Extractor: Extracting Attack Behavior from Threat Reports

Kiavash Satvat
University of Illinois at Chicago
ksatva2@uic.edu

Rigel Gjomemo
University of Illinois at Chicago
rgjome1@uic.edu

V.N. Venkatakrishnan
University of Illinois at Chicago
venkat@uic.edu

531.00 ©2021 IEEE | DOI: 10.1109/EUROSP51992.2021.00046

***Abstract*—The knowledge on attacks contained in Cyber Threat Intelligence (CTI) reports is very important to effectively identify and quickly respond to cyber threats. However, this knowledge is often embedded in large amounts of text, and therefore difficult to use effectively. To address this challenge, we propose a novel approach and tool called** EXTRACTOR **that allows precise automatic extraction of concise attack behaviors from CTI reports.** EXTRACTOR **makes no strong assumptions about the text and is capable of extracting attack behaviors as provenance graphs from unstructured text. We evaluate** EXTRACTOR **using real-world incident reports from various sources as well as reports of DARPA adversarial engagements that involve several attack campaigns on various OS platforms of Windows, Linux, and FreeBSD. Our evaluation results show that** EXTRACTOR **can extract concise provenance graphs from CTI reports and show that these graphs can successfully be used by cyber-analytics tools in threat-hunting.**

I. INTRODUCTION

Cyber Threat Intelligence (CTI), as commonly reported in technical reports, whitepapers, blogs, and newsgroups, is

of the malware's geographical origin, though interesting, does not contribute to the description of the malware behavior in a system.

Challenge 2. CTI text complexity An important assumption of the previous approaches is that the text structure of CTI reports is (a) relatively simple [52] or (b) that it follows a specific grammatical structure [45] or (c) assuming some patterns in describing concepts [88] or (d) considering stable grammatical relations in the presentation of the sentence in the form of subject, verb and object [52], [45]. While these assumptions do not interfere with the goal of these works to extract IOCs and threat action in isolation, in fact, the majority of CTI reports contain much more complex domain-specific contexts (see Section II), which makes the extraction of attack behavior and causal inference more challenging. The CTI reports' syntactic and semantic complexities, the prevalence of technical terms, and lack of proper punctuation in these reports [62] can easily impact the interpretation of the report