Towards a Classification Model for Identifying Risky IFTTT Applets

Bernardo Breve[1], Gaetano Cimino[1] and Vincenzo Deufemia[1]

[1]University of Salerno, via Giovanni Paolo II, Fisciano (SA), 84084, Italy

Abstract

With the rapid growth of Internet-of-Things (IoT) devices, especially in the context of smart homes, we witnessed the rise of different services aimed at providing end-users with tools for the definition of custom behaviors. Among these, If-This-Than-That (IFTTT) became the most used end-user programming tool for creating *event-condition-action* (ECA) rules. However, while defining such rules, end-users might expose both their smart devices and personal information to security and privacy threats. This paper presents the progress achieved in the definition of a classification model based on neural networks for the identification of possible security and privacy issues within an IFTTT applet.

Keywords

End-user programming, IFTTT service, Security and Privacy threats, Internet of Things (IoT)

1. Introduction

Internet-of-Things (IoT) platforms and devices are being widely used in industrial and domestic contexts. Recently, we witnessed the rise of services that facilitate the interoperability between different smart devices and cloud services, offering tailoring to end-users through the usage of specifically studied paradigms. Such functionalities allow users to define custom behaviors for both the devices and the services by means of simple conditional rules [1, 2].

In the IoT domain the behavioral rules, which are commonly defined as Event-Condition-Action (ECA) rules, express how and when IoT devices have to be activated in reaction to detected events [1]. The end-user development (EUD) platforms allow for the definition of such rules through appropriate levels of abstraction, helping the user to define an ECA rule in a simple way. In particular, while interacting with these platforms, the end-users are not required to deal with technical details that may be complex to understand, simplifying the definition process by means of basic semantic concepts [3]. Among the released services for designing ECA rules, IFTTT (If-This-Then-That) affirmed itself as one of the most used for its simplicity in building rules. Furthermore, IFTTT provides a variety of pre-composed rules that the end-user can use in his/her smart environment.

The behaviors defined through ECA rules can concern different aspects of a smart environment, for example, it is possible to define the switching on or off behavior of smart lights relative

EMPATHY: Empowering People in Dealing with Internet of Things Ecosystems. Workshop co-located with INTERACT 2021, August 30, 2021, Bari, Italy

bbreve@unisa.it (B. Breve); gcimino@unisa.it (G. Cimino); deufemia@unisa.it (V. Deufemia)

0000-0002-3898-7512 (B. Breve); 0000-0001-8061-7104 (G. Cimino); 0000-0002-6711-3590 (V. Deufemia)

to the position of the user, as well as the publication or the gathering of social network posts following the occurrence of certain conditions. In addition, through ECA rules, it is also be possible to define behaviors triggered in response to the occurrence of certain events which could jeopardize the security of the smart environment [4].

However, an important aspect to tackle when dealing with ECA rule definition is related to the possibility of introducing inconsistencies, which might lead to security and privacy risks [5, 6]. Indeed, the low technical knowledge of end-users might not allow them to notice immediately the existence of such inconsistencies, and the consequences that a *risky* ECA rule could entail. As an example, let us consider the following ECA rule: *"When I go to the gym, make a post on Facebook"*. Although nothing wrong could be immediately deducted from this rule, we can identify a serious privacy disclosure concern through painstaking analysis. In fact, such a rule could provide useful information to a malicious individual, e.g., a thief might know when you are out the house. Thus, it is crucial to support end-users during the definition of ECA rules by providing information about their potential risks.

In the literature, Artificial Intelligence approaches aimed at extracting knowledge from ECA rules have been applied. In fact, in [7] authors proposed a neural network architecture based on the attention mechanism, aiming at inferring the morphology on an IFTTT rule from textual descriptions. In particular, such a mechanism allows each element of the text, namely token, to be assigned to a weight determining its importance within the description. Each weight depends not only on the single token but also on the context it belongs to, i.e. the tokens preceding and succeeding it. Finally, by analyzing the weights calculated on each token, the authors can extract details regarding what elements of the text constitute the event and the action of an IFTTT rule. Furthermore, many other works in literature propose security- and privacy-preserving solutions in several application domains [8, 9]. Unfortunately, few efforts have been devoted to the identification of possible security and privacy risks underlying ECA rules. This is a serious shortcoming if taking into consideration the large number of users who are approaching these new technologies.

This position paper highlights the preliminary assessments obtained in the definition of a classification model for identifying risky IFTTT applets. The proposed solution exploits neural networks for analyzing the IFTTT applets' elements and predicting whether they entail privacy or security threats for both the users and the smart environment.

2. The IFTTT Applet Structure

An IFTTT applet is built of two components: a *trigger* that causes the applet to run, and an *action* that define what should happen. The trigger represents the "**this**", while the action represents the "**that**" of the statement "*if* ***this***, *then* ***that***" [6]. Each trigger and action belongs to a "channel", which specifies the service provider that the trigger or the action is associated with, (e.g. the IoT device manufacturer, or a social media company) called *trigger channel* and *action channel*, respectively. The triggers associated to the channels can be either simple as "when a certain time is reached" or more complex, such as "when the camera detects motion". Similarly, the actions depending on the selected channels can range from sending an email to arming a security system. Figure 1 shows an example of ECA rule and highlights the components

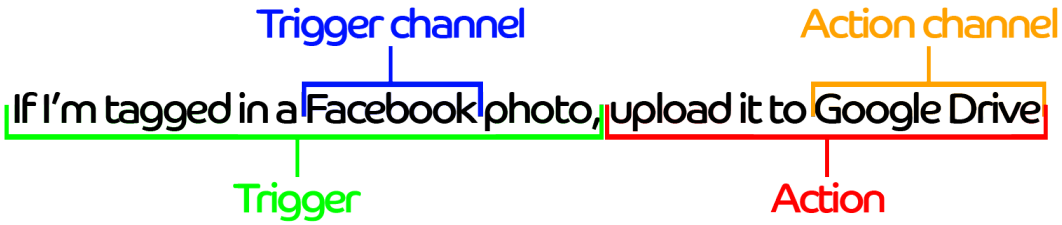

Figure 1: An example of IFTTT applet highlighting the trigger and action components

of the corresponding IFTTT applet.

An applet has also associated a *title*, which is a short text summarizing what is expected from the applet, and a *description*, which should provide details about its triggers, its actions, and any other useful information users should be aware of.

3. Identifying Risky IFTTT Applets

To deal with the problem of classifying IFTTT applets concerning underlying security or privacy risks, we considered the application of a Long Short-Term Memory (LSTM) model. However, the latter require large amount of data to achieve effective predictions. To this end, we considered a dataset of 320k applets published by Mi *et al.*, obtained through the application of crawling activities on the IFTTT website from November 2016 to April 2017 [10]. The dataset preserves the structure of the IFTTT rules, i.e., the decomposition of the applet in trigger, action, and the corresponding channels.

To evaluate the supervised learning technique for classifying risky IFTTT applets, we generated a training set of applets derived from the initial dataset. In particular, we labeled the applets concerning the risk they could represent for the end-user's security and privacy, according to the four classes defined in [5]:

- **Harmless**: apparently harmless applet (to which class 0 has been associated).
- **Personal**: causes loss of sensitive data (to which class 1 has been associated). This damage is self-inflicted since any damage is the result of user behavior.
- **Physical**: damage to physical health or goods (to which class 2 has been associated). This damage is external as a third party can potentially inflict the damage.
- **Cybersecurity**: interruption of an online service or distribution of malware (to which class 3 has been associated). This damage is external as a third party can potentially inflict the damage.

Initially, a small number of applets were selected and labeled. Later, to speed up the manual labeling process, we decided to adopt a sentence embedding technique. These types of techniques can represent entire sentences and their semantic meaning in the form of vectors, thus helping computers to understand the context, the intent, and other nuances within the texts. In particular,

we employed SentenceBERT, a model introduced in 2018, based on a Siamese network-like architecture that accepts a sentence as input and generates its vector representation. The sentence is supplied as input to a BERT model and a pooling layer to generate its 512-dimensional embedding. This model was applied to the already labeled applets, taking into consideration, for each of them, both the title and the description. To filter the applets of a given class, for every single already labeled applet, we employed the cosine similarity on the sentence embeddings between the labeled applet and all the remaining unlabeled ones. For each of these operations, an Excel file was generated in which the unlabeled applets were arranged in descending order based on their similarity to the considered labeled applet. Even in this case, we pursued with manual labeling, but this strategy allowed us to derive other applets "similar" to those already labeled and to discover new types of applets.

Therefore, we manually labeled a subset of 2000 applets taken from the complete dataset[1], obtaining an initial training set, which was given in input to several neural network models. For each applet, we selected the following textual information: the title of the applet, its description, and the textual description of both the trigger and the action of the applet. Before applying LSTM, the texts were preprocessed, and features were extracted and represented into a 50-dimensional vector using GloVe model. Initial results showed an average accuracy score of around 82% with LSTM. Currently, we are evaluating the effectiveness of other classification models by varying the considered features.

At the workshop, we will present the results achieved with the trained model, and discuss the challenges related to the creation of larger labeled datasets of ECA rules and to the communication of the identified risks to end-users.

Acknowledgments

This work has been supported by the Italian Ministry of Education, University and Research (MIUR) under grant PRIN 2017 "EMPATHY: EMpowering People in deAling with internet of THings ecosYstems" (Progetti di Rilevante Interesse Nazionale – Bando 2017, Grant 2017MX9T7H).

References

[1] G. Desolda, C. Ardito, M. Matera, Empowering end users to customize their smart environments: Model, composition paradigms, and domain-specific tools, ACM Trans. Comput.-Hum. Interact. 24 (2017).

[2] G. Ghiani, M. Manca, F. Paternò, C. Santoro, Personalization of context-dependent applications through trigger-action rules, ACM Trans. Comput.-Hum. Interact. 24 (2017).

[3] F. Corno, L. De Russis, A. M. Roffarello, A high-level semantic approach to end-user development in the internet of things, International Journal of Human-Computer Studies 125 (2019) 41–54.

[1] Note that the second author of the paper was in charge of manually classifying the applet, while the other two cross-checked the data.

[4] B. Breve, G. Desolda, V. Deufemia, F. Greco, M. Matera, An end-user development approach to secure smart environments, in: D. Fogli, D. Tetteroo (Eds.), Proceedings of 8th International Symposium on End-User Development IS-EUD 2021, Limassol, Cyprus, July 6-8, Springer, Berlin, 2021.

[5] M. Surbatovich, J. Aljuraidan, L. Bauer, A. Das, L. Jia, Some recipes can do more than spoil your appetite: Analyzing the security and privacy risks of IFTTT recipes, in: Proceedings of the 26th International Conference on World Wide Web, WWW '17, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 2017, p. 1501–1510.

[6] C. Cobb, M. Surbatovich, A. Kawakami, M. Sharif, L. Bauer, A. Das, L. Jia, How risky are real users' IFTTT applets?, in: Proceedings of the Sixteenth Symposium on Usable Privacy and Security (SOUPS 2020), USENIX Association, San Diego, CA, 2020, pp. 505–529.

[7] C. Liu, X. Chen, E. C. Shin, M. Chen, D. Song, Latent attention for if-then program synthesis, Advances in Neural Information Processing Systems 29 (2016) 4574–4582.

[8] N. Mehdy, C. Kennington, H. Mehrpouyan, Privacy disclosures detection in natural-language text through linguistically-motivated artificial neural networks, in: International Conference on Security and Privacy in New Computing Environments, Springer, Cham, 2019, pp. 152–177.

[9] Z. Yang, M. Yang, Y. Zhang, G. Gu, P. Ning, X. S. Wang, Appintent: Analyzing sensitive data transmission in android for privacy leakage detection, in: Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security, Association for Computing Machinery, New York, NY, USA, 2013, pp. 1043–1054.

[10] X. Mi, F. Qian, Y. Zhang, X. Wang, An empirical characterization of IFTTT: ecosystem, usage, and performance, in: Proceedings of the 2017 Internet Measurement Conference, Association for Computing Machinery, New York, NY, USA, 2017, pp. 398–404.