

On Privacy Disclosure from User-Generated Content of Automation Rules^{*}

Bernardo Breve^{1,*,\dagger}, Gaetano Cimino^{1,*,\dagger}, Vincenzo Deufemia^{1,\dagger} and Annunziata Elefante^{1,*,\dagger}

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

Abstract

Trigger-Action Platforms (TAPs) are systems that enable users to automate routine tasks, such as turning off lights at a specific time, without requiring technical skills. In the process of creating automation rules, users are prompted to provide descriptions in natural language, which are referred to as User-Generated Content (UGC), such as the title that explains the intended behavior of the rule. However, UGC may contain sensitive information that could expose users to unwanted situations or be exploited by cyber attackers. This position paper provides an initial assessment of the risks associated with UGC in TAPs and discusses the use of NLP techniques to mitigate these risks. Additionally, the paper highlights the need for further research to better understand the impact of UGC on privacy and to develop effective privacy-preserving mechanisms for TAPs.

Keywords

Trigger-action platforms, Privacy leakage, User-generated content, Automation rules, Smart homes

1. Introduction

The blossoming of smart technology in contemporary society has significantly impacted every aspect of everyday life. Terms such as “smart cities”, “smart houses”, “smart mobility”, and “smart health” are now well-known within our vocabulary. The Internet of Things (IoT) [1], has revolutionized the way end-users interact with technology-injected variants of everyday objects, allowing for unprecedented control and management over the Internet.

In order to simplify the way end-users can interact and customize smart devices, the End-User Development (EUD) [2, 3] paradigm has become in our days increasingly popular, enabling individuals to access and utilize IoT technology throughout various domains, from business to healthcare (eHealth) [4]. Smart Houses, in particular, represent a rapidly growing area of interest and application for IoT technology, enabling users to control all aspects of their home, such as lighting, television, air conditioning, and garage. To simplify the automation of all these household tasks, users can utilize Trigger-Action Platforms (TAPs) to create automation rules

IS-EUD 2023: 9th International Symposium on End-User Development, 6-8 June 2023, Cagliari, Italy

*Corresponding author.

^{\dagger}These authors contributed equally.

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

🆔 0000-0002-3898-7512 (B. Breve); 0000-0001-8061-7104 (G. Cimino); 0000-0002-6711-3590 (V. Deufemia); 0009-0001-7141-6105 (A. Elefante)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

that incorporate triggers, conditions, and actions [5]. These rules connect online services that represent both digital and physical resources and are executed when the conditions associated with triggers are satisfied, leading to the completion of the action. For example, a user may program a rule that automatically turns on lights at sunset or a rule for activating the air-conditioning at a specific time of the day.

The most popular TAP to date is If-This-Then-That (IFTTT)¹, a web-based platform that boasted more than 18 million users as of 2020 [6]. The popularity of IFTTT can be attributed to its user-friendly interface, which allows even novice users to easily create new rules from scratch or use pre-existing ones from its catalog. Searching for rules in the catalog is also a straightforward process, as each rule is defined by a specific natural language textual title and description. These two components, formally addressed as User-Generated Content (UGC) [7], offer significant benefits, as they assist the rule creator in remembering the behavior of their automatism and provide a means of support for new users to comprehend how the rule operates. For example, a user can readily activate a rule such as:

IF I PUBLISH A PHOTO ON FACEBOOK THEN SHARE IT ON INSTAGRAM

which by means of a UGC could be described as so:

Keep your Instagram followers updated! This rule allows you to automatically synchronize any new photo you upload on Facebook with your Instagram profile.

At first glance, this rule appears to provide users with significant benefits as it saves individuals from having to upload their photos manually on both social networks. However, automation rules can intrinsically raise privacy and security concerns either for the smart environment or the users, especially when such rules are defined and used by inexperienced users [6, 8, 9, 10, 11]. With regard to the previous example, there might be scenarios where a user would not want to share his or her photos with followers of one social network over another one, leading to unwilling uploads of photos that could cause embarrassment.

In addition, UGC employed by users to describe their rules may provide further damage to their privacy. In fact, users might mistakenly disclose sensitive information when explaining the intended behavior of their rules. Alternatively, a user may choose to allow the platform to automatically complete the fields with relevant information. However, in either case, a user may inadvertently publish a rule with private and personal data (e.g., the user's real email), as shown by the following description:

*When Lautaro Martinez publishes a photo on Instagram, then send an email to **EMAIL ADDRESS***

Therefore, end-users might thus publicly share their sensitive information, particularly since the typical user of these platforms lacks technical background and may be unaware of the potential privacy risks implied by the degree of freedom when typing UGC.

¹<https://ifttt.com>

This position paper outlines a viable solution to mitigate the sensitive information leakage issue within the context of TAPs.

2. Identifying Privacy Leakage from UGC in the TAP domain

In recent years, several studies have highlighted the sensitive information that is inadvertently disclosed by users of automation platforms, such as TAPs [12]. In particular, researchers have investigated the possibility of inferring and constructing a complete profile of the user from the release of personal data on the Internet, without the user being aware of the harm involved [9, 13, 14, 15, 16].

Identifying vulnerabilities in the domestic environment, particularly in the smart devices that are utilized by millions of users on a daily basis in their houses, is a related area of concern regarding privacy and data leakage. In fact, if an attacker gains knowledge of all the rules published by a user on a TAP, s/he could potentially descend to the level of individual devices [17] and deduce private information about the user. In such cases, it is imperative to conduct an analysis of the IoT infrastructure to identify and mitigate these security risks.

The information pertaining to personal data and IoT devices is derived from the unregulated usage of TAPs by users who may not possess a comprehensive understanding of the internal mechanisms of these systems. As a result, when users divulge information through UGC, they may not fully contemplate the ramifications that even a solitary piece of sensitive information could have on their privacy. UGC in the TAP domain has the potential to cause privacy breaches in various ways. For example, UGC may inadvertently contain personal information, such as location data or personal identifiers, which can be easily accessed by third parties, including attackers and data brokers. Additionally, UGC may be utilized to uncover personal information by identifying patterns of behavior or preferences. For instance, a user who frequently posts about their workout routine may be inferred to be health-conscious, potentially making them a target for health-related advertisements or offers. It is crucial for users to be cognizant of the potential risks associated with UGC in the TAP domain and to take necessary steps to safeguard their privacy.

One promising strategy for addressing the problem of privacy leakage in UGC is the application of Natural Language Processing (NLP) techniques to analyze and comprehend human language used by online users [10]. These techniques can be employed in multiple ways to help identify any sensitive information being shared. For instance, NLP can detect personal identifiers like names, addresses, and phone numbers, as well as sensitive data such as financial information, health data, or passwords. Additionally, it can recognize patterns of behavior or preferences that may reveal sensitive information about a person. Finally, NLP can also scrutinize metadata linked to UGC, including timestamps, locations, and devices used to post the content. While these elements may seem meaningless when considered alone, they could potentially provide malicious individuals with useful information to plan attacks. For instance, if a thief is aware that a user has activated the rule “Turn off living room lights when I leave home”, they could examine the rule-targeted device (the lights) and its location (the living room) to determine the right moment to carry out a theft.

An NLP-based methodology for achieving such goals is the employment of Named Entity

Recognition (NER) techniques, which focus on extracting and classifying from texts different types of entities according to the domain of interest [18]. In the TAP domain, the entities should refer to the users' information and the smart devices and online services they use within automation rules. Specifically, it is necessary to define specific labels, such as **PERSON** to indicate a person's first and/or last name, **ORG** to denote an online service, and **SENS** to highlight sensitive data. Below is an example demonstrating the application within the rule description shown in Section 1:

When Lautaro Martinez **PERSON** publishes a photo on Instagram **ORG**, then send an email to l.martinez@unisa.com **SENS**

In conclusion, through the application of Natural Language Processing (NLP), we can gain a deeper understanding of the potential privacy risks associated with UGC in the TAP domain and take measures to mitigate them.

At the workshop, we will discuss how the involvement of NLP techniques can benefit the achievement of the discussed goals.

Acknowledgments

This work has been supported by the Italian Ministry of University and Research (MUR) 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] L. Atzori, A. Iera, G. Morabito, The internet of things: A survey, *Computer networks* 54 (2010) 2787–2805.
- [2] P. Markopoulos, J. Nichols, F. Paternò, V. Pipek, End-user development for the internet of things, *ACM Transactions on Computer-Human Interaction (TOCHI)* 24 (2017) 1–3.
- [3] B. R. Barricelli, F. Cassano, D. Fogli, A. Piccinno, End-user development, end-user programming and end-user software engineering: A systematic mapping study, *Journal of Systems and Software* 149 (2019) 101–137.
- [4] S. S. Mishra, A. Rasool, IoT health care monitoring and tracking: A survey, in: *Proceedings of 3rd International Conference on Trends in Electronics and Informatics (ICOEI)*, IEEE, 2019, pp. 1052–1057.
- [5] G. Ghiani, M. Manca, F. Paternò, C. Santoro, Personalization of context-dependent applications through trigger-action rules, *ACM Transactions on Computer-Human Interaction (TOCHI)* 24 (2017) 1–33.
- [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 16th USENIX Conference on Usable Privacy and Security*, USENIX Association, 2020, pp. 505–529.

- [7] X. Chen, X. Song, R. Ren, L. Zhu, Z. Cheng, L. Nie, Fine-grained privacy detection with graph-regularized hierarchical attentive representation learning, *ACM Transactions on Information Systems (TOIS)* 38 (2020) 1–26.
- [8] B. Breve, G. Cimino, V. Deufemia, Towards explainable security for ECA rules, in: *Proceedings of the 3rd International Workshop on Empowering People in Dealing with Internet of Things Ecosystems (EMPATHY '22)*, volume 3172 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2022, pp. 26–30.
- [9] Y.-H. Chiang, H.-C. Hsiao, C.-M. Yu, T. H.-J. Kim, On the privacy risks of compromised trigger-action platforms, in: *Proceedings of 25th European Symposium on Research in Computer Security (ESORICS 2020)*, Springer, 2020, pp. 251–271.
- [10] B. Breve, G. Cimino, V. Deufemia, Identifying security and privacy violation rules in trigger-action IoT platforms with NLP models, *IEEE IoT J* 10 (2023) 5607–5622.
- [11] 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*, ACM, 2017, p. 1501–1510.
- [12] R. Xu, Q. Zeng, L. Zhu, H. Chi, X. Du, M. Guizani, Privacy leakage in smart homes and its mitigation: IFTTT as a case study, *IEEE Access* 7 (2019) 63457–63471.
- [13] X. Chen, X. Song, G. Peng, S. Feng, L. Nie, Adversarial-enhanced hybrid graph network for user identity linkage, in: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 1084–1093.
- [14] A. Abbas, J. Holmberg, Information extraction from short text messages, *LU-CS-EX 2019-18* (2019).
- [15] F. Erlandsson, M. Boldt, H. Johnson, Privacy threats related to user profiling in online social networks, in: *Proceedings of International Conference on Privacy, Security, Risk and Trust and International Conference on Social Computing*, IEEE, 2012, pp. 838–842.
- [16] X. Song, X. Wang, L. Nie, X. He, Z. Chen, W. Liu, A personal privacy preserving framework: I let you know who can see what, in: *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018, pp. 295–304.
- [17] S. Rizvi, R. Pipetti, N. McIntyre, J. Todd, I. Williams, Threat model for securing internet of things (IoT) network at device-level, *Internet of Things* 11 (2020) 100240.
- [18] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, C. Dyer, Neural architectures for named entity recognition, *arXiv preprint arXiv:1603.01360* (2016).