and interpret these contextual factors and use these to discern the plausibility of actions in a given sequence. We can then adapt the technology to provide cues that are intended to encourage the user to perform these actions. In this manner, coaching becomes a matter of creating the situations in which users of objects can learn to associate an object with an action, and an action with an intention, and an intention with a desirable outcome.

ACKNOWLEDGMENTS

The work in this paper was supported, in part, by EU STREP Project CogWatch [FP7-ICT-288912].

REFERENCES

1. Op den Akker, H., Jones, V. M., & Hermens, H. J. (2014) Tailoring Real-Time Physical Activity Coaching Systems: A Literature Survey and Model, *User Modeling And User-Adapted Interaction*, 24, 351–392.
2. Schwartz, M. F. (2006). The cognitive neuropsychology of everyday action and planning. *Cogn Neuropsychol.*, 23(1), 202-221.
3. Hughes, C.M.L., Baber, C., Bienkiewicz, M., Worthington, A., Hazell, A. & Hermsdörfer, J. (2015) The application of SHERPA (Systematic Human Error Reduction and Prediction Approach) in the development of compensatory cognitive rehabilitation strategies for stroke patients with left and right brain damage, *Ergonomics*, 58, 75-95.
4. Hermsdörfer, J., Bienkiewicz, M., Cogollor, J.M., Russel, M., Jean-Baptiste, E., Parekh, M., Wing, A.M., Ferre, M. & Hughes, C. (2013) CogWatch—Automated Assistance and Rehabilitation of Stroke-Induced Action Disorders in the Home Environment. In *International Conference on Engineering Psychology and Cognitive Ergonomics*, Springer Berlin Heidelberg 343-350.
5. Pastorino, M., Fioravanti, A., Arredondo, M., Cogollor, J., Rojo, J., Ferre, M., Bienkiewicz, M., Hermsdörfer, J.; Fringi, E. & Wing, A. (2014) Preliminary Evaluation of a Personal Healthcare System Prototype for Cognitive eRehabilitation in a Living Assistance Domain, *Sensors*, 14, 10213-10233.
6. Giachritsis, C. & Randall, G. (2012) CogWatch: Cognitive Rehabilitation for Apraxia and Action Disorganization Syndrome Patients. In *The Seventh International Workshop on Haptic and Audio Interaction Design August 23-24 2012 Lund, Sweden*.
7. Jean-Baptiste, E.M., Nabieci, R., Parekh, M., Fringi, E., Drozdowska, B., Baber, C., Jancovic, P., Rotschein, P. & Russell, M. (2014) Intelligent assistive system using real-time action recognition for stroke survivors, *2014 IEEE International Conference on Healthcare Informatics (ICHI)*, IEEE, 39-44.
8. Jean-Baptiste, E.M., Rotshtein, P. & Russell, M (2015) POMDP based action planning and human error detection, *IFIP International Conference on Artificial Intelligence Applications and Innovations*, Springer International Publishing, 250-265
9. Jean-Baptiste, E.M., Rotshtein, P. & Russell, M. (2016) CogWatch: Automatic prompting system for stroke survivors during activities of daily living. *Journal of Innovation in Digital Ecosystems*, 3(2), 48-56.
10. Kortuem, G., Kawsar, F., Fitton, D. & Sundramoorthy, V. (2010) Smart objects as building blocks for the Internet of Things, *IEEE Internet Computing*, Jan/Feb., 44-51.
11. Ishii, H. (2008) The tangible user interface and its evolution. *Communications of the ACM*, 51(6), 32-36.
12. Ishii, H., Leithinger, D., Follmer, S., Zoran, A., Schoessler, P., & Counts, J. (2015) TRANSFORM: Embodiment of Radical Atoms at Milano Design Week.

- In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems* (pp. 687-694). ACM
13. Zuckerman, O. (2015) Objects for change: A case study of a tangible user interface for behavior change. In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction* (pp. 649-654). ACM.
 14. Sung, M-H. & Chian, C-W. (2016) The Research of Using Magnetic Pillbox as Smart Pillbox System's Interactive Tangible User Interface, *International Conference on Human-Computer Interaction*, Springer International Publishing, 451-456.
 15. Reeder, B., Chung, J., Le, T., Thompson, H. J., & Demiris, G. (2014) Assessing older adults' perceptions of sensor data and designing visual displays for ambient assisted living environments: An exploratory study. *Methods of information in medicine*, 53, 152.
 16. Dobkin, B. H. (2016) A Rehabilitation-Internet-of-Things in the Home to Augment Motor Skills and Exercise Training, *Neurorehabilitation and Neural Repair*.
 17. Bonanni, L., Lee, C.H. & Selker, T., 2005. CounterIntelligence: Augmented reality kitchen. In *Proc. CHI* (Vol. 2239, p. 45).
 18. Cervantes-Solis, J.W., Baber, C., Khattab, A. & Mitch, R., 2015, July. Rule and theme discovery in human interactions with an 'internet of things'. In *Proceedings of the 2015 British HCI Conference* (pp. 222-227). ACM.
 19. Young, W.R., Shreve, L., Quinn, E.J., Craig, C. and Bronte-Stewart, H. (2016) Auditory cueing in Parkinson's patients with freezing of gait. What matters most: Action-relevance or cue continuity? *Neuropsychologia*, 87, 54-62.
 20. Bienkiewicz, M. N., Gulde, P., Schlegel, A., & Hermsdörfer, J. (2014) The Use of Ecological Sounds in Facilitation of Tool Use in Apraxia, 2nd International Conference on NeuroRehabilitation (ICNR2014). Paper presented at the Replace, Repair, Restore, Relieve – Bridging Clinical and Engineering Solutions in Neurorehabilitation, Aalborg, 24-26 June, 2014.
 21. Hermsdörfer, J., Li, Y., Randerath, J., Goldenberg, G., & Johannsen, L. (2012) Tool use without a tool: kinematic characteristics of pantomiming as compared to actual use and the effect of brain damage, *Experimental Brain Research*, 218, 201-214
 22. Randerath, J., Li, Y., Goldenberg, G., & Hermsdörfer, J. (2009) Grasping tools: Effects of task and apraxia. *Neuropsychologia*, 47, 497-505.
 23. Graham, N. L., Zeman, A., Young, A. W., Patterson, K., & Hodges, J. R. (1999) Dyspraxia in a patient with cortico-basal degeneration: the role of visual and tactile inputs to action, *Journal of Neurology, Neurosurgery and Psychiatry*, 67, 334-344.
 24. Matheson, H., Newman, A.J., Satel, J. & McMullen, P., (2014) Handles of manipulable objects attract covert visual attention: ERP evidence, *Brain and cognition*, 86, 17-23.
 25. Poupyrev, I., Nashida, T. & Okabe, M., 2007, February. Actuation and tangible user interfaces: the Vaucanson duck, robots, and shape displays. In *Proceedings of the 1st international conference on Tangible and embedded interaction* (pp. 205-212). ACM....
 26. Norman, D. A. (1990) *The Design of Everyday Things*, New York: Doubleday.
 27. Gaver, W. (1991). Technology affordances, *Proceedings of CHI 1991*, New York: ACM Press, 79-84.
 28. Gibson, J.J. (1986) *The Ecological Approach to Visual Perception*, Houghton Mifflin Company.
 29. Blevins, E., Bødker, S., Flach, J., Forlizzi, J., Jung, H., Kaptelinin, V., Nardi, B. & Rizzo, A. (2015) Ecological perspectives in HCI: Promise, problems, and potential, *Proceedings of CHI 2015*, New York: ACM Press, 2401-2404.
 30. Gellersen, H-W., Beigl, M. & Krull, H. (1999) The MediaCup: awareness technology embedded in an everyday object, *Handheld and Ubiquitous Computing 1st International Symposium HUC'99*, Berlin: Springer, 308-310.
 31. Jarrassé, N., Kuhne, M., Roach, N., Hussain, A., Balasubramanian, S., Burdet, E. and Roby-Brami, A. (2013) Analysis of grasping strategies and function in hemiparetic patients using an instrumented object, *Proceedings of the 13th International Conference on Rehabilitation Robotics (ICORR)*, 1-8.