

Controlling Security Rules Using Natural Dialogue: an Application to Smart Home Care

Stanley Goffinet, Donatien Schmitz, Igor Zavalyshyn, Axel Legay, Etienne Rivière

ICTEAM, UCLouvain, Belgium

{stanley.goffinet,donatien.schmitz}@student.uclouvain.be,{igor.zavalyshyn,axel.legay,etienne.riviere}@uclouvain.be

ABSTRACT

Smart home systems have revolutionized the way we interact with our living environment, but the concerns over sensor data privacy and security have become one of the major barriers to their widespread adoption. Despite a considerable research effort in designing secure access control mechanisms, the end users are still reluctant to use those either due to the complexity of the interfaces or because they do not have sufficient skills. For the elderly users the problem gets worse. Making the right security choices is increasingly more difficult for this group of users. To assist the elderly users in defining their access control policies, we design a dialogue-based system which allows to create new security rules or update existing ones in a simple and intelligible way using natural language and familiar terms.

CCS CONCEPTS

• Security and privacy → Access control.

KEYWORDS

privacy, smart home, elderly, access control, voice assistant

ACM Reference Format:

Stanley Goffinet, Donatien Schmitz, Igor Zavalyshyn, Axel Legay, Etienne Rivière. 2021. Controlling Security Rules Using Natural Dialogue: an Application to Smart Home Care. In *Adjunct Proceedings of the 2021 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2021 ACM International Symposium on Wearable Computers (UbiComp-ISWC '21 Adjunct)*, September 21–26, 2021, Virtual, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3460418.3479331>

1 INTRODUCTION

Smart home technologies have seen an unprecedented growth over the last five years. Numerous Internet-connected devices, like smart TVs or smart speakers, can now be found in almost every home. However, the more smart devices people put in their living rooms and bedrooms the more concerned they become over the data these devices collect. In fact, a simple motion sensor can reveal the times when a home is empty, and a smart TV's watch history can accurately predict a user's political views and religion. Unfortunately, such concerns have been supported by actual cases of sensitive data

abuse [5, 8], unauthorized sharing [9, 10, 12], eavesdropping [4], and massive data leaks [13].

Various access control mechanisms have been proposed to improve the security and privacy properties of existing and future smart home systems [1–3, 6, 11]. However, common across all these efforts is the assumption that users have sufficient technical experience and skills to configure access control policies and to clearly understand the consequences of the choices they make. This, however, is in stark contrast with the reality: there are significant gaps in user perceptions of IoT devices and services which result in poor security choices and lead to potentially dangerous situations [14, 15]. The situation is even worse for older smart home users. Due to their limited technological literacy older people tend to be unaware of and susceptible to various security and privacy risks posed by smart home systems [7], which results in either acting recklessly or avoiding a smart technology use completely.

There is an urgent need for technological and social solutions that specifically target elderly smart home users and offer adequate tools and mechanisms that consider their mental models and expectations regarding sensitive data privacy and security, as well as their limited technical knowledge and potential usability issues with new technologies.

In this paper we describe our experience in building a dialogue-based voice assistant that helps older smart home users to define their security and privacy preferences in an assisted living scenario. Given a set of events generated by various devices monitoring the health and well-being of an elderly person and a list of contacts that need to be notified of those events, e.g. family members, professional caregivers or emergency services, the assistant will automatically suggest appropriate security rules and ask for approval from the elderly. With this voice assistant one can define rules of various complexity using a natural language free from technical jargon. For instance, a positive answer to a "Should I send a message to your family members when you go out?" question will create a security rule that restricts notifications about this specific event to relatives only. The assistant constantly monitors the usage statistics of existing policy rules and suggests modifications if some of the rules become irrelevant or are rarely used. The latter mechanism stimulates the user to revisit their past security choices and ensures that the security policy always reflects current user preferences.

We evaluated a prototype of the voice assistant on a small set of potential users in a round of individual interviews. The results proved that usability-wise voice assistants are much more easier and pleasant to use than traditional web- or mobile-based interfaces, especially when configuring access control policy rules.

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 the author(s) 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.
UbiComp-ISWC '21 Adjunct, September 21–26, 2021, Virtual, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8461-2/21/09...\$15.00
<https://doi.org/10.1145/3460418.3479331>




Figure 1: An architecture of the proposed system.

2 SYSTEM OVERVIEW

A general architecture of the proposed system is shown at Figure 1. It consists of four main parts:

- **Smart hub** collects sensor data from the various connected smart home devices and wearables and notifies a predefined list of trusted contacts when certain events happen. For instance, it can send an alert to an emergency service when a fall is detected, or send a message to a family member when a pill dispenser is empty. The hub provides a web interface through which a system administrator or any of the family members can make an initial configuration of the system and define the events of interest and a list of trusted contacts who can be notified about these events via a text message or an audio call issued by a hub’s Alert API.
- **Voice assistant** activates when the initial system configuration is completed. Given a list of events and trusted contacts it will generate a set of security rules that restrict event notifications to selected contacts only. The assistant will then ask the elderly person in a form of a dialogue to approve or modify these rules. Based on the received answers it will activate corresponding access control policy rules that allow or block sending certain notifications and alerts to specific contacts. With a powerful speech recognition engine the assistant is able to detect and recognize a wide variety of user responses ranging from simple “yes” or “no”, to a more complex ones featuring a name of a selected contact, e.g., “David”, or its relation to the elderly, e.g., “my son” or “my neighbor”.
- **User** is an elderly person whose health and well-being is being monitored. The user defines her privacy and security preferences by simply replying to the questions posed by the voice assistant.
- **Trusted contacts** comprise a social network around the elderly person. These can be family members, friends, neighbors, professional caregivers, such as nurses and family doctors, as well as emergency and social service workers.

2.1 Policy rule format

The proposed system was carefully designed to be compatible with existing smart home solutions, e.g., SmartThings, and their access control policies. The rules generated by the assistant follow a well-defined format that specifies a trigger event (TRIGGER) which

```

TRIGGER: {
  match (hrmonitor).(heartrate)
  satisfy (>)-(60)
}

ACTION: [
  {perform (message).(son)},
  {perform (call).(nurse)},
  {perform (message).(doctor)}]

```

Listing 1: A policy rule generated for a heart rate monitor

generates an alert or a notification (ACTION) sent by the hub to a specific contact (ACTOR). Initial trigger events are inferred from the type of available sensor devices or from their corresponding software drivers installed at the hub. A door lock, for instance, can generate a ‘door lock’ or a ‘door closed’ events. The actions describe various ways of notifying people about these events. For instance, a hub can send an SMS message or an email for non-critical events, or make an audio call in case of an emergency. Finally, the list of contacts is initially defined by the system administrator and can later be adjusted by adding or removing a specific contact. Each contact in this list is accompanied by the description of its relation to an elderly person, e.g., a neighbor, a son, or a visiting nurse, a full name and the contact details, e.g., a phone number or an email address.

To construct a policy rule, the assistant uses the following template and replaces event, actor and action placeholders based on user responses:

- Should I {{ACTION}} {{ACTOR}} if {{EVENT}}?
- If {{EVENT}}, should I {{ACTION}} {{ACTOR_1}} or {{ACTOR_2}}?

Listing 1 shows an example of a policy rule processing a heart rate monitor data and sending alerts to various contacts, a family member (son), a nurse and a doctor, when a heart rate rises above a certain level. Both son and a doctor will get a text message, but a nurse will receive a call prompting her to check on the elderly in a timely manner.

2.2 Policy rule types

The number of policy rules that need to be approved by the elderly can quickly become overwhelming even with just a few device events and a small list of trusted contacts. To prevent a *decision fatigue* voice assistant internally implements a special heuristics that aims to prioritize certain events over the others. To this end, we classify each event type based on its criticality on a scale from 0 to 3 where a higher value constitutes a higher criticality level. For instance, a detected fall event (criticality score: 3) is more important than an event of a light switch (criticality score 0). Furthermore, for each event type the system selects a subset of contacts based on their social ties with the elderly and their ability to react accordingly. For instance, physical activity events would be more relevant for family members and a doctor than for a neighbor. Similarly, emergency service contacts are preferred when processing potentially life-threatening events from a gas or smoke detector.

2.3 Policy prompts

Using the aforementioned heuristics the voice assistant creates a set of security policy rules that are ordered based on the criticality of events they depend on. Rules with higher criticality scores are prompted for approval first, followed by the less critical ones that can be processed later to avoid overwhelming the user. For a policy rule shown in Listing 1 the voice assistant will ask the following question: *"Should I notify your son, your doctor and a nurse when your heart rate is higher than 60 bpm?"*. An elderly person can then select a specific contact out of the proposed ones, for example, by responding with a name of that contact, or authorize all three of them. The assistant does not require elaborate long answers and intentionally structures its prompts in a way that should provoke a short and definite answer. We provide several examples of potential policy prompts and the list of expected answers below:

- *Q: Should I notify your son and your neighbor if you forget to take your pills?*
Expected answers: {"No", "Yes", "My neighbor", "My son", "Both"}
- *Q: If you fall down should I call your son?*
Expected answers: {"No", "Yes"}

The actors or actions suggested by the assistant may sometimes be in conflict with the expectations of the elderly person. In this case, upon receiving a negative response the assistant will ask if the user is willing to change the action type. If the user agrees, the assistant will find the second most relevant action for the current event, for instance, based on other previously approved policy rules, and offer this new choice. The user can decline the offered choices until she finds an appropriate action or until there are no more actions available. The process is similar for changing the suggested actors list until, eventually, the user defines a new policy rule that fully satisfies her preferences.

If an answer to a given prompt was not recognized successfully, the assistant will repeat the question three more times before marking the prompt as unsuccessful and postponing it. To minimize the risks of mistakes, the assistant will repeat a recognized answer and ask for confirmation. Finally, all the policy rules that were successfully configured and confirmed by the elderly are saved and activated at the hub.

Policy rules may become outdated or irrelevant when certain devices or contacts get removed. To this end, the voice assistant will issue a prompt to update the deprecated rules accordingly. Update prompts have higher priority than those suggesting to create a new policy rule. The assistant will also occasionally remind the user about the existing rules and ask to confirm if they are still in line with user expectations.

3 IMPLEMENTATION

We implemented a prototype of the system using a collection of open source tools. We use Rhasspy as a base for our voice assistant with its built-in text-to-speech engine, Kaldi speech recognition engine and the Fsticuffs intent detection toolkit.

We chose to implement the heuristics and question generator in PHP within a Laravel project. Our implementation exposes a REST API endpoint and a web interface for initial system configuration, e.g., specify the list of trusted contacts. We use a MySQL database

User	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8	Q9	Q10	SUS score
1	4	1	5	2	3	2	4	1	4	1	82,5
2	4	1	5	1	4	1	4	1	5	1	92,5
3	4	1	4	1	5	1	5	1	5	3	90
4	5	1	5	2	5	1	4	1	5	2	92,5
5	5	2	4	2	4	2	5	1	5	2	85

Table 1: Usability assessment results based on SUS

to store policy rules, contact details and information about the connected devices and supported event types.

4 EVALUATION

To evaluate our system we conducted a set of individual interviews with 5 participants with an average age of 72 years and a minimum age of 59. As part of these interviews we asked each participant to define two security policy rules using our voice assistant, and then describe their experience by filling in a form. For the latter, we rely on a SUS (System Usability Scale) 10-item questionnaire that allows the participants to rate the usability of a given system. Each item has a score ranging from 1 (strongly disagree) to 5 (strongly agree). A total score computed with a small formula provides a subjective assessment of system usability: the higher the value the easier it is for the user to use the system and perform a given task.

Table 1 presents the results of our usability analysis. The right-most column displays a SUS score computed based on each participant's answers. Scores higher than 80 are generally considered to be the indicators of a great usability. The two participants that interest us the most are the two oldest, number 2 and 3, as they are the typical target for such a system. The SUS score computed for these participants is higher than 90 which is an encouraging result.

5 CONCLUSIONS

In this paper we describe our experience in designing, developing, and evaluating an alternative interface for access control systems in smart homes tailored for elderly. It allows to define security rules using natural language in a dialogue-like fashion using a voice assistant prompts. We evaluated the prototype on a small set of users that fit the target audience. The results were quite encouraging with all of the participants agreeing that such a dialogue-based system is a much more user-friendly alternative to the more common web interfaces.

REFERENCES

- [1] Z Berkay Celik, Leonardo Babun, Amit Kumar Sikder, Hidayet Aksu, Gang Tan, Patrick McDaniel, and A Selcuk Uluagac. 2018. Sensitive information tracking in commodity IoT. In *27th USENIX Security Symposium (USENIX Security 18)*. 1687–1704.
- [2] Z. Berkay Celik, Patrick McDaniel, and Gang Tan. 2018. Soteria: Automated IoT Safety and Security Analysis. In *Proceedings of USENIX Annual Technical Conference (USENIX ATC)*.
- [3] Z. Berkay Celik, Gang Tan, and Patrick McDaniel. 2019. IoTGuard: Dynamic Enforcement of Security and Safety Policy in Commodity IoT. In *Proceedings of Network and Distributed System Security Symposium (NDSS 2019)*.
- [4] Ry Crist. 2019. *Amazon and Google are listening to your voice recordings. Here's what we know about that*. CNET. <https://cnet.co/3jdHbcN> Accessed: June 2021.
- [5] Adam Clark Estes. 2018. *Yes, Your Amazon Echo Is an Ad Machine*. Gizmodo. <https://gizmodo.com/yes-your-amazon-echo-is-an-ad-machine-1821712916> Accessed: June 2021.
- [6] Earlene Fernandes, Justin Paupore, Amir Rahmati, Daniel Simionato, Mauro Conti, and Atul Prakash. 2016. Flowfence: Practical data protection for emerging iot application frameworks. In *25th USENIX security symposium (USENIX Security 16)*. 531–548.

- [7] Alisa Frik, Leysan Nurgalieva, Julia Bernd, Joyce Lee, Florian Schaub, and Serge Egelman. 2019. Privacy and security threat models and mitigation strategies of older adults. In *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*.
- [8] FTC. 2017. *VIZIO to Pay 2.2 Million to FTC, State of New Jersey to Settle Charges It Collected Viewing Histories on 11 Million Smart Televisions without Users' Consent*. US Federal Trade Commission. <https://www.ftc.gov/news-events/press-releases/2017/02/vizio-pay-22-million-ftc-state-new-jersey-settle-charges-it> Accessed: June 2021.
- [9] Caroline Haskins. 2018. *Amazon Sent 1,700 Alexa Recordings to the Wrong Person*. Vice. <https://www.vice.com/en/article/pa54g8/amazon-sent-1700-alexa-recordings-to-the-wrong-person> Accessed: June 2021.
- [10] Christine Hauser. 2018. *Police Use Fitbit Data to Charge 90-Year-Old Man in Stepdaughter's Killing*. New York Times. <https://www.nytimes.com/2018/10/03/us/fitbit-murder-arrest.html> Accessed: June 2021.
- [11] Yunhan Jack Jia, Qi Alfred Chen, Shiqi Wang, Amir Rahmati, Earlence Fernandes, Zhuoqing Morley Mao, Atul Prakash, and SJ Unviersity. 2017. ContextIoT: Towards Providing Contextual Integrity to Appified IoT Platforms.. In *Proceedings of Network and Distributed System Security Symposium (NDSS 2017)*, Vol. 2. 2–2.
- [12] Alfred Ng. 2019. *You shared Ring footage with police. They may share it, too*. CNET. <https://www.cnet.com/home/security/you-shared-ring-footage-with-police-they-may-share-it-too/> Accessed: June 2021.
- [13] Amanda Yeo. 2019. *Data leak by IoT device maker Wyze exposes personal information of 2.4 million people*. Mashable. <https://mashable.com/article/wyze-smart-home-data-leak-breach> Accessed: June 2021.
- [14] Eric Zeng, Shrirang Mare, and Franziska Roesner. 2017. End user security and privacy concerns with smart homes. In *thirteenth symposium on usable privacy and security ({SOUPS} 2017)*. 65–80.
- [15] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User perceptions of smart home IoT privacy. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–20.