

Exploiting IFTTT and Usage Control Obligations for Smart Home Security and Management

Giacomo Giorgi¹ | Antonio La Marra² | Fabio Martinelli¹ | Paolo Mori¹ | Athanasios Rizos³ | Andrea Saracino¹

¹Istituto di Informatica e Telematica,
Consiglio Nazionale delle Ricerche, Pisa,
Italy

²Security Forge srl, Pisa, Italy

³Department of Computer Science,
University of Pisa, Pisa, Italy

Correspondence

*Andrea Saracino, Istituto di Informatica e Telematica, Consiglio Nazionale delle Ricerche, Via G. Moruzzi n.1, 56124, Pisa, Italy. Email: andrea.saracino@iit.cnr.it

Abstract

In this paper we present an application of the Usage Control paradigm to a Smart Home infrastructure, based on a model extension and structured use of obligations. In the proposed extended model obligations are exploited to enforce two different access revocation time, namely revoke and suspend. This increases the policy expressiveness and enable to optimize the resource usage. Furthermore, obligations are exploited to send commands via IFTTT to different interconnected Smart Home devices, to impose safety-relevant behaviors, or to act on policy attributes to implement a self-healing paradigm for revoked sessions. The paper is motivated by a parental control use case where deep learning is used in combination with Usage Control to regulate dynamically viewing rights of a Smart TV and interactions with interconnected devices. Accuracy and performance experiments show the effectiveness and feasibility of the proposed work.

KEYWORDS:

Access Control, Deep Learning, IFTTT, Parental Control, Usage Control

1 | INTRODUCTION

In the last years, smart devices are increasingly becoming pervasive in all aspects of our daily life. For instance, traditional devices and appliances in our houses are being replaced by their smart versions (e.g., smart-TV, fridge, washing machine, etc.), which provide additional functionalities and improved effectiveness because they are connected to the network. This enables the *Smart Home* paradigm, where distinct and different devices in the same house communicate among them and with other entities through the local network and the Internet, thus being able to support novel services. As a matter of fact, the standalone capabilities of these devices are exponentially increased when they are able to communicate. Each device has, in fact, different sensors and actuators, and has access to different kind of information which can be used collaboratively to provide more complex services. Hence, a number of novel and advanced applications can be envisaged for the Smart Home scenario, spanning from efficient energy consumption management, to comfort enhancement, elderly and disabled people assistance, home safety and surveillance, and parental control.

Besides introducing the aforementioned advantages in terms of availability of a new generation of innovative and advanced services, the automatization of the devices deployed in a house also poses some new security and privacy concerns and challenges. For instance: is a child authorized to switch on the smart-TV or the smart oven? If the smart-TV has been switched on by an adult and it is playing content for adults only, what we would like the Smart Home management system to do if a child comes in front of the smart-TV?

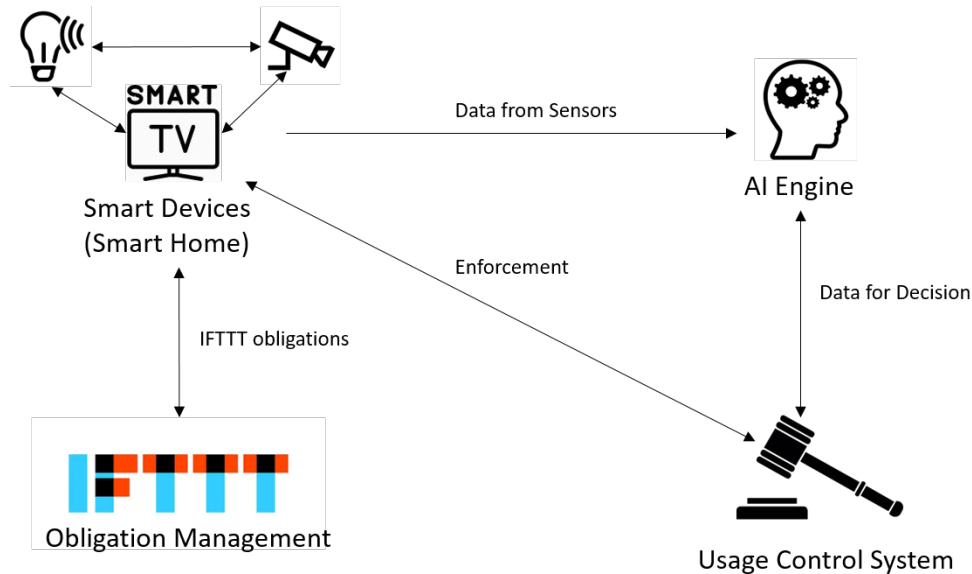


FIGURE 1 Logical Architecture of a Smart-TV use case.

To address the previously mentioned issues, the Usage Control System¹ has been already exploited in the Smart Home environment as a control tool able to handle seamless device access management and data protection^{2,3}, as well as to detect policy violations occurring while the access is in progress in order to react by executing mitigating actions. The Usage Control System is an extension of the eXtensible Access Control Markup Language (XACML) standard which defines both a language and a reference architecture to perform traditional Attribute-Based Access Control (ABAC). The extension concerns both the XACML language and reference architecture, and it emphasizes the role of obligations, that are mandatory actions that have to be performed in conjunction with the policy enforcement, and the continuity of access decision by evaluating and enforcing policies based on mutable attributes⁴, following the Usage Control (UCON) model⁵. Hence, coming back to the previous example, the Usage Control System, once having given to an adult the permission to switch on the smart-TV and to select content for adults, would be also able to detect when a child comes in front of the TV as a policy violation and, consequently, to ask the execution of an obligation as a countermeasure, such as the smart-TV switch off.

However, although both the XACML language and its extended version, U-XACML, include a basic construct to declare the obligation section within a policy, they do not provide a standard for the description of obligation semantic. While the intention is the one of non imposing constraints on the format of the represented information, the direct consequence is that obligation management engines have to be developed ad-hoc. Instead, having a standard semantic for obligation representation which is meaningful at least in an application macro-environment, would push developers and policy editors to use common expressions to represent and enforce obligations. In particular, in the Smart Home environment, which is a subcase of the Internet of Everything (IoE) environment, it is possible to express obligations as commands for an inter-operable platform used by a multitude of Internet of Things (IoT)/IoE devices, called If This Then That (IFTTT)¹. IFTTT is a free web-based platform that triggers the execution of actions on IoT devices as a consequence of changes that happen in the specified web-services. Exploiting IFTTT triggers to express obligations will allow the devices receiving such obligations from the Usage Control System as a consequence of the policy evaluation decision to easily execute them, without hard-coding the actual obligation interpretation in the device, or demand the obligation interpretation to a specific applet.

Before our proposal, some effort to extend IFTTT for managing security features have been already performed^{6,7}, still not reaching the full functionality of an Access Control system.

In this paper, we present an architecture that exploits the Usage Control System and IFTTT to manage resources and operations in a smart home environment. This framework aims at managing security and safety in IoT environment through a single standard interface and policy language. As an example, in Figure 1 we depict a high level view of the proposed framework, applied in a Smart Home environment for a parental control application that will be detailed in the following. As shown, the Usage Control

¹<https://www.ifttt.com/>

System interacts directly with the smart devices, acquiring data from sensors, possibly elaborated through mechanisms such as artificial intelligence (e.g. for user age estimation), while the IFTTT services are used for obligation management, to control specific actions of the connected smart devices.

We introduce extensions to the standard UCON model to make the paradigm more suitable for a dynamic environment, by handling automated resume of sessions revoked because of policy violations through self-healing policies exploiting IFTTT obligations to force attribute changes needed to satisfy the policy again.

Our proposal aims at enhancing the capabilities of the Usage Control System applied in the Smart Home environment leveraging on the obligation concept, and focusing on security and safety aspects. As a matter of fact, although the Usage Control System already supports obligations, the lack of a standard to express them in Usage Control policies is a relevant obstacle which limits their usage.

In the Smart Home environment we overcame this issue and we significantly enhance the capabilities of the Usage Control System thanks to the integration of the IFTTT platform which brings two major advantages: it allows to easily express in the Usage Control policy obligations operating on the IoT devices installed in the Smart Home, and it also simplifies the interactions of the Usage Control System with such devices when asking the execution of actions implementing obligations. This would allow the Usage Control System users to exploit the Smart Home devices to execute countermeasures to deal with unauthorized (i.e., undesired and/or dangerous) situations. This integration strengthens the Usage Control System simplifying its adoption in a large number of interesting scenarios.

In the following of the paper, we will use the previously introduced example concerning a parental control scenario, where a Smart-TV interacts with other smart devices in the Smart Home in order to enforce safety policies for smart-TV usage, as reference application to validate the proposed approach.

The contribution of this work are the following:

- We present an architecture based on Usage Control for advanced policy management in a Smart Home environment with interconnected devices.
- We introduce a methodology to express and enforce Access Control and Usage Control obligations by exploiting the IFTTT service.
- We propose an extension to the standard Usage Control model which introduces two new operations to suspend and resume accesses instead of revoking them, which increases the policy expressiveness and allows to optimize the resources usage time.
- We exploit the novel proposed model and the IFTTT-based obligations to implement self-healing policies, where obligations are used after an access revocation (or suspension) to force actions and behaviors meant to bring the system back in a situation compliant with usage policies.
- We present the implementation of a parental control use case, based on the one presented in⁸, which exploits the proposed framework to implement a safe and secure usage of a smart-TV and interconnected devices.
- We discuss the design and performance of a novel deep-learning classifier for age estimation, which enables the implementation of the parental control use case by identifying with high accuracy and with a limited performance overhead the number of children and adults currently watching the smart-TV: essential attributes for parental control policies enforcement.
- We evaluate the performance and effectiveness of the proposed approach in both a simulated and in a real use case, demonstrating its feasibility.

This work strongly extends and complete the work presented in⁹ and in⁸, by proposing a novel application environment and exploiting a novel extension of the UCON model. In particular, the novel contributions w.r.t. such previous papers are (i) the extended usage control model which considers the coexistence of the revoke and suspend workflows for a more flexible management of continuous control; (ii) the usage of obligations to implement a self-healing paradigm, to force the execution of actions needed to resume the access in case of policy violations; (iii) the application of IFTTT obligations to a parental control use case with implementation of specific IFTTT applets; (iv) a unique deep learning classifier that simultaneously detects faces and estimates their ages, improving accuracy and performances.

The rest of this article is organized as follows: In Section 2 we report some background information on Usage Control, IFTTT platform, and TV Parental Guidelines. Section 3 details the proposed architecture and operative workflow. Section 4 details our implementation and discusses the results of the performance analysis. In Section 5 we present a set of related works about security applications of UCON in IoT and XACML obligation expressions. Finally, in Section 6 we offer some conclusions and hint at future directions.

2 | BACKGROUND

In this section, we provide some background knowledge regarding Usage Control, IFTTT platform, and TV parental guidelines.

2.1 | Usage Control

The UCON model^{5,10} extends traditional access control models mainly introducing *obligations* and *mutable attributes* in policy evaluation.

Obligations are decision factors that impose the execution of actions as a consequence of the access decision process. For instance, an usage control policy could impose that a smart-light must be switched on when a person enters a given room and it must be kept on as long as there are people in such room.

Traditional access control models perform the decision process taking into account *immutable attributes* only, i.e., attributes that do not change their values as a consequence of the normal operation of the system, and that is modified only through administrative actions. Mutable attributes, instead, represent features of subjects, resources, and environment that change their values as a consequence of the normal operation of the system⁴. For instance, the number of people in a room is a mutable attribute of the room which changes when a person enters or leaves the room.

Since mutable attributes might change their values during the usage of an object, the UCON model allows the definition of policies (*pre-policies*) which are evaluated before the access (*pre-decision*), and policies (*ongoing-policies*) which are evaluated continuously during the access to the object (*ongoing-decision*). The *ongoing-decision* phase is meant to detect policy violations that occur while the usage of an object is in progress because the new value of a mutable attribute causes the policy to be not satisfied anymore. For this reason, UCON introduces the concept of access revocation, i.e., the capability of stopping ongoing access.

This paper takes into account the Usage Control System (UCS) presented in a previous work¹, whose architecture, shown in Figure 2, is based on the XACML reference architecture¹¹. As policy language, the UCS exploits an extension of the XACML language, called U-XACML, for dealing with Usage Control features.

In the UCS architecture, the Policy Enforcement Points (PEPs) embedded in the controlled system intercept the execution of security-relevant operations, and invoke the Context Handler (CH), which is the front-end of the UCS. The Policy Information Points (PIPs) are the components invoked by the CH to retrieve the current values of the attributes required by the Policy Decision Point (PDP) for the execution of the decision process, i.e., to evaluate the policy retrieved from the Policy Administration Point (PAP). Each specific scenario where the UCS is exploited has its own set of Attribute Managers (AMs) which manage the attributes required for the policy evaluation. Consequently, a number of PIPs must be properly configured for each specific scenario in order to be able to interact with such AMs.

The phases of the Usage Control decision process are regulated by the interactions between the PEPs and the UCS as follows (derived from¹²):

TryAccess: is the *pre-decision* phase, which begins when the *TryAccess* message is sent by the PEP to the UCS because a subject requests to execute an access. The *pre-decision* phase finishes when the UCS sends the response to the PEP. The possible responses are: *PermitAccess*, to allow the access, or *DenyAccess*, to reject the access request;

StartAccess: this interaction triggers the *ongoing-decision* phase, which begins when the *StartAccess* message is sent by the PEP to the UCS because the access has just started, and causes a continuous evaluation of the *ongoing-policy*. In case of policy violation, a *RevokeAccess* message is sent to the PEP, causing the access to be stopped, otherwise the *ongoing-decision* phase is executed until the access is not terminated by an *EndAccess* message;

RevokeAccess: this is part of the *ongoing-decision* phase. Every time an attribute changes its value the policy is evaluated again and, in case of violation, a *RevokeAccess* message is sent by the UCS to the PEP, causing the access to be stopped;

EndAccess: sent by the PEP to the UCS when ongoing access terminates normally.

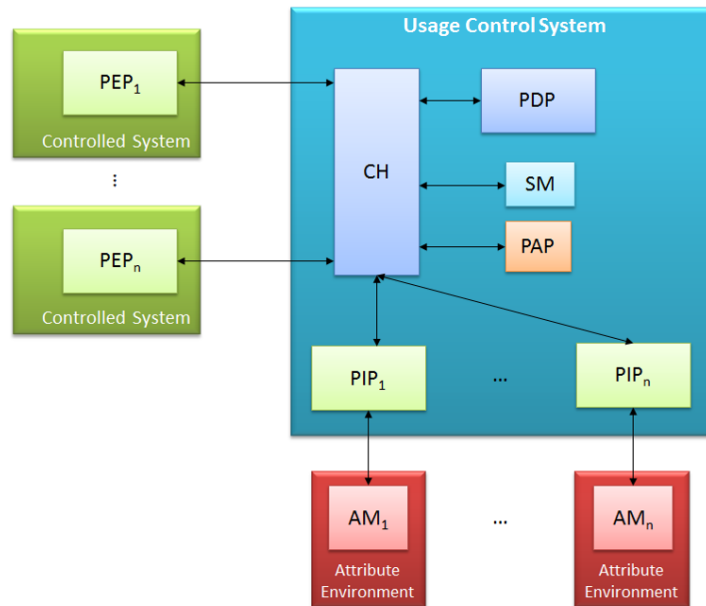


FIGURE 2 Usage Control System Architecture.

The workflow of the Usage Control decision process is the following. When the subject s tries to execute a security relevant action a , the PEP suspends the execution of a and retrieves the information related to the access context (subject and resource IDs, etc.). The PEP sends the *TryAccess* message with the data previously collected to the UCS, which performs the *pre-decision* process and returns the result of the *pre-policy* evaluation to the PEP, which enforces it. If the execution of a is permitted, the PEP sends the *StartAccess* message to the UCS as soon as a is started, and the UCS executes the first part of the *on-decision* phase, thus performing a first evaluation of the *ongoing-policy*. If the policy is violated, the UCS sends the *RevokeAccess* message to the PEP, in order to take proper countermeasures. Instead, if the policy is satisfied, the execution of a continues, and each time an attribute changes its value, the second part of the *on-decision* phase is performed, thus evaluating the *ongoing-policy* again. If the policy is violated, the UCS sends the *RevokeAccess* message to the PEP, in order to take proper countermeasures. Otherwise, no message is sent and the execution of a continues.

2.2 | IFTTT

IFTTT derives its name from the programming conditional statement “If This, Then That.”. IFTTT provides a software platform that connects apps, devices, and services from different developers in order to trigger one or more automation involving those apps, devices, and services.

IFTTT creates chains of simple conditional statements that are called *Applets* and are triggered by changes that happen within web-services.

IFTTT consists of the following structure:

Services: The basic building block of IFTTT. They describe data and actions controlled by an Application Programming Interface (API), and each service has a specific set of *Triggers* and *Actions*.

Triggers: The “*This*” part of the *Applet*, causing thus the triggering of the *Action*.

Actions: The “*That*” part of the *Applet* which is the result of a *Trigger*.

Applets: The composition of a *Trigger* and an *Action*.

An advantage of IFTTT is that it can work with various platforms and devices in IoT. A drawback of IFTTT is that it allows only a single *Trigger* and a single *Action* in the same *Applet*. The same *Trigger* can be used for other *Actions* but not inside the same *Applet*.

There are various competitors of IFTTT and in Table 1 we see a summary of the pros and cons of them. More in detail, the biggest competitor of IFTTT is Microsoft Flow. It is free for up to 750 runs/month, it can allow multi-step connection and works with many apps including Gmail, Facebook, etc. and can be accessed either via apps or browsers. The cons are that,

| Name | Cost | Connections | Apps/ Devices | Number of Services | Access |
|----------------|-----------|-------------|------------------|-----------------------|-------------|
| IFTTT | Free | One-to-One | Yes/Yes | >600 | App/Browser |
| Microsoft Flow | Free/Paid | Multiple | Yes/No | 226 | App/Browser |
| Zapier | Free/Paid | Multiple | Yes/No | >1500 | Browser |
| Yonomi | Free | Multiple | No/Yes | 200 | App |
| Stringify | Free | Multiple | Yes/Yes | 70 | App/Browser |
| Workflow | Free | Multiple | Yes(iOS)/No | N/A | App |

TABLE 1 Comparison of IFTTT and its Alternatives.

up to now, the apps are around 226 and it does not work with physical devices. But, it also supports Do/While and For/Each loop whereas IFTTT supports only If conditionals. Hence, it is more difficult and complicated to operate. Another example is Zapier that works only with apps and not physical devices. It allows multi-step connection between devices and has a simple free usage up to 100 runs/month and also paid plans of use. On the contrary, Yonomi works only with physical devices and not with apps, but allows multi-step connection between the about 100 compatible devices, it is free and focuses more on home applications. Furthermore, Stringify is also free and can host multiple connections but has only about 70 services and cannot be accessed via a browser like IFTTT but only via an app. Finally, there is also Workflow that works only with iOS apps and allows multiple-step systems. It is again free but there is not a list of compatible services but only the most famous iOS apps are used like Safari, Photo Gallery, Facebook, etc.

2.3 | TV Parental Guidelines

Multimedia content, such as movies and TV programs, are classified according to standards in order to advise the intended audience about features that might be considered dangerous or offensive to a set of spectators. In this work we refer to the American standard for parental guidelines ², where the guidelines are based on two parameters, namely the *Audience* and the *Content Label*. The *Audience* parameter may assume one of the following values:

- (Y) for all children: This means that the program is not expected to frighten younger (2–6-year-old) children.
- (Y7) older children: This program is designed for children at the age of 7 and above. Some of these programs may also have the Fantasy Violence (FV) *Content Label* in case that fantasy violence is more combative than in other programs.
- (G) General Audience: Program for all ages. Even if this program is not specifically designed for children, most parents may let younger children watch this program unattended.
- (PG) Parental Guidance: many parents may find this program unsuitable for younger children (the theme may call for parental guidance).
- (14) Parents Strongly Cautioned: many parents may find this program unsuitable for children under 14. They are also urged against letting children unattended.
- (MA) Mature Audience: Programs specifically designed to be viewed by adults. *Content Label* indicates if a show might contain violence, sexually explicit content, mature language, etc.

Each program may have assigned one or more *Content Labels* as follows:

- (D): Suggestive dialogue
- (L): Coarse or crude language

²http://www.tvguidelines.org/resources/TV_Parental_guidelines_Brochure.pdf

- (S): Sexual situations
- (V): Violence
- (FV): Fantasy Violence (children’s programming only)

The presence of one or more *Content Labels* among D, S, and L, defines whether a program or a content should have intended audience is PG, 14 or MA.

3 | METHODOLOGY

In this section, we report the description of the proposed framework for enforcing UCON policies in the Smart Home scenario implementing obligations via IFTTT. As a matter of fact, as shown in Section 2.1, the UCON model covers obligations, but the extension of the XACML policy language that has been defined for dealing with UCON features, U-XACML¹³, does not cover a language to express such obligations. The main reason is that the obligations that are enforced by the UCON policies are typically very dependent on the application scenario where the policy is enforced, and hence it is difficult to define a standard language for dealing with all kinds of obligations. Instead, in some application scenarios, such as the Smart Home one, specific languages already exist for expressing scenario-specific action execution commands. For this reason, in the Smart Home scenario, we decided to exploit IFTTT as a language to define obligations in UCON policies.

3.1 | Extended Usage Control Model

In the original UCON model (described in Section 2.1), when a policy violation is detected in the *ongoing-decision* phase, the corresponding access is revoked. This work considers an extended version of the UCON model¹⁴, which pairs the standard revoke operation with a softer usage revocation, named *SuspendAccess*. In fact, some scenarios only need to suspend the access as long as the Usage Control policy is violated, and the access should be resumed as soon as the Usage Control policy is satisfied again.

To describe the proposed extension we follow a similar approach to¹², where a set of *Usage Control actions* are used to describe the actions performed by the subject and the ones performed by the UCON system to implement the Usage Control decision process. Hence, given that the triple (s, o, r) which represents an access request where a subject s wants to access an object o through an operation that requires the right r , we extend the UCON model as follows:

- *TryAccess* (s, o, r) : performed when the subject s requests a new access (s, o, r) . This triggers the execution of the *pre-decision* phase.
- *PermitAccess* (s, o, r) : performed by the system as a result of the *pre-decision* phase, when the evaluation of the policy allows the access request (s, o, r) .
- *DenyAccess* (s, o, r) : performed by the system as a result of the *pre-decision* phase, when the evaluation of the policy forbids the access request (s, o, r) .
- *StartAccess* (s, o, r) : performed when the access (s, o, r) has been allowed by the *PermitAccess* (s, o, r) and it has actually started. This triggers the execution of the *ongoing-decision* phase.
- *End Access* (s, o, r) : performed when the subject s terminates an ongoing access (s, o, r) .
- *RevokeAccess* (s, o, r) : performed by the system as a result of the *ongoing-decision* phase when the policy is violated and it requires to block the usage of a resource.
- *Suspend Access* (s, o, r) : performed by the system as a result of the *ongoing-decision* phase when the policy is violated and it requires to suspend ongoing access (s, o, r) .
- *ResumeAccess* (s, o, r) : performed by the system as a result of the *ongoing-decision* phase when the policy requires to resume an access (s, o, r) previously suspended by the *Suspend Access* (s, o, r) action.

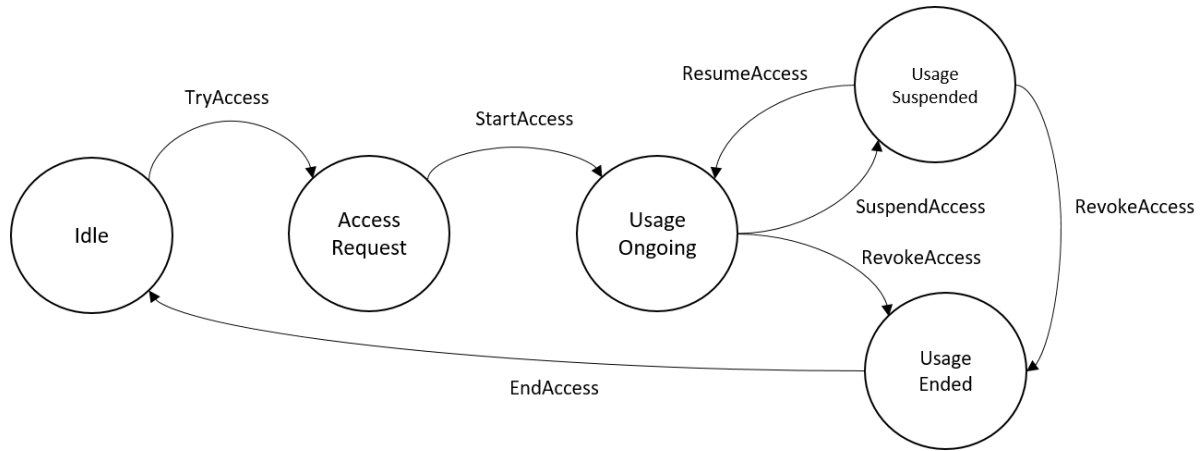


FIGURE 3 State machine of the extended UCON model.

As shown in Figure 3, with respect to the traditional UCON model, two new actions have been added to the Usage Control decision process: *SuspendAccess* and *ResumeAccess*. The *SuspendAccess* action, instead of definitely terminating ongoing access, simply suspends it. This means that the *ongoing-decision* phase is still performed for that access, until either the subject decides to terminate it, or the UCON system decides to resume or definitely terminate such access. The *ResumeAccess* is issued for suspended access when, during the *ongoing-decision* phase, the evaluation of the policy determines that the access is authorized again and, consequently, it must be resumed.

This model helps in maximizing the resource usage time when an *ongoing-policy* violation only requires a temporary suspension of the access until the policy is satisfied again. In fact, when the *ongoing-policy* is violated, the standard UCON model terminates the resource usage through a *RevokeAccess*. Consequently, the access can be continued only if the user explicitly issues a new *TryAccess* request, given that the request is matching the *pre-policy* and the *ongoing-policy*. However, it is possible that mutable attributes already were back in a configuration matching the policy before the actual time in which the new *TryAccess* is performed. In such a case, the time elapsed between the attribute changes and the moment in which the new *TryAccess* is performed by the user is wasted usage time. Formally, naming t_c the time in which the mutable attributes values match again the ongoing policy conditions, and naming $t_{ta} > t_c$ the time in which the *TryAccess* is performed, the time $\Delta_w = t_{ta} - t_c$ is the amount of usage time that goes unused and wasted. The introduction of the *ResumeAccess* reduces Δ_w .

Differently, from what proposed in¹⁴, the extended model we propose in this paper still includes the *RevokeAccess* action to increase the expressiveness of the policies. In fact, this model allows to differentiate among (i) violations that require that the whole usage control policy is re-evaluated (i.e., both the *pre-policy* and the *ongoing-policy*) to continue the access (e.g., active authentication is required each time the resource is accessed), and (ii) violations which require that only the *ongoing-policy* is re-evaluated to decide whether to resume the access (e.g., monitoring the temperature in the room, suspending the usage if it goes over the threshold and resuming it as it changes back to the acceptable range). For (i), the access revocation has to be handled with the *RevokeAccess*, which triggers the *EndAccess* that definitely terminates the access. In (ii), a *SuspendAccess* can be used instead, thus mutable attributes will be continuously monitored and the *ongoing-policy* re-evaluated. When the re-evaluation returns a permit, the *ResumeAccess* will be invoked enabling the usage again.

At the implementation level, the management of the two different strategies to react to *ongoing-policy* violations can be handled through policy obligations. In fact, in U-XACML an access revocation is enforced as the result of a DENY decision returned by the PDP as a result of the evaluation of an *ongoing-policy*. Thus, an obligation which specifies if this deny decision should imply an access revocation or an access suspension should be included in the policy. This obligation is thus returned to the PEP together with the DENY decision, driving the PEP behavior. In case the obligation is forcing a *RevokeAccess*, the PEP will terminate the resource usage and will invoke an *EndAccess* as for the usual UCON flow, which in turns blocks the attribute monitoring and policy re-evaluation for that access. In case the obligation implies a *SuspendAccess*, the resource access is also terminated or put on hold (depending on the resource type). The *EndAccess* is not invoked, and attributes are still monitored in such a way that, every time they change, the *ongoing-policy* is re-evaluated for that access, until a permit decision is returned, or in alternative a definitive revoke due to a different policy rule returning deny.




FIGURE 4 UCON obligation via IFTTT in a smart-greenhouse installation.

3.2 | Use Cases

This section presents three different applications of our framework to show the flexibility of our model as it adapts to several different context.

UCON obligation via IFTTT in a smart-greenhouse

The first example consists of advanced management through policy enforcement for remote urban farming in smart-greenhouses⁹. The representation of the scenario is shown in Figure 4. The greenhouse is composed of a set of smart-sensors, which in our specific scenario are humidity and temperature sensors, The UCS by monitoring the greenhouse attributes, can provide access to the smart-heating device to perform scheduled heating of the plants. The obligations are targeting the watering schedule of the smart-watering device which is responsible for watering the plants. The goal in this use case is to maintain the temperature and humidity levels of the greenhouse to the desired values via controlling the smart-heating with UCS and change the schedule of the smart-watering device with IFTTT via the obligations coming from UCS. The PEP component is embedded in the smart-heating device and it asks for permission to operate on the smart-greenhouse. It communicates with the UCS, initially with a StartAccess action, in order to request access to operate. The UCS monitors the values of the humidity, temperature, and it takes a decision on the basis of the request and the policy set. The response to the StartAccess sent by the UCS of a successful request contains an obligation including a schedule for the smart-watering device with the operative settings. The PEP will enforce the obligation by sending the web-request to the IFTTT platform that acts directly on the smart-watering device, setting thus the received schedule. The UCS performs a continuous re-evaluation of the policy monitoring the attributes received by the smart-sensors. Thus, when occurring a policy violation, e.g the weather is too hot or the smart-heating device is operating for too long, the UCS performs the RevokeAccess to the PEP to stop operating. Such action includes an obligation containing the new schedule for the smart-heating.

UCON obligation via IFTTT in a smart-office

Another use case is smart management through policy enforcement for controlling a smart-office⁹. The goal of this scenario is to continuously monitor the presence of people inside an office, the air quality of the room while they are inside and the temperature. The representation of this scenario is shown in Figure 5, where there are shown two cases. On the left, there is




FIGURE 5 UCON obligation via IFTTT in a smart-office installation for monitoring air quality.

an example with a few people inside the office. The PEP is installed in the smart-heating device and request for access from the UCS to operate as scheduled. In the response of the UCS after the StartAccess there is an obligation for the IFTTT to force stop the ventilation system. The other example depicted In Figure 5 represents a situation with too many people inside the office. Thus, due to the not comfortable condition inside the office, there is a policy violation. Then, independently from the schedule of the smart-heating, UCS issues the RevokeAccess action to force the smart-heating device to stop. The obligation that comes together with the RevokeAccess, forces, through the execution of the IFTTT Applet, the smart-ventilation system to start operating. In this example, the first obligation comes with the response after a successful StartAccess and forces through the IFTTT the smart-ventilation system to stop. The second obligation comes with the RevokeAccess and forces through the IFTTT the smart-ventilation system to operate.

UCON obligation via IFTTT in a Smart Home

The third use case, which is used as a reference scenario to validate the proposed framework, concerns a Smart Home where the usage of a smart-TV is regulated through our framework. As a matter of fact, the smart-TV allows the people in the house to easily access a large number of multimedia contents, being some of them addressed to adults only (the ones that fall in the category MA according to the classification reported in Section 2.3), or requiring the presence of an adult for a child to safely watch them (categories 14 and PG). Hence, an issue that arises with smart-TVs is that children or minors can very easily watch contents (e.g., live programs or on-line videos) that are not meant for them. The relevance of the problem is confirmed by the fact that some smart-TVs and some apps or content providers already include some solutions to enforce parental control. However, these solutions are mainly based on manual settings of channels, programs, or time slots that should be accessible or not to the children, where the forbidden contents are generally protected through a PIN or a password. This makes the parental control service cumbersome to use since adults are not willing to insert a password/PIN any time they wish to watch a program or content not meant for children. Moreover, these approaches have limited flexibility, because they don't consider, for example, those contents which are allowed to minors if watched in the presence of an adult. For this scenario, is also not considered the case in which the adult can temporarily leave the minor to watch the content, without pausing it. Figure 6 shows the smart-parental control use case. The attributes that must be taken into account by the policy are represented by the ages of the people in front of the smart-TV which are detected by the artificial intelligence engine that analyzes the video stream of the camera embedded in the smart-TV and the TV program recommendation provided by the TV content metadata. The PEP is embedded in the smart-TV to enforce the decision provided by the UCS through IFTTT obligation. The UCS, monitoring the aforementioned attributes, takes a decision on the bases of the request and the policy set. If the decision result returns a PERMIT, the UCS response does not contain any IFTTT obligation, whereas if a violation occurs, the UCS response returns a DENY and embed an IFTTT obligation to be executed on the PEP that performs a revocation or a suspension of the TV content.

3.3 | Relevant Usage Control Policies

Considering the aforementioned Smart Home scenario, we define the following set of Usage Control policies to regulate the usage of the smart-TV:

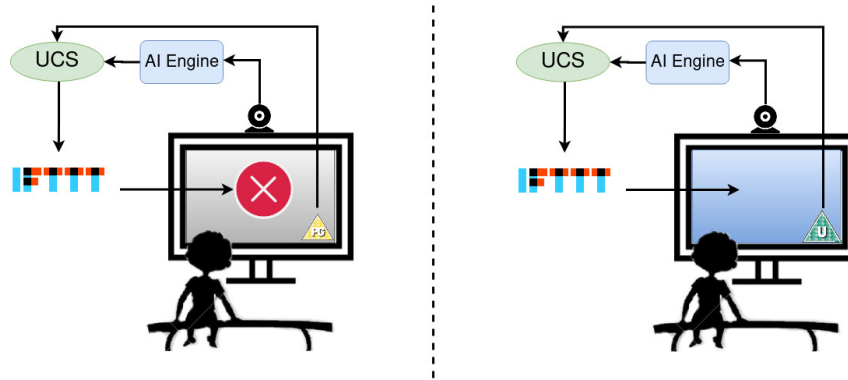


FIGURE 6 UCON obligation via IFTTT in a smart-TV installation for parental control.

Policy 1: The user can watch the TV IF the *program classification* is one-of {Y, Y7, G} OR the *program classification* is PG and the *number of adults* is greater than 0 OR *program classification* is one of {14, MA} and the *number of children* is equal to 0. On PERMIT switch on the light if ambience light is not sufficient.

This is a *pre-policy* to ensure that the right of watching the TV is granted according to the parental guidelines on the streaming program. Furthermore, the policy contains an obligation used to enforce the right lighting conditions in the room to not fatigue the eyesight of spectators, by interacting with the smart-lights in the room, if needed.

Policy 2: AS LONG AS the program classification is one-of {Y, Y7, G}, the vision can continue.

Policy 3: AS LONG AS the program classification is PG and the *number of adults* is greater than 0, the vision can continue. In case the program classification is PG and the *number of adults* is equal to 0, the vision is *suspended* and the adult has to choose one of: (i) change program, (ii) wait for adult presence, (iii) switch off the TV.

Policy 4: AS LONG the program classification is one-of {14, MA}, and the *number of children* is equal to 0, the vision can continue. In case the program classification is one-of {14, MA}, and the *number of children* is greater than 0, the vision is *suspended* and the adult has to choose one of: (i) change program, (ii) wait for children absence, (iii) switch off the TV.

Policy 2 is an *ongoing-policy* which simply allows to continue the vision of the program as long as its the classification is Y, Y7, or G.

Policy 3 is an *ongoing-policy* which allows to continue the vision of a PG classified program only if at least an adult is watching such program. This policy suspends the right visualize the program when no adults are in front of the TV to avoid that a child is left alone watching a PG classified program. In such a case, through the obligation enforced via IFTTT, a notification is sent to the TV owner smartphone, who can either change the program, switch off the TV, or do not act. Changing the program will imply the mutation of the *program classification* attribute value, hence a policy re-evaluation. For instance, if the new program is classified Y, Y7 or G and no other attributes have changed, the re-evaluation will allow the streaming of the new program, hence the usage of the smart-TV will be granted again.

Policy 4 is an *ongoing-policy* which suspends the right to watch the TV if a child pops up while the adult is watching a program classified 14 or MA. Similarly to the previous case, through the obligation enforced via IFTTT, a notification is sent to the TV owner smartphone, who either change the program, switch off the TV, or do not act.

Policy 5: IF the *current time* is different from allowed viewing hours and the *number of children* is at least 1, the vision of the TV is *revoked* or *denied*.

This policy contains a pre-authorization and ongoing condition to define ranges of time in which the child/children is allowed to watch the TV, either forbidding the TV usage out of allowed time ranges or stopping the usage once the allowed watching time range passes.

3.4 | Architecture

The aim of this paper is to evaluate the adoption of the UCON framework to implement a parental control system applied on a smart-TV. The main idea is to react to policy violations through UCON obligations enforced via IFTTT. To this aim, we propose the architecture in Figure 7, which is composed by four components described in the following.

- **Smart-TV:** the main component of the architecture is a smart-TV offering internet connectivity and a range of applications (apps) for the TV program managing. In our scenario, the smart-TV embeds a camera used to acquire the scene in front of the TV. Such camera is used by the *video capture* component to transmit the live video to the AI Engine. The built-in app that manages the smart-TV retrieves the TV program recommendation (*Audience* and *Content Label*) directly from the content metadata when available (natives in Android TVs), or querying external databases with content information such as the Internet Movie Database (IMDb). Such metadata are continuously sent to the PIP_{Adv} located in the UCON framework. Finally, a *PEP* is embedded in the smart-TV to interact with the UCS, i.e., to send access requests to the UCS and to receive the related responses. If the response of the UCS includes an obligation, the PEP enforces it directly on the smart-TV or by activating the Trigger of the corresponding IFTTT Applet.
- **AI Engine:** the aim of the AI Engine server is to analyze the received video stream related to the environment in front of the smart-TV. The analysis is performed to detect the faces of the people in front of the TV and estimate their ages. Such activity exploits a real-time Deep Learning object detector trained to simultaneously detect faces and their ages on a video stream. The output of the network is modified to inform in a continuous way the PIP_{child} and PIP_{adult} components about the number of children and adults detected in front of the smart-TV camera.
- **UCS:** the Usage Control System is invoked by the PEP embedded in the smart-TV to trigger the *pre-decision* phase every time the smart-TV is switched on. To evaluate the *pre-policy*, PIP_{Adv} extracts the TV program recommendation (*Audience* and *Content Label*) from the message received from the smart-TV, while PIP_{Child} and PIP_{Adult} extract the number of children and adults, respectively, from the AI Engine. These PIPs also notify the UCS when the value of these attributes change, in order to trigger the *ongoing-decision* phase. Exploiting these attributes, the PDP evaluates the *pre-policy* and the *ongoing-policy* according to the Usage Control decision workflow, and the result is sent to the PEP located to the smart-TV. This result could embed obligations to be executed by the PEP, which specify actions to be triggered through the IFTTT, e.g., switch on the TV back light. However, when the evaluation of the *ongoing-policy* returns DENY, obligations are also used to communicate to the PEP whether the access must be revoked or simply suspended.
- **IFTTT Component** The IFTTT component is the service that manages the execution of actions controlled through web based requests. It includes a *Trigger* component, which is invoked by the PEP for the enforcement of obligations. Each web request done by the PEP calls such trigger that invoke a specific *Action*, that in our scenario can be a message notification on the TV owner smartphone or a lighting control in the room.

3.5 | Workflow

The complete workflow of our framework is presented in Figure 8 . The figure describes the communication between the PEP and UCS implementing the Usage Control process workflow, exploiting the Usage Control actions defined in Section 3.1.

As shown in Figure 8 , the PEP primarily initiates communication with the UCS by performing the *TryAccess* action for the evaluation of the request (1:TryAccess(request)). The CH component of the UCS receives the request and, consequently, collects the current values of the attributes through the PIPs (PIP_{child} , PIP_{adults} , $PIP_{advisory}$, PIP_{time}). The request and the current attribute values are sent to the PDP which evaluates the *pre-policy* (2:Policy Evaluation). Let us suppose that the result of the policy evaluation is PERMIT. In this case, a unique session ID is assigned to this access and the PEP is informed that the access is granted (3:PermitAccess(sessionID,obligation)). The PEP executes the obligations included in the UCS response (4:Obligation execution) by triggering the IFTTT action (5:Trigger). The IFTTT server forwards the command to the applet which interacts with the related smart object. For instance, following Policy 1, the IFTTT server must enforce the right lighting condition by interacting with the smart-lights in the room, if needed (6:lighting control). After the obligation enforcement, the PEP starts the actual usage of the resource, and notifies the UCS by performing the *StartAccess* action for the session with the specific ID (7:StartAccess(SessionID)). Consequently, the UCS starts the continuous evaluation of the *ongoing-policy* (8.1:Policy re-evaluation/8.2:Policy re-evaluation). While the usage session is in progress, the PIPs monitor the value of attributes and notify the UCS when the value of an attribute changes, thus triggering the PDP for a re-evaluation of the *ongoing-policy*. As long as the policy is satisfied, the UCS does not interact with the PEP. In case of policy violation, the policy itself specifies to the UCS to perform either the *SuspendAccess* suspending the ongoing access to the resource (9.1:SuspendAccess(SessionID, Obligation) or the *RevokeAccess*, definitely terminating such access (9.2:RevokeAccess(SessionID, Obligation)). Let us suppose to enforce Policy 3. In case of policy violation, the access is suspended and the obligation is to notify the TV owner. Hence, the PEP

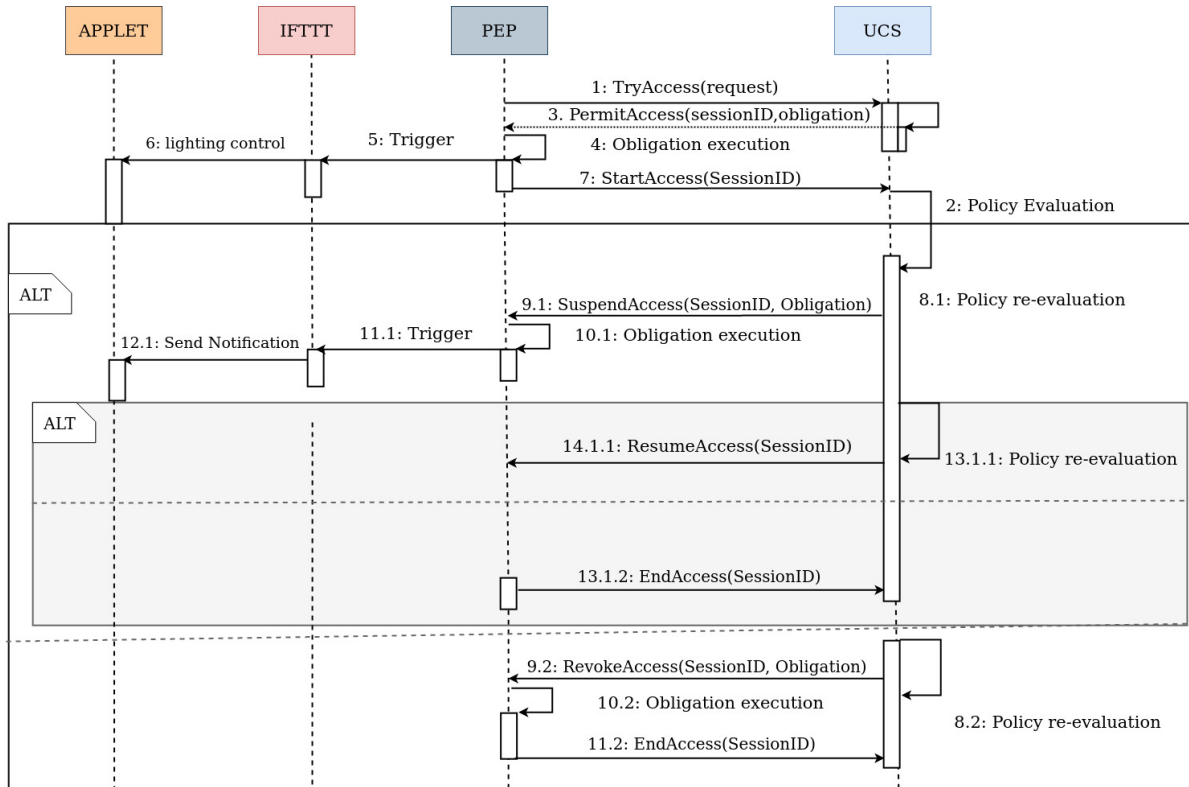


FIGURE 8 Policy enforcement workflow.

connections like the residual network in ResNet¹⁷. To simultaneously obtain face localization and age estimation, we trained our classifier exploiting the concept of transfer learning, following the same methodology explained in ⁸. As shown in figure 9 ,

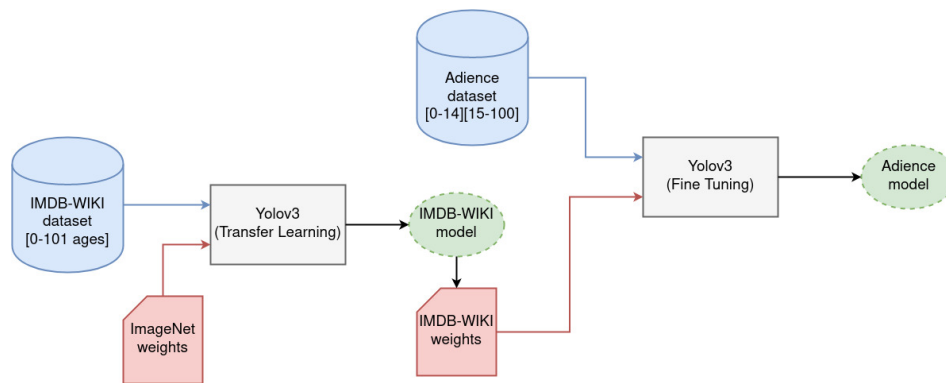


FIGURE 9 Training methodology.

we started with a darknet-53 pre-trained model that exploits the weights learned on the imageNet dataset which contains 1000 classes. Starting from the fact that the first layers of the network (convolutional layers) learn lower-level image features and the last fully connected layers evaluate the features to determine to which class they belong, the main idea is to start with some knowledge about object detection (imagenet dataset), then retrain the last fully connected layer using the specific dataset suitable for estimating ages and detect faces. As a transfer learning dataset, we exploited IMDb-WIKI dataset¹⁸. It is the largest public dataset of face images with age and gender labels. It is gathered using the most popular 100,000 celebrity images from IMDb and Wikipedia websites. In total the dataset is composed of 460,723 face images from 20,284 celebrities from IMD and 62,328




FIGURE 10 IMDb-WIKI age distribution.

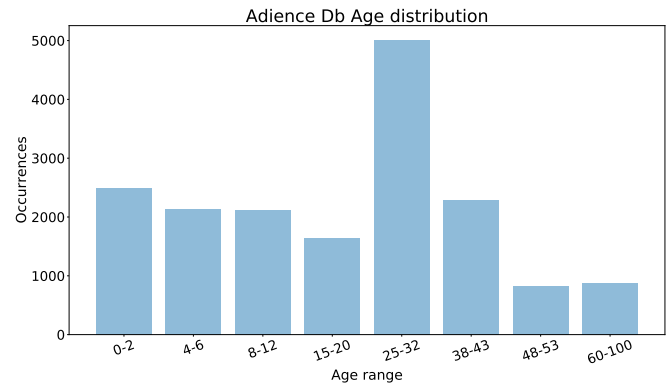


FIGURE 11 Adience Db age distribution.

from Wikipedia, thus 523,051 in total. As shown in Figure 10 and explained in⁸, the dataset is unbalanced with missing data in the range [0-15] and [70-100] ages. However, the dataset contains a large amount of data useful to learn lower level features to detect faces and estimate ages in an image frame. For that reason, starting from the Yolo V3 pre-trained on ImageNet, we transfer its learning on the IMDb-WIKI dataset to adapt the network on the new tasks (face detection and age estimation of 101 classes). Thus, we fine-tuned the obtained model on a smallest but balanced dataset in order to get higher accuracy on the minority classes. As balanced dataset we exploited the Adience DB^{19,20} dataset, which contains 26,580 face images from 2,284 subjects divided in 8 different age ranges, equally distributed as shown in figure 11. The dataset has been filtered, removing images with multiple faces belonging to different age ranges or images wrongly labeled, obtaining thus a dataset consisting of 5.949 faces in the age range [0-14] and 5.937 faces in the age range [15-100].

Testbed implementation and smart-TV simulation

The smart-TV can be considered by definition a component of a Smart Home environment. In fact, together with offering a large set of advanced entertainment features, being connected to the home network, it can easily communicate with other devices connected that are part of the Smart Home. Hence, it is possible to exploit the different functionalities and computational power of the device for providing additional joint functionalities and improving the whole system performance.

A preliminary version of this architecture has been presented in⁸, our previous work more focused on the deep learning aspects. The architecture has been refined, by upgrading the functionalities of the components and by introducing the IFTTT-based obligations in the workflow. More in details, as anticipated, the FEL and ACL have been completely re-implemented with a new, single network logic, the training strategy has been completely revised and this brought better performance in terms of accuracy and overhead. In addition, we have considered the interaction with more external devices (e.g. the Smart-Light to ensure lighting conditions which do not harm the viewer sight). Furthermore, the presence of IFTTT-based obligation, strongly improves the functionalities of the smart-TV by (i) considered the interaction with more external devices (e.g. the Smart-Light to ensure lighting conditions which do not harm the viewer sight), and (ii) by enabling self healing policy for a smarter reaction and better control of the parental functionalities.

In particular, since from the point of view of performance the smart-TV should be considered a relatively constrained device, running the FEAC and UCS components directly on it, might impose a considerable overhead on the smart-TV and might result in poor performance. To this end, we consider the possibility of outsourcing these functionalities to other nodes of the Smart Home network. In our testbed implementation, we simulated the smart-TV with an NVIDIA Shield where we implemented a simple live TV streaming application written for Android TV. The application is able to manage the live video streaming allowing the suspend and the resume video actions, and it extracts the video information rating directly from the content metadata when available (natives in Android TVs), or querying external Internet Movie Database (IMDb). The PEP is embedded in the simulated smart-TV as an Android application. It can communicate with the live TV streaming application allowing the enforcement on the streaming. The last component embedded in the smart-TV is the video capture, implemented as Android application as well. It exploits an external camera to capture the live video streaming in front of the TV with frame resolution 650x650, and forwards it to the AI Engine server, which provides a secure endpoint for the transmission.

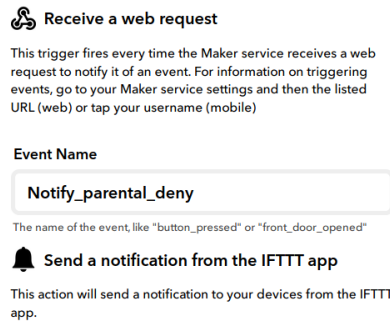


FIGURE 12 IFTTT Trigger part.

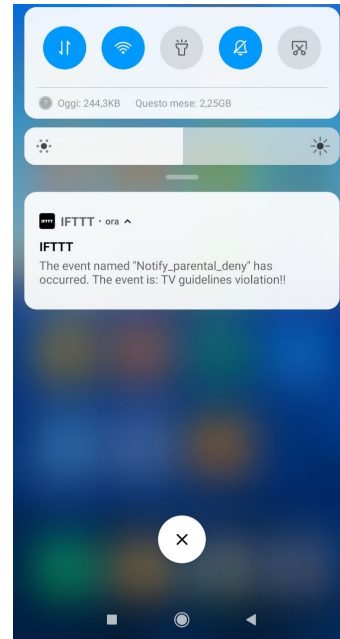


FIGURE 13 IFTTT Notification.

In our implementation the UCS is installed directly in the (simulated) smart-TV, still it is possible to envision different configurations with the UCS running on a separate device. The UCS itself does not require high performance devices; it is a service that can be installed in any device given that such device can be customized at a certain level. Android devices, as the one used in our testbed, make a perfect fit since they allow installing third party applications. Another possible set-up envisions the UCS as placed in a remote device (e.g. a server) and only the PEP installed in the controlled device. Again, a certain level of customization is required, which makes devices where it is possible to install third-party app perfect candidates. More closed devices might require a deeper modification to deploy the PEP and/or the UCS, such as installing kernel module, or modifying the device firmware. This approach is seldom taken, as it requires a deployment update every time the kernel/firmware version is updated. If these modifications are not possible, a further approach consists of in-lining called external primitives to hijack the operation and enforce UCON policies, or installing the controlling component on a connected device that interacts in a read/write manner with the device that cannot be modified.

Policy Enforcement

The policy enforcement is performed by the PEP, which is embedded in the smart-TV. In our simulated scenario, the PEP is installed into the smart-TV and it interacts with the live TV streaming application to enforce the suspend and the resume actions through a REST API provided by the streaming application. The IFTTT server is implemented using Webhooks³, that allows to create an HTTP endpoints (hooks) that can be used as a Trigger of the Applet. The *Action* is a service of the IFTTT platform that sends a notification to the smartphone of the TV owner exploiting the IFTTT app. Figure 12 shows the Trigger part which provides a box for specifying the unique Event Name of the Webhooks service.

The action can be invoked through an HTTP request to the webhooks service specifying the event name, the service key and the values to transmit in the payload. Such web request is called by the PEP when it extracts the appropriate obligation from the policy with the values to transmit them on the payload (policy violation message or lighting control). The web request is as follows:

`https://maker.ifttt.com/trigger/A/with/key/B, A ← EventName, B ← Key`

Figure 13 shows the IFTTT notification received on a smartphone.

³https://www.ifttt.com/maker_ebhooks

4 | RESULTS

This section presents a set of experiments to evaluate the performance overhead introduced by the proposed framework and to demonstrate its viability.

4.1 | Classification Results

In this section are reported the classification results, in terms of accuracy of face detection and age estimation obtained evaluating the FEAC component on the public datasets used in the training phase.

Face Detection and Age Estimation results

The FEAC component has been tested using the 10% of the total images contained in the Audience DB dataset (testing set) to evaluate both the accuracy in detecting the face and the accuracy in estimating the ages range. As usual, we define:

- **True Positive:** The number of correctly classified items belonging to each class;
- **False Negative:** For each class, the number of items belonging to that class that have been wrongly classified as belonging to another class;
- **False Positive:** For each class, the number of items belonging to a different class, classified by mistake as belonging to that class.

Table 2 concerns the face detection, and it shows the results computed on the testing set. In particular for each age range are reported: (i) the number of faces contained in the testing set (#Faces), (ii) the number of faces correctly detected (True Positive), (iii) the total number of faces undetected (False Negative), (iv) the number of non-faces objects detected as faces (False Positive), and (v) the Face Detection Rate, computed as the true positive over the number of faces in the testing set. The total face detection rate reached is 97.21% with a low level of false-positive and negative. The majority of errors are related to the classes with few training samples [48-63] and [60-100]. Table 3 concerns the age estimation, and for each age range are reported: (i) the number of faces detected correctly by the classifier (#Faces detected), (ii) the number of faces with a correct estimation of the age range (True Positive), (iii) the number of faces belonging to a specific age range that have been wrongly classified as belonging to another age range (False Negative), (iv) the number of faces belonging to another age range that have been wrongly classified as belonging to that age range (False Positive), and (v) the Age Estimation Rate for each age range computed as the true positive over the number of faces detected. The overall age estimation rate obtained in the performed experiments is 95.05%. The majority of errors fall in the [8–12] category reaching 85% of age estimation rate. However, it is worth noting that adolescents in this age range might be difficult to categorize even for a human eye, due to different growth speed. Moreover, it is worth noting that the problem that has been considered in this set of experiments is more complex than the one related to the enforcement of the majority of parental control policies. As a matter of fact, in our experiments, we will consider the ability to discern people with more than 14 years old from people with less than 14 years old.

| | [0-2] | [4-6] | [8-12] | [15-20] | [25-32] | [38-43] | [48-63] | [60-100] | Total |
|----------------------------|--------|--------|--------|---------|---------|---------|---------|----------|--------|
| #Faces | 219 | 158 | 129 | 86 | 326 | 147 | 59 | 60 | 1184 |
| True Positive | 215 | 157 | 125 | 80 | 320 | 144 | 54 | 55 | 1151 |
| False Negative | 4 | 1 | 4 | 6 | 6 | 3 | 5 | 5 | 34 |
| False Positive | 2 | 3 | 4 | 0 | 5 | 0 | 0 | 3 | 17 |
| Face Detection Rate | 98.17% | 99.36% | 96.89% | 93.02% | 98.15% | 97.95% | 91.52% | 91.66% | 97.21% |

TABLE 2 Face detection on Audience Db.

| | [0-2] | [4-6] | [8-12] | [15-20] | [25-32] | [38-43] | [48-63] | [60-100] | Total |
|---------------------|--------|--------|--------|---------|---------|---------|---------|----------|--------|
| #Faces detected | 215 | 157 | 123 | 81 | 315 | 147 | 59 | 58 | 1152 |
| True Positive | 210 | 136 | 111 | 70 | 314 | 146 | 59 | 58 | 1095 |
| False Negative | 5 | 21 | 18 | 12 | 1 | 1 | 0 | 0 | 57 |
| False Positive | 6 | 8 | 18 | 5 | 1 | 0 | 0 | 0 | 38 |
| Age Estimation Rate | 97.67% | 90.24% | 85% | 86.41% | 99.68% | 99.31% | 100% | 100% | 95.05% |

TABLE 3 Age Estimation on Adience Db trained on [0-14] [15-100] age ranges.

4.2 | Simulation Scenario Results

The simulating experiment has been used to evaluate the capability of the framework to estimate the age of spectators from a video stream, simulating what the smart-TV camera is continuously monitoring. To this aim, we simulate the spectators in front of a smart-TV using a video portraying a family in front of a webcam. The video is composed of (i) shares in which the whole family is in front of the camera, (ii) shares with the child alone, and (iii) shares with parents only. The whole video has been reproduced three times, supposing that the spectators will be watching at first content with classification Y, afterward a content with classification PG, and finally one with classification MA. The policies are shown in Section 3.3 have been applied. The duration of the video is 5 minutes and 3 seconds with a resolution of 1280x720 at 25 fps. Setting an interval extraction frames of 5 s, 60 frames have been extracted, of which 24 portraying only adults, 11 representing the child alone, and 25 showing the adults and the child together. Out of 115 faces globally present in the 60 frames, the FEAC component was able to correctly detect 113 faces, thus achieving a detection rate of 98.26%. As we can see from Figure 15 the faces undetected, surrounded by red boxes, are peculiar cases such as face cut, people not looking in the direction of the TV, or people too far from the camera. The FEAC was able to correctly estimate the age range of 103 out of 113 faces, thus achieving an age estimation rate of 91.15%. In particular, considering a single frame, 7 times an adult face (out of 74) has been classified as a child and only 2 times the child has been recognized as an adult. Thus the impact on the user experience would have been minimal since the TV program would have only been suspended one time for a few seconds, during the reproduction, according to the policy. Figure 14 shows the confusion matrix of the age estimation process produced by the FEAC component.

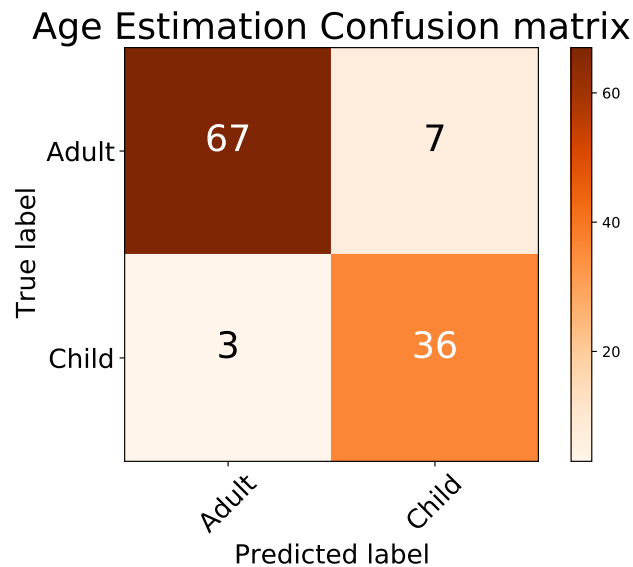


FIGURE 14 Age estimation confusion matrix.




FIGURE 15 Faces misclassified by the FEAC.

4.3 | Real Scenario Results

A real scenario has been reproduced in a house setting. One of the captured frames is reported in Figure 16 , portraying a father with his daughter in front of the TV. The frames have been collected at the same rate of the simulated use case with a resolution 650x650. The frames which have been collected are the following: 3 frames with only one adult, 12 frames with only one child, and 13 frames with the adult and the child. The FEAC was able to detect 41 faces over 41, thus achieving an accuracy of 100%, whereas it was able to correctly classify 39 over 41 ages with 95% of age estimation rate, with the adult misclassified 2 times and the child only once.



FIGURE 16 Real experiments frame.

4.4 | Result comparison

This section summarizes the results obtained in face detection and age estimation. In particular, is shown the comparison among the approach using a single deep learning network (FEAC component) and the deep learning pipeline presented in⁸ (FEL + ACL). Table 4 reports the face detection rate, the age estimation rate, and the computation time of the networks both in the simulation and in the real scenario. The FEAC component that exploits the implementation with a single deep learning network based on Yolo v3, shows an improvement of detection and estimation rate of 2-3% respect to the FEL and ACL components based respectively on the CNN MMOD and the Deep EXpectation architectures. The proposal to use a single deep learning network reduces the overhead cost in computation time. In fact, we obtained a time computation reduction of 32 ms in the real scenario (frame size 650 x 650) and 38 ms in the simulation scenario (frame size 1280 x 720).

| Metric/Network | FEAC | FEL + ACL |
|---|--------|-----------|
| Face Detection Rate Simulation Scenario | 98.26% | 96% |
| Age Estimation Rate Simulation Scenario | 91.15% | 89% |
| Face Detection Rate Real Scenario | 100% | 100% |
| Age Estimation Rate Real Scenario | 95% | 92% |
| Computation Time Real Scenario | 40 ms | 72 ms |
| Computation Time Simulation Scenario | 85 ms | 123 ms |

TABLE 4 Comparison among single network and pipeline approach.

4.5 | Performance Evaluation

This section evaluates the performance of our framework by monitoring the timing actions of each component. Considering the configuration described in Section 3.6 We measured the timings related to:

- UCON Actions
- Face Detection and Age estimation

UCON performances

To perform an extensive evaluation of the performances of the UCON framework, we tested the system starting from a simple case with one attribute until a case with 40 attributes (we recall that in the reference scenario we used only 4 attributes). We repeated each experiment five times and we report the average. As shown in Table 5, the number of attributes increases the time that the UCS needs to evaluate each request. The *SuspendAccess* and *ResumeAccess* have a similar trend. In fact these actions have a similar execution flow: they perform the policy evaluation and, respectively, they perform the operations to suspend or to resume the session.

| UCON action / attribute no. | 1 | 4 | 5 | 10 | 15 | 20 | 25 | 30 | 35 | 40 |
|-----------------------------|-------|-------|-------|-------|------|-------|-------|-------|-------|-------|
| UCON TryAccess | 312.6 | 375.8 | 387.8 | 389.6 | 358 | 380 | 393.6 | 433.4 | 476.2 | 477.8 |
| UCON StartAccess | 79 | 87.4 | 94.2 | 76.8 | 98.2 | 90.6 | 127 | 106.2 | 131.4 | 169.8 |
| UCON RevokeAccess | 52 | 68.3 | 70.2 | 72.8 | 92.6 | 103.8 | 121.4 | 115.6 | 170.4 | 146.4 |
| UCON SuspendAccess | 36.3 | 40.5 | 42.4 | 49.6 | 69.8 | 80.2 | 94.5 | 101.2 | 114.3 | 120.8 |
| UCON ResumeAccess | 74.8 | 81.5 | 84.3 | 90.1 | 99.6 | 116.7 | 135.5 | 151.1 | 178.2 | 190.8 |

TABLE 5 UCON Actions performances (execution times in ms).

Face detection and Age estimation

The time required to extract faces and estimate ages is reported for three different cases: (i) the time to extract faces from a frame on the simulated testbed (1280x720 resolution frame), (ii) the time to extract faces from a frame on the real use case (650x650 resolution frame), and (iii) the time to extract faces from a frame contained a number of faces in the interval [5-30] (650x650 resolution frame). Table 6 shows the described performances computed executing the FEAC component on a GeForce GTX1080 graphic card. The results shown as the time needed to extract faces and estimate ages from a frame is

| Scenario | Time (ms) |
|------------------------------|-----------|
| T simulation FEAC | 85 |
| T real time FEAC | 40 |
| T multiple faces FEAC | 40 |

TABLE 6 FEAC performances.

extremely dependent from the image resolution and the number of faces in the frame. In the simulation scenario we consider as resolution 1280x720 that is higher respect to the resolution in the other two cases. Such aspect, increase the time to extract the faces but can increase the accuracy in the detection and the estimation.

Framework time

Considering the performances evaluation described in the previous paragraphs, we can estimate the total response time of the framework taking into account our reference scenario. In case of policy violation, we are interested to minimize the time that a specific not permitted action can be performed. Thus considering the TV parental guideline violation (case in which a child remains alone during the streaming of a PG program), we can estimate the reaction time as the overall Suspend Response Time of the proposed framework, $SRT_{ExtendedFramework}$, as reported in Formula 1.

$$SRT_{ExtendedFramework} = T_{FEAC} + T_{UCSuspendAccess} + T_{StopStreaming} \quad (1)$$

Considering the real time scenario, the suspend response time of the framework is the sum of the AI engine to detect and estimate the ages (40 ms), the time for policy re-evaluation to decide the suspension of the resource access (40.5 ms), and the time that the PEP needs to perform the stop enforcement on the video streaming (evaluated in 63 ms). Thus, the overall amount of time necessary to stop the reproduction of a video in case of policy not respected in a real scenario is 143.5 ms.

As well as the evaluation of the suspension time, is interesting analyze the Resume Reproduction Time (RRT), i.e. the time necessary to resume the reproduction of the video when the policy is once again matched. Such time is expressed in Formula 2 and it is composed by the AI Engine time to detect the people in front of the camera and their ages (40 ms), the time for policy re-evaluation to decide to resume the access to the resource (81.5 ms) and the time that the PEP employs to resume the streaming (evaluated in 75 ms). Thus the overall amount of time necessary to resume the reproduction of a video due to the attribute changing is 196.5 ms.

$$RRT_{ExtendedFramework} = T_{FEAC} + T_{UCSResumeAccess} + T_{ResumeStreaming} \quad (2)$$

Considering the same operations (stop and resume video execution) using the traditional UCON model, the suspend and resume access would have been replaced by the RevokeAccess in case of policy violation, and a TryAccess in case of resume conditions. In addition, has to be considered the time to start the video (90 ms). The two time are reported in Formula 3 and Formula 4.

$$RRT_{OriginalFramework} = T_{FEAC} + T_{UCSTryAccess} + T_{StartStreaming} \quad (3)$$

$$SRT_{OriginalFramework} = T_{FEAC} + T_{UCSRevokeAccess} + T_{StopStreaming} \quad (4)$$

The estimation time computed considering 4 attributes is 505.8 ms and 171.3 ms respectively for the $RRT_{OriginalFramework}$ and $SRT_{OriginalFramework}$. Such evaluation shows the benefits of the extended UCON framework that, through the introduction of the suspend and resume access, allows to decrease the response time as shown in Table 7 .

| Framework type | SRT time (ms) | RRT time (ms) |
|------------------|------------------|------------------|
| Extended UCON | 143.5 | 196.5 |
| Traditional UCON | 171.3 | 505.8 |

TABLE 7 UCON framework comparison.

5 | RELATED WORK

An application of UCON in IoT environments has been presented in³. The authors propose a distributed model of the standard UCON framework, discussing a smart home use case. An implementation specific for IoT communication protocols is presented in^{2,21}, where the authors present the integration of the UCON paradigm in the MQTT and CoAP workflow. The focus of these works, however, is not centered on obligations, they only exploit in an appropriate way the authorization constructs. On the side of IFTTT an effort to create simple policy algorithms for IoT is presented in²². This work focuses more on the policy specification and how it could be predicted according to the user features and the specific IoT domain. This work, though, does not study the way policies are written, or the obligation format and context but focuses more on the part of how to create a policy according to what is happening in a specific IoT domain. Furthermore, enforcement mechanisms are not considered in their analysis. In⁶, the author describes the development of a specific framework that can modify the concepts of the IFTTT platform so as to provide services that can be used in order to make home automation systems more secure. This work, though, creates specific services in the IFTTT platform and also depends only on the security mechanisms of IFTTT without any control on the installed environment and, also, no continuous way of evaluation or policy enforcement. In⁷, the authors describe also some possible security and privacy risks in the *Applets* of IFTTT.

Concerning the XACML model for policies, in²³, the authors present a description of obligation expression based on examples in Grid and networking security. The work is very specific to this limited environments and does not provide a way to formalize the obligations semantic. Finally, in²⁴, the authors present a variation of XACML specific for obligation extraction. Firstly, they propose a method of creating a new file for obligations in XACML which is produced by the initial file but they do not give a clear description of how the obligation part should be constructed so as to be easily readable by both the PDP and the PEP. Though interesting, their approach is not based on standard XACML, differently from our solution that integrates in XACML 3.0 policies without requiring any modification to the standard architecture and workflow. Usage control for data management has been used for controlling dynamically the right to use files on mobile devices and in the cloud respectively in^{25,1}. These works are mainly focused on the description of the usage control framework infrastructure, presenting generic policies. The present work describes instead a real application of the usage control in a real environment, with effects not only on security, but safety as well.

A framework which exploits the embedded camera in a smart-TV for user authentication is presented in²⁶. By authenticating users the framework also performs parental control enforcement, however, differently from Smart Parental Advisory the age recognition is based on the manual assignment of the age to the recognized face. On the other hand, our framework performs age recognition automatically, without exploiting the identity. Hence, Smart Parental Advisory is able to recognize new adults and children without the need of retraining it for any new spectator, being thus more general. A similar machinery to the one of²⁶ has been patented for commercial use in²⁷. Deep learning for face recognition has been used with high accuracy in²⁸. However, this work only focuses on classification aspects, without considering parental control aspects. A recent work focused instead only on age estimation is the one presented in²⁹, where deep learning algorithms have been used to automatically infer the age of people from face pictures. In our work instead the process of age recognition is more challenging since the faces on which the age is estimated are extracted from live streams. Another effort towards exploiting smart device capabilities is shown in³⁰, where the authors focusing on extending parental control capabilities through rooting on smart-TVs. This approach though needs specifically modified smart-TVs and software installed to them, that may cause instability of the operating system and break of the warranty. Instead our approach does not require such modifications to smart-TVs. Finally, in³¹, the authors suggest a blockchain framework that monitors the actions over resources in a smart home case, rewarding good behaviors and penalizing bad ones. But, the authors do not implement any verification mechanisms on the tasks needed to be fulfilled. They only suggest the usage of surveillance cameras for such a task. In our work, though, we present a face detection and age estimation algorithm, that is able to identify the presence or not of kids and adults when they are in front of a camera.

6 | CONCLUSION

Controlling access rights to operations and devices in Smart Home settings enables the definition of security, safety and management policies able to improve the level of customization and quality of provided services. In this paper we have presented an architecture exploiting UCON and IFTTT to manage safety and security in a smart home setting, focusing on a parental control use case. The presented use case strongly motivate our work, showing the enforcement of safety policies in an interconnected smart home environment, which involves three cooperating and interconnected devices. The proposed approach and the novel introduced models show improvement for what concerns policy expressiveness and overall system effectiveness and performance. By exploiting the IFTTT webservice, this work also provides a method to better exploit the UCON obligations, by using an actual protocol with standardized expressions.

As future extensions, we plan to design more structured IFTTT-based obligations with dedicated tags possibly to be proposed as standard extensions. Furthermore, we plan the design and testing of more complex use cases, with devices having a dedicated UCS and IFTTT capable applications to enforce their own and third parties obligations.

ACKNOWLEDGEMENTS

The activities leading to this research have been partially funded by the EU H2020 project SPARTA GA 830892 and the EU H2020 project SIFIS-Home GA952652.

References

1. Carniani E, D'Arenzo D, Lazouski A, Martinelli F, Mori P. Usage Control on Cloud systems. *Future Gener. Comput. Syst.* 2016; 63: 37–55. doi: 10.1016/j.future.2016.04.010
2. La Marra A, Martinelli F, Mori P, Rizos A, Saracino A. Improving MQTT by Inclusion of Usage Control. In: Wang G, Atiquzzaman M, Yan Z, Choo KR., eds. *Security, Privacy, and Anonymity in Computation, Communication, and Storage - 10th International Conference, SpaCCS 2017, Guangzhou, China, December 12-15, 2017, Proceedings*. 10656 of *Lecture Notes in Computer Science*. Springer; 2017: 545–560
3. La Marra A, Martinelli F, Mori P, Saracino A. Implementing Usage Control in Internet of Things: A Smart Home Use Case. In: *2017 IEEE Trustcom/BigDataSE/ICSS, Sydney, Australia, August 1-4, 2017* IEEE Computer Society, Trustcom/BigDataSE/ICSS, Sydney, Australia, August 1-4, 2017; 2017: 1056–1063
4. Park J, Zhang X, Sandhu RS. Attribute Mutability in Usage Control. In: Farkas C, Samarati P., eds. *Research Directions in Data and Applications Security XVIII, IFIP TC11/WG 11.3 Eighteenth Annual Conference on Data and Applications Security, July 25-28, 2004, Sitges, Catalonia, Spain*. 144 of *IFIP*. Kluwer/Springer; 2004: 15–29
5. Park J, Sandhu RS. The UCON_{ABC} usage control model. *ACM Trans. Inf. Syst. Secur.* 2004; 7(1): 128–174. doi: 10.1145/984334.984339
6. Vorapojpisut S. A Lightweight Framework of Home Automation Systems Based on the IFTTT Model. *JSW* 2015; 10(12): 1343–1350. doi: 10.17706/jsw.10.12.1343-1350
7. Surbatovich M, Aljuraidan J, Bauer L, Das A, Jia L. Some Recipes Can Do More Than Spoil Your Appetite: Analyzing the Security and Privacy Risks of IFTTT Recipes. In: Barrett R, Cummings R, Agichtein E, Gabrilovich E., eds. *Proceedings of the 26th International Conference on World Wide Web, WWW 2017, Perth, Australia, April 3-7, 2017* ACM; 2017: 1501–1510
8. Giorgi G, La Marra A, Martinelli F, Mori P, Saracino A. Smart Parental Advisory: A Usage Control and Deep Learning-Based Framework for Dynamic Parental Control on Smart TV. In: Livraga G, Mitchell CJ., eds. *Security and Trust Management - 13th International Workshop, STM 2017, Oslo, Norway, September 14-15, 2017, Proceedings*. 10547 of *Lecture Notes in Computer Science*. Springer; 2017: 118–133

9. La Marra A, Martinelli F, Mori P, Rizos A, Saracino A. Using IFTTT to Express and Enforce UCON Obligations. In: Heng S, López J., eds. *Information Security Practice and Experience - 15th International Conference, ISPEC 2019, Kuala Lumpur, Malaysia, November 26-28, 2019, Proceedings*. 11879 of *Lecture Notes in Computer Science*. Springer; 2019: 213–231
10. Pretschner A, Hilty M, Basin DA. Distributed usage control. *Commun. ACM* 2006; 49(9): 39–44. doi: 10.1145/1151030.1151053
11. OASIS . eXtensible Access Control Markup Language (XACML) Version 3.0. 2013.
12. Zhang X, Parisi-Presicce F, Sandhu RS, Park J. Formal model and policy specification of usage control. *ACM Trans. Inf. Syst. Secur.* 2005; 8(4): 351–387. doi: 10.1145/1108906.1108908
13. Colombo M, Lazouski A, Martinelli F, Mori P. A Proposal on Enhancing XACML with Continuous Usage Control Features. In: Desprez F, Getov V, Priol T, Yahyapour R., eds. *Grids, P2P and Services Computing [Proceedings of the CoreGRID ERCIM Working Group Workshop on Grids, P2P and Service Computing, 24 August 2009, Delft, The Netherlands]* Springer; 2009: 133–146
14. Martinelli F, Mori P, Saracino A. Enhancing Android permission through usage control: a BYOD use-case. In: Ossowski S., ed. *Proceedings of the 31st Annual ACM Symposium on Applied Computing, Pisa, Italy, April 4-8, 2016* ACM; 2016: 2049–2056
15. Redmon J, Farhadi A. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* 2018.
16. Deng J, Dong W, Socher R, Li LJ, Li K, Fei-Fei L. Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition* Ieee. ; 2009: 248–255.
17. Szegedy C, Ioffe S, Vanhoucke V, Alemi AA. Inception-v4, inception-resnet and the impact of residual connections on learning. In: *Thirty-first AAAI conference on artificial intelligence*; 2017.
18. Rothe R, Timofte R, Gool LV. DEX: Deep EXpectation of Apparent Age from a Single Image. In: *2015 IEEE International Conference on Computer Vision Workshop, ICCV Workshops 2015, Santiago, Chile, December 7-13, 2015* IEEE Computer Society; 2015: 252–257
19. Eidinger E, Enbar R, Hassner T. Age and gender estimation of unfiltered faces. *IEEE Transactions on Information Forensics and Security* 2014; 9(12): 2170–2179.
20. Rothe R, Timofte R, Gool LV. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision* 2018; 126(2-4): 144–157.
21. Rizos A, Bastos D, Saracino A, Martinelli F. Distributed UCON in CoAP and MQTT Protocols. In: Katsikas SK, Cuppens F, Cuppens N, et al., eds. *Computer Security - ESORICS 2019 International Workshops, CyberICPS, SECPRE, SPOSE, and ADIoT, Luxembourg City, Luxembourg, September 26-27, 2019 Revised Selected Papers*. 11980 of *Lecture Notes in Computer Science*. Springer; 2019: 35–52
22. Nadkarni A, Enck W, Jha S, Staddon J. Policy by Example: An Approach for Security Policy Specification. *CoRR* 2017; abs/1707.03967.
23. Demchenko Y, Koeroo O, Laat dC, Sagehaug H. Extending XACML authorisation model to support policy obligations handling in distributed application. In: Schulze B, Fox GC., eds. *Proceedings of the 6th International Workshop on Middleware for Grid Computing (MGC 2008), held at the ACM/IFIP/USENIX 9th International Middleware Conference, December 1-5, Leuven, Belgium* ACM; 2008: 5
24. Chadwick David LM. Obligation Standardization. In: *W3C Workshop on Access Control Application Scenarios*; 2009: 1–5.
25. Lazouski A, Martinelli F, Mori P, Saracino A. Stateful Data Usage Control for Android Mobile Devices. *Int. J. Inf. Sec.* 2017; 16(4): 345–369. doi: 10.1007/s10207-016-0336-y

26. Lee S, Sohn M, Kim D, Kim B, Kim H. Smart TV interaction system using face and hand gesture recognition. In: *IEEE International Conference on Consumer Electronics, ICCE 2013, Las Vegas, NV, USA, January 11-14, 2013*IEEE; 2013: 173–174
27. Kwoh LCCMRJMDEC. Apparatus and method for total parental control of television use. <http://www.freepatentsonline.com/5382983.html>; 1995.
28. Taigman Y, Yang M, Ranzato M, Wolf L. DeepFace: Closing the Gap to Human-Level Performance in Face Verification. In: *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*IEEE Computer Society; 2014: 1701–1708
29. Dong Y, Liu Y, Lian S. Automatic age estimation based on deep learning algorithm. *Neurocomputing* 2016; 187: 4–10. doi: 10.1016/j.neucom.2015.09.115
30. Majchrowicz M, Kapusta P, Faustryjak D, Jackowska-Strumillo L. System for remote parental control and management of rooted smart TVs. In: *2018 International Interdisciplinary PhD Workshop (IIPhDW)*; 2018: 357-360.
31. Suhaad SA, Mashiko K, Ismail NB, Abidin MHZ. Blockchain Use in Home Automation for Children Incentives in Parental Control. In: *Proceedings of the 2018 International Conference on Machine Learning and Machine IntelligenceMLMI2018*. Association for Computing Machinery; 2018; New York, NY, USA: 50–53

