Enabling Personalisation of Remote Elderly Assistance

Luca Corcella[1], Marco Manca[1], Jan Egil Nordvik[2], Fabio Paternò[1], Anne-Marthe Sanders[2], Carmen Santoro[1]
(1) CNR-ISTI HIIS Laboratory, Italy, (2) SUNNAS RH, Norway

ABSTRACT

One of the goals of Ambient Assisted Living (AAL) solutions is to extend the time that elderly people can live independently in their preferred environments by using ICT technologies for personal healthcare. However, in order to be optimal, remote monitoring services and health-related interventions should be strongly personalised to specific individuals' requirements, preferences, abilities and motivations, which can vary among the elderly, and even dynamically evolve over time for the same person, depending on changing user needs and conditions associated with the current context of use.

In this paper we present an End User Development (EUD) tool for the personalisation of context-dependent assistance by non-technical users in the AAL domain. In particular, we have considered applications for remotely monitoring and assisting elderly people at home. The design and development of the tailoring environment has been carried out in an iterative manner, informed by the feedback that was gathered through empirical evaluations done with some elderly people and caregivers.

KEYWORDS

End-user development, Ambient Assisted Living, Personalisation Rules.

INTRODUCTION

Nowadays, context-dependent adaptation is acquiring increasing importance in many applications and services covering a plethora of different domains. With the advent of the Internet of Things (IoT) its relevance is even more prominent given the increasing number of user environments characterised by the dynamic presence of sensors, objects and devices, according to which applications need to dynamically tailor their behaviour. In addition, applications that are targeted for highly assorted populations, also need to be dynamically customized according to the needs of specific end users. All these aspects are especially relevant in the AAL domain, particularly for solutions targeting the elderly population. Indeed, older adults are a highly heterogeneous target group in terms of familiarity with technology, activity level, level of social relationships, and physical and mental well-being. This translates into highly diverse requirements, making it challenging for developers to design a technology which fits for all of them. In addition, elderly also have specific aging-related requirements that are likely to evolve individually over time. In such environments it could be very difficult for developers to foresee all the possible context-dependent scenarios (and associated customizations to support in the software), because there could be some unanticipated (at design time) need that should be incorporated at runtime, when the application is actually used. Fortunately, the increasing affordability and availability of technology has also promoted new types of participation by end users in the creation process of software, such as End User Development (EUD) approaches [15], to improve the flexibility and acceptability of technological solutions by final users who at some point might want to incorporate new behaviour in their applications. This can be obtained through tools that do not assume specific technical background from their users.

In this paper we present an environment enabling end users to customize the behaviour and appearance of web applications and associated appliances in a context-dependent manner, by using an intuitive trigger-action paradigm. The solution can be applied to remote assistance applications, since in-house monitoring of elderly using intelligent ubiquitous sensors has emerged as a useful AAL service due to its potential of increasing the independence, safety and quality of life of the elderly while minimizing the risks of living alone and avoiding the costs of more expensive hospitalization solutions.

Combining the data originated by the sensors installed in the elderly's house and exploiting the tool for dynamically tailoring applications, both elderly (having some familiarity with technology) and people providing informal and formal care to them (e.g. family members or medical staff) can be empowered to, for example, set up reminders, alarms and messages for promoting a healthy lifestyle, check medication adherence, support monitoring functionalities, so adding new personalization possibilities not foreseen at design time, specified in a context-dependent manner. The personalisable actions supported by the platform range from user interface modifications, to sending messages, to the possibility of changing the state of appliances and devices available in the surrounding context. The resulting tool has also been tested by a set of elderly people and caregivers.

In this paper we mainly focus on the evolution of the editor which is used by end users to define their personalization rules for elderly assistance. The editor has been designed and developed through iterative cycles and by validating it through multiple user tests. After each test, the editor appearance and its user experience has been improved taking into account the feedback received from the participants.

RELATED WORK

In the AAL domain, the emergence of tools for older adults based on ambient intelligence paradigm has been identified and reported [19], together with the need to provide them with personalised services [25]. Previous work in this area include the contribution by Carmien and Fischer [2] who present the Memory Aiding Prompting System (MAPS), an environment in which caregivers can create scripts for people with cognitive disabilities to support them in carrying out daily tasks. Tetteroo et al. [23] propose TagTrainer, a physical rehabilitation technology that supports physiotherapists in the creation of customized rehabilitation exercises for people with neurological impairments.

One further relevant aspect concerns tools helping elders reach their goals with IT products by working with their caregivers. One example of the latter is the work of Zhao et al. [26] who present the CoFaçade approach helping elderly people to reach their goals using digital artefacts by working collaboratively with helpers. In this approach, the elder uses a simple interface having a small number of customizable triggers, which are mapped to procedures that accomplish high-level goals with any IT product. The caregiver uses a customization interface to link triggers to procedures that accomplish recurring high-level goals with IT products. To demonstrate the effectiveness of their approach, the authors implemented a prototype using a handheld physical trigger interface and a desktop customization interface for defining procedures for both computer applications and consumer electronics. While this approach goes in the direction of supporting the elderly in their everyday tasks, differently from our solution they do not consider context-dependent aspects that can modify the execution of procedures.

Also less recent work considered the need of enabling the elderly to customise their applications and environments. In [10] the authors present a tangible interaction technique using magnetic cards to empower the elderly in augmenting their ambient environment with software driven personalized behaviour. In particular, NFC-enabled magnetic cards allow elderly users to create personalized behaviour for their ambient environment and a digital memo board acts as a place holder for active behaviour ensuring the metaphoric resemblance of a memo board and post cards. However, this work is deeply centred on tangible interaction, while we consider Web applications because they can operate through a plethora of different devices.

In [9] the authors provide means to let end users create rule-based smart behaviour through the notion of object augmentation. However, in that work the developers still play a central role since they have to define, implement and install an augmentation module, while end users are expected just to configure the augmentation once it is installed by developers.

The trigger action approach is emerging as a useful paradigm to allow people without experience to connect and integrate devices and services. Lucci and Paternò [16] reported on a user test with three Android apps (Tasker, Atooma, and Locale) supporting it. It revealed that the one able to support most features (Tasker) was the less usable in terms of task performance and errors, thus revealing the need for more careful design of such environments.

IFTTT (IF This Then That, https://ifttt.com/) is a common tool that allows people without programming experience to create simple applications according to the pattern IF <something happens> THEN <do some action>. IFTTT allows for connecting widely used Web services. However, with respect to our approach, which allows for combining multiple triggers, it has lower expressiveness, since it does not allow users to create more structured rules, i.e., those combining multiple events and actions. In addition, previous work [24] found that inexperienced users can quickly learn to create programs containing multiple triggers or actions. Another contribution [1] reported on a test with IFTTT and Atooma and in the end provided some indications for the design of new tools, which included that they should allow the combination of more than one trigger and more than one action in the same rule. This shows that the trigger-action approach seems suitable to support EUD of context-dependent applications, but needs to be improved in order to allow users to express various desired combinations of events and corresponding actions. For this purpose, some authors have considered the 5W model [11], in which they consider five standard questions when specifying rules. This allows users to provide complete descriptions but may make the specification phase too elaborated, especially for people without technical background. Corno and others [6] have explored the use of semantic Web ontologies to support the creation of more complex rules but it is difficult to model them in such a way to be easily understood by end users.

In addition, as highlighted in [13], rule-based approaches, and in particular trigger-action rules, could raise some ambiguity in their interpretation due to potential discrepancies in end users' mental models. In [18] an established theory of mental models is used to guide the design of interfaces for EUD so that people can easily comprehend and manipulate logical expressions. According to such theory, people find it easier to conceptualize logical statements as a disjunction of conjunctions (an OR of ANDs), as opposed to other logically equivalent forms. Thus, the authors propose a paradigm to facilitate the specification of complex logical expressions that however is still far from providing general solutions. Coutaz and Crowley [8] also point out the importance of supporting deployment under real-world conditions, incremental installation of devices and services, and meaningful feedback and feedforward.

In this work, we consider the TARE platform [12], a platform for supporting adaptation of Internet of Things applications. It has been used for tailoring the behaviour of a students' home [5] in order to investigate the feasibility and convenience of using its strategy for tailoring applications, however that work considered young people (students), which is a rather different population from the one considered in this work. In particular, in this paper we present how the tailoring environment has been designed to better support customisation of context-dependent applications supporting elderly based on feedback received in various iterations. Thus, the contribution of the paper is to present the evolutionary design that a technology platform supporting EUD has undertaken –through a number of evaluation cycles– in the specific domain of AAL.

THE PERSONALISATION ENVIRONMENT

Requirements

An initial requirement elicitation activity was carried out [4], in which we addressed mainly two classes of stakeholders: elderly users and informal caregivers. We involved both user groups from the earliest stages of the design phase, so that their views, knowledge and feedback could be fed into the design and development of the platform. For each type of stakeholder separate processes were conducted to understand current practices, as well as opportunities for meaningful customizations in the AAL domain.
We submitted a questionnaire to 71 older adults. The large majority of respondents was in the age group of 65-74 years (64%), followed by 21% in the group between 75-85 years, and 11% between 55 and 64 years. A minority was over 85 years old, none of the respondents was younger than 55 years old. Since the considered sample was mainly composed of active and overall healthy elderlies, on the one hand it came out that continuous monitoring due to health issues is scarcely required. On the other hand, there is the need for not only having features tailored to specific individual abilities, but also to specific lifestyles and living environments significantly came out. Transparency, data control, and social aspects were identified as other relevant user needs.
Two interactive workshops, one survey and three personas were carried out with the aim to scan the types of daily routines and health-related activities which informal caregivers are confronted with daily, to understand how elderly people interact with their caregivers and whether opportunities for personalisation can be identified. In the end, a fully gender-balanced group of eight people aged 55-80 years participated over the two workshops. Caregivers report that elderly people having strong dependency needs require help with their hygiene. Moreover, healthy nutrition (including a sufficient water intake) is another important factor for elderly, who may need support in properly selecting food that gives them the best nutritional value: indeed, with ageing, the metabolism slows down, so while fewer calories are needed then before, at the same time elderly need more of certain nutrients. In addition, informal caregivers highlighted the need of providing support with mild changes in memory and other thinking skills that are common as people age e.g. forget to switch off appliances, especially important for e.g. stoves and/or irons. Moreover, several older people mentioned the need of being helped by a "medication plan": as individuals get older, comorbidity becomes more frequent, thus there is the need of supporting them in properly managing multiple medications. In addition, some older people use to do physical exercises (to stay fit and work against impairments) as well as cognitive ones (e.g. they solve Sudoku and crossword puzzles regularly).

We also analysed the state of the art in the area of AAL solutions for elderly, to identify further user needs as well as key requirements of AAL systems. An interesting work was [21], where authors analysed the issue of older people's trust in various characteristics of AAL technology. Fifty (50) participants aged between 60 and 90 years (average age 71.3 years) were surveyed about the perceived positives and negatives of using technological support in everyday life. The analysis was based on data collected in semi-structured, face-to-face interviews. Considering a 10-point Likert scale, in the study men had distinctly higher levels of trust in sensor technology than women (7.6 vs. 6.8). In addition, people living together with another person showed higher trust values than people living in a single household (7.8 vs. 6.4). Regarding which characteristics formed the basis for trust in technology, reliability and ease of use (both scoring above 9.5) were assessed as the most important aspects, followed by visibility (i.e. the exclusive use of invisible sensors may lead to lower trust levels in AAL) which scored 8.45, whereas costs only slightly influenced (3.90) reliance in AAL technology. In a recent work, Cesta and others [3] report the results of a systematic work devoted to the elicitation and validation of users' expectations on AAL intelligent services. In this work the authors involved about 135 persons among elderly and caregivers for gathering their feedback. Results show that, users appreciated the potentiality of AAL solutions in supporting independent living and in improving quality of life, although privacy concerns have been raised especially by elderly people. The importance to make users understand the meaning and the value of what the system can do clearly emerged, highlighting that in this way a service that is considered useful is also more acceptable. The obtained results have then been used to derive a list of priorities associated to services, which can provide advice for system developers, mainly pointing out the importance of building intelligent technology that can be adapted and personalized according to the varying users' requirements. These requirements can lead to systems that better respond to the individual needs of supporting the elderly live independently in their own homes. Along the same line, in [14] the author proceeded from the assumption that elderly people have for the most negative attitude toward technologies, but also that correctly informing potential users, raising their awareness, and understanding of the usefulness of remote home care are important for elderly acceptance of ICT innovation in home settings. They interviewed 114 users and from their analysis it came out that services must, in the first place, be in line with the needs, habits, desires, and opinions of those who will use the new technologies and related services. Implementation of these innovations must therefore emerge from a model that actively involves users and identifies them as the main actors in this process. To sum up, the need of having solutions able to take into account the wide spectrum of seniors' characteristics, preferences and routines was acknowledged as a relevant need for smart AAL solutions.

An Example Scenario

In this section we describe a scenario in which we envisage a possible use of the considered platform and highlight how seniors can benefit from it. Luisa is 80 years old. She is a retired teacher and lives in a large city. Luisa lives alone since many years, as she lost her husband several years ago. She has been healthy for almost all of her life, but in recent years she got diabetes, then she has to do regular blood tests. In addition, due to diabetes, she is also experiencing some vision issues and then she stopped driving, which further worsened her social isolation and her mood. In addition to such physical impairments, recently Luisa has started to experience some sporadic and tiny memory slips. For instance, she sometimes forgets when specific door-to-door separate collection of waste should be done (specific garbage should be put outside her house in specific days and time intervals) or she forgets turning off appliances and lights in empty rooms. Luisa herself realises that she is somewhat struggling with some cognitive tasks she easily managed to do in the past, for instance e.g. to track monthly bills is taking a bit longer than usual.

Also her daughter Giulia has noticed that her mother is recently showing some subtle signs of cognitive difficulties. Thus, she convinces Luisa to see a specialist, to assess her situation and better evaluate if some kind of interventions could be needed. Luisa acknowledges this. The specialist says that there is no need to worry about this for now, as there is no specific cognitive problem apart from sporadic difficulties connected with normal ageing. Luisa and her daughter feel very relieved about this. However, the specialist suggests using a new system that could support Luisa in her daily activities, provide reminders and notifications, more easily control appliances in her house as well as monitoring some life parameters whose variation can be a sign of the onset of more serious frailty.

Using this platform Luisa can easily manage some daily tasks at her will. For instance, the system can be set up to detect whether Luisa is not performing sufficient physical activity, and motivates her, based on her current status (e.g. physical activity preferences, health status) in performing further physical exercises. In addition, the system is able to guide her while undergoing her physical exercises e.g. by providing short videos demonstrating the exercises suggested by the physiotherapist that he/she can perform at home and can change according to her current health status, following a possible worsening (or improvement) of her conditions. Moreover, the system provides Luisa with evidence about the effectiveness of activity monitoring, by providing her with easily understandable data: getting feedback about her health status often contributes to reassuring Luisa about her current health situation.

In addition, using the suggestions of the specialist, Giulia sets up many rules for better managing lights in her mother's house. For instance, she sets up lights so that Luisa is more exposed to bright light especially during the winter, which should positively improve her mood. In addition, she also sets up another rule that reduces the intensity of lights and changes the light temperature to a warmer value the last two hours before going to bed, to prepare the body of Luisa for a good night sleep. On the contrary, another rule is set up to boost and 'activate' the body at the beginning of the day: from the morning the light gradually rises to high intensity and high colour temperatures (cold light).

The system is able to support Luisa in managing their medications and habits, especially when there are some temporary changes affecting them. For instance, when Luisa has to do the blood test, she has to remember to postpone the usual medication intake (done early in the morning) to after the test. The same happens whenever the doctor decides for a change of her therapy: the system helps Luisa to remember the changes in her intakes. If Luisa has not taken a medicine and it is time to do it, the system, using the Luisa's most preferred device and according to her current position, shows a reminder with the image of the medicine to take, and its expected dosage.

In addition, Giulia sets for her mother a rule for better managing the energy. For instance, when the heater or air conditioner is running, all the external windows and doors should be closed. In addition, another rule is created for managing the air conditioner, aiming at using it only on really hot or humid days. In particular, when a hot day is expected, the air conditioning should be turned on earlier than usual, avoiding to wait until the home becomes too hot. Similarly, another rule will start heating early when a cold day is foreseen. In addition, using the system, it is possible to automatically adjust the temperature according to Luisa's habits and activities currently going on. For instance, if Luisa is doing her physical exercises, the temperature in her room will be decreased a bit. On the contrary, over winter weekends, when the nephews of Luisa visit their grandmother, the temperature is set a bit higher than usual to improve the comfort of the children.

The Personalisation Rule Editor

The research reported in this paper used the TARE, Trigger-Action Rule Environment [12] as starting point, which has been re-designed in various versions developed according to an iterative refinement process that leveraged on a series of evaluations that will be reported afterwards. The tool allows users to define adaptation rules following a trigger-actions paradigm. We have considered it because it is part of a more comprehensive platform, that also includes other two modules, the Adaptation Engine and the Context Manager.

The Adaptation Engine stores and manages personalisation rules, associating them to the available applications and users; it is also responsible for deciding which rule apply when multiple rules are triggered at the same time.

The Context Manager is a middleware module that receives the state of sensors installed in the context and informs the Adaptation Engine when a change in the context would trigger the execution of a rule. When it happens, the Context Server notifies the Adaptation Engine, which extracts the list of actions from the concerned personalization rules, and sends them to the application for interpreting and executing them. For this purpose, the application contains some JavaScript scripts able to understand the requests that correspond to the actions to be done on the application in order to realise the meant personalisation changes.

The starting point for the setup of the Rule Editor for the AAL domain considered is a generic definition of the context model (the set of entities which compose the context of use), which includes common attributes that can be shared between all application domains (e.g. personal information such as age, gender, education, environment attributes such as temperature, humidity, light level, etc.), thus it is domain-independent. Then, it has been customised with the support of domain experts (caregivers), to identify domain-specific context entities that will compose the triggers. Below are listed some of the context attributes we identified as relevant for the AAL domain:

- Physical information: walking ability, heart rate, daily steps, body temperature, respiration rate, posture, weight);
- User position inside the house;
- Cognitive information: attention, memory, language;
- Medication: planned and occurred (medicine name, dosage, notification time);

- Motivation (wellness, fitness, health, social);
- Environment attributes (temperature, humidity, gas presence, motion, light and noise level);
- State of the devices and smart physical objects installed in the elderly's home;

All the context attributes compose the domain-specific context model and are available in the Personalization Rule Editor for specifying personalization rules.

When an instance of the platform is created for being used in a specific environment, it needs to be customised according to two aspects: the considered context of use, and the application to personalise. For the context of use, the customization depends on the available information and sensors, because they will provide the input to determine whether triggers can be fired. The application is considered to customise the possible actions that can be actually triggered. For our trials we considered a number of sensors. For such sensors we developed some software (context delegates) that take as input the data from sensors and communicate such data to the Context Manager server. For our user tests we have used the PLUX BITalino chest band (http://bitalino.com/en/), which includes a number of bio-signal sensors, so providing data associated with heart rate, number of steps, current position of the user (including e.g. 'supine', 'prone', 'standing'), respiration rate and body temperature. We also use a Fitbit Charge 2 fitness wristband, which provides information about physical activity of the user (steps taken, distance covered). Furthermore, we exploit a number of Estimote Proximity beacons to derive the user position. The Proximity delegate application defines three proximity zones: immediate, near, and far, according to the strength of the signal received from the beacons (from a few centimetres to some meters) and it informs the Context Manager that the user is inside a room (each beacon is associated to a room). The configuration of the environment sensors is composed of an Arduino Uno (enhanced by an Ethernet shield to provide the Arduino board with Internet access) which acts as a master since it receives the data from multiple Arduino Micro (Nodes) connected to it and then it sends those data to the Context Manager. Each Arduino unit is equipped with a NRF24L01 module, which is a transceiver that allows to make a wireless communication between the Arduino micro and the master. As for the nodes, we have an Arduino-compatible sensor for gas (MQ5), another one for detecting temperature and relative humidity (DHT11) and one for detecting motion.

In the context of the scenario previously described, when Giulia wants to set up a rule for her mother, the Personalization Rule Editor should first load the relevant context model, which provides her with a representation of relevant contextual elements (e.g. events and conditions) that can be associated with the triggers that she wants to consider in her rules. For instance, for the personalization rule aimed to improve Luisa's emotional state by acting on lights, the relevant *triggers* would be the current level of illumination in the house, the time of the day and the current emotional status of Luisa, while the relevant *actions* would be acting on some parameters of lights (e.g. colour, temperature).

In particular, *triggers* refer to elements identified in the contextual model. They are organized according to a logical hierarchy, by grouping together related context entities. At the highest level, the hierarchy of triggers consider the following aspects: user, environment, technology, social aspects. *Actions* can be associated with: appliances commands (to change the state of some actuator); UI modifications (to change the presentation, content or navigation of the application UI); functionalities (to access externals services e.g. weather forecast service); alarms (to highlight some potentially dangerous situations); and reminders (to indicate tasks that should be accomplished).

After Giulia finishes specifying the rule through the editor, the rule is sent to the Adaptation Engine, a software service that subscribes to the Context Manager for being notified about the occurrence of events associated with the rules received in the current context. Afterwards, when the application is activated, it subscribes to the Adaptation Engine in order to receive the adaptation actions when the rule is triggered. When the application receives from the Adaptation Engine the actions, through some JavaScript scripts included in it, it is able to understand the requests corresponding to such actions and then realise the meant personalisation changes.

An Example Application

One of the applications that we have developed for our trials addresses mainly elderly people with a good degree of independence, with the goal of improving their quality of life.

In particular, the web application (see Figure 1) allows the elderly to control the home environment. Moreover, it supports remote monitoring, thus it includes a number of functionalities providing information about e.g. elderly's health-related data, wellbeing goals settings, planning of the activities, personal profile and main contacts, environment sensors status. In addition, within the application it is also possible to control some lights installed in the elderly's home. In particular, the lights considered were Philips Hue lamps equipped with Philips Hue Bridge 1.0. The Bridge is connected both to the Internet Network (Ethernet connection) and to the lights (ZigBee connection), thus it allows to remotely control the light system. In addition, real time values of steps performed, heart and respiration rate, user body position and temperature are detected through BITalino sensors embedded in a chest band worn by the older person. Weather and news contents are incorporated as RSS feed from external providers based on location information.

The original application had limited possibilities in terms of personalization. We want to show how the approach proposed allows end users to dynamically change its possibilities by editing relevant rules. The goal is to dynamically adapt the application according to rules, e.g. by changing the state of appliances when some events occur, or by receiving motivational messages for improving the current lifestyle of the elderly when an unhealthy behaviour is detected, so as to reinforce the motivation of specific users and to effectively encourage them to exercise.

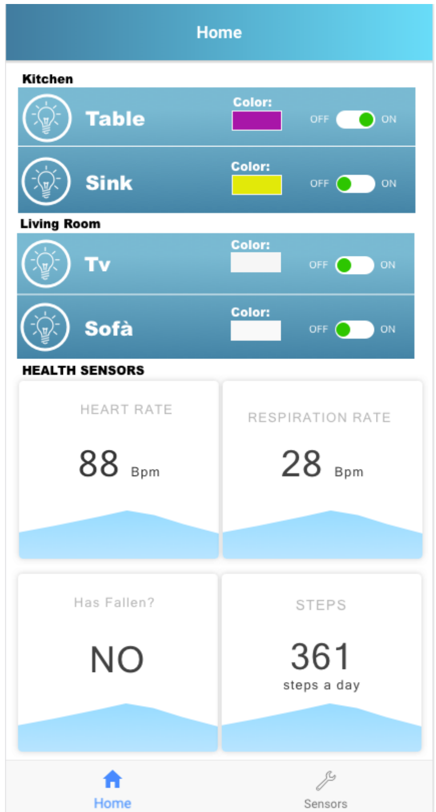

Figure 1:Remote monitoring application

Adaptation through Personalization Rules

In this section we report examples of the personalization rules defined with stakeholders and created using the Personalization Rule Editor.

- If the notification time of a planned medication expired and it does not exist a corresponding occurred medication, send the elderly a reminder (also visualising the name, shape and colour of the pill to take)
- When the user sits on the sofa in the afternoon, turn on the music player and change the colour of the ambient lighting to green to make the elderly relax
- If the user has not yet completed the number of physical exercises planned for that day send her a reminder encouraging her to finish the exercises
- If motion in the corridor is detected during the night switch on the lights
- If the steps performed during the last week are less than the goal steps and the weather is sunny and the TV in the living room is on, send a motivational message on the TV inviting the user to exit and have a walk
- When the number of steps has decreased of a specific percentage over the last month compared to the previous month send a message to the caregiver and send to the elderly an encouraging message inviting the elderly to do more exercise
- If an appointment is scheduled in a few hours and the weather forecasts are good and the destination is at elderly's house walking distance, suggest the elderly to have a walk
- When the elderly enters the living room in the afternoon, set ideal temperature in the living room and switch to an energy saving mode in the other rooms
- When there are only two hours left before going to bed (sleep scheduled time minus current time is less than two hours), set the light colour to red (the red colour is the least disruptive for the sleep quality because it has the lowest wavelength).

USER TESTS

As we used an iterative approach in which multiple evaluation sessions informed the design, in this section we report on the evaluations that we carried out on the rule editor prototype, which was subject to multiple refinements through a number of cycles.

First User Test

Figure 2 shows the initial version of the Rule Editor that was assessed during the first user test. It was carried out with a sample of 7 participants that included three older people (aged 74-80) and four informal caregivers (aged 45-67). A laptop PC (Lenovo Z570) with a 15.6-inch display was used. People were recruited by a Swiss foundation operating as a representative body for mature people and as a service provider in the market for the elderly [4]. The test session was divided into two main parts. First, the users had to rate the exhaustiveness of the way in which the user context of use and the possible actions were modelled. During the first part of the test users were asked (on a 1-7 scale; 1=very bad; 7=very good) about the exhaustiveness of the set of triggers that can be specified using the tool (min:3, max:6, median:5), and of the set of actions that can be specified (min:4, max:5, median:5).

The second part focused on the usability of the rule editor. The test leader gave the users some simple rules written in natural language (in German), covering the majority of the contextual aspects of the Personalization Rule Editor. The participants had to specify such rules using the tool. The same Likert scale as before was used to rate some aspects of the Personalisation tool, for which the following ratings were obtained:

- Usability of the action selection mechanism (min:2, max:6, median:4)
- Usability, in general, of the rule-based approach (min:3, max:6, median:4)
- Exhaustiveness of the set of events that can be specified (min:3, max:6, median:5)
- Exhaustiveness of the set of actions that can be specified (min:4, max:5, median:5)
- Usability of the tool support for reusing previously saved rules (min:3, max:7, median:5)
- Usefulness of describing the rules in natural language (min:6, max:7, median:7)

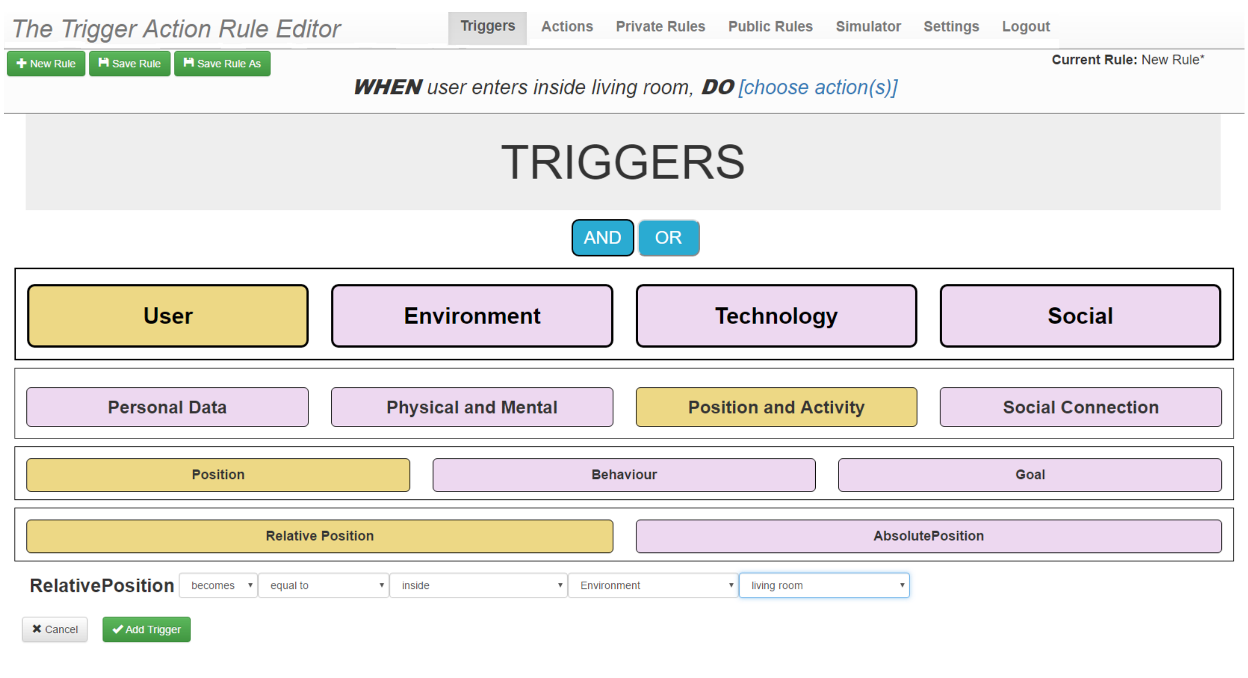

Figure 2: First version of the Rule Editor

The results of the usability test show that the structure of the application was clear and well-structured at a first glance. However, once deep in the tree structure, users might get lost. For instance, some of them reported having difficulties to remember where they can find the next steps, because the whole structure is no longer visible. To overcome this problem, we developed a search function helping users to more easily find a specific hierarchy element without exploring the whole tree (of triggers). In addition, users could not easily remember where to find some functions of the tool. Other times users could not use the system properly due to e.g. too small font sizes, or because the selection fields were too small.

Furthermore, some fields associated with a trigger or an action needed to be filled in with free values. However, users preferred having UIs that provide more control over data values and types: for instance, instead of directly editing values, they would prefer selecting values through drop-down menus, or virtual number keypads if numbers are needed.

Most users found it difficult to get the whole overview and understand the full potentiality of the system by just doing a few exercises. For further development, a more consistent and clear structure of the UI was suggested: in particular, consistent use of buttons, symbols and colours is important for smooth handling by end users (e.g. consistently use green colour for “ok” or “done” and red colour for “not already done” or “does not work”). During the tests, users reported some difficulties in navigating the hierarchy representing the Context Model, especially when it is unfolded in multiple levels. As it can be seen in **Errore. L'origine riferimento non è stata trovata.**, the only difference between folded and unfolded elements was the background colour; some users reached the lowest level in the hierarchy without finding the desired element, at this point they did not understand that it is possible to unfold other elements on top of the current selection. In order to overcome this usability issue we differentiated the folded and unfolded elements by adding respectively a “-“ and a “+” symbols beside each element (see Figure 4).

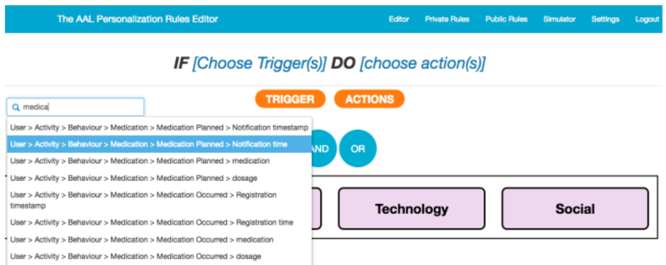

Figure 3: The new support for identifying context attributes, identified as result of the first test

As a result of this study a new user interface for the rule editor has been designed and implemented. It supports users in more easily finding the relevant aspects of the context of use in order to specify the triggers in the personalization rules (see Figure 3).

For this purpose, a specific search functionality was added: after specifying the desired concept in the text field, the tool shows the paths of the elements where such a concept is considered within the context model. Then, the user can easily select the most relevant one amongst those listed. When the user selects the desired concept, the user interface shows it in the logical structure of the trigger classification, and provides users with the possibility to edit the corresponding attributes (see Figure 4).

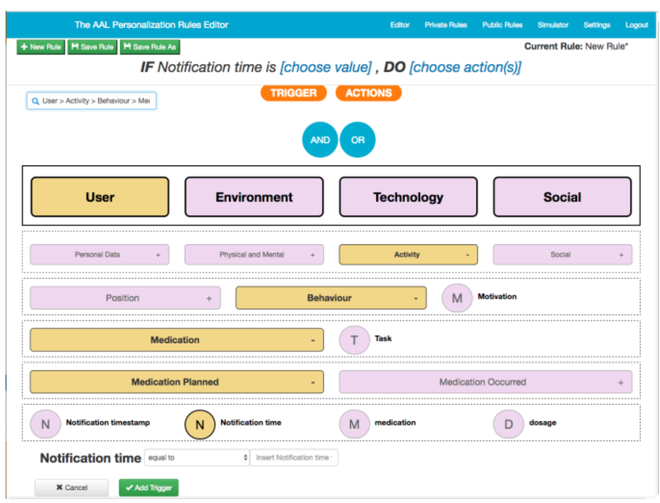

Figure 4: The selected element is presented in the context structure

After including such improvements in the Rule Editor, the tool underwent an additional test in order to assess how the new version was perceived by users.

Second User Test

Participants

Twenty participants were involved in this second test. They were recruited by the institution moderating the test, Sunnaas, a Norwegian hospital specialised in the field of physical medicine and rehabilitation, by placing an advertisement in a newspaper, asking for people over the age of 65. Participants were expected to know how to interact with a browser, either using a tablet or a personal computer, but no further expertise was required. The participants' responsibilities were to complete a set of representative task scenarios in an as efficient and timely manner as possible, and to provide feedback regarding the usability and acceptability of the Rule Editor. The participants were directed to provide honest opinions regarding the usability of the application, and to participate in post-session subjective questionnaires and debriefing. Participants were not previously aware of the application being tested. Regarding the age distribution, almost all the respondents belong to the age group of 65-75 years (95%). A single participant was 90+. No respondent was younger than 65 years old. Regarding gender, the distribution was overall balanced: 55% of the respondents were male, 45% female. When asked about health condition, all the participants reported that they did not suffer from any severe impairments. All the participants characterized themselves as older adults, without having a caregiver role.

Procedure

Participants took part in the usability test at Sunnas Rehabilitation Hospital in Oslo. The tasks were completed using a laptop (ASUS ZenBook UX305, Intel(R) Core(TM) i7-6500U CPU @ 2,50GHz, 3200x1800, Windows 10 Home, Firefox Browser). The participant's interaction with the applications was monitored by a facilitator seated in the same room. Note takers and data logger(s) also monitored the sessions. The facilitator briefed the participants on the application. Participants signed an informed consent that acknowledged that the participation is voluntary and it can cease at any time, and their privacy of identification will be safeguarded (i.e. the performance of any test participant must not be individually attributable). Participants completed a demographic pre-test and a background information questionnaire. At the start of each task, the participant reads aloud the task description from a printed copy and begins the task. The facilitator instructed the participant to 'think aloud' so that a verbal record exists of their interaction with the application. The facilitator observed and entered user behaviour, user comments, and system actions in the data logging application. After each task, the participant completes the post-task questionnaire and elaborates on the task session with the facilitator. After all task scenarios are attempted, the participant completes the post-test satisfaction questionnaire for each application**.** Before carrying out the assigned tasks, no specific training was provided to participants, just an overview about the test procedure, equipment and the Rule Editor. The concepts that were required to understand the tasks to be executed were explained as necessary before the corresponding tasks. The test was organized in two main stages: a pre-test questionnaire and the real interaction with the rule editor. In the pre-test phase, for each participant we collected their age, gender, if they are a caregiver for a senior, if they have any impairments that might impair the use of the application, and, in case they do, which impairment. Then, after an introduction to the editor, the participants were asked to carry out three tasks concerning creating some rules with the Rule Editor. The difficulty and effectiveness of each task was registered. Workload was measured with the NASA TLX questionnaire.

Tasks

Due to the short time in which each participant was available, we selected tasks with low complexity. Data introduced by the participants during the test was deleted after the test was completed to make sure that all users begun their test in the same state.

Task1: Users were required to create a rule having a simple trigger. In particular, they had to create a rule (Rule1) to turn on the living room lights when the user is inside the living room, then save the rule. This task required to specify a trigger involving a condition.

*Task2***:** Create a rule to turn on the TV when the user enters the living room (Rule 2), then save the rule. This task required to specify a trigger involving an event.

*Task3***:** This task was composed of two parts:

Sub-Task3.1: Create a rule that increases the font size when the user is outside the home. Save the rule.

Sub-Task 3.2: Go back to the Remote Monitoring Application. Imagine now that you are seeing the application in a tablet and you go outside your home. [Moderator triggers rule to change the font size, from 16px to 21px, a 30% increase]

Results

Task1 (Create a rule having a simple trigger involving a condition)

As for Task1, the interviewees were asked how difficult the completion of the task was (Figure 7). A clear majority (15 interviewed persons) reported that the task was easy or very easy. Only one person rated the task as difficult. However, during the completion of the task, 80% needed help. At the same time, the error rate was high, as 50% of the respondents could not complete the task without errors. However, it turned out that the participants did not notice the possibility to search for triggers.

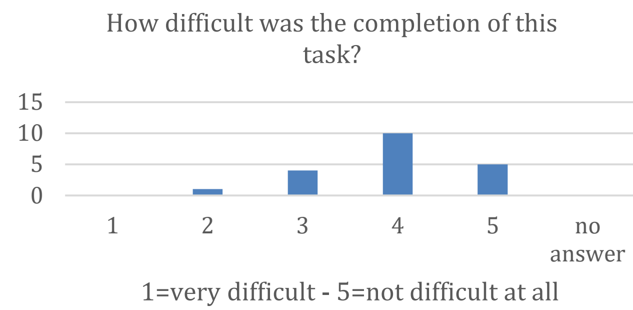

Figure 5: Perceived difficulty of completing Task1

Task2 (Create a rule having a simple trigger involving an event)
As for Task2, the majority of the respondents rated this task again as easy or very easy (14 persons). One respondent considered the task as difficult (see Figure 8). About 63% of the interviewees needed help with the completion of the task and only 26% could complete it without error. Still, these number represent an improvement over the previous task, as can be expected.

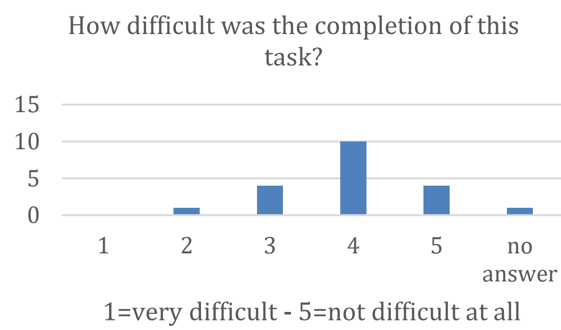

Figure 6: Perceived difficulty of completing Task2

As for the usefulness of the provided Rule2, the majority of the test persons think that such a rule is very useful or useful. Only 3 persons rated such a rule as not useful. When asked about the difference between "is" and "becomes", it turned out that only 41% understood the difference, but almost all participants (85%) managed to use it correctly.

Task3 (Create a rule changing the aspect of the target application and then simulate the rule execution)

As for the first subtask of Task3, overall the respondents did not find it complicated. The majority of 12 persons considered the task as not difficult. However, three persons considered the task as difficult, which is an increase if compared to the previous two tasks. 72% of the respondents could not complete the task without help, but also 72% were able to complete it without errors. As for the results of subtask2 of Task3: most of the respondents (71%) did notice the change in the font size. The majority of 14 respondents considered such a rule as useful or very useful. A minority of 2 persons rated it as not useful.

Evaluation of the Cognitive Workload (NASA TLX)

After completion of the tasks, participants answered the NASA-TLX questionnaire, which consists of six subscales: Mental, Physical, and Temporal Demands, Frustration, Effort, and Performance, with the assumption that the combination of these dimensions can represent the "workload" experienced by most people performing most tasks. As for the evaluation of workload, all items are generally considered not demanding. Nevertheless, some respondents reported a high mental workload, consistent with the higher complexity of this application. This can be expected due to the nature of the application, and the coverage it entails (rules can be created to deal with many situations, ranging e.g. from appliance and home control, to user interface aspects). The overall raw (non-weighted) NASA-TLX score was 36,3, which is on the low end of the scale, representing a low requested workload.

Further Qualitative Remarks

The Rule Editor mainly targets caregivers, and in the second place elderly. From the gathered data and the overall perception of the user tests we can say that, although the usability of the Rule Editor can be improved, interacting with the Rule Editor does not require a high workload. It is also worth pointing out that the participant sample of this second test is mainly characterized as active older adults (represented by the majority of the participants), who are not the primary target users of this application. Nonetheless, they were also able to perceive the usefulness of the application and to use it to set simple rules. Active older adults could understand the Rule Editor concepts and use it to set rules that can adapt the user interface of the considered applications and automate tasks in the older adults' homes and in the applications they use (e.g. setting reminders for taking medication). One of the biggest concerns of the users was about user orientation within the tool: users complained about the need of considerably using the scrolling within the Rule Editor (which was often required to navigate within the hierarchies of triggers and of actions). Thus, long scrolling needed for navigating the hierarchies of triggers and actions (by e.g. folding elements) revealed to be a bit problematic for users because when they were deep within the page they could missed the top part of the hierarchy), so losing important information about the current context.

In addition, users had difficulties in understanding the point (and the progress) currently achieved in the process of building a rule (e.g. what they have already specified and what still remained for completing the specification of a rule, e.g. "it is difficult to understand what I've done and what should I do next"). Moreover, as for the user interface of the tool, users complained about the use of too technical terms, some of them rather resembling a "programming language" style, with concepts difficult to understand. Thus, they asked for adopting a more user-friendly language in the tool user interface so as not to intimidate users. Some participants also suggested having definitions in the text available.

Some of them complained that it was not immediately intuitive how some elements were organised in the hierarchy (e.g. finding "position" under "activity"). Moreover, some of them said that saving was difficult to find. Finally, some of them highlighted the need of expressing, within triggers, events and/or conditions that do not hold, as – they said – many times, especially in remote monitoring applications, checking whether some events have *not* occurred could be particularly relevant e.g. for detecting abnormal situations.

The New Tool Design

We redesigned the tool to address some of the issues that were identified in the last usability test. In particular, we included a sidebar (see left part of the user interface shown in Figure 7) to improve user's awareness of the progress achieved in creating rules. From the sidebar it is possible to see what the user has already done in creating a rule, and what else should still be done (see Figure 7). In particular, in Figure 7, on the one hand the "Actions" label has a green background with a tick sign to render that at least an action has been completely specified in the rule, thus it could be considered completed. On the other hand, the "Trigger" part has a grey background colour to render that the specification of triggers is incomplete, since at least one trigger needs to be defined. Furthermore, from the sidebar it is possible to directly carry out some actions supported by the tool (e.g. modify and delete rules). We also moved the button for saving rules in a more visible position in the UI. Figure 4 shows the earlier version of the tool, where the "Save" button was placed in a part of the UI that was judged not very visible by users, as it was included in a menu in the top-left part of the rule editor. The "Save" button (see Figure 7) is disabled or enabled according to the current state reached in specifying the rule: as soon as a user specify all the main parts of a rule (e.g. both the trigger part and the action part), the "Save" button will be enabled, otherwise it will be disabled. The same applies to

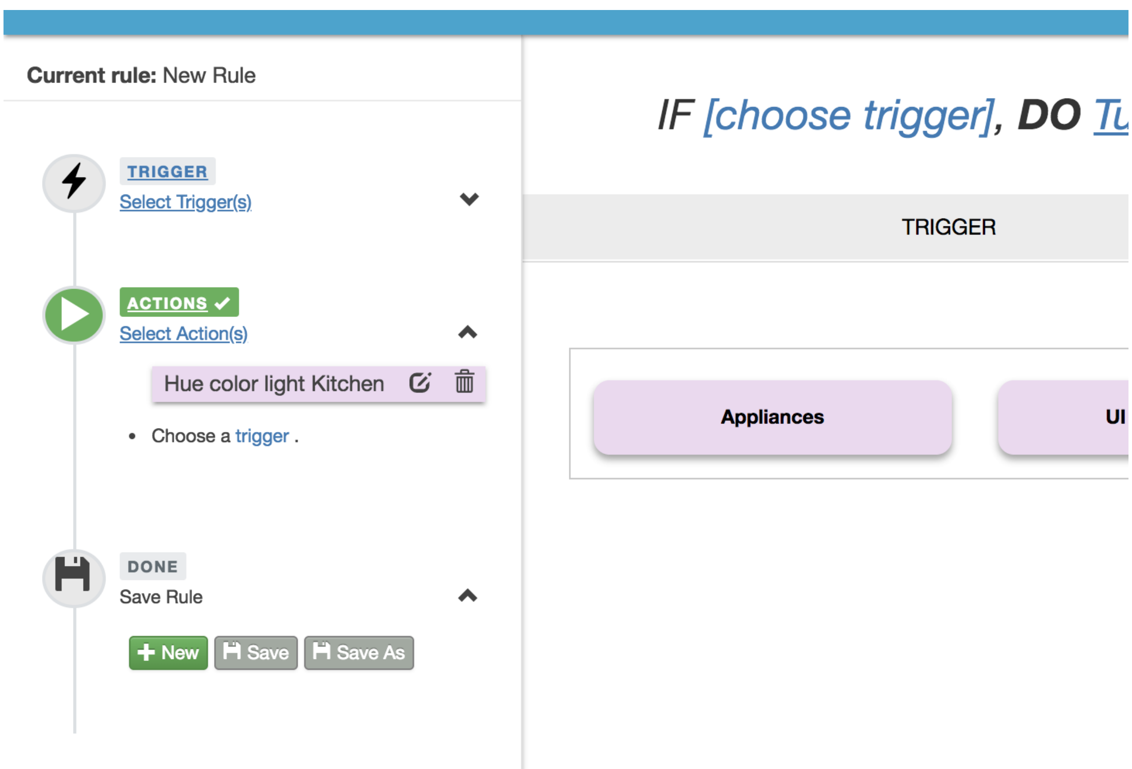

Figure 7: Sidebar added after the first User test

Moreover, we improved the style used in the natural language description. Indeed, in the new version of the tool we exploited a simpler, more conversational and less technical language to better cater for different user skills/literacy. For instance, in the natural language description of rules, instead of using the technical expression "is equal to" we just used "is"; instead of "Please Set Value" we just used "choose value". Furthermore, still in the natural language section, instead of expressing the sentence representing the rule following the pattern "trigger name + operator + value" (e.g. when the user posture becomes equal to standing), we used some heuristics to provide a clearer, more conversational/natural style of the language used for describing the rule. Thus, the previous sentence becomes "When the user stands up".

In addition, we further checked that a sensible use of colour was followed (i.e. we changed the colour for highlighting the current elements selected in the visualized hierarchy, opting for a brighter yellow (see the difference between **Errore. L'origine riferimento non è stata trovata.** and Figure 7) and further providing an even more simplified design of the UI. In particular, we redesigned the tool in such a way that the need of vertical scrolling (which was activated e.g. when the user navigated down of several levels in the hierarchy) would be highly reduced in the new version. With the new tool, the user interface automatically folds the visualization of the tree (of triggers or of actions) as soon as a trigger (respectively: an action) is added to the current rule, in order to reduce the visualization of unnecessary information and dynamically organize the UI so that at any moment the number of UI elements is limited only to the needed ones. Along the same line, there is the possibility for the user to fold the various sections of the sidebar, if the user does not want or does not need them and then he wants to hide them from her current view. Moreover, in the previous version of the tool the part dedicated to editing the attributes of the selected trigger was placed under the tree representation of the trigger hierarchy and this was a bit burdensome because the user was forced to continuously scroll the page (see Figure 2 and 4). In the new version of the Rule Editor the parameters editing of the selected trigger is supported through a modal dialog that appears on top of the main window (thus avoiding the scroll), and temporarily disables the main window (see Figure 8).

In addition, we supported users in better controlling the user interface, by providing more suitable UI widgets for more effectively supporting specific tasks and prevent errors (e.g. in Environment > Date and in Environment > Time, now we have input fields of type respectively "date" and "time"). We also changed some elements in the hierarchy of triggers, to better reflect their content. For instance, the former "Activity" trigger was renamed into "Activity and Position" trigger since this category also included triggers activated according to user position.

Moreover, during the tests, some users complained that it is only possible to define triggers describing if or when something happens: they suggested that it would be useful to support the possibility of specifying rules that should be triggered when some associated events or conditions do *not* occur in a specific time interval. Thus, we introduced the NOT Boolean operator, to be used in the specification of triggers. An example of its use is visualized in Figure 8, which shows the editing of the rule IF user does not stand up between 8 and 9 am D= send a message to caregiver and turn on the bedroom light.

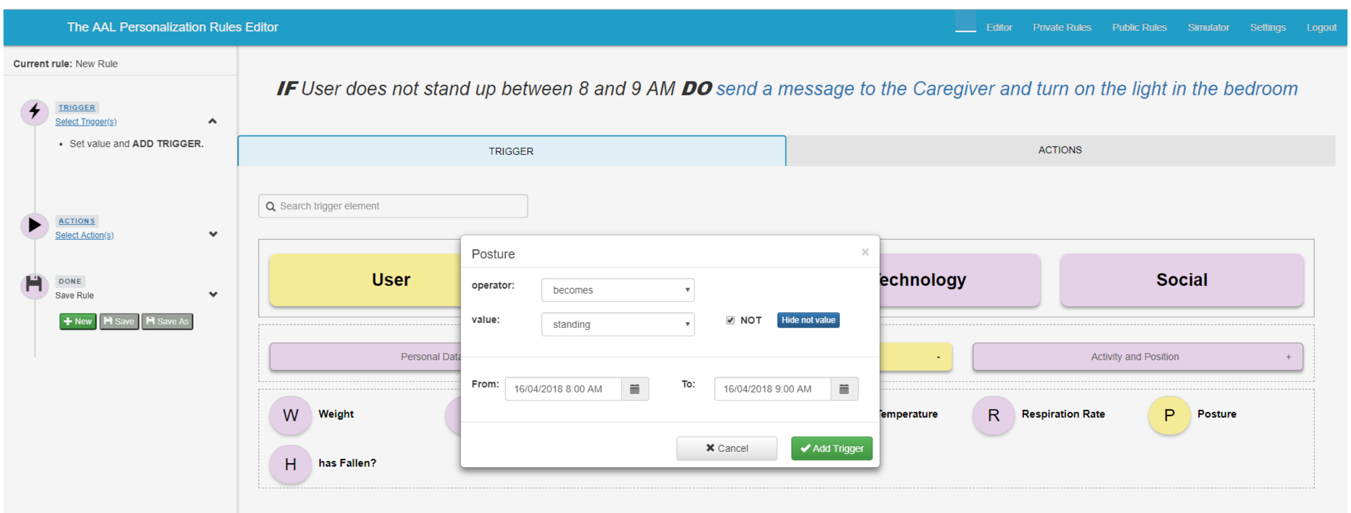

Figure 8: The new version of the tool, with an example of use of the NOT operator

In this example, the user wants to specify the situation when a user does not stand up within a particular time interval. First of all, the user has to specify the event which models the change of the user posture (e.g.: *When user posture becomes standing*) and then she adds the NOT operator which allows to define the time interval in which the system has to check if the event occurred or not. If the event does not occur within the specified time interval, then the rule is triggered: a message is sent to the caregiver and the bedroom light is turned on.

Implications for Design

We would like to derive some more general implications for good design of EUD tools in AAL scenarios.

First of all, tools should use a simple and conversational language referring to users' familiar concepts, so as to cater to various user backgrounds and education, and avoid intimidating and confusing less tech-savvy persons with complicated and convoluted expressions. As a consequence, the representation of intended behaviour (the rules in our case) should be easily understandable and visualized in a clear way so that even non-professional users can make sense of it. Connected to this point there is the need of involving end users when structuring the logical organisation of the elements referred by triggers, so that not only such terms refer to a vocabulary familiar to users, but also the relationships between elements make sense to them: this will result in facilitating users in finding relevant triggers within the logical structure proposed by the tool. Moreover, by using the EUD tool, end users are asked to assume responsibilities that are traditionally intended for developers in that they can take control and co-create solutions that fit their own needs and contexts. Thus, the tool needs to be designed in a way that allows end users to be effectively included in the software lifecycle and actively participate in the application development with full awareness of the supported features and the impact that they can have in the real environment. Therefore, it is imperative that users feel in control of the application (e.g. be aware of the point currently reached in the interaction process), which could also imply the need to use multiple, even redundant representations catering to different users and different tasks to convey the same key information. Although the workflow in editing the personalisation rules is not particularly complex, EUD tool users without a technical background may still need some support to better understand what the relevant options are to progress towards their completion at any given step. Moreover, especially in the AAL domain, the user interface of the EUD tool should guide the user in capturing the most relevant information, thus restricting users' attention only to relevant data in order to limit the cognitive effort and avoid distracting them by unnecessary information by which they can feel overwhelmed. In addition, it is useful to provide users with the possibility to specify triggers not only associated with situations that occur, but also with situations that in some period of time do not occur. Finally, the tool design should encourage progressive exploration of used concepts and it should support quick learning. Users should be able to specify needed behaviour with minimal training, and be able to easily incorporate the technology into their everyday lives at their own pace.

CONCLUSIONS AND FUTURE WORK

In this paper we present the iterative design and development of a personalisation rule editor for end user development of context-dependent applications, which has been customized for applications aiming at remotely monitoring and assisting elderly at home.

After an initial test with some elderly people and caregivers the platform was redesigned to address their main concerns. The new version was then assessed in a new test done with only elderly, which informed a redesign of the solution, expected to solve the identified usability issues. As future work we plan to validate the new version of the system in longitudinal studies in the elderly homes to investigate how the use and appropriation of the tool vary over time.

ACKNOWLEDGMENTS

This work was partly supported by the Ambient Assisted Living Project PersonAAL (http://www.personaal-project.eu).

REFERENCES

[1] Federico Cabitza, Daniela Fogli, Rosa Lanzilotti, Antonio Piccinno: Rule-based tools for the configuration of ambient intelligence systems: a comparative user study. Multimedia Tools Appl. 76(4): 5221-5241 (2017)

[2] Stefan Parry Carmien and Gerhard Fischer. 2008. Design, adoption, and assessment of a socio-technical environment supporting independence for persons with cognitive disabilities. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08). ACM, New York, NY, USA, 597-606. DOI: https://doi.org/10.1145/1357054.1357151

[3] Amedeo Cesta, Gabriella Cortellessa, Francesca Fracasso, Andrea Orlandini, Marcello Turno: User needs and preferences on AAL systems that support older adults and their carers. JAISE 10(1): 49-70 (2018)

[4] Cristina Chesta, Luca Corcella, Stefan Kroll, Marco Manca, Julia Nuss, Fabio Paternò, Carmen Santoro: Enabling Personalisation of Remote Elderly Assistant Applications. CHItaly 2017: 6:1-6:9

[5] Luca Corcella, Marco Manca and Fabio Paternò. 2017. Personalizing a Student Home. Proceedings of the 6th International Symposium on End User Development (IS-EUD'17). 18-33. http://dx.doi.org/10.1007/978-3-319-58735-6_2

[6] Fulvio Corno, Luigi De Russis, Augusto Monge Roffarello, A Semantic Web Approach to Simplifying Trigger-Action Programming in the IoT, Computer 50 (11), 18-24, 2017.

[7] Fulvio Corno, Luigi De Russis, and Augusto Monge Roffarello, 2017. A High-Level Approach Towards End User Development in the IoT. In: CHI 2017: The 35th Annual CHI Conference on Human Factors in Computing Systems, Denver, CO (USA), May 6–11, 2017. pp. 1546-1552

[8] Joelle Coutaz and James Crowley, 2016. A first person experience with end-user development for smart home. IEEE Pervasive Computing, vol. 15, no 2, May-June 2016: 26:39

[9] Johan Criel, Laurence Claeys, and Lieven Trappeniers. 2011. Deconstructing Casensa: The CAEMP context-aware empowering platform. Bell Lab. Tech. J., 16, 1 (June 2011), 35-53.

[10] Johan Criel, Marjan Geerts, Laurence Claeys, and Fahim Kawsar. 2011. Empowering Elderly End-Users for Ambient Programming: The Tangible Way. International Conference on Grid and Pervasive Computing (GPC'11). 94-104. http://dx.doi.org/10.1007/978-3-642-20754-9_11

[11] Giuseppe Desolda, Carmelo Ardito, Maristella Matera, Empowering End Users to Customise their Smart Environments: Model, Composition Paradigms, and Domain-Specific Tools, ACM Trans. Comput.-Hum. Interact. 24(2): 14:1-14:33

[12] Giuseppe Ghiani, Marco Manca, Fabio Paternò, and Carmen Santoro. 2017. Personalization of Context-Dependent Applications Through Trigger-Action Rules. ACM Trans. Comput.-Hum. Interact. 24(2): 12:1-14:52

[13] Justin Huang, and Maya Cakmak. 2015. Supporting mental model accuracy in trigger-action programming. In Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15). ACM, New York, NY, USA, 215-225.

[14] Boštjan Kerbler (2018). Using Information and Communication Technology in Home Care for the Elderly, Caregiving and Home Care, Prof. Mukadder Mollaoglu (Ed.), InTech, DOI: 10.5772/intechopen.72083. Available from: https://www.intechopen.com/books/caregiving-and-home-care/using-information-and-communication-technology-in-home-care-for-the-elderly

[15] Henry Lieberman, Fabio Paternò, Markus Klann, Volker Wulf. 2006. End-User Development: An Emerging Paradigm. In End User Development, Henry Lieberman, [7] Fabio Paternò, and Volker Wulf (eds.). Springer, The Netherlands, 1-8.

[16] Gabriella Lucci, F Paternò, Understanding end-user development of context-dependent applications in smartphones, International Conference on Human-Centred Software Engineering, 182-198

[17] David Markland, and Lew Hardy. 1993. The Exercise Motivations Inventory: Preliminary Development and Validity of a Measure of Individuals' Reasons for Participation in Regular Physical Exercise: Personality & Individual Differences, 15, 289-296

[18] Georgios Metaxas, & Panos Markopoulos (2017). Natural Contextual Reasoning for End Users. ACM Transactions on Computer-Human Interaction, 24(2), 1–36.

[19] Parisa Rashidi and Alex Mihailidis. 2013. A Survey on Ambient-Assisted Living Tools for Older Adults: IEEE Journal of Biomedical and Health Informatics, 17, 3, 579-590

[20] Jeff Sauro, Joseph S. Dumas. 2009. Comparison of Three One-Question, Post-Task Usability Questionnaires. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09). 1599-1608. http://dx.doi.org/10.1145/1518701.1518946

[21] Frederick Steinke, Tobias Fritsch, Daniel Brem, and Svenja Simonsen. 2012. Requirement of AAL systems: older persons' trust in sensors and characteristics of AAL technologies. In Proceedings of the 5th International Conference on PErvasive Technologies Related to Assistive Environments (PETRA '12). ACM, New York, NY, USA, , Article 15 , 6 pages. DOI=http://dx.doi.org/10.1145/2413097.2413116

[22] Pedro J. Teixeira, Eliana V. Carraça, David Markland, Marlene N. Silva and Richard M. Ryan. 2012. Exercise, physical activity, and self-determination theory: a systematic review: Int J Behav Nutr Phys Act., 9, 78, 1-30

[23] Daniel Tetteroo, Paul Vreugdenhil, Ivor Grisel, Marc Michielsen, Els Kuppens, Diana Vanmulken, Panos Markopoulos: Lessons Learnt from Deploying an End-User Development Platform for Physical Rehabilitation. CHI 2015: 4133-4142

[24] Blase Ur, Elyse McManus, Melwyn P. Yong Ho, Michael L. Littman. 2014. Practical Trigger-Action Programming in the Smart Home. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14). 803-812. http://dx.doi.org/10.1145/2556288.2557420

[25] Dimitrios D. Vergados. 2010. Service Personalization for Assistive Living in a Mobile Ambient Healthcare-Networked Environment. Personal and Ubiquitous Computing, 2010, 14, 6, 575-590.

[26] Jason C. Zhao, Richard C. Davis, Pin S. Foong, and Shengdong Zhao. 2015. CoFaçade: A Customizable Assistive Approach for Elders and Their Helpers. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '15), 1583-1592. http://dx.doi.org/10.1145/2702123.2702588