

Insight into Insiders and IT: A Survey of Insider Threat Taxonomies, Analysis, Modeling, and Countermeasures

IVAN HOMOLIAK, STE-SUTD Cyber Security Laboratory, Singapore University of Technology and Design and Faculty of Information Technology, Brno University of Technology

FLAVIO TOFFALINI and JUAN GUARNIZO, STE-SUTD Cyber Security Laboratory, Singapore University of Technology and Design

YUVAL ELOVICI, STE-SUTD Cyber Security Laboratory, Singapore University of Technology and Design

MARTÍN OCHOA, STE-SUTD Cyber Security Laboratory, Singapore University of Technology and Design and Department of Applied Mathematics and Computer Science, Universidad del Rosario

Insider threats are one of today's most challenging cybersecurity issues that are not well addressed by commonly employed security solutions. Despite several scientific works published in this domain, we argue that the field can benefit from our proposed structural taxonomy and novel categorization of research that contribute to the organization and disambiguation of insider threat incidents and the defense solutions used against them. The objective of our categorization is to systematize knowledge in insider threat research, while using existing grounded theory method for rigorous literature review. The proposed categorization depicts the workflow among particular categories that include: 1) *Incidents and datasets*, 2) *Analysis of incidents*, 3) *Simulations*, and 4) *Defense solutions*. Special attention is paid to the definitions and taxonomies of the insider threat; we present a structural taxonomy of insider threat incidents, which is based on existing taxonomies and the 5W1H questions of the information gathering problem. Our survey will enhance researchers' efforts in the domain of insider threat, because it provides: a) a novel structural taxonomy that contributes to orthogonal classification of incidents and defining the scope of defense solutions employed against them, b) an overview on publicly available datasets that can be used to test new detection solutions against other works, c) references of existing case studies and frameworks modeling insiders' behaviors for the purpose of reviewing defense solutions or extending their coverage, and d) a discussion of existing trends and further research directions that can be used for reasoning in the insider threat domain.

Additional Key Words and Phrases: Insider threat, malicious insider threat, unintentional insider threat, masqueraders, traitors, grounded theory for rigorous literature review, 5W1H questions.

1 INTRODUCTION

Insider threats are one of the most challenging attack models to deal with in practice. According to a recent survey, 27% of all cyber crime incidents were suspected to be committed by insiders, and 30% of respondents indicated that the damage inflicted by insiders was more severe than the damage caused by outside attackers [Trzeciak 2017]. Similar numbers are reported by [Collins et al. 2016]: "23% of electronic crime events were suspected or known to be caused by insiders," while 45% of the respondents thought that the consequences were worse than the consequences of outsiders. In another survey investigating economic crime [PWC 2017], internal fraudsters acted as the main perpetrator in 29% of cases. According to a survey conducted by Vormetric, Inc. [Kellett 2015], only 11% of respondents felt that their organization was not vulnerable to insider attacks, while 89% felt at least somewhat vulnerable to insider attacks.

In recent years, famous whistle-blowers have made their way into the media, such as the high profile data leakage cases involving Edward Snowden or Chelsea Manning (see [Collins et al. 2016]

0000. 0360-0300/0000/9-ART00 \$15.00

<https://doi.org/0000001.0000001>

for a collection of famous insider threat cases). Although such cases may be viewed as security issues breaking the confidentiality of secret information, they may also be viewed as human loyalty manifested to the country or the society (see Section 3 for further discussion).

In general, insiders are authorized users that have legitimate access to sensitive/confidential material, and they may know the vulnerabilities of the deployed systems and business processes. Many attacks caused by malicious insiders are more difficult to detect compared to those of external attackers whose footprints are harder to hide [Moore 2016]. In addition, there has been an increasing trend of unintentional insider threat in recent years [Collins et al. 2016]. Therefore, the motivation for dealing with insider threat is very high and is likely to grow.

Concerns about insider threat in the literature are not new, and there is an impressive body of knowledge in this broad field. In the last decade there have also been several attempts to survey this field. However, after reviewing such works, we encountered various shortcomings and identified the need for an up-to-date, more comprehensive survey. For instance, some surveys focus exclusively on detection approaches [Bertacchini and Fierens 2008; Gheyas and Abdallah 2016; Salem et al. 2008; Sanzgiri and Dasgupta 2016] or lack a systematical approach to the categorization of the literature [Azaria et al. 2014]. The objective of this work is thus to conduct an extensive literature review in the field of insider threat, while aiming at systematization of the knowledge and research conducted in this area. We view this as important for both the researchers that design experimental defense solutions, as well as the security practitioners who seek to understand the problem and are tasked with selecting or implementing appropriate defense solutions for their specific needs.

1.1 Survey Approach

Our main objective for this work is to address the identified gaps and consolidate the information contained in existing surveys, by including a more exhaustive and up-to-date literature set, emphasizing the review and unification of taxonomies, and using a systematical approach for the categorization of the literature. To this end, we have revised and updated the bibliography of previous surveys, and to ensure that we have not inadvertently left out any relevant work, we manually checked the top 100 best ranked papers in the field according to Google Scholar, as well as the 100 most cited papers in the Web of Science databases, by querying the term *insider threat*. Overall, our sample set contained 322 works, and 108 of them were filtered out based on our inclusion and exclusion criteria (explained below). Note that given the vast amount of work in the field, the goal of this work was not to exhaustively cover all of the literature, but to select a large enough literature set to review the state-of-the-art and to propose a reasonable categorization of it.

Survey Scope. The scope of our survey is based on the following criteria: **a)** The articles included in this survey were selected based on a widespread view of the insider threat problem, ranging from definitions and taxonomies of insider threat, and analysis and modeling of incidents, to conceptual defense solutions and their proofs of concept; **b)** We focused on studies for which the insider threat problem was the main subject; therefore, we excluded papers that only mentioned the insider threat problem tangentially. We did, however, include several examples of masquerade detection approaches, which are related to the identity theft problem but are also considered part of the insider threat problem; **c)** We excluded non-unique papers from the same group of authors who presented a concept/approach across multiple papers. The older version of a study was usually superseded by the newer one, except in cases in which the newer version contained fewer details or also focused on another idea.

After specifying the scope of the survey, we applied an iterative process to construct a literature categorization based on grounded theory for rigorous literature review [Wolfswinkel et al. 2013], which serves as guidelines on an analysis and presentation of findings in a particular field of

research. While processing the set of papers, we identified several abstract concepts, proposed a workflow-based categorization, and rearranged the bibliography according to it. The proposed categorization consists of a review of datasets, case studies, analysis and modeling of incidents, simulations, conceptual and practical defense solutions, and best practices. Note that our main categories are not meant to be disjoint, as there are works that address aspects of various proposed categories. However, we believe that our categorization offers useful dimensions with which to classify works in the literature, while enabling researchers to identify relevant related work.

1.2 Contributions

In sum, this work presents a novel insider threat survey that aims to be comprehensive, yet succinct and easy to access by researchers looking for new avenues to explore or to learn about the subject. The main contributions of this survey can be summarized as follows: **a)** To the best of our knowledge, this is the first work that systematically categorizes heterogeneous insider threat studies and thereby enables readers to obtain a panoptic view on this disparate topic. **b)** We survey existing taxonomies of the insider threat problem, and based on them, we propose a practical and unified taxonomy that can be used to classify: 1) an insider threat incident, or 2) specialization/coverage of a defense solution. **c)** We aggregate information about publicly available datasets that can be utilized for testing a detection solution and comparison with other works included in this survey that have previously used the datasets. **d)** We identify open questions and challenges in insider threat detection, prevention, and mitigation, while proposing further research directions.

The rest of the paper is organized as follows. Section 2 briefly reviews existing surveys in the field and also mentions their contributions, shortcomings, and how our work differs from them. In Section 3 we discuss the scope of insider threat and provide a survey of definitions and existing taxonomies; based on existing taxonomies, we propose a structural taxonomy of insider threat incidents. After providing this background, we change our focus to the categorization of research conducted in this domain, which is discussed in Section 4. A detailed description and the subcategorization of specific major categories is presented as follows: incidents and datasets are dealt within Section 5, an analysis of incidents is provided in Section 6, simulation research is presented in Section 7, and defense solutions are covered in Section 8. The last part – Section 9 – concludes the paper and proposes further directions in the field.

2 EXISTING SURVEYS

In this section we provide a brief summary of existing surveys involving insider threats, and then we describe how our proposed taxonomy differs from them. Salem et al. [2008] introduced a taxonomy of malicious insiders, dividing them into two categories according to the knowledge they have about the target system: traitors and masqueraders. The authors reviewed the literature on insider detection works and divided the works into three types of approaches: a) “host-based user profiling approaches,” b) “network level approaches,” and c) “integrated approaches.” Network level and host-based profiling may have, according to the authors, a high chance of detecting traitors, while host-based user profiling may be successful in identifying masqueraders. Also, the authors posited that malicious actions of insiders occur at the application and business process levels. Hunker and Probst [2011] proposed a categorization of the research based on a combination of psycho-social input data with technical data. The resulting categories consist of three types of approaches to insider threat detection: 1) “sociological, psychological, organizational,” 2) “socio-technical,” and 3) “technical.” The authors emphasized that successful insider threat detection techniques require a combination of various approaches. Azaria et al. [2014] divided related works into six categories, likely based on the most significant trends in the field: 1) “psychological and social theories,” 2)

“anomaly-based approaches,” 3) “honeypot-based approaches,” 4) “graph-based approaches,” 5) “game theory approaches,” and 6) “motivating studies.” However each category draws from a different dimension of possible categorizations, and thus a new approach may not fit any of the proposed categories. The authors aimed to reinterpret principles and results of the approaches included in their survey. Ophoff et al. [2014] conducted a literature review and classified insider threat research based on “50 top ranked MIS journals from the ISI world, as well as top security journals in the IS domain.” The search term “*insider threat*” yielded over 600 results, however after removing irrelevant research and filtering duplicates they had 90 papers, which they divided into six categories and 13 subcategories using the grounded theory approach [Wolfswinkel et al. 2013]. The categorization is composed of: 1) “insider threat mitigation,” 2) “theoretical perspectives,” 3) “insider threat management,” 4) “insider threat behavior,” 5) “insider threat overview”, and 6) “miscellaneous.” This work focused on categorization, rather than focusing on providing an overview of relevant papers in each category. Gheyas and Abdallah [2016] applied “a systematic literature review and meta-analysis” defined by PRISMA [Moher et al. 2009] for a review of malicious insider threat detection approaches and their concepts. The authors ranked research papers considering the theoretical properties and the clarity of the contributions presented. They also proposed a categorization of studies according to the input dataset, features used, detection strategy implemented, and the underlying machine learning algorithm. The authors discussed key challenges in malicious insider threat detection from the big data perspective, stated trends in the field, and provided best practice recommendations for future research. Sanzgiri and Dasgupta [2016] proposed classification of insider threat detection methods based on strategies and features used in the detection itself, introducing the following nine classes: 1) “anomaly-based,” 2) “role-based access control,” 3) “scenario-based,” 4) using “decoys and honeypots,” 5) “risk analysis using psychological factors,” 6) “risk analysis using workflow,” 7) “improving defense of the network,” 8) “improving defense by access control,” and 9) “process control to dissuade insiders.” Bertacchini and Fierens [2008] conducted a literature review on masquerader detection approaches in the Unix command domain and divided them into several categories: 1) “information-theoretic-based,” 2) “text mining-based,” 3) “HMM-based,” 4) “Naïve Bayes-based,” 5) “sequence-based and bioinformatics-based,” 6) “SVM-based,” and 7) “other approaches.” In addition, the authors summarized the properties and the Unix command datasets used in specific works.

2.1 Comparison with Our Survey

In sum, the existing surveys deal with either categorization of *heterogeneous studies* [Azaria et al. 2014; Hunker and Probst 2011; Ophoff et al. 2014] or *homogeneous studies* that are constrained to detection approaches [Bertacchini and Fierens 2008; Gheyas and Abdallah 2016; Salem et al. 2008; Sanzgiri and Dasgupta 2016]. A characteristic of homogeneous studies is that they are limited to specific application domains: masquerader detection approaches [Bertacchini and Fierens 2008], insider threat detection approaches [Gheyas and Abdallah 2016; Sanzgiri and Dasgupta 2016], and host-based/network-based profiling of insiders [Salem et al. 2008]. Since we deal with a broad scope of research related to insider threat, our survey involves heterogeneous studies, and thus, is more similar to [Azaria et al. 2014; Hunker and Probst 2011; Ophoff et al. 2014]. However, in contrast to existing surveys, we first perform coarse-grained categorization based on the workflow of research efforts, and then we make fine-grained categorization within each category.

3 DEFINITIONS AND TAXONOMIES

This section focuses on definitions and taxonomies of insider threats. First, we start by emphasizing the difficulty of determining the border of distinction for insider threats and also intent of whistleblowers. Then, we survey several disparate definitions of insiders and insider threats, followed by

brief descriptions of state-of-the-art taxonomies, which we split into three types according to an insider's intention: *malicious*, *unintentional*, and special case of taxonomies involving both types of intentions, denoted as *combined* taxonomies. Afterwards, based on the taxonomies described, we propose a structural taxonomy of insider incidents by providing answers to questions regarding the information gathering problem.

Scope of Insider Threat. An inherent problem associated with defining the scope of insider threats is that often it is difficult to distinguish between insiders and outsiders of an organization once they are operating within an internal network. Further complicating matters is the fact that there are insiders that *attack from outside*. One representative example of such an insider is an ex-employee who recently left the organization.

According to Neumann [2010], outsiders who have successfully penetrated into an IS or network are considered outsiders unless they have obtained enough knowledge to make them indistinguishable from the insiders they are masquerading as; the author considers a Turing test to make such a differentiation. Moreover, insider threat has also been defined as part of broader topics that have emerged throughout the history of computer systems, such as intrusion attempts [Anderson 1980], threats to IS security [Loch et al. 1992], and counterproductive workplace behavior (CWB) [Sackett 2002]. Currently, the most relevant scope of insider threat is maintained by the CERT division of Software Engineering Institute at Carnegie Mellon University (referred to in this paper as CERT) and is derived from a database of more than 1000 real case studies. We refer the reader to a discussion of different perspectives on insider threats in Appendix A.

Ethical Question of Whistle-Blowers. As we have already mentioned, whistle-blowers usually act due to their loyalty to the society and to the country, while at the same time they break official secret act or prior consent about the confidentiality of information. We do not take a side of this question, instead, we present this kind of insider threat as part of the state-of-the-art taxonomies that put it under malicious insider threat. We utilize such an assignment in our proposed structural taxonomy of insider threat incidents purely for incident classification purposes.

3.1 Definitions

Most of the following definitions distinguish between insiders and insider threat. Insider definitions usually refer to static descriptions of individuals, using terms such as access, knowledge, trust, or security policy, while insider threat definitions refer to corresponding actions such as misusing access or knowledge which insiders have or violation of a security policy. Note that the distinction between malicious and unintentional insider types only makes sense for definitions of insider threat, as they require that some action be performed. Most existing definitions of insider threat implicitly assume a malicious intent of this threat [Einwechter 2010; Schultz and Shumway 2001; Theoharidou et al. 2005]. However, other existing definitions do not distinguish between a malicious and unintentional insider and may cover both [Bishop 2005; Greitzer and Frincke 2010; Hunker and Probst 2011; Pfleeger et al. 2010; Predd et al. 2008]. Only one of the studies explicitly define unintentional insider threat [Collins et al. 2016].

3.1.1 Term Insider. Pfleeger et al. [2010] define an insider as “a person with legitimate access to an organization's computers and networks.” The report from RAND Corp. [Brackney and Anderson 2004] defines an insider as “an already trusted person with access to sensitive information and information systems” (IS). Aimed at database security, Garfinkel et al. [2002] consider an insider “as a subject of the database who has personal knowledge of information in the confidential field.” Chinchani et al. [2010] attribute the term insider to “legitimate users who abuse their privileges, and given their familiarity and proximity to the computational environment, can easily cause significant

damage or losses.” According to Althebyan and Panda [2007], an insider is “an individual who has the knowledge of the organization’s IS structure for which he/she has authorized access and who knows the underlying network topologies of the organization’s IS.” Sinclair and Smith [2008] define an insider as “any person who has some legitimate privileged access to internal digital resources, i.e., anyone who is allowed to see or change the organization’s computer settings, data, or programs in a way that arbitrary members of the public may not.” A trust-based definition of an insider by the Dagstuhl seminar on Countering Insider Threat [Probst et al. 2008] concluded that an insider “is a person that has been legitimately empowered with the right to access, represent, or decide about one or more assets of the organization’s structure.” Greitzer and Frincke [2010] define the insider as “an individual currently or at one time authorized to access an organization’s IS, data, or network.” Bishop [2005] determines the insider definition with regard to security policy that contains specified rule set. He defines an insider as “a trusted entity that is given the power to violate one or more rules in a given security policy.” A non-binary approach, indicating “*degrees of insiderness*” with access control rules used to develop these degrees, was proposed by Bishop et al. [2009a], who defined someone as “an insider with respect to access to some well-defined data or resource.” According to Predd et al. [2008], an insider “is someone with legitimate access to an organization’s computers and networks.” Note that the authors deliberately do not define the meaning of word *legitimate*, hence they do not separate insiders from outsiders; instead, they state that “both legitimate access and the system’s perimeter are a function not only of system-specific characteristics but also of a given organization’s policies and values.” Therefore, an insider might also be represented by an external entity such as contractor, ex-employee, business partner, etc. [Predd et al. 2008].

3.1.2 Term Insider Threat. Pfleeger et al. [2010] defines insider threat as “an insider’s action that puts at risk an organization’s data, processes, or resources in a disruptive or unwelcome way.” According to Greitzer [2010], the insider threat refers to “harmful acts that trusted insiders might carry out; for example, something that causes harm to an organization, or an unauthorized act that benefits the individual.” According to Theoharidou et al. [2005], insider threat refers to “threats originating from people that have been given access rights to an IS and misuse their privileges, thus violating the IS security policy of the organization.” Insider threat is defined by Hunker and Probst [2011] as “an individual with privileges who misuses them or whose access results in misuse.” Schultz and Shumway [2001] define an insider attack as “the intentional misuse of computer systems by users who are authorized to access those systems and networks.” Bishop [2005] defines insider threat as an event occurring when “a trusted entity abuses the given power to violate one or more rules in a given security policy.” According to Predd et al. [2008], insider threat is “an insider’s action that puts an organization or its resources at risk.”

3.1.3 Term Unintentional Insider Threat. According to Collins et al. [2016], an unintentional insider threat is defined as “a current or former employee, contractor, or other business partner” who: 1) “has or had authorized access to an organization’s network, system, or data,” and 2) “had no malicious intent associated with his/her action (or inaction) that caused harm or substantially increased the probability of future serious harm to the confidentiality, integrity, or availability of the organization’s information or IS.” Liu et al. [2009] consider inadvertent insider threat to be defined as “inattentive, complacent, or untrained people who require authorization and access to an IS in order to perform their jobs.” Raskin et al. [2010] introduce the notion of “*unintended inference*, which addresses what is not explicitly said in the public text made by an insider,” and thus he/she may inadvertently reveal some private information.

3.2 Unintentional Insider Threat Taxonomies

Taxonomy by CERT. Greitzer et al. [2014] define four types of unintentional insider threat (derived from the Privacy Rights Clearinghouse): “**malicious code**” – sensitive information social engineered (e.g., planted USB drive, phishing attack) in combination with malware or spyware; “**disclosure**” – “sensitive information posted publicly on a Web or sent to unauthorized recipients via fax, mail, or email;” “**improper/accidental disposal of physical records**” – “lost, discarded, or stolen non-electronic records, such as paper documents;” and “**portable equipment no longer in possession**” – “lost, discarded, or stolen data storage device, such as a laptop, PDA, smart phone, portable memory device, CD, hard drive, or data tape.”

Typologies of Insiders. Considering motivation, Wall [2013] states that unintentional insider threats are divided according to previous research into **well-meaning** and **negligent**. Then, the author presents the typology of four risk types of employees who can cause data spillage: **underminers**, (those that make their lives easier by not respecting security policies), **over-ambitious**, (those that purposely bypass security measures in order to be more effective), **socially engineered**, and **data-leakers**. The author outlines various ways in which data can be leaked, such as accidentally (losing USB sticks); for user’s convenience (copy data to home PCs); among others (data present on the hard drives of discarded PCs, sharing with third parties, leakage through public post, etc.).

Taxonomy of Human Errors. Reason [1990] defines human error as “the failure to achieve the intended outcome in a planned sequence of mental or physical activities when failure is not due to chance.” The author proposes a generic error modeling system (GEMS). GEMS divides errors into: “**slips**”, which represent “the incorrect execution of a correct action sequence” (execution failures), and “**mistakes**”, which represent “the correct execution of an incorrect action sequence” (planning failures).

3.3 Malicious Insider Threat Taxonomies

Inside Misusers. One of the earliest classifications of internal misuse of computer systems was proposed by Anderson [1980] who distinguishes among three types of illicit inside users, ordered by the ascending difficulty of their detection from audit trails: **masqueraders** – who can be either external attackers that bypassed security controls and penetrated into a computer system, or internal users who intend to exploit another user’s credentials in order to perform some malicious action; **misfeasors** – users who do not masquerade, but instead abuse their own privileges in order to misuse the system; and **clandestine** users – who represent superusers with the capability of staying under the radar of security controls, which they manage and thoroughly know, and making them most difficult to detect. Note that internal misuse of a computer system may also include some activities that fall into the category of counterproductive workplace behavior, which are not covered by any of the insider threat definitions. Therefore, we consider insiders misusing computer systems as the superset of malicious insider threat.

Insider Attack Types. Bellovin [2008] distinguishes among three types of insider attacks: **the misuse of access** that is, according to the author, the hardest type of incident to detect, because an insider uses a legitimate access for the improper purpose; **defense bypass** where insiders have already passed some lines of defense (e.g., firewalls) and can also bypass others; and **access control failure** that represents a technical problem, where either vulnerabilities are present in an access control element or such an element is misconfigured.

Insider Threat with Various Types of Knowledge. Salem et al. [2008] divide malicious insider threats into two groups: *traitors* and *masqueraders*. These two classes can be distinguished based

upon the amount of knowledge they have. **Traitors** have full knowledge about the systems they work with on a daily basis, as well as the actual security policies. Traitors usually act on their own behalf, and therefore use their own credentials for malicious actions. On the other hand, **masqueraders** may have far less knowledge than traitors. They are attackers who steal the credentials of another legitimate user, and then use the stolen credentials for executing a malicious act on behalf of another user. Another example is obtaining an access to the account of a victim by exploiting some system vulnerability. Note that these two classes are not necessarily disjoint; traitors can first use their legitimate access for obtaining another user's access and then perform damage using that access.

Categorization of Insider Threat by Network Monitoring. Myers et al. [2009] divide malicious insider threat into two categories from the perspective of network event monitoring: **“the unauthorized use of privileges,”** representing insiders who access data that they are not authorized to access (i.e., data that is not related to the project or role of an insider), or misuse their authorized access in an inappropriate way (e.g., sharing classified data with an unauthorized person); and **“automated insiders”** that represent either bots performing reconnaissance of the internal network, or probes identifying vulnerabilities and misconfigurations in the internal systems.

Categorization of Insider Threat from the Cloud Computing Perspective. Claycomb and Nicoll [2012] categorize insider threat with regard to cloud computing into three categories: **the rogue cloud provider administrator** who can access a potential victim's data or other leased resources that can be leaked or misused for fraud, IP theft, or sabotage of the cloud provider's reputation; **insiders who exploit cloud vulnerabilities for unauthorized purposes**, e.g., exploitation of replication lag among replicas for fraud; and **insiders who exploit the cloud services** in order to conduct nefarious activity against the company (e.g., cracking password files, DDoS attacks against an employer, exfiltration of data when leaving the company). In contrast to previously mentioned categories, targeted systems or data are not necessarily directly related to cloud services. Another categorization of insider threat from cloud computing perspective is presented by Kandias et al. [2011] who distinguish between a) an **“insider employed by the cloud provider”**, and b) an **“insider employed by an organization”** that outsources IT to the cloud.

Classification of Insiders. Cole and Ring [2005] classify insiders into four classes: 1) A class of **pure insiders**, which contains regular employees that only have privileges necessary to perform their jobs (door keys, access cards, access to a specified list of network services, etc.). A special case is *“an elevated pure insider”* who has privileged access. 2) A class of **inside associates** (a.k.a. external insiders [Franqueira et al. 2010]) consists of third party personnel, such as contractors and suppliers, as well as internal employees with limited authorized access to various compartments inside an organization, for example, security guards, servicemen, and cleaners. The limited authorized access of inside associates is usually represented by physical access to the department/facility, as opposed to access to the IT infrastructure of an organization. Inside associates, such as cleaners, have physical access to the workspace of other employees, and therefore they may find privacy-sensitive information on employees' desks or in trash bins. Moreover, they may plant key-loggers for the purpose of sniffing credentials or other sensitive information entered by employees using keyboards. In comparison to the previous classes, 3) **inside and 4) outside affiliate** classes represents people that do not have any justified and legitimate reason to enter the building of an organization. An inside affiliate class involves family members, friends, or clients of an employee. They have no access to the organization's facility but can steal and misuse the access cards/credentials of an employee to obtain such access and further perpetrate malicious acts. The class of outside affiliates includes untrusted people external to an organization, who can obtain internal access to the organization's

network by, for example, using unprotected WiFi, bypassing weakly protected WiFi, or social engineering the credentials of authorized employees (e.g., phishing attacks).

High-End and Low-End Profiles of Insiders. Cole and Ring [2005] suggest profiling insiders into two categories, taking into account the actions, appearances, and instincts of co-workers: **1) low-end insiders** – this profile includes the characteristics of the insiders that have been caught, accused, and potentially convicted, while also keeping in mind that the majority of high-end insiders are smarter and have not been caught. The basic characteristics of low-end insiders are as follows: they have no or minimal technical education/knowledge; they have already worked in a variety of positions; the attacks they perpetrate involve theft of IP; their primary motivation is personal gain; they are unaware of potential negative consequences of their acts; their wrong doing causes an increased attention and suspicion of their colleagues; they are influenced by their emotional state (e.g., anger or grievance). **2) high-end insiders** – this profile of insider involves people who usually view their malicious mission as their career decision, in contrast to low-end insiders whose actions are usually short-term. Therefore, the authors consider high-end insiders as more dangerous than low-end insiders. The high-end insider profile has, according to the authors, several common characteristics: they are diligent employees, they want to achieve high career positions quickly, while evincing very good leadership skills, job dedication, and trustworthiness.

Levels of Insider Threats. Based on the various potential consequences and harm to an organization, Cole and Ring [2005] distinguish among three levels of insider threats: **Level 1: Self-motivated insiders** are not motivated by any third party; instead they decide to act on their own due to some personal reason, such as revenge, need for correction of organization's actions (e.g., whistle-blowers), or complacency. **Level 2: Recruited insiders** are those who do not independently decide to act maliciously but are convinced and hired by third party. Usually, these insiders are successfully recruited due to financial problems or other personal weakness that can be exploited by a third party. Because this type of recruited insider can be risky for both the insider and the recruiter, the preferred way for malicious third parties is to plant their own insider (e.g., spy) into a target victim's company. **Level 3: Planted insiders** are placed by malicious organizations that find someone suitable for an insider job, train them, get them hired at the target victim's company, allow them time to earn trust within the company, and then start exploiting the insider for data exfiltration/espionage.

Types of Motivations. According to Cole and Ring [2005], there are plenty of factors that may contribute to the motivation of an insider who "turns to the dark side," however there are three main motivations that appear repeatedly: **financial** – when organizations are recruiting people to perform inside attacks, they may target people with financial difficulties or those who want to earn some extra money, which may also be the reason for a self-motivated insider; **political** – some employees have strong political opinions, and if their employers have substantially different opinions and actions, then these employees are motivated to cause harm when opportunity comes or to commence collaboration with malicious entities outside the organization; and **personal** – this type of motivation can occur in one of two ways: 1) the recruiter digs into the victim's past and tries to find out the victim's deepest, darkest secrets and use them to threaten the victim (e.g., blackmail), or 2) if the victim has no secrets accessible to the recruiter, then the recruiter tries to orchestrate a trap scenario that will create a new secret (e.g., using an attractive human bait).

Insider Profiling by CERT. Cappelli et al. [2012] propose profiling malicious insider threats into three types: **1) IT sabotage** in which "an insider uses IT to direct specific harm at an organization or an individual." Such insiders are usually disgruntled employees with technical background who

have administrative privileges. An example of this category is the installation of a logic bomb that is activated after an employee's termination. **2) Theft of Intellectual Property.** Common cases involve espionage, and they are usually committed by technical staff (e.g., engineers and developers), as well as non-technical staff (e.g., clerks and salesmen). Perpetrators may steal information that they daily access, and take it with them when they leave the organization (e.g., using the IP for their own business, taking it to a new employer, or passing it to another organization). **3) Fraud** in which "an insider uses IT for the unauthorized modification, addition, or deletion of an organization's data for personal gain or theft." Insider fraud is usually perpetrated by low-level staff with non-technical backgrounds, such as human resources or help desk staff. The motivation for this is often greed or financial difficulties, and this type of crime is typically long-term. Recruitment of such fraudsters by external entities is also very common. The case studies that do not fall under these three profiles are denoted as miscellaneous [Collins et al. 2016].

3.4 Combined Taxonomies of Insider Threat

Categorization of Insiders from Hayden. Hayden [1999] categorizes insider threat into four groups: "traitor, zealot, browser, and well intentioned," where the former two are related to malicious insider threat and the latter two to unintentional insider threat. The **traitor** group consists of people "who have malevolent intent to damage, destroy, or sell out their organization." The **zealot** group refers to people "who believe strongly in the correctness of one's position or feel the organization is not on the right side of a certain issue." In this case, the insider may try to "correct the organization," for example, by disclosure or removal of confidential data, or providing access to unauthorized parties. The **browser** category refers to people "who are overly curious in nature, often violating the need-to-know principle." The **well-intentioned** category of insiders involve people who "commit violations through ignorance, e.g., disabling anti-virus protection, using unapproved thumb drives."

Policy-Based Taxonomy of Insider Threat. A policy-based taxonomy of insider threat can be derived from the definitions and corresponding model of insider threat proposed by Bishop and Gates [2008]. The authors expanded their former definition of insider threat [Bishop 2005] by arguing "that a security policy is inherently represented by the access control rules employed by an organization." Then, an insider threat can be categorized considering two primitive actions: "**violation of a security policy using legitimate access,**" involving actions that do not respect security policy in force, for example, confidential data is leaked to an unauthorized person; and "**violation of an access control policy by obtaining unauthorized access**" – in this case, insiders misuse their legitimate access in order to extend the scope of their actual privileges, while breaking the security policy in force as well as the access control mechanism, e.g., gaining superuser access by exploiting some system vulnerability. Note that authors do not explicitly mention an intent of the insider threat, therefore it may cover malicious as well as unintentional threats. With this in mind, an example of gaining superuser access can be made for the purpose of stealing confidential data or for the purpose of making the work more efficient, while disobeying policies.

Categorization of Inside Misusers by Physical Presence. Considering physical presence, Neumann [2010] classifies insiders into two categories: logical and physical. A **logical insider** executes his/her actions physically outside of an organization's workspace, while a **physical insider** acts inside of the physical boundaries of the organization's workspace, including external trusted networks. Note that this taxonomy does not specify the intent of the insiders, and therefore can be applied to both unintentional and malicious insider threats.

Human-Centric Taxonomy of Inside Misusers. Magklaras and Furnell [2002] propose a human-centric taxonomy of inside misusers considering three dimensions: *insider's role in the*

system, reason for misuse, and consequences on the system. 1) **System role** – the first dimension classifies people by “the type and level of system knowledge they possess.” This dimension consists of: *system masters* who have full control over most of the IS resources (e.g., system administrators); *advanced users* who, although do not have privileged access, acquired a large amount of knowledge about the systems and networks of the organization, and also are often capable of revealing system vulnerabilities (e.g., programmers, database administrators); and *application users* who use certain standard applications, like Internet browsers, office suites, and email clients, but usually do not have extra privileges to access resources other than those required by their applications. 2) **Reason for misuse** – this dimension describes attributes of insider threat incidents. Considering this dimension, the authors categorize insiders into two groups: *intentional* misfeasors who act for various reasons (e.g., deliberate ignorance, revenge) and *accidental* misusers who can also be classified by the actual reason that negatively influenced the behavior of a legitimate user (e.g., lack of training, excessive workload, personal problems). 3) **System consequences** – this dimension distinguishes among the various ways a misuse act occur, which is manifest by certain traces in the IT infrastructure at the system level. The authors describe three levels that are attributed to these consequences: *OS* – modifications to the structure of a file system, the installation of unauthorized software, etc.; *network* – network packets may contain unauthorized content, data exfiltration of confidential data may be perpetrated through email or file sharing services, etc.; and *hardware* – vandalism or removal of hardware components, installation of key-loggers, modifications of default configurations to critical hardware components (e.g., for the purpose of sabotage or IP theft).

Detection-Oriented Taxonomy of Inside Attacks. Similar to the system consequences dimension proposed by Magklaras and Furnell [2002], Phyo and Furnell [2004] propose a taxonomy of inside attacks, which distinguishes among four monitoring levels of a target system at which an attack may be detected: 1) **network**, 2) **operating system**, 3) **application**, and 4) **data**. This taxonomy is based on the assumption that one inside attack may be manifest at specific levels of the system, while traces of another inside attack may be present at different levels (e.g., fraud breaching integrity of data may be obvious at application and data levels, while data exfiltration may be manifest at network and OS levels). Note that this taxonomy does not discern the intent of the insiders, hence it can be inherently applied to both malicious and unintentional insider threats.

Subtypes of Perpetrators. The typology of insider perpetrator subtypes is presented by Shaw and Fischer [2005] and includes two subtypes representing unintentional insider threat and six subtypes representing malicious insider threat. This typology includes: **explorers** who are inquisitive people (largely with benign intent) that violate policies while they explore the system and its components; **samaritans** who are individuals that do not follow approved procedures for fixing some issues, but instead “hack into a system” in a more efficient or faster way; **hackers** who are individuals that, despite having records of previous hacking activities, are employed and continue in these activities. They may also install logic bombs that serve as employment insurance in cases in which their malicious activities are discovered; **machiavellians** who are individuals that perpetrate sabotage, industrial espionage, intellectual property theft, or other types of activities that potentially lead to their job promotion (e.g., intentionally introduce critical bugs that are discovered by these *triumphant* insiders); **proprietors** who are individuals with a strong feeling of possession of the systems they are managing and are ready to defend “the control and power over this territory” – in many cases these insiders prefer to destroy some critical components of the system or the system itself instead of giving up; **avengers** who are disgruntled individuals that act impulsively due to perceived injustice against them; **career thieves** who are individuals that start their job exclusively with the intention of perpetrating malicious acts that they can benefit from,

such as financial fraud or the theft of intellectual property; and **moles** who are individuals that start their job exclusively with the intention of stealing intellectual property for the benefit of a foreign country, organization, or competing company.

Insider Taxonomy from Sinclair. Sinclair and Smith [2008] divide the insider threat into three classes: insiders with an intention to commit malicious action; insiders acting for their personal profit; and insiders who accidentally or unwittingly act in a harmful way.

Categorization of Insider Threats using Trust. Probst and Hunker [2010] consider trust and categorize insider threats into two categories: **those not representing a violation of trust**, which consist of *accidental insider* threats and *disobeying a security policy* in force; and **insider threats representing a violation of trust**, which include *simple insider* and *high profile insider* threats (similar to the low-end and high-end insiders in [Cole and Ring 2005]).

3.5 Structural Taxonomy of Insider Incidents by 5W1H

Taking into account the definitions and taxonomies introduced above, in this section we present our structural taxonomy of insider threat. In order to provide a unified view regarding the existing taxonomies, we utilized the *who, what, where, when, why, and how* (5W1H) methodology. The 5W1H are elementary questions addressing the information gathering problem, which were originally used to report a news story, but have had other applications as well (e.g., building domain ontologies [Yang et al. 2011]). Since the insider incident investigation problem can also be viewed an instance of information gathering, we selected this approach for a structural taxonomy of this type of incident.

First, we selected several sample taxonomies of insider threat answering 5W1H questions with regards to new information brought. In the selection, we targeted taxonomies that are mostly orthogonal, while we skipped ones that are specific to a particular domain (e.g., [Claycomb and Nicoll 2012]) or that overlap with other more broad taxonomies (e.g., [Sinclair and Smith 2008]). Then, according to the remaining relevant answers to the 5W1H questions, we supplemented our selection with additional further subcategorization that can enhance the description of a particular insider threat incident in a systematic manner. The proposed schema of structural taxonomy is depicted in Figure 1 (the left side of the figure is related to malicious insider threat, and the right side is related to unintentional insider threat). In addition to assigning some of described taxonomies to 5W1H questions, we subcategorized malicious insider threat incidents by adding two ways of addressing the question of *when*; the first one describes duration, and the second differentiates between whether the incident was performed before or after an insider left the place of his/her employment (i.e., job termination). For the sake of simplicity, this inherently assumes that every malicious insider is fired by an organization or decides to leave a place of employment on his/her own. The next two additional subcategorizations were added to the unintentional insider threat incidents; the first one relates to the question of *why*, and the second one relates to the question of *when*. Note that after distinguishing between malicious and unintentional insider threat, some of the taxonomies can be based on an abstract level reduced into one (e.g., [Salem et al. 2008] and “malicious” part of taxonomy derived from [Bishop and Gates 2008]). In the case of subcategorization by a violation of policy, we depicted the requirement of obtaining unauthorized access based on the use of legitimate access with a red dashed arrow. By using the proposed structural taxonomy, all important information about an insider threat incident is kept in an easily maintained and clear format, which can be extended or modified in a straightforward manner in order to integrate future case studies. Two examples that demonstrate the application of our proposed structural taxonomy are presented in Appendix B.

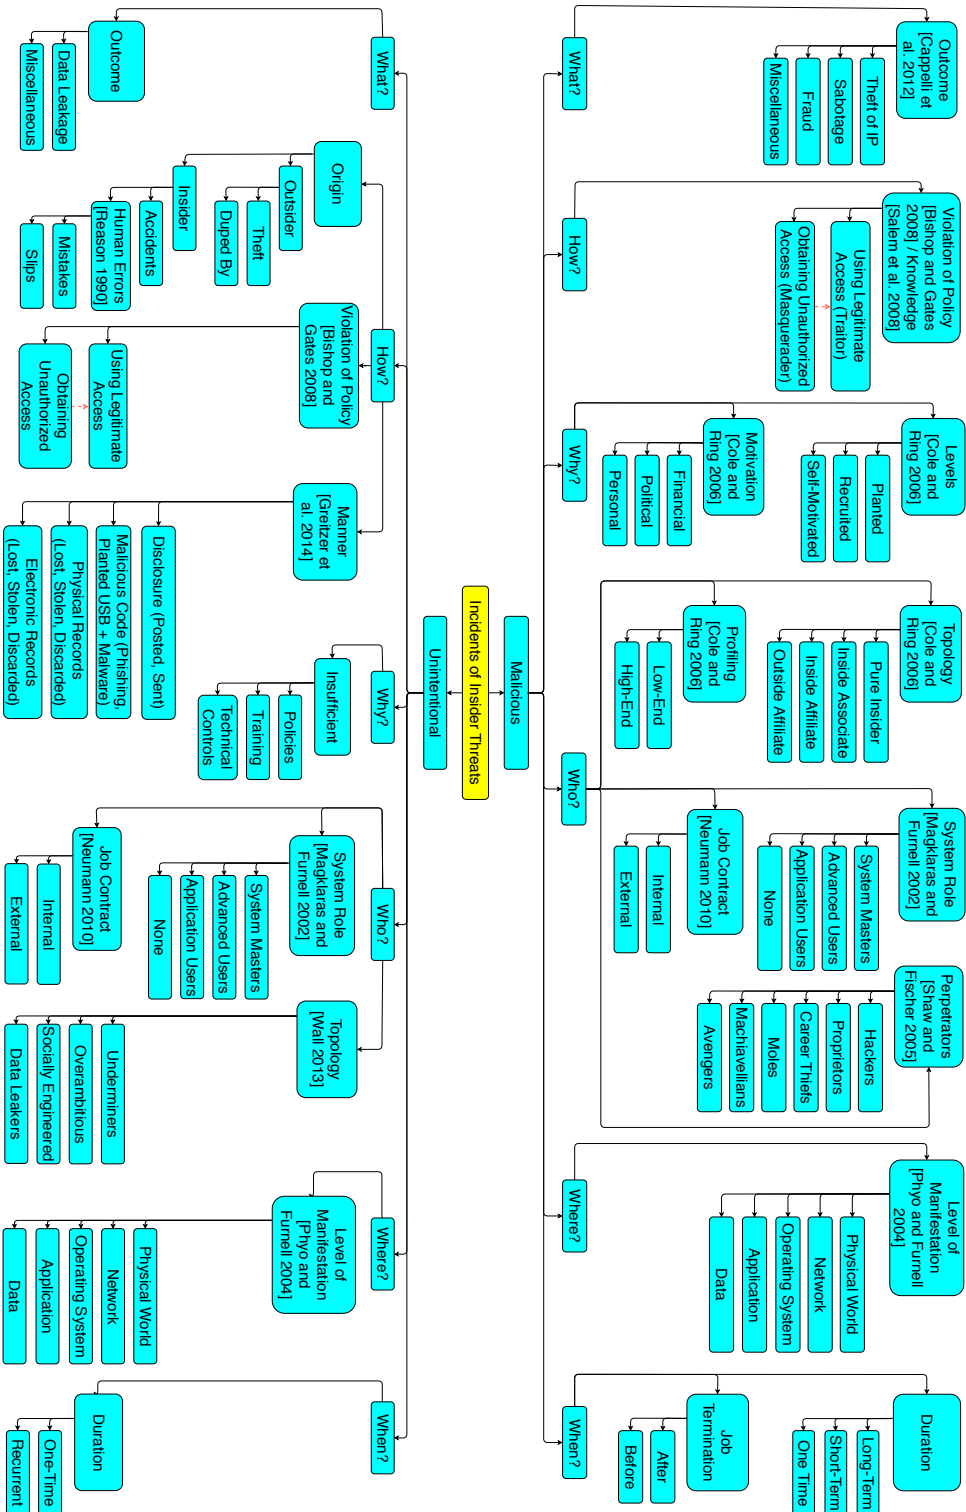


Fig. 1. Structural taxonomy of insider threat (based on previous research and 5W1H questions)

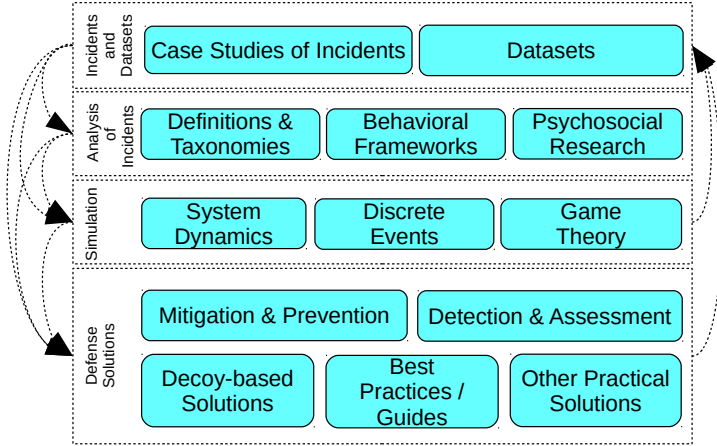


Fig. 2. Workflow of research contributions

4 PROPOSED CATEGORIZATION

After surveying definitions and taxonomies related to insider threat, we now change our focus to the second main contribution of this work – our proposed novel categorization of existing insider threat research. In this section we present an overview of the four main categories, and we subsequently expand on each of the categories in dedicated sections below.

4.1 Workflow of Research Contributions

For the categorization and review of all of the studies included in our survey (and respectively their contributions), we applied the grounded theory approach for rigorous literature review [Wolfswinkel et al. 2013], which consists of five stages. In the initial stage, several inputs are specified: the criteria for inclusion/exclusion, field of research, appropriate data sources, and search terms. Searching is performed in the second stage, followed by the stage of selecting, which includes filtering out doubles, refining samples based on titles and abstracts (later based on full texts), and inclusion of forward and backward references. The fourth stage is aimed at analysis, and it applies the key principles of the grounded theory approach: identification of high level categories from found concepts (open coding); identification of subcategories among the high level categories (axial coding); and refinement and integration of existing categories and their subcategories (selective coding). The last stage deals with presenting the findings and insights in the area.

In addition to the application of grounded theory, our intention was to depict the workflow among particular categories of the contributions that would follow the direction from incidents to solutions or vice versa. In this process, we identified step-by-step *defense solutions* for the detection, assessment, prevention, and mitigation of insider threats; these solutions were further subdivided to discern *simulation* approaches as an independent category. Afterward, we identified research contributions that either analyze and model the behavior of an insider threat, or study its psychosocial precursors; then we created a category of such contributions and called it *analysis of incidents*. Finally, we established a separate category for insider threat *incidents and datasets*, which involves collections of miscellaneous datasets of cyber observable data and real-world case studies. Thus, the proposed workflow of insider threat contributions consists of four main categories:

- The **Incidents and Datasets** category contains reference *datasets* applicable for the evaluation of insider threat detection approaches, as well as collections describing real-world

insider *incident case studies*, which can be utilized for the analysis and modeling of insider threat or the design of defense solutions aimed at specific types of incidents. More information about this category can be found in Section 5.

- The **Analysis of Incidents** category aims at generalization and modeling of all related aspects and behaviors of insider threat incidents and includes *definitions and taxonomies* (previously described in Section 3), *behavioral frameworks* modeling the insider threat lifecycle, observed indicators and critical pathways from the security perspective, as well as contributions of studies from *psychological and social* areas that also involve criminology theories. This category has significant importance in understanding the behavior of the malicious insider and his/her trails and observations, which are indications that malicious activities have begun. This information should be taken into account when designing defense solutions. Detailed information about this category can be found in Section 6.
- The **Simulations** category includes research contributions that perform experiments with programmed models of simulated insider environments, either for the purpose of investigating the impact of various simulation settings on the execution of a simulation, or for the purpose of synthetic dataset generation that may be used for testing defense solutions. The simulation category contains three subcategories: *system dynamics*, *discrete event simulation*, and *game theory* works. Further information regarding this category is available in Section 7.
- The **Defense Solutions** category is the largest category and includes contributions that propose a solution for the insider threat *detection assessment, prevention, or mitigation*; we include a special subcategory for *decoy-based solutions* (such as honeypots and honeytokens) and procedural defense solutions in the form of *best practices and guides*. Moreover, for the sake of comprehensiveness, this category also contains various *other practical defense solutions* involving several commercial tools. In sum, this category provides knowledge about the spectrum of the defense options and reveals trends and ideas in the development of defense solutions. Section 8 provides a detailed description of this category.

The workflow-based relations of insider threat categories, along with their subcategories, are depicted in Figure 2. In this figure, the top-down direction represents the direction from incidents to solutions (arrows on the left side of the diagram), and the bottom-up direction represents the workflow from solutions to incidents (arrows on the right side of the diagram). The direction from incidents to solutions (top-down) represents a goal-oriented workflow with the ultimate goal of the design and development of defense solutions. In this paper, we followed this order of the workflow in order to model and present the categorization. More specifically: all categories are dependent on the incidents and datasets category; the analysis of incidents can be used for the purpose of the simulation and design of defense solutions; defense solutions can also utilize the results of the experiments from simulation, and thereby improve the settings of the detection solutions.

On the other hand, in some cases the workflow may go in the opposite direction – from the bottom to top. In this case, simulations may contribute to the generation of synthetic datasets. Best practices related to detection may extend the knowledge about real case studies (as more incidents can be detected), while detection can produce real datasets as a testbed for other defense solutions. A more detailed categorization of the research contributions is presented in Appendix I.

5 INCIDENTS AND DATASETS

A dataset is a key element for designing and evaluating new ideas and solutions in every applied research field. In this section we present relevant data about case studies of incidents of insider threat, as well as existing datasets gathered either from laboratory experiments or from the real




Fig. 3. Detailed categorization of the *Incidents and Datasets* category

world. We believe that the following information will help readers acquaint themselves with various case studies and available datasets in this field. See Figure 3 for an overview of this category.

5.1 Case Studies of Incidents

This subcategory includes several examples of various case studies of insider threat incidents, and we also include two state-of-the-art assignments of such incidents into clusters. The following states research that covers various types of insider threat case studies.

The majority of the research aimed at obtaining information about insider incidents was conducted by CERT and the US Secret Service (e.g., [Kowalski et al. 2008; Randazzo et al. 2005]). For example, Randazzo et al. [2005] examined 23 case studies of insider threat occurring in the finance sector between 1996 and 2002; 15 of the incidents involved fraud, four involved IP theft, and four involved sabotage of the IS/network. In addition, Kowalski et al. [2008] described 36 insider threat case studies in the government sector that occurred between 1996 and 2002; 21 of the incidents involved various types of fraud, nine involved sabotage, three involved IP theft, and three involved a combination of sabotage and IP theft. Fischer described several case studies of insider computer abuse of US defense systems in [2003]. Shaw et al. [1998] described several case studies of critical information technology insiders (CITI), including system administrators, programmers, and network professionals. Sabotage using IT in critical infrastructure sector was also addressed by Keeney et al. [2005]. Shaw and Fischer [2005] thoroughly examined 10 case studies of malicious insiders within national critical infrastructure in Washington, DC. Furthermore, all versions of the CERT guides (e.g., [Collins et al. 2016]) provide several descriptions of insider threat case studies; the authors highlight that all insider actions could be prevented by respecting the best practices proposed in the guides.

Interesting case studies are also described in the works of Magklaras and Furnell [2002], Jabbour and Menascé [2009a], Probst et al. [2010], Bishop et al. [2009a], and Predd et al. [2008]. All of these authors dealt with either data exfiltration, IP theft, or sabotage, particularly that found in the financial and military sectors. Moreover, some of them also focused on the unintentional insider threat, e.g., Probst et al. [2010] dealt with phishing attacks, and Predd et al. [2008] described an episode of unintentional denial of service.

Clustering of Case Studies. In their book, Cappelli et al. [2012] described a sample of 51 case studies, which were assigned to five clusters with respect to proposed profiling: 24 involved *sabotage*, three cases contained *fraud with sabotage*, six cases involved *IP theft*, 12 cases involved *fraud* alone, and finally, six cases exhibited a *miscellaneous* character.¹ In another book, Cole and Ring [2005] divided the presented case studies into two wide clusters based on the environment of their occurrence: *government* and *corporations*. The government cluster contains a description of four examples from *state and local government* and 17 examples from the *federal government*, while

¹Note that this clustering of case studies corresponds to the taxonomy proposed by these authors – see Section 3.3.

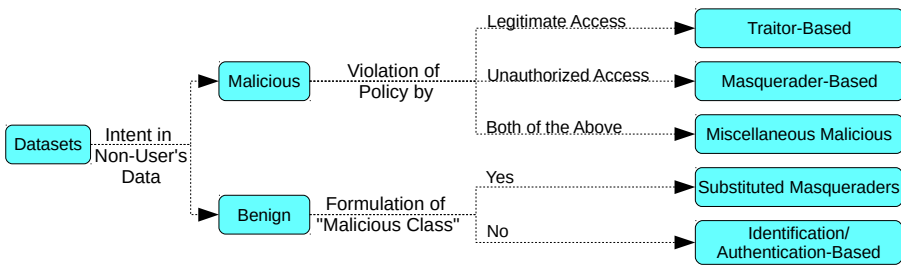


Fig. 4. Categorization of datasets used in insider threat research

in the corporation cluster, the authors present 15 *commercial* examples, eight examples from the *banking and financial sector*, and two examples of *government subcontractors*.

5.2 Datasets

After a review of the insider threat literature, and defense solutions in particular, we divided commonly used datasets into five categories: 1) *masquerader-based*, 2) *traitor-based*, 3) *miscellaneous malicious*, 4) *substituted masqueraders*, and 5) *identification/authentication-based*. All of these categories are depicted in Figure 4, and are also described in our previous work [Harilal et al. 2017]. As can be seen in the figure, these categories can be obtained by applying the following decision steps: a) by discerning the user’s intent in non-user’s data (i.e., data not belonging to a user), which yields *malicious* and *benign* branches; b_1) for the malicious intent branch, by the manner in which the policy violation was executed – either by the use of a legitimate user’s access (*traitor-based*), by obtaining unauthorized access (*masquerader-based*), or when both of the cases are included in a dataset separately (*miscellaneous malicious*); and b_2) for the benign intent branch, by discerning whether the malicious class was formulated by authors of a dataset or not, where the *substituted masqueraders* category include dataset with samples containing labels of such an explicitly built “malicious class” and the *identification/authentication-based* category does not – samples contain only labels of user identification. Splitting datasets of the benign intent branch into two subcategories enables us to isolate datasets that specify equal conditions on a detection/classification task, and thereby the evaluation of an approach on these datasets is always reproducible with the same setting. In contrast to them, identification/authentication-based datasets enable researchers to select various mixtures of samples for a malicious class and thus potentially simplify the classification task. Also, note that each of the five dataset classes we proposed can be sub-categorized according to the origin of data into: 1) real-world datasets and 2) laboratory datasets. However, we are aware only about a single real-world dataset [CALO Project 2015], therefore we use this division criteria only tangentially.

5.2.1 Masquerader-Based Datasets. Despite the fact that there has been a significant amount of research conducted dealing with the masquerader detection problem, only a few studies have used datasets specifically designed for this purpose. The following examples utilize datasets that contain malicious intent in data entries with malicious labels, and the corresponding malicious scenarios are aimed at violations of policy by obtaining unauthorized access.

RUU Dataset. Are You You (RUU) is a masquerader dataset that was introduced by Salem and Stolfo [2009; 2011b]. The dataset was collected from the PCs of 34 normal users and consists of host-based events derived from file system access, processes, the Windows registry, dynamic library loadings, and the system GUI. The dataset contains masquerade sessions performed by 14 humans based on the specified task of locating any piece of data that has a direct or indirect financial value; the users were not limited to any particular means or resources.

WUIL Dataset. The Windows-Users and Intruder simulations Logs (WUIL) dataset was designed and implemented by Camiña et al. [2011] and contains generic file system interactions regardless of their type (e.g., open, write, read). The WUIL dataset contains records from 20 volunteer users (increased to 76 in [Camiña et al. 2016]), who were monitored at different periods of time during their routine daily activities. Although, some users produced approximately an hour’s worth of logs, others produced logs spanning several weeks. The data was collected using an internal tool for file system auditing of Windows machines of various versions (i.e., XP, 7, 8, and 8.1). While the legitimate users’ data was collected from real users, the masquerade sessions were simulated using batch scripts considering three skill levels of users: *basic*, *intermediate*, and *advanced*.

DARPA 1998 Dataset. The DARPA 1998 Intrusion Detection Evaluation dataset was synthesized by the MIT Lincoln Laboratory based on statistical parameters of a “government site containing 100’s of users on 1000’s of different hosts” [Lippman et al. 2000], and its primary purpose was to evaluate and improve intrusion detection systems. However, it was also used in research on the insider threat detection problem [Parveen et al. 2011b]. The DARPA 1998 dataset consists of network traces and system call logs captured on attacked machines, and the attacks performed are divided into four groups: 1) “*denial of service*”, 2) “*remote to user*”, 3) “*user to root*”, and 4) “*surveillance*”. From the insider threat perspective, the only interesting group of attacks is the “*user to root*” group, which may be viewed as masquerade attacks. Nevertheless, only system call logs are relevant, as direct consequences of these attacks can only be monitored on the victim hosts. The DARPA 1998 dataset received a critique [McHugh 2000] and currently is considered outdated.

5.2.2 Traitor-Based Datasets. In the malicious datasets branch, we identified the masquerader research trend, while research dedicated to traitor detection has been more limited. This difference can be explained by the assumption that masquerader detection is simpler and more straightforward than traitor detection, as argued by Salem et al. [2008] who mentioned that “a masquerader is likely to perform actions inconsistent with the victim’s typical behavior”. On the other hand, given that an attacker is a moving target, he/she may imitate a behavior of a victim to some extent, along malicious actions he/she performs. The following research utilizes datasets that include malicious intent in data considered as malicious, and is aimed at policy violations using legitimate access.

Enron Dataset. The Enron dataset [CALO Project 2015] consists of a collection of 500,000 real-world emails (from 1998 to 2002) associated with 150 users, mainly senior management of Enron Corporation. Although, some of the emails were deleted as they contained attachments or confidential information, the dataset contains interesting information that can be used for the analysis of text in emails and social network analysis aimed at the detection of insider threat involving collaborating traitors.

APEX 2007. The APEX ’07 dataset was collected, according to Santos et al. [2008], by the National Institute of Standards and Technology (NIST) with an intention to simulate the tasks of analysts in the intelligence community. The APEX ’07 dataset consists of actions and research reports of eight benign analysts, while the malicious insider threat was simulated by five analysts whose tasks were based on the tasks of benign analysts in order to make a detection more challenging.

5.2.3 Miscellaneous Malicious Datasets. Datasets composed of both malicious insider subtypes (masqueraders and traitors) belong to this category. Therefore, this subcategory of datasets may serve as general testbed for the detection of malicious insider threat.

CERT Datasets. CERT, along with other partners, has generated a collection of synthetic insider threat datasets. The approach employed for the dataset generation is described in [Glasser and Lindauer 2013]; the datasets were generated using scenarios containing traitor instances, as well as

other scenarios involving masquerade activities. The logs collected contain logon data, browsing history, file access logs, emails, device usage, psychometric information, and LDAP data.

TWOS Dataset. The TWOS dataset has been collected by Harilal et al. [2017] as the outcome of a multi-player game designed to reproduce interactions in real companies while stimulating existence of masqueraders and traitors. The game involved 24 users, organized into 6 teams that played for one week. Masquerade sessions were performed by “temporarily” malicious users, who, once in a while, received credentials of other users (victims) and were able to take control over victim’s machines for a period of 90 minutes. Traitor sessions were collected when a few participants were fired from their original team. The dataset consists of miscellaneous data types such as a mouse, keyboard, network, and host monitor logs of system calls. Furthermore, the authors presented applicable state-of-the-art features and demonstrated the potential use of the TWOS dataset in multiple areas of cyber-security that relate to the insider threat area, such as authorship verification and identification, continuous authentication, and sentiment analysis [Harilal et al. 2018].

5.2.4 Substituted Masqueraders from Benign Users. In this category of datasets, data considered as malicious has been explicitly substituted by legitimate data that was never seen before (e.g., data from other users). We created a dedicated category of such datasets, although these dataset can be viewed as related to the authentication problem – *does an input sample belong to the particular user?* Note that previous research has indicated that such datasets are less suitable for testing masquerader detection solutions than *masquerader-based* datasets [Salem et al. 2008].

Schonlau Dataset. The Schonlau dataset (a.k.a. SEA dataset) was introduced by Schonlau et al. [2001] and contains around 15,000 Unix commands per user; the dataset was generated by 50 individuals who had various roles inside of an organization. In this dataset, masqueraders are approximated by randomly blending the data collected from unknown users (i.e., users not among the previously mentioned 50 users), and thus the data used in masquerade sessions does not contain any malicious intent. Maxion and Townsend [2002] showed that the Schonlau dataset is not appropriate for the detection of masqueraders because of the following: 1) its data was collected during different time periods for each user; 2) each user performed a different number of login sessions; 3) all Unix commands were captured in the order of their termination, instead of the order of their commencement. Nevertheless, according to the findings of Salem et al. [2008], until 2008, SEA served as a common benchmark dataset, leading to a large amount of research aimed at the detection of substituted masqueraders in Unix commands.

5.2.5 Authentication/Identification-Based Datasets. This category of datasets can be used for the purpose of identification or authentication of any user, regardless of his/her intent, although benign intent is implicitly assumed. Therefore, this sort of dataset may be used for addressing the identification and authentication questions: 1) *which user does the input sample belong to?* and 2) *does an input sample belong to the particular user?* The following datasets in this category were used in various insider threat detection works.

Greenberg’s Dataset. Greenberg’s dataset [Greenberg 1988] is the first known collection of authentication-based data. The author collected a dataset comprised of full command-line entries (including arguments and timestamps) from 168 Unix users of the *cs* shell. The dataset maintains anonymity of the users, while it retains the semantical information of the executed commands. The dataset is split into four groups according to the level of user knowledge and skills, and consists of 52 scientists with programming skills, 55 novices, 36 advanced users, and 25 non-technical users, respectively. The researchers that use Greenberg’s dataset usually adapt two different configurations: *plain* and *enriched*. The plain configuration provides only commands and aliases, and therefore is

equivalent to the configuration of the SEA dataset. In contrast, the enriched configuration contains commands/aliases together with their full arguments.

Purdue University Dataset. The Purdue University (PU) dataset was introduced by Lane and Brodley in [1997]. The PU dataset consists of eight subsets of preprocessed Unix command data that were taken from the *tcs* shell histories of eight computer users at Purdue University during a two year period (initially it contained data from only four users). Commands included in this dataset are enriched, and thus contain command names, arguments, and options; nevertheless, filenames are omitted.

MITRE OWL Dataset. The Organization-Wide Learning (OWL) dataset from MITRE was designed for collecting application usage statistics in order to provide individual feedback and tutoring to users of an organization [Linton et al. 2000], however it was also used for the analysis of human interactions with GUI-based applications for the purpose of user authentication [El Masri et al. 2014]. During a period of two years (from 1997 to 1998), the data was collected from 24 employees using Microsoft Word on the Macintosh operating system. The participants were employees of an artificial intelligence group, researchers, and technical and support staff.

6 ANALYSIS OF INCIDENTS

In this section we aim to generalize all related aspects and behaviors of a malicious insider before, during, and after conducting an incident. We present the research that deals with the analysis of incidents by considering the insider’s lifecycle, observed indicators, and the critical pathway. First, we consider this mainly from the security perspective and describe behavioral frameworks of insider incidents, and this is followed by consideration of psychological and social theory works. Later we will see that significant portion of the research contributions related to the analysis of the incidents serves as a basis for defense solutions (see Section 8). Although definitions and taxonomies are also part of this category, these were discussed in Section 3, and thus we will not include them here. See Figure 5 for an overview of this category.

6.1 Behavioral Frameworks

The security action cycle for inside abuse is a general way in which inside attackers can be modeled in time, which was presented by Straub and Welke [1998] and later extended by Willison and Warkentin [2013] who enriched it with pre-kinetic events, which represent temporal antecedents of an act of abuse and occur prior to the deterrence stage. This framework is depicted in Figure 6, in which the vertical axis represents the discrete time relative to the occurrence of computer abuse. The middle column represents *Actions & Countermeasures* that an organization may take in order to address (potential) insider abuse at various stages of its lifetime. The successful outcome of such actions and countermeasures appears in the left column (*Desired Events & States*), while the negative outcome appears in the right column (*Undesired Events & States*). Therefore, the objective of an organization should be to perform *Actions & Countermeasures* that maximize




Fig. 5. Detailed categorization of the *Analysis of Incidents* category

Desired Events & States, and simultaneously minimize Undesired Events & States. Like [Willison and Warkentin 2013], the remedies category and deterrence feedback loops are not included in the figure for the sake of clarity. Working from the top to the bottom of the figure, interaction between the organization and the employee is modeled; pre-kinetic events can be perceived by the employee. Positive pre-kinetic events include a friendly and fair working environment, while negative events include organizational injustice, disgruntlement, justification of committing an incident, etc. Furthermore, negative pre-kinetic events form negative attitudes that may lead to the intention to perform computer misuse, however this might be deterred by policies, training, or other means. If deterrence attempts are insufficient, a misuse attempt is executed, and then it is the turn of prevention controls and implemented policies. If prevention controls do not stop the misuse act, the act is successfully perpetrated. At this point, the organization has a chance of detecting it by technical controls or social means.

The following include several examples of the behavioral framework subcategory. Wood [2000] contributes to this area by elaborating the attributes of malicious insiders: *knowledge, access, privilege, risk, tactics, skills, motivation, and process*. Andersen et al. [2004] presents the system dynamics scheme of the malicious insider threat, based on six case studies involving disgruntlement and financial gain as the motivation factors; this research formed the summary of the 2nd Workshop on System Dynamics Modeling for Information Security. Based on 10 case studies in the national critical infrastructure industry, Shaw and Fischer [2005] outlined a conceptual framework describing events on the critical pathway of malicious insider attacks, including: 1) *personal/professional stressors*, 2) *emotional and maladaptive behavioral reactions*, 3) *results in official attention*, 4) *ineffective intervention*, and 5) *attack*. Pfleeger et al. [2010] proposed a conceptual framework for modeling any insider threat, which is aimed at the risk factors and based on four components and their interactions: *organization, environment, system, and individual*. Claycomb et al. [2012] performed a chronological examination of 15 sabotage case studies, and as a result, the authors identified six consecutive common events: 1) *tipping point* (first observed disgruntlement event), 2) *malicious act* (e.g., installing logic bomb), 3) *occurrence of the attack*, 4) *detection of the attack*, 5) *end of the attack*, and 6) *action on insider* (response). Nurse et al. [2014] proposed a conceptual framework for characterizing any form of insider attack, constructed by using grounded theory on 80 case studies. The framework models consecutively connected aspects of an insider incident: 1) *catalyst*, 2) *actor characteristics*, 3) *attack characteristics*, and 4) *organization characteristics*. Farahmand and Spafford [2013] investigated accepted models of risk-taking behaviors in the insider threat field and then introduced ordinal scales to a selected model that represents perceived risk as a function of *consequence* and *understanding*. Maasberg et al. [2015] proposed a theoretic insider threat model based on the MOC (motive, opportunity, capability) concept, the theory of planned behavior, and dark triad personality traits. CERT's MERIT project contains, among other parts, system dynamics models whose purpose is to reveal patterns in the course of insider threat cases over time. Particular profiles of insider threats with corresponding system dynamics models are addressed in research as follows: IT sabotage in [Moore et al. 2008], IP theft in [Moore et al. 2011], fraud involving financial services in [Cummings et al. 2012], and IT sabotage and espionage in [Band et al. 2006].

Formalization Frameworks. Probst et al. [2006] proposed a formal model for static analysis of insiders based on two parts: 1) "a high-level system model based on graphs," and 2) "a process calculus called acKlaim," which extends μ Klaim calculus with access control features. Dimkov et al. [2010] presented PORTUNES, a framework that represents malicious scenarios present in the digital, social, and physical domains; the authors designed and applied abstract language based on Klaim calculus. Chen et al. [2015] formally modeled malicious insider threat at micro and macro levels that are represented by an *intentional analysis* and a *behavioral analysis*, respectively. The

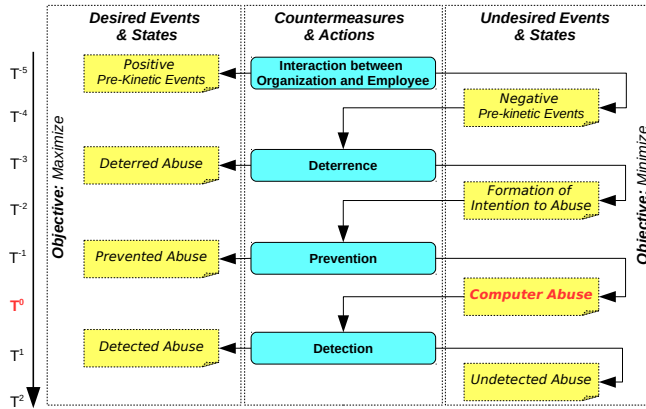


Fig. 6. Security action cycle (based on [Straub and Welke 1998] and [Willison and Warkentin 2013]).

intentional analysis models likelihood of an employee to be an insider threat by Bayesian network, and once an employee intends to attack, the behavioral analysis models the probability of success using Markov decision processes, considering the formal description of an environment in acKlaim calculus. Kammüller et al. [2016] formally modeled both malicious and inadvertent insider threats in the IoT using attacks trees and the higher order logic interactive theorem prover, Isabelle.

6.2 Psychological and Social Theory

Prior psychological and social research found correlations between the insider incidents committed and personality factors, current emotional states, predispositions to malicious behavior, and mental disorders [Greitzer et al. 2013], [Shaw 2006], [Cappelli et al. 2012], [Cole and Ring 2005], [Ambrose et al. 2002]. The following works examine psychological or sociological aspects of insider threat.

Shaw et al. [1998] focused on critical information technology insiders (such as system administrators, programmers, network professionals) and explored personal and cultural vulnerabilities consisting of: 1) “introversion,” 2) “dependency on computers,” 3) “personal and social frustrations,” 4) “ethical flexibility,” 5) “low loyalty,” 6) “feeling of entitlement,” and 7) “lack of empathy.” Potential indicators for the prediction of malicious insider attacks were proposed by Schultz [2002] and consist of “deliberate markers, meaningful errors, preparatory behaviors, correlated usage patterns, verbal behavior, and personality traits.” Leach [2003] analyzed several factors influencing employees’ secure behavior, while emphasizing that three of them improve such secure behavior: 1) *the behavior of other personnel, especially managers*, 2) *the employee’s “common sense of security” and decision-making experiences*, and 3) *the employee’s “psychological contract with the company.”* Ho [2008] investigated changes in trustworthiness as indicators of insider threat, with trustworthiness defined as “the degree of correspondence between communicated intentions and behavioral outcomes that are observed over time.” Farahmand and Spafford [2009] investigated well-known models of risk perception in relation to the benefits of performing malicious actions by insiders. The authors argue that cognitive understanding and potential consequences are the most significant features influencing a risk perceived by insiders. Willison and Warkentin [2009] addressed how perceptions of fairness are formed inside the workplace environment; denoted as *Organizational Justice*, which consists of four constructs: a) *distributive* (e.g., differences in the rewards of employees with the same responsibilities), b) *procedural* (e.g., in disputes), c) *interpersonal* (e.g., managers treat their subordinates respectfully and with dignity), and d) *informational justice* (e.g., discussion of demotion decisions). Later, Willison and Warkentin [2013] reviewed existing *neutralization techniques* of insiders that rationalize their malicious behavior. The level of trust and its consequences to insider threat as part of the risk analysis process are examined in [Probst and Hunker 2010]. Ethical and

social issues of insider threat monitoring are examined in [Greitzer et al. 2010]. Martinez-Moyano et al. [2011] focused on behavior-based detection of terrorists, considering the level of technical expertise, history of suspicious activities, and intensity of religious radicalism. An investigation of causal reasoning of insiders after the implementation of new security measures was performed by Posey et al. [2011]. The relative risk of insider attack based on *dynamic environmental stressors, static/dynamic personal characteristics, capability, and counterproductive behavior* was modeled with Bayesian networks in [Axelrad et al. 2013]. Based on the analysis of the game they conducted, Ho et al. [2016] supported the hypothesis that language incentives in group communication change significantly once an insider has turned to the malicious side.

Criminology Theories. Behavior pertaining to insider misuse, which can be also considered as a type of misbehavior in the workplace, has been studied in criminology research as well. Related studies aim at analyzing insider misuse based on various criminology theories. Theoharidou et al. [2005] discussed major criminology theories with regards to ISO 17799 (a standard in IS security management); they focus on theories applicable to insider misuse – general deterrence theory (GDT), social learning theory (SLT), social bond theory (SBT), theory of planned behavior (TPB), and theory of situational crime prevention (SCP). Lee and Lee [2002] adopted the framework of TPB for computer abuse within organizations and assessed the influences of GDT, SBT, and SLT on TPB. Willison and Siponen [2009] showed how SCP techniques can be utilized with the universal script framework of insider threat; they demonstrated its use on an example of fraud.

7 SIMULATION RESEARCH

This section includes papers related to the modeling and simulation domain, as well as some game theory approaches used for the purpose of simulation. According to Banks [1998], “simulation is the imitation of a real system’s operation over time,” and it contains creation of “*an artificial history*” of a system’s model and “drawing inferences concerning the operational characteristics of the real system that is represented,” based on that artificial history. All studies in this section contain experiments based on the execution of programmed models as simulations. We have identified three subcategories: 1) *approaches working with discrete events*, 2) *system dynamics approaches*, and 3) *game theory approaches*. See Figure 7 for an overview of this category.

Discrete Events. According to Banks [Banks 1998], the goal of discrete event simulation is “to portray the activities in which entities of the system engage, and thereby learn something about the system’s dynamic behavior”; this is accomplished “by defining the states of the system and constructing activities that move it from state to state.” The start and end of each activity is represented by events and the state of the simulated system remains unchanged between two consecutive events. The following simulation studies belong to this subcategory.

Althebyan and Panda [2008] utilized this type of simulation as a performance evaluation technique for an insider threat risk assessment model. In a similar vein, Alghamdi et al. [2006] evaluated the performance of proposed multi-entity Bayesian networks, while using two networks: a *generative network* for simulating the observations and an *inference network* for inferring whether a user is an insider threat or not. For the purpose of performance evaluation, a discrete time Markov chain was used as a database transaction simulator in [Panigrahi et al. 2013]. Simulation in OMNet++ for the purpose of data generation was utilized in [Eberle and Holder 2009]; the data had a graph-based




Fig. 7. Detailed categorization of the *Simulations* category

representation and was related to business transactions and processes inside of organizations. Using colored Petri nets, Nasr and Varjani [2014] performed threshold-based anomaly detection of simulated malicious insiders in a supervisory control and data acquisition (SCADA) system.

System Dynamics. According to Richardson et al. [2001], system dynamics is “a computer-aided approach to policy analysis and design that applies to dynamic problems arising in complex social, managerial, economic, or ecological systems,” meaning that any dynamic system is represented by dependencies among its components, mutual interactions, feedback loops, and circular causalities. Thus, system dynamics enables us to analyze complex systems containing various causalities among their components, and more specifically, it allows us “to investigate the effect of changes in one variable on the other variables over time” [Martinez-Moyano et al. 2008]. Note that system dynamics simulation may run either in continuous or discrete time, and the state of a simulated dynamics can be evaluated at any point in time. The following studies belong to this domain.

Simulations aimed at the inspection of four implementation levels of formal security controls (*absent, poor, normal, and high*) and their impact on the perception of security by management, security level, and incident cost were performed in [Melara et al. 2003]. In the paper, the authors simulated the real sabotage case of Tim Lloyd/Omega. Rich et al. [2005] applied system dynamics simulation in long-term fraud detection by evaluating the alignment training of “information workers” and “security officers” compared to the base run. During the base run, information workers adjusted their decision thresholds solely on their own, while in the alignment training, some experiences of the security officers were transferred to the information workers. Martinez-Moyano et al. [2008] inspected a similar fraud detection scenario, however they evaluated three policy strategies against the base run: *perfect information, consistency training, and alignment training*. In the base run, the firm guarded itself against inside attackers unsuccessfully in comparison to: 1) perfect information, where judges (security officers and information workers) did not change their judging strategies but used better information cues; 2) consistency training in which the defenders were trained to become better judges who responded consistently; and 3) alignment training that used the same tactics as [Rich et al. 2005], enriched by training the security officers who should ensure adherence to security standards by all employees. In the latter paper, Martinez-Moyano et al. [2011] also focused on the convergence of the decision thresholds of security officers for the detection of terrorist activity.

Game Theory. Myerson [1997] defines game theory as “the study of mathematical models of conflicts and cooperation between intelligent and rational decision-makers.” Game theory consists of methods that can be used for the analysis of problems in which two or more entities make decisions influencing one another’s benefit. The term *game* represents a social situation that involves two or more entities who are called *players*. The following studies contribute to this domain.

Using Nash equilibrium, Liu et al. [2008] proposed a two-player stochastic game between a system administrator and an insider who perpetrates a fraud. Aimed at anomaly detection of insider threat, Zhang et al. [2010] proposed a few algorithms for “the establishment of the defender’s reputation,” which resulted in improvement of the trade-off between false positives and true positives. Kantzavelou and Katsikas [2010] used quantal response equilibrium (QRE), which adjusts players’ preferences, and the Von Neumann-Morgenstern utility function that assigns numbers reflecting players’ preferences. In a comparison to Nash equilibrium, QRE was capable of capturing players’ bounded rationality, and thus enabled the players to select even not necessarily the best action. Tang et al. [2011] applied dynamic Bayesian networks consisting of various user-behavioral variables as part of a defender/insider game that used QRE.

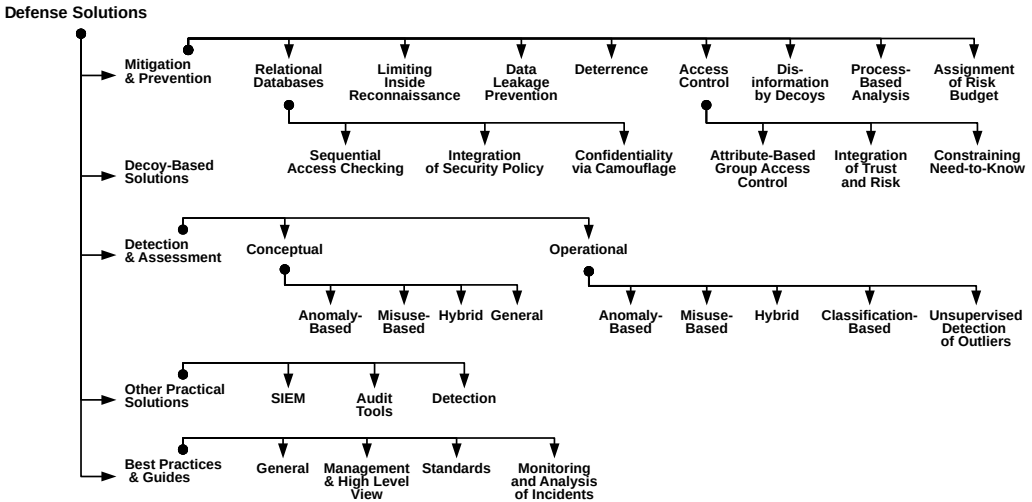


Fig. 8. Detailed categorization of the *Defense Solutions* category

8 DEFENSE SOLUTIONS

This section categorizes and briefly describes defense solutions for the insider threat problem. We emphasize that many of the solutions presented may be considered as general IT security mechanisms, and we survey them here because of the relevancy and the importance to insider threats. First, means-based categorization of **mitigation/prevention approaches** is presented, and this is followed by intrusion-detection-derived categorization of insider threat **assessment and detection** research. As complementary defense solutions, we identified **best practices and guidelines**, **decoy-based solutions**, and **other practical solutions**. However, due to space constraints, we refer the reader to appendices D, E, F for descriptions and examples of these subcategories. A detailed overview of the defense solutions category is depicted in Figure 8. It is important to state that because of the broad nature of insider threats, some of the defense solutions discussed span across other topics that are out of the scope of this survey. More specifically, we refer the reader to the surveys of Shabtai et al. [2012] and Alneyadi et al. [2016] in the area of data leakage, and recommend the works of Lazouski et al. [2010], Servos and Osborn [2017] in the area of access control. Additionally, some related works have focused on the detection and mitigation of specific kinds of fraud attacks, such as financial fraud, telecommunication fraud, Internet marketing fraud, or insurance fraud, which may be perpetrated by outsiders as well as by insiders (depending on the type of fraud). The defense solutions against fraud perpetrators are specific to the type of fraud. For details about the defense techniques used against various kinds of fraud, we refer the reader to the surveys of Abdallah et al. [2016] and Edge et al. [2009]. In this section, we will only include the works that have a significant intersection of insider threat with other topics.

8.1 Mitigation and Prevention

Studies dealing with the mitigation and prevention of insider threat are divided into several categories according to the type of prevention technique or specific application domain.

8.1.1 Deterrence. Vance et al. [2012] presented several subconstructs of accountability as deterrent factors in policy access violation and conducted a factorial survey to support their assumptions. The authors showed that convenient features of the user interface of an IS contribute to awareness about accountability, which as a result, cause deterrence of policy violations.

8.1.2 Data Leakage Prevention. Wu et al. [2011] proposed the concept of “active data leakage prevention” employing an encrypted secure data container (SDC) and ensuring that data is accessed by authorized users in a trusted environment, while data access and manipulation respect expected patterns. Johnson et al. [2016] proposed SimpleFlow, “an information flow-based access control system” that delays action “until a process tries to write confidential data outside of the local network.” Pramanik et al. [2004] proposed context-based security policies that prevent illegal information flow among documents; this is based on the idea of prohibiting modification of a file when another inappropriate file is open. Bertino and Ghinita [2011] proposed an approach for the prevention of data exfiltration, which uses “provenance tracking” through the watermarking and “the confine & mark method.” The confine step is achieved using network segmentation and virtualization, whereas the mark step is handled by tokens.

8.1.3 Process-Based Analysis. Bishop et al. [2014] proposed a process-based approach for the identification and elimination of places where data exfiltration and sabotage can be perpetrated, by applying fault tree analysis and finite-state verification techniques.

8.1.4 Assignment of Risk Budget. Liu et al. [2009] proposed a mitigation technique for unintentional insiders by assigning each user a risk budget that specifies the maximum amount of accumulated risk an employee can cause during the execution of a task; employees may be rewarded for staying within their budget or punished for depleting the budget.

8.1.5 Disinformation by Decoys. Stolfo et al. [2012] proposed a combination of behavioral monitoring and offensive decoy technology for the prevention of unauthorized access by masqueraders in cloud-based environments. When unauthorized access is suspected, the system triggers a “disinformation attack” that confuses a potential attacker by a large volume of decoy data.

8.1.6 Relational Databases.

Sequential Access Checking. Chagarlamudi et al. [2009] proposed a concept that prevents execution of malicious users’ tasks by checking the partial order of database transactions assigned to each particular application’s task against predefined Petri network models. Yaseen and Panda [2011] proposed an insider mitigation strategy for sequential access to particular data in the domain of relational databases, using threat prediction graphs.

Confidentiality via Camouflage. Gopal et al. [2002] proposed a concept denoted as confidentiality via camouflage (CVC) that performs deterministically numerical interval-based responses of ad hoc queries to relational databases, while preserving confidentiality. Garfinkel et al. [2002] proposed an extension of this approach for binary fields and later proposed an improvement for CVC, aimed at insiders that have knowledge of some data in the confidential fields [2006].

Integration of Security Policy. Considering various role access, Jabbour and Menascé [2009b] proposed a concept for protecting a database environment against insider threat by integrating a security policy mechanism as an inseparable part of the protected system.

8.1.7 Access Control.

Attribute-Based Group Access Control. Bishop et al. [2008] proposed attribute-based group access control (ABGAC) utilized with Carlson’s unifying policy hierarchy as a concept for mitigation of malicious and inadvertent insider threats.

Integration of Trust and Risk. Crampton and Huth [2010] extended the concept of access control by providing support for awareness of trustworthiness and risk. In a similar vein, Baracaldo and Joshi [2012] proposed an extension of the role-based access control (RBAC) concept by applying a risk assessment of roles and the system’s trust in its users.

Constraining Need-to-Know. Focused on the document access problem, Aleman-Meza et al. [2005] presented an ontological approach based on measuring the distance of access requested documents from the domain of the insider’s need-to-know. Desmedt and Shaghghi [2016] proposed the concept of “function-based access control (FBAC),” inspired by functional encryption, using operations that access only particular parts of a document (i.e., atoms). Shalev et al. [2016] proposed a Linux container-based solution for isolating the system administrators from resources irrelevant to their current ticket’s task, while enabling them to obtain additional permissions when approved by the permission broker. As part of their work, Eom et al. [2011] proposed the concept of a misuse monitor that can control access to resources by matching the actual processing pattern to the expected processing patterns. Althebyan and Panda [2007] proposed a conceptual model for the prevention of data exfiltration by insiders, which is based on an insider’s knowledge base and dependency graphs among documents. The proposed model prevents access when the insider’s knowledge exceeds an access limit for a particular document cluster. The concept of capability acquisition graphs was presented by Mathew et al. [2008], who proposed performing periodical evaluations of the privileges accumulated by users with respect to critical information assets.

8.1.8 Limiting Inside Reconnaissance. Achleitner et al. [2016] proposed a network deception system based on software-defined networking (SDN), which defends against reconnaissance (e.g., advanced persistent threat – APT) conducted by insider adversaries. The system simulates virtual network topologies that are able to thwart network reconnaissance by delaying the scanning results of attackers, and moreover, the system deploys honeypots that enable the identification of attackers. Markham and Payne [2001] proposed a second generation hardware firewall called network edge security (NES), which is placed on the host’s NIC and contains dedicated memory and CPU to which the host does not have access. NES can only be accessed remotely.

8.2 Detection and Threat Assessment Approaches

Works in this section are organized into two subcategories: *conceptual* and *operational*. A paper is considered conceptual if its contribution is theoretical (i.e., no relevant empirical results are presented), while it is considered operational if the authors include a proof-of-concept (i.e., evaluate their solution against a relevant dataset). These two groups are further subclassified based on criteria derived from intrusion detection, which is described below. Note that we only include several of the most significant and cited works here; the rest of the studies included in our survey is presented in Appendix G.

8.2.1 Conceptual Works. Considering two general classes of intrusion detection – misuse-based and anomaly-based detection – we have divided the conceptual techniques for the detection of insider threat into four subcategories that have specific characteristics: *anomaly-based*, *misuse-based*, *hybrid*, and *general*. Anomaly-based approaches model legitimate behavior as a baseline profile, and with new input, they compute a score that represents the distance from the baseline profile. In contrast, misuse-based approaches model malicious behaviors and measure the similarity or conformity of new input with them. The anomaly and misuse-based categories can be viewed as one-class techniques, as they model just one type of behavior. On the other hand, there are other types of approaches that simultaneously model both kinds of behavior and therefore can be considered two-class techniques; we denote them as hybrid techniques. Finally, we include a subcategory for approaches that aim at general reasoning in detection of insider threat, and denote them as general techniques.

Anomaly-Based. In the relevant literature, anomaly-based conceptual techniques consider a wide range of observables, such as host-based observables, psychosocial indicators, cyber-environmental properties, etc. Early conceptual works in our database just tended to propose the use of cyber observables. For example, Magklaras and Furnell [2002] developed “evaluated potential threat” (EPT) metrics that model user behavior by features, such as role, file system knowledge, access to critical components, previous intrusion records, and network activity. Later, these authors proposed a method for the estimation of end user sophistication [2005], presuming it to be a potential factor influencing the capability of users to commit insider misuse. Further trends in the insider threat detection research brought *role-based monitoring*, *natural language processing* (NLP), *assessing the environment*, *business process monitoring*, and *psychological indicators*. Considering **role-based monitoring**, Ali et al. [2008] proposed host-based user profiling with policies defined by a role-based trust matrix, where access decisions are made according to user-specific thresholds. Focusing on the **database domain**, Bertino and Ghinita [2011] proposed to use pattern matching techniques for the detection of data exfiltration anomalies of database users whose baseline profiles are created during normal periods of activity. Considering **psychological factors**, Kandias et al. [2010] used psychometric tests based on the MOC model for the computation of a per user threat score. **NLP** is a branch of techniques usually used to infer various psychological or emotional indicators (sentiment analysis), which can be used for other purposes as well. For example, Raskin et al. [2010] addressed an unintended inference by identifying hidden information from social media and conversations. In terms of **assessing the environment**, Althebyan and Panda [2007] adapted their previous approach to Bayesian networks (BN), creating the knowledge Bayesian attack graph, which enabled them to estimate risk values for various objects in a system by Bayesian inference. Considering **business process monitoring**, Gritzalis et al. [2014] focused on detecting performance deviations of users, which were correlated with other indicators from social media and technical controls.

Misuse-Based. These approaches do not usually occur in the well-known signature matching form, but rather incorporate softer forms of matching represented by similarity measurement. Ray and Poolsapassit [2005] utilized concepts of *attack trees* for online monitoring and assessment of insider threat; their assessment is based on a comparison of the minimal attack tree of a system, which is generated according to the initially specified user’s intent, and the user’s runtime actions on a system. Utilizing active directory services, Bhilare et al. [2009] proposed a *rule-based* concept for the detection of insider threat violating policies in an academic campus environment. Agrafiotis et al. [2016] proposed *tripwire grammar* capable of capturing abstraction of policies that organizations adopted, as well as signatures of insider misbehaviors. As part of this technique alerts are generated when policy is violated or signature of insider misbehavior is matched.

Hybrid. Approaches simultaneously combining anomaly and misuse-based detection belong to this category. Liu et al. [2009] proposed a conceptual model called *sensitive information dissemination detection* (SIDD) for the detection of insider threats involving the exfiltration of sensitive data to external networks of an organization. SIDD is a network device that is placed at the edge of the network and transparently performs three tasks: identification of applications from the payload of network packets, matching of content signatures, and detection of covert channels.² Considering observables outside of cyberspace (situation-aware observables), Buford et al. [2008] described an architecture for insider threat detection based on *belief-desire-intention* (BDI) *agents* that model behaviors of user roles and malicious insiders and compares these behaviors to the predefined set of plans. Aimed at insiders in the intelligence community, Park and Ho [2004] proposed the

²We consider this work as a conceptual one, as the evaluation was not performed using data clearly related to insider threat.

composite role-based monitoring (CRBM) approach as an extension to RBAC, having separate role structures for organizations, operating systems, applications, and their mutual mappings. In CRBM, a user's behavior is monitored in three separated sessions (OS, application, and organization), and it is compared with expected and unexpected behaviors.

General. A general framework of a tiered conceptual detector was presented by Legg et al. [2013] who designed the framework on the bases of *top/down and bottom/up reasoning in hypothesis trees*, while incorporating three layers: hypothesis, measurement, real world.

8.2.2 Operational Works. In this section we briefly survey insider threat detection and assessment approaches that contain a proof-of-concept evaluated on relevant data. Similar to the previous section, we grouped the approaches into five categories reflecting intrusion detection classification, which is in addition to the previous section enriched by machine learning-based perspective. The resulting categories are: *anomaly-based, misuse-based, hybrid, classification-based, and unsupervised outlier detection*. Here, the difference between hybrid and classification-based approaches is that the former independently merges misuse-based and anomaly-based types, while the latter does that simultaneously using two-class (or multi-class) classifications techniques. In contrast to all of the remaining categories, unsupervised outlier detection category does not require labeled training data.

In addition to intrusion detection-derived categorization, we propose using two other categorizations that can also be applied on operational works dealing with the detection and assessment of insider threat: 1) categorization based on ***the dataset setting used for evaluation*** (in accordance with Section 5.2), and 2) categorization based on ***the feature domains*** (in accordance with [Gheyas and Abdallah 2016]). Further information regarding these additional categorizations are present in appendices C and H, which also demonstrates their application.

Anomaly-Based. Similar to the conceptual works, the most widely adopted approach of operational works is anomaly-based detection. There are works that deal with specific types of cyber observables, as well as studies that combine such cyber observables. In a set of works dealing with **specific cyber observables, Unix/Linux command histories** are the most widespread data source. The most well-known work in this domain was done by Schonlau et al. [2001] who evaluated six anomaly-based methods (sequence match, uniqueness, Markov models, compression and incremental probability action modeling (IPAM)); they also provided the research community with the SEA dataset (referred to as SEA in this paper). Another type of data source used in this field is **file system interaction**. As an example, Camiña et al. [2014] proposed detection systems for masqueraders utilizing SVM and k-NN as one-class techniques that were evaluated on the WUIL dataset. Another example aimed at detection of masqueraders but using graph partitioning is presented by Toffalini et al. [2018] who evaluated their approach on WUIL and TWOS datasets. Considering **system calls**, Liu et al. [2005] performed supervised anomaly detection based on k-NN, in which input features were aggregated by n-grams, histograms, and parameters of system calls. Companies' **databases** represent an attractive target for insiders, and therefore some research has addressed this area specifically. For example, Panigrahi et al. [2013] performed user profiling for the detection of suspicious transactions by using the extended Dempster-Shafer theory to combine multiple pieces of evidence, using inter and intra-transactional features. The possibilities of user identification also increased with the introduction of **GUI**. For example, Sankaranarayanan et al. [2006] used IPAM and compared it with a classification-based Naïve Bayes approach for the detection of masqueraders in Microsoft Word user behavior. In contrast to the previous anomaly-based techniques, an example dealing with **general cyber observables** was presented by Salem and Stolfo [2011b] who proposed an approach using one-class SVM and seven features modeling

search behavior that is aimed at detecting anomalies in the RUU dataset. In addition to finding correlations between insider threat and **psychosocial observables**, operational kinds of research involving such observables have also increased. For example, Brdiczka et al. [2012] proposed a traitor assessment using Bayesian techniques that combined “structural anomaly detection from information and social networks with psychological profiling,” and then evaluated the approach on the World of Warcraft dataset.

Misuse-Based. Aimed at insiders performing data exfiltration resulting into underlying changes to the integrity of directory services, Claycomb and Shin [2010] proposed a *combination of policy with monitoring*, which leverages the capabilities of directory virtualization. Hanley and Montelibano [2011] demonstrated the utilization of *signature alerts in the SPLUNK* logging engine for the detection of data exfiltration. Aimed at high privileged system users, Sibai and Menascé [2011] proposed a system for insider threat detection as a network element that is based on *rule-based policies* (in Snort format) defined for different categories of applications; the system executes decryption of network traffic for payload inspection. For specifying insider misuse signatures, Magklaras and Furnell [2012] designed the *insider threat prediction and specification language* (ITPSL), which has markup features and utilizes logical operators; they evaluated ITPSL’s application on a game containing several malicious and accidental scenarios.

Hybrid. Aimed at the detection of data exfiltration from **network data**, Maloof and Stephens [2007] proposed the ELICIT system, which is based on 76 binary detectors that examine volumetric anomalies, suspicious behaviors, etc.; the outputs of these detectors are passed into a Bayesian networks (BN) to perform threat assessment. Considering cyber observables enriched by **psychosocial indicators**, Legg et al. [2015] proposed a hybrid approach based on known attacks and policy violations combined with threshold and deviation-based anomalies. Considering MOC-based features and using a high level of abstraction, AlGhamdi et al. [2006] applied a multi-entity BN for the assessment of insider threat and legitimate behavior in the **document relevance problem**. Maybury et al. [2005] presented results of the ARDA NRRC (MITRE) workshop for the US intelligence community, where the authors only considered cyber observables for the fusion of three approaches: 1) honeytokens used as Web pages, 2) stealth watch sensors that monitor abnormal network and host activities, and 3) structured top-down analysis modeling pre-attack indicators.

Classification-Based. Aimed at data exfiltration, Azaria et al. [2014] employed SVM and Naïve Bayes as part of their BAIT framework that considers **cyber data** from actions such as transfer (print, copy to thumb drive) and send (by email, HTTP/HTTPS services), while distinguishing between external and internal actions. As part of their work, Mathew et al. [2010] proposed a data-centric approach to role-based masquerader detection in **relational databases**, which applied supervised techniques such as SVM, J.48 decision tree, and Naïve Bayes classifiers. The detection of masquerades in **Unix commands** was addressed by Macion and Townsend [2002], who used a Naïve Bayes classifier with an updating scheme, which considered frequency of particular commands for each user. Inspired by cache memory term locality, Camiña et al. [2016] addressed masqueraders in **file system access** based on temporal and spatial locality features processed by TreeBagger – an ensemble of decision tree classifiers from MATLAB. Garg et al. [2006] proposed the detection of masqueraders in the **GUI** environment by SVM, considering only mouse derived features such as speed, distance, angles, and their statistical properties; evaluation was performed on data collected from three users. For the purpose of assessing the **trustworthiness of entities** such as actors or documents, Mayhew et al. [2015] proposed behavior-based access control (BBAC), which is based on a sequential combination of k-means clustering and SVM. Dealing with **NLP** in comments of YouTube users, Kandias et al. [2013] employed SVM, logistic regression, and Naïve Bayes classifiers

in order to predict users with negative/radical political attitudes, assuming these attitudes to be precursors of insider threat.

Unsupervised Detection of Outliers. Dealing with masqueraders in **Unix commands**, Lane and Brodley [1998] presented an online learning system of user behavior that supports the *concept drift*. Considering **general cyber observables**, Senator et al. [2013] presented the results of the PRODIGAL team, in which several outlier detection methods were designed; the evaluation was performed on data collected by SureView from approximately 5,500 real users and red team scenarios injected by CERT. As part of their work, Mathew et al. [2010] dealt with cluster-based detection of masqueraders in **relational databases**. Aimed at indirect **psychological observables**, Kandias et al. [2013] detected outliers based on social network graphs containing more than one million Greek Twitter users, while the authors assumed that narcissism/extroversion is a precursor of insider threat. Considering aspects of **mutual interactions of users**, Okolica et al. [2008] expanded probabilistic latent semantic indexing (PLSI) by including users in order to extract an individual's interests from emails, while they modeled sharing of interests with individuals' co-workers versus those interests only shared with people external to the organization. More general outlier detection was performed by Eberle and Holder [2010] who applied graph-based anomaly detection for the identification of traitors, which was evaluated on: 1) the Enron dataset, 2) cell phone calls from VAST grant, and 3) business processes containing fraud scenarios.

9 CONCLUSION

The objective of this study was to provide a systematization of knowledge in insider threat research, while leveraging existing grounded theory method for rigorous literature review. In so doing, we identified four main categories of research and development efforts (see Figure 2): (1) In the *incidents and datasets* category we provided references to several sources of insider threat case studies, as well as categorization and details about related datasets. (2) In the *analysis of incidents* category we provided generalization of aspects and behaviors of insider threat and aimed at including research contributions that addressed an insider attack's lifecycle, indicators, and critical pathways, as well as psychosocial point of view. (3) In the *simulations* category we described research utilizing modeling and simulation approaches for experiments with programmed detection methods or for the purpose of data generation. We identified three groups of approaches: a) approaches working with discrete events, b) system dynamics approaches, and c) game theory approaches. (4) The *defense solutions* category, which constitutes the majority of the research included in our survey, was divided into four subcategories: a) mitigation and prevention, b) detection and threat assessment, c) best practices and guidelines, d) decoy-based solutions, and e) other practical solutions. In addition to the categorization itself, our intention was to illustrate workflow between categories, which followed the direction from incidents to defense solutions. Special attention was paid to definitions and taxonomies of insider threat – we proposed a structural taxonomy of insider threat incidents (see Figure 1), which is based on existing taxonomies and the 5W1H questions of the information gathering problem. The proposed taxonomy contributes to orthogonal classification of incidents and defining the scope of defense solutions employed against them.

Finally, we made the following observations aimed at highlighting directions and challenges for future research:

- As indicated by Salem et al. [2008], “a major challenge of insider threat detection research is the lack of real data” for assessing defense solutions; this fact has not changed since then. Moreover, only a few synthetic datasets contain samples of malicious insider attacks. It is important to note that these synthetic datasets are not validated as correctly modeling real

environments. Therefore, the challenge of absence of datasets still holds, and we encourage researchers to create and share related datasets with the community.

- Assuming that an insider is a moving target, we emphasize that detection methods built upon existing datasets may not work in practice. Therefore, we recommend to pursue adversarial classification for improvement of detection approaches as well as for the enrichment of datasets (e.g., using Generative Adversarial Networks [Goodfellow et al. 2014]).
- New datasets should be also designed and created with collusion attacks and several variations of concept drift in mind; taking such natural phenomena into account would improve the datasets and enable more realistic and challenging testing of detection methods.
- There is a trend in the development of detection approaches based on data collected from single-player and multi-player games. The majority of these games are deterministically guided according to scenarios given to the players. The challenge associated with this centers on providing the players of multi-player games with ethical dilemmas, i.e., whether to perform malicious actions for self-benefit or bolster the team and refrain from performing malicious acts (such as in [Brdiczka et al. 2012; Ho et al. 2016]).
- In Table 1 of Appendix H we see a trend of detection approaches that utilize opportunity-based features and do not make use of motive and capability-based features. The reasons for the absence of motive-based features (attributed to psychosocial observables and decoys) may involve privacy issues or the fact that the underlying datasets do not contain suitable data. In contrast, capability-based features can often be derived from existing data (e.g., similar to Magklaras and Furnell [2005]). Therefore, another challenge is to include this kind of information in datasets as well as the detection approaches working with them.
- Social engineering testing or decoy-based solutions may help defend against malicious insider threat and may also generate valuable data for the design of detection methods.
- We observed a trend toward anomaly-based and unsupervised outlier approaches, which can be attributed to class imbalance in datasets and fear of zero-day malicious attacks. However, we believe that a good and robust insider threat defense program should contain a combination of several independent solutions. In the first line of defense, procedural and other best practices involving mitigation and prevention techniques (such as those mentioned in Section 8.1) should be present. In the second line of defense, there should be misuse-based detection that covers existing insider threat scenarios. Finally, in the third line of defense, anomaly-based and unsupervised outlier detection should be deployed. Alerts from the latter two layers should be correlated together in a dedicated view.
- The last recommendation for future research in defense solutions is to specify the MOC-based feature domains of the approach (as presented in Appendix H) and enumerate the detection scope covered by the defense solution (e.g., using the taxonomy in Figure 1).

REFERENCES

- A. Abdallah, M. A. Maarof, and A. Zainal. 2016. Fraud detection system: A survey. *Journal of Network and Computer Applications* 68 (2016), 90–113.
- S. Achleitner, T. La Porta, P. McDaniel, S. Sugrim, S. V. Krishnamurthy, and R. Chadha. 2016. Cyber Deception: Virtual Networks to Defend Insider Reconnaissance. In *Int. Workshop on Managing Insider Security Threats*. ACM, 57–68.
- I. Agrafiotis, A. Erola, M. Goldsmith, and S. Creese. 2016. A Tripwire Grammar for Insider Threat Detection. In *Int. Workshop on Managing Insider Security Threats*. ACM, 105–108.
- B. Aleman-Meza, P. Burns, M. Eavenson, D. Palaniswami, and A. Sheth. 2005. An ontological approach to the document access problem of insider threat. In *Int. Conference on Intelligence and Security Informatics*. Springer, 486–491.
- G. AlGhamdi, K. B. Laskey, E. J. Wright, D. Barbará, and K. Chang. 2006. Modeling Insider Behavior Using Multi-Entity Bayesian Networks. In *Int. Command and Control Research and Technology Symposium*.

- G. Ali, N. A. Shaikh, and Z. A. Shaikh. 2008. Towards an automated multiagent system to monitor user activities against insider threat. In *Int. Symposium on Biometrics and Security Technologies*. IEEE, 1–5.
- S. Alneyadi, E. Sithirasenan, and V. Muthukkumarasamy. 2016. A survey on data leakage prevention systems. *Journal of Network and Computer Applications* 62 (2016), 137–152.
- Q. Althebyan and B. Panda. 2007. A knowledge-base model for insider threat prediction. In *Information Assurance and Security Workshop, 2007. IAW'07*. IEEE SMC. IEEE, 239–246.
- Q. Althebyan and B. Panda. 2008. Performance analysis of an insider threat mitigation model. In *Int. Conference on Digital Information Management*. IEEE, 703–709.
- M. L. Ambrose, M. A. Seabright, and M. Schminke. 2002. Sabotage in the workplace: The role of organizational injustice. *Organizational behavior and human decision processes* 89, 1 (2002), 947–965.
- D. F. Andersen, D. Cappelli, J. J. Gonzalez, M. Mojtahedzadeh, A. Moore, E. Rich, J. M. Sarriegui, T. J. Shimeall, J. Stanton, E. Weaver, and others. 2004. Preliminary system dynamics maps of the insider cyber-threat problem. In *Int. Conference of the System Dynamics Society*. 25–29.
- J. P. Anderson. 1980. *Computer security threat monitoring and surveillance*. Technical Report. James P Anderson Company.
- R. H. Anderson. 1999. *Research and Development Initiatives Focused on Preventing, Detecting, and Responding to Insider Misuse of Critical Defense Information Systems*. Technical Report. RAND Corporation.
- E. T. Axelrad, P. J. Sticha, O. Brdiczka, and J. Shen. 2013. A Bayesian network model for predicting insider threats. In *Security and Privacy Workshops*. IEEE, 82–89.
- A. Azaria, A. Richardson, S. Kraus, and V. S. Subrahmanian. 2014. Behavioral Analysis of Insider Threat: A Survey and Bootstrapped Prediction in Imbalanced Data. *Transactions on Computational Social Systems* 1, 2 (2014), 135–155.
- S. R. Band, D. M. Cappelli, L. F. Fischer, A. P. Moore, E. D. Shaw, and R. F. Trzeciak. 2006. *Comparing insider IT sabotage and espionage: A model-based analysis*. Technical Report. DTIC Document.
- J. Banks. 1998. *Handbook of simulation: principles, methodology, advances, applications, and practice*. John Wiley & Sons.
- N. Baracaldo and J. Joshi. 2012. A trust-and-risk aware RBAC framework: tackling insider threat. In *Symposium on Access Control Models and Technologies*. ACM, 167–176.
- S. M. Bellovin. 2008. The insider attack problem nature and scope. In *Insider Attack and Cyber Security*. Springer, 1–4.
- V. H. Berk, G. Cybenko, I. Gregorio-de Souza, and J. P. Murphy. 2012. Managing malicious insider risk through bandit. In *Hawaii Int. Conference on System Science*. IEEE, 2422–2430.
- M. Bertacchini and P. Fierens. 2008. A survey on masquerader detection approaches. In *Congreso Iberoamericano de Seguridad Informática, Universidad de la República de Uruguay*. 46–60.
- E. Bertino and G. Ghinita. 2011. Towards mechanisms for detection and prevention of data exfiltration by insiders: keynote talk paper. In *Symposium on Information, Computer and Communications Security*. ACM, 10–19.
- D. Bhilare, A. Ramani, and S. Tanwani. 2009. Protecting intellectual property and sensitive information in academic campuses from trusted insiders: leveraging active directory. In *SIGUCCS fall conference*. ACM, 99–104.
- M. Bishop. 2005. Position: Insider is relative. In *Workshop on New Security Paradigms*. ACM, 77–78.
- M. Bishop, H. M. Conboy, H. Phan, B. I. Simidchieva, G. S. Avrunin, L. A. Clarke, L. J. Osterweil, and S. Peisert. 2014. Insider threat identification by process analysis. In *Security and Privacy Workshops*. IEEE, 251–264.
- M. Bishop, S. Engle, S. Peisert, S. Whalen, and C. Gates. 2008. We have met the enemy and he is us. In *Workshop on New Security Paradigms*. ACM, 1–12.
- M. Bishop, S. Engle, S. Peisert, S. Whalen, and C. Gates. 2009a. Case studies of an insider framework. In *Hawaii Int. Conference on System Sciences*. IEEE, 1–10.
- M. Bishop and C. Gates. 2008. Defining the insider threat. In *Workshop on Cyber security and Information Intelligence Research*. ACM, 15.
- M. Bishop, C. Gates, D. Frincke, and F. L. Greitzer. 2009b. AZALIA: an A to Z Assessment of the Likelihood of Insider Attack. In *Technologies for Homeland Security, 2009. HST'09*. IEEE Conference on. IEEE, 385–392.
- C. Blackwell. 2009. A security architecture to protect against the insider threat from damage, fraud and theft. In *Workshop on Cyber Security and Information Intelligence Research*. ACM, 45.
- B. M. Bowen, S. Hershkop, A. D. Keromytis, and S. J. Stolfo. 2009. Baiting inside attackers using decoy documents. In *Int. Conference on Security and Privacy in Communication Systems*. Springer, 51–70.
- R. C. Brackney and R. H. Anderson. 2004. *Workshop on Understanding the Insider Threat*. Technical Report. RAND Corporation, DTIC document.
- O. Brdiczka, J. Liu, B. Price, J. Shen, A. Patil, R. Chow, E. Bart, and N. Ducheneaut. 2012. Proactive insider threat detection through graph learning and psychological context. In *Security and Privacy Workshops*. 142–149.
- C. R. Brown, A. Watkins, and F. L. Greitzer. 2013. Predicting insider threat risks through linguistic analysis of electronic communication. In *Hawaii Int. Conference on System Sciences*. IEEE, 1849–1858.

- J. F. Buford, L. Lewis, and G. Jakobson. 2008. Insider threat detection using situation-aware MAS. In Int. Conference on Information Fusion. IEEE, 1–8.
- CALO Project. 2015. Enron Email Dataset. (2015). <http://www.cs.cmu.edu/~enron/> Accessed on May/2017.
- B. Camiña, C. Hernández-Gracidas, R. Monroy, and L. Trejo. 2014. The Windows-Users and Intruder simulations Logs dataset (WUIL): An experimental framework for masquerade detection mechanisms. Expert Systems with Applications 41, 3 (2014), 919–930.
- B. Camiña, R. Monroy, L. A. Trejo, and M. A. Medina-Pérez. 2016. Temporal and spatial locality: an abstraction for masquerade detection. IEEE transactions on information Forensics and Security 11, 9 (2016), 2036–2051.
- B. Camiña, R. Monroy, L. A. Trejo, and E. Sánchez. 2011. Towards building a masquerade detection method based on user file system navigation. In Mexican Int. Conference on Artificial Intelligence. Springer, 174–186.
- J. B. Camiña, J. Rodríguez, and R. Monroy. 2014. Towards a Masquerade Detection System Based on User’s Tasks. In Int. Workshop on Recent Advances in Intrusion Detection. Springer, 447–465.
- D. M. Cappelli, A. P. Moore, and R. F. Trzeciak. 2012. The CERT guide to insider threats: how to prevent, detect, and respond to information technology crimes (Theft, Sabotage, Fraud). Addison-Wesley.
- Carbon Black. 2017. Cb Endpoint Security Platform. (2017). <https://www.carbonblack.com/products/> Accessed on August/2017.
- CERT. 2014. Unintentional insider threats: social engineering. Software Engineering Institute (2014).
- M. Chagarlamudi, B. Panda, and Y. Hu. 2009. Insider Threat in Database Systems: Preventing Malicious Users’ Activities in Databases. In Int. Conference on Information Technology: New Generations. IEEE, 1616–1620.
- T. Chen, F. Kammüller, I. Nemli, and C. W. Probst. 2015. A probabilistic analysis framework for malicious insider threats. In Int. Conference on Human Aspects of Information Security, Privacy, and Trust. Springer, 178–189.
- Y. Chen, S. Nyemba, and B. Malin. 2012. Detecting anomalous insiders in collaborative information systems. IEEE transactions on dependable and secure computing 9, 3 (2012), 332–344.
- R. Chinchani, D. Ha, A. Iyer, H. Q. Ngo, and S. Upadhyaya. 2010. Insider threat assessment: Model, analysis and tool. In Network Security. Springer, 143–174.
- W. R. Claycomb, C. L. Huth, L. Flynn, D. M. McIntire, T. B. Lewellen, and C. I. T. Center. 2012. Chronological Examination of Insider Threat Sabotage: Preliminary Observations. JoWUA 3, 4 (2012), 4–20.
- W. R. Claycomb and A. Nicoll. 2012. Insider threats to cloud computing: Directions for new research challenges. In Annual Computer Software and Applications Conference. IEEE, 387–394.
- W. R. Claycomb and D. Shin. 2010. Detecting insider activity using enhanced directory virtualization. In Workshop on Insider threats. ACM, 29–36.
- E. Cole and S. Ring. 2005. Insider threat: Protecting the enterprise from sabotage, spying, and theft. Syngress.
- L. Coles-Kemp and M. Theoharidou. 2010. Insider threat and information security management. In Insider threats in cyber security. Springer, 45–71.
- M. L. Collins, M. C. Theis, R. F. Trzeciak, J. R. Strozer, J. W. Clark, D. L. Costa, T. Cassidy, M. J. Albrethsen, and A. P. Moore. 2016. Common sense guide to prevention and detection of insider threats 5th edition. Published by CERT, Software Engineering Institute, Carnegie Mellon University (2016).
- J. B. Colombe and G. Stephens. 2004. Statistical profiling and visualization for detection of malicious insider attacks on computer networks. In Workshop on Visualization and data mining for computer security. ACM, 138–142.
- C. Colwill. 2009. Human factors in information security: The insider threat—Who can you trust these days? Information security technical report 14, 4 (2009), 186–196.
- S. Coull, J. Branch, B. Szymanski, and E. Breimer. 2003. Intrusion detection: A bioinformatics approach. In Computer Security Applications Conference. IEEE, 24–33.
- J. Crampton and M. Huth. 2010. Towards an access-control framework for countering insider threats. In Insider Threats in Cyber Security. Springer, 173–195.
- A. Cummings, T. Lewellen, D. McIntire, A. P. Moore, and R. Trzeciak. 2012. Insider threat study: Illicit cyber activity involving fraud in the US financial services sector. Technical Report. CERT.
- DarkTrace. 2017. Darktrace (Core). (2017). <https://www.darktrace.com/products/> Accessed on August/2017.
- Y. Desmedt and A. Shaghghi. 2016. Function-Based Access Control (FBAC): From Access Control Matrix to Access Control Tensor. In Int. Workshop on Managing Insider Security Threats. ACM, 89–92.
- T. Dimkov, W. Pieters, and P. Hartel. 2010. Portunes: representing attack scenarios spanning through the physical, digital and social domain. In Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security. Springer, 112–129.
- G. Doss and G. Tejay. 2009. Developing insider attack detection model: a grounded approach. In Int. Conference on Intelligence and Security Informatics. IEEE, 107–112.

- W. Eberle, J. Graves, and L. Holder. 2010. Insider threat detection using a graph-based approach. *Journal of Applied Security Research* 6, 1 (2010), 32–81.
- W. Eberle and L. Holder. 2009. Mining for insider threats in business transactions and processes. In *Computational Intelligence and Data Mining, 2009. CIDM'09. IEEE Symposium on*. IEEE, 163–170.
- M. E. Edge and P. R. F. Sampaio. 2009. A survey of signature based methods for financial fraud detection. *computers & security* 28, 6 (2009), 381–394.
- N. Einwechter. 2010. Preventing and Detecting Insider Attacks Using IDS. (2010). <https://www.symantec.com/connect/articles/preventing-and-detecting-insider-attacks-using-ids> Accessed on August/2017.
- Ekran System. 2017. Ekran System. (May 2017). <https://www.ekransystem.com> Accessed on May/2017.
- A. El Masri, H. Wechsler, P. Likarish, and B. B. Kang. 2014. Identifying users with application-specific command streams. In *Int. Conference on Privacy, Security and Trust*. IEEE, 232–238.
- H. Eldardiry, E. Bart, J. Liu, J. Hanley, B. Price, and O. Brdiczka. 2013. Multi-Domain Information Fusion for Insider Threat Detection. *Security and Privacy Workshops* (2013), 45–51.
- J. Eom, M. Park, S. Park, and T. Chung. 2011. A framework of defense system for prevention of insider's malicious behaviors. In *Int. Conference on Advanced Communication Technology*. IEEE, 982–987.
- F. Farahmand and E. H. Spafford. 2009. Insider behavior: an analysis of decision under risk. In *Int. Workshop on Managing Insider Security Threats*. 22.
- F. Farahmand and E. H. Spafford. 2013. Understanding insiders: An analysis of risk-taking behavior. *Information systems frontiers* 15, 1 (2013), 5–15.
- L. F. Fischer. 2003. Characterizing information systems insider offenders. In *Conference of the Intl. Military Testing Association*. Citeseer.
- V. N. Franqueira, A. van Cleeff, P. van Eck, and R. Wieringa. 2010. External insider threat: A real security challenge in enterprise value webs. In *Int. Conference on Availability, Reliability, and Security*. 446–453.
- R. Garfinkel, R. Gopal, and P. Goes. 2002. Privacy protection of binary confidential data against deterministic, stochastic, and insider threat. *Management Science* 48, 6 (2002), 749–764.
- R. Garfinkel, R. Gopal, and D. Rice. 2006. New approaches to disclosure limitation while answering queries to a database: protecting numerical confidential data against insider threat based on data or algorithms. In *Hawaii Int. Conference on System Sciences*, Vol. 6. IEEE, 125a–125a.
- A. Garg, R. Rahalkar, S. Upadhyaya, and K. Kwiat. 2006. Profiling users in GUI based systems for masquerade detection. In *2006 IEEE Information Assurance Workshop*. IEEE, 48–54.
- C. Gates, N. Li, Z. Xu, S. N. Chari, I. Molloy, and Y. Park. 2014. Detecting Insider Information Theft Using Features from File Access Logs. In *European Symposium on Research in Computer Security*, Vol. 8713 LNCS. Springer, 383–400.
- I. A. Gheyas and A. E. Abdallah. 2016. Detection and prediction of insider threats to cyber security: a systematic literature review and meta-analysis. *Big Data Analytics* 1, 1 (2016), 6.
- J. Glasser and B. Lindauer. 2013. Bridging the gap: A pragmatic approach to generating insider threat data. In *Security and Privacy Workshops*. IEEE, 98–104.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- R. Gopal, R. Garfinkel, and P. Goes. 2002. Confidentiality via Camouflage: The CVC Approach to Disclosure Limitation When Answering Queries to Databases. *Operations Research* 50, 3 (2002), 501–516.
- S. Greenberg. 1988. *Using Unix: Collected traces of 168 users*. Technical Report.
- F. L. Greitzer and D. A. Frincke. 2010. Combining traditional cyber security audit data with psychosocial data: towards predictive modeling for insider threat mitigation. In *Insider Threats in Cyber Security*. Springer, 85–113.
- F. L. Greitzer, D. A. Frincke, and M. Zabriskie. 2010. Social/ethical issues in predictive insider threat monitoring. *Information Assurance and Security Ethics in Complex Systems: Interdisciplinary Perspectives* (2010), 132–161.
- F. L. Greitzer and R. E. Hohimer. 2011. Modeling Human Behavior to Anticipate Insider Attacks. *Journal of Strategic Security* 4, 2 (2011), 25.
- F. L. Greitzer, L. J. Kangas, C. F. Noonan, C. R. Brown, and T. Ferryman. 2013. Psychosocial modeling of insider threat risk based on behavioral and word use analysis. *e-Service Journal* 9, 1 (2013), 106–138.
- F. L. Greitzer, L. J. Kangas, C. F. Noonan, A. C. Dalton, and R. E. Hohimer. 2012. Identifying at-risk employees: Modeling psychosocial precursors of potential insider threats. In *Hawaii Int. Conference on System Science*. IEEE, 2392–2401.
- F. L. Greitzer, J. Strozer, S. Cohen, J. Bergey, J. Cowley, A. Moore, and D. Mundie. 2014. Unintentional insider threat: contributing factors, observables, and mitigation strategies. In *Hawaii Int. Conference on System Sciences*. IEEE, 2025–2034.
- D. Gritzalis, V. Stavrou, M. Kandias, and G. Stergiopoulos. 2014. Insider threat: enhancing BPM through social media. In *Int. Conference on New Technologies, Mobility and Security*. IEEE, 1–6.

- M. D. Guido and M. W. Brooks. 2013. Insider threat program best practices. In Hawaii Int. Conference on System Sciences. IEEE, 1831–1839.
- D. Ha, S. Upadhyaya, H. Ngo, S. Pramanik, R. Chinchani, and S. Mathew. 2007. Insider threat analysis using information-centric modeling. In IFIP Int. Conference on Digital Forensics. Springer, 55–73.
- M. Hanley and J. Montelibano. 2011. Insider threat control: Using centralized logging to detect data exfiltration near insider termination. Technical Report. DTIC Document.
- A. Harilal, F. Toffalini, J. Castellanos, J. Guarnizo, I. Homoliak, and M. Ochoa. 2017. TWOS: A Dataset of Malicious Insider Threat Behavior Based on a Gamified Competition. In Int. Workshop on Managing Insider Security Threats. ACM, 35–46.
- Athul Harilal, Flavio Toffalini, Ivan Homoliak, John Castellanos, Juan Guarnizo, Soumik Mondal, and MartiÀrn Ochoa. 2018. The Wolf Of SUTD (TWOS): A Dataset of Malicious Insider Threat Behavior Based on a Gamified Competition. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA) 9, 1 (March 2018), 54–85.
- Hawlett Packard. 2017. ArcSight Express. (2017). <http://www8.hp.com/us/en/software-solutions/arcSight-express-siem-appliance/index.html> Accessed on August/2017.
- M. Hayden. 1999. The insider threat to US government information systems. Technical Report. DTIC Document.
- S. Ho, J. Hancock, C. Booth, M. Burmester, X. Liu, and S. Timmarajus. 2016. Demystifying insider threat: Language-action cues in group dynamics. In Hawaii Int. Conference on System Sciences. IEEE, 2729–2738.
- S. M. Ho. 2008. Attribution-based anomaly detection: trustworthiness in an online community. In Social Computing, Behavioral Modeling, and Prediction. Springer, 129–140.
- E. Humphreys. 2008. Information security management standards: Compliance, governance and risk management. information security technical report 13, 4 (2008), 247–255.
- J. Hunker and C. W. Probst. 2011. Insiders and Insider Threats: An Overview of Definitions and Mitigation Techniques. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications 2, 1 (2011), 4–27.
- W. B. Jaballah and N. Kheir. 2016. A Grey-Box Approach for Detecting Malicious User Interactions in Web Applications. In Int. Workshop on Managing Insider Security Threats. ACM, 1–12.
- G. Jabbour and D. Menascé. 2009a. Stopping the insider threat: the case for implementing autonomic defense mechanisms in computing systems. In Int. Conference of Information Security and Privacy.
- G. G. Jabbour and D. A. Menascé. 2009b. The insider threat security architecture: a framework for an integrated, inseparable, and uninterrupted self-protection mechanism. In Int. Conference on Computational Science and Engineering. IEEE, 244–251.
- R. V. Johnson, J. Lass, and W. M. Petullo. 2016. Studying Naive Users and the Insider Threat with SimpleFlow. In Int. Workshop on Managing Insider Security Threats. ACM, 35–46.
- F. Kammüller, J. R. C. Nurse, and C. W. Probst. 2016. Attack tree analysis for insider threats on the iot using isabelle. In Int. Conference on Human Aspects of Information Security, Privacy, and Trust. Springer, 234–246.
- M. Kandias, K. Galbogini, L. Mitrou, and D. Gritzalis. 2013. Insiders trapped in the mirror reveal themselves in social media. In Int. Conference on Network and System Security. Springer, 220–235.
- M. Kandias, A. Mylonas, N. Virvilis, M. Theoharidou, and D. Gritzalis. 2010. An Insider Threat Prediction Model. In TrustBus - Int. Conference. Springer, 26–37.
- M. Kandias, V. Stavrou, N. Bozovic, L. Mitrou, and D. Gritzalis. 2013. Can we trust this user? Predicting insider’s attitude via YouTube usage profiling. In Int. Conference on Ubiquitous Intelligence and Computing. IEEE, 347–354.
- M. Kandias, N. Virvilis, and D. Gritzalis. 2011. The insider threat in cloud computing. In Int. Workshop on Critical Information Infrastructures Security. Springer, 93–103.
- I. Kantzavelou and S. Katsikas. 2010. A game-based intrusion detection mechanism to confront internal attackers. computers & security 29, 8 (2010), 859–874.
- M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, S. Rogers, and others. 2005. Insider threat study: Computer system sabotage in critical infrastructure sectors. Technical Report. National Threat Assessment Centre, Washington DC.
- Andrew Kellett. 2015. Trends and Future Directions in Data Security – 2015 Vormetric Insider Threat Report. Technical Report. Vormetric Data Security.
- Kibana. 2017. Kibana. (2017). <https://www.elastic.co/products/kibana> Accessed on May/2017.
- K. S. Killourhy and R. A. Maxion. 2008. Naive bayes as a masquerade detector: Addressing a chronic failure. In Insider Attack and Cyber Security. Springer, 91–112.
- H. Kim and S. Cha. 2005. Empirical evaluation of SVM-based masquerade detection using UNIX commands. Computers & Security 24, 2 (2005), 160–168.
- E. Kowalski, T. Conway, S. Keeverline, M. Williams, D. Cappelli, B. Willke, and A. Moore. 2008. Insider threat study: Illicit cyber activity in the government sector. US Secret Service, SEI CMU (2008).

- T. Lane and C. E. Brodley. 1997. An application of machine learning to anomaly detection. In National Information Systems Security Conference, Vol. 377. Baltimore, USA, 366–380.
- T. Lane and C. E. Brodley. 1998. Approaches to Online Learning and Concept Drift for User Identification in Computer Security. In KDD. 259–263.
- M. Latendresse. 2005. Masquerade detection via customized grammars. In Int. Conference on Detection of Intrusions and Malware, and Vulnerability Assessment. Springer, 141–159.
- A. Lazouski, F. Martinelli, and P. Mori. 2010. Usage control in computer security: A survey. Computer Science Review 4, 2 (2010), 81 – 99. <http://www.sciencedirect.com/science/article/pii/S1574013710000146>
- J. Leach. 2003. Improving user security behaviour. Computers & Security 22, 8 (2003), 685–692.
- J. Lee and Y. Lee. 2002. A holistic model of computer abuse within organizations. Information management & computer security 10, 2 (2002), 57–63.
- P. Legg, N. Moffat, J. R. C. Nurse, J. Happa, I. Agrafiotis, M. Goldsmith, and S. Creese. 2013. Towards a Conceptual Model and Reasoning Structure for Insider Threat Detection. JoWUA 4 (2013), 20–37.
- P. A. Legg, O. Buckley, M. Goldsmith, and S. Creese. 2015. Automated insider threat detection system using user and role-based profile assessment. (2015).
- D. Liginlal, I. Sim, and L. Khansa. 2009. How significant is human error as a cause of privacy breaches? An empirical study and a framework for error management. computers & security 28, 3 (2009), 215–228.
- F. Linton, D. Joy, H. Schaefer, and A. Charron. 2000. OWL: A recommender system for organization-wide learning. Educational Technology & Society 3, 1 (2000), 62–76.
- R. P. Lippman, D. J. Fried, I. Graf, J. W. Haines, K. R. Kendall, D. McClung, D. Weber, S. E. Webster, D. Wyschogrod, R. K. Cunningham, and others. 2000. Evaluating intrusion detection systems: The 1998 DARPA off-line intrusion detection evaluation. In DARPA Information Survivability Conference and Exposition, DISCEX'00, Vol. 2. IEEE, 12–26.
- A. Liu, C. Martin, T. Hetherington, and S. Matzner. 2005. A comparison of system call feature for insider threat detection. In Proc. of the 6th Annual IEEE Systems, Man & Cybernetics, Information Assurance Workshop. 341–347.
- D. Liu, X. Wang, and J. Camp. 2008. Game-theoretic modeling and analysis of insider threats. Int. Journal of Critical Infrastructure Protection 1 (2008), 75–80.
- D. Liu, X. Wang, and L. J. Camp. 2009. Mitigating inadvertent insider threats with incentives. In Int. Conference on Financial Cryptography and Data Security. Springer, 1–16.
- Y. Liu, C. Corbett, K. Chiang, R. Archibald, B. Mukherjee, and D. Ghosal. 2009. SIDD: A framework for detecting sensitive data exfiltration by an insider attack. In Hawaii Int. Conference on System Sciences. IEEE, 1–10.
- K. D. Loch, H. H. Carr, and M. E. Warkentin. 1992. Threats to information systems: today's reality, yesterday's understanding. Mis Quarterly (1992), 173–186.
- M. Maasberg, J. Warren, and N. L. Beebe. 2015. The dark side of the insider: detecting the insider threat through examination of dark triad personality traits. In Hawaii Int. Conference on System Sciences. IEEE, 3518–3526.
- G. Magklaras and S. Furnell. 2002. Insider threat prediction tool: Evaluating the probability of IT misuse. Computers & Security 21, 1 (2002), 62–73.
- G. Magklaras and S. Furnell. 2005. A preliminary model of end user sophistication for insider threat prediction in IT systems. Computers & Security 24, 5 (2005), 371–380.
- G. Magklaras and S. Furnell. 2012. The Insider Threat Prediction and Specification Language. In INC. 51–61.
- G. Magklaras, S. Furnell, and M. Papadaki. 2011. LUARM—An audit engine for insider misuse detection. In WDFIA. 133–148.
- M. A. Maloof and G. D. Stephens. 2007. Elicit: A system for detecting insiders who violate need-to-know. In Int. Workshop on Recent Advances in Intrusion Detection. Springer, 146–166.
- T. Markham and C. Payne. 2001. Security at the network edge: A distributed firewall architecture. In DARPA Information Survivability Conference & Exposition, Vol. 1. IEEE, 279–286.
- I. Martinez-Moyano, E. Rich, S. Conrad, D. Andersen, and T. Stewart. 2008. A behavioral theory of insider-threat risks: A system dynamics approach. ACM Trans. on Modeling and Computer Simulation 18, 2 (2008).
- I. J. Martinez-Moyano, S. H. Conrad, and D. F. Andersen. 2011. Modeling behavioral considerations related to information security. computers & security 30, 6 (2011), 397–409.
- S. Mathew, M. Petropoulos, H. Q. Ngo, and S. J. Upadhyaya. 2010. A Data-Centric Approach to Insider Attack Detection in Database Systems. In RAID. Springer, 382–401.
- S. Mathew, S. Upadhyaya, D. Ha, and H. Q. Ngo. 2008. Insider abuse comprehension through capability acquisition graphs. In Int. Conference on Information Fusion. IEEE, 1–8.
- R. Maxion. 2003. Masquerade Detection Using Enriched Command Lines. In Int. Conference on Dependable Systems and Networks, Vol. 3. 5–14.
- R. A. Maxion and T. N. Townsend. 2002. Masquerade detection using truncated command lines. In Int. Conference on Dependable Systems and Networks. IEEE, 219–228.

- M. Maybury, P. Chase, B. Cheikes, D. Brackney, S. Matzner, T. Hetherington, B. Wood, C. Sibley, J. Marin, and T. Longstaff. 2005. Analysis and detection of malicious insiders. Technical Report. DTIC Document.
- M. Mayhew, M. Atighetchi, A. Adler, and R. Greenstadt. 2015. Use of machine learning in big data analytics for insider threat detection. In Military Communications Conference, MILCOM 2015-2015 IEEE. IEEE, 915–922.
- M. McCormick. 2008. Data Theft: A Prototypical Insider Threat. Springer US, Boston, MA, 53–68.
- J. McHugh. 2000. Testing intrusion detection systems: a critique of the 1998 and 1999 darpa intrusion detection system evaluations as performed by lincoln laboratory. ACM Transactions on Information and System Security (TISSEC) 3, 4 (2000), 262–294.
- C. Melara, J. M. Sarriegui, J. J. Gonzalez, A. Sawicka, and D. L. Cooke. 2003. A system dynamics model of an insider attack on an information system. In Int. Conference of the System Dynamics Society. Citeseer, 20–24.
- D. Miller, S. Harris, A. Harper, S. VanDyke, and C. Blask. 2010. Security information and event management (SIEM) implementation. McGraw Hill Professional.
- R. Miller and M. Maxim. 2015. I have to trust someone... don't I? Dealing with insider threats to cyber-security. Technical Report. CA Technologies.
- MKinsight. 2017. MKinsight. (May 2017). <http://www.mkinsight.com> Accessed on May/2017.
- D. Moher, A. Liberati, J. Tetzlaff, and D. G. Altman. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. Annals of internal medicine 151, 4 (2009), 264–269.
- A. P. Moore, D. M. Cappelli, T. C. Caron, E. Shaw, D. Spooner, and R. F. Trzeciak. 2011. A preliminary model of insider theft of intellectual property. Technical Report. CERT.
- A. P. Moore, D. M. Cappelli, and R. F. Trzeciak. 2008. The “big picture” of insider IT sabotage across US critical infrastructures. Technical Report. Carnegie Mellon University.
- M. Moore. 2016. Cybersecurity Breaches and Issues Surrounding Online Threat Protection. IGI Global.
- J. P. Murphy, V. H. Berk, and I. Gregorio-de Souza. 2012. Decision support procedure in the insider threat domain. In Security and Privacy Workshops. IEEE, 159–163.
- J. Myers, M. Grimaila, and R. Mills. 2009. Towards insider threat detection using web server logs. In Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies. ACM.
- R. B. Myerson. 1997. Game theory. Harvard university press.
- Y. Naghmouchi, N. Perrot, N. Kheir, R. Mahjoub, and J. Wary. 2016. A New Risk Assessment Framework Using Graph Theory for Complex ICT Systems. In Int. Workshop on Managing Insider Security Threats. ACM, 97–100.
- P. M. Nasr and A. Y. Varjani. 2014. Alarm based anomaly detection of insider attacks in SCADA system. In Smart Grid Conference (SGC), 2014. IEEE, 1–6.
- P. G. Neumann. 2010. Combatting insider threats. In Insider Threats in Cyber Security. Springer, 17–44.
- N. T. Nguyen, P. L. Reiher, and G. H. Kuenning. 2003. Detecting Insider Threats by Monitoring System Call Activity. In IAW. Citeseer, 45–52.
- J. R. C. Nurse, O. Buckley, P. A. Legg, M. Goldsmith, S. Creese, G. R. T. Wright, and M. Whitty. 2014. Understanding Insider Threat: A Framework for Characterising Attacks. In Workshop on Research for Insider Threat. IEEE, 214–228.
- ObserveIt. 2017. Observe It. (May 2017). <http://www.observeit.com> Accessed on May/2017.
- M. Oka, Y. Oyama, H. Abe, and K. Kato. 2004. Anomaly detection using layered networks based on eigen co-occurrence matrix. In Int. Workshop on Recent Advances in Intrusion Detection. Springer, 223–237.
- J. S. Okolica, G. L. Peterson, and R. F. Mills. 2008. Using PLSI-U to detect insider threats by datamining e-mail. Int. Journal of Security and Networks 3, 2 (2008), 114–121.
- J. Ophoff, A. Jensen, J. Sanderson-Smith, M. Porter, and K. Johnston. 2014. A Descriptive Literature Review and Classification of Insider Threat Research. Technical Report.
- S. Panigrahi, S. Sural, and A. K. Majumdar. 2013. Two-stage database intrusion detection by combining multiple evidence and belief update. Information Systems Frontiers 15, 1 (2013), 35–53.
- J. S. Park and J. Giordano. 2006. Role-based profile analysis for scalable and accurate insider-anomaly detection. In Int. Performance Computing and Communications Conference. IEEE.
- J. S. Park and S. M. Ho. 2004. Composite Role-Based Monitoring (CRBM) for Countering Insider Threats. Springer, 201–213.
- P. Parveen, J. Evans, B. Thuraisingham, K. W. Hamlen, and L. Khan. 2011a. Insider threat detection using stream mining and graph mining. In Int. Conference on Privacy, Security, Risk and Trust. IEEE, 1102–1110.
- P. Parveen, Z. R. Weger, B. Thuraisingham, K. Hamlen, and L. Khan. 2011b. Supervised learning for insider threat detection using stream mining. In Int. Conference on Tools with Artificial Intelligence. IEEE, 1032–1039.
- S. L. Pfleeger, J. B. Predd, J. Hunker, and C. Bulford. 2010. Insiders behaving badly: addressing bad actors and their actions. IEEE Transactions on Information Forensics and Security 5, 1 (2010), 169–179.

- A. H. Phyo and S. M. Furnell. 2004. A detection-oriented classification of insider it misuse. In Third Security Conference. Citeseer.
- A. H. Phyo, S. M. Furnell, and F. Portilla. 2004. A Framework for Role-Based Monitoring of Insider Misuse. Springer US, Boston, MA, 51–65.
- R. Posadas, C. Mex-Perera, R. Monroy, and J. Nolasco-Flores. 2006. Hybrid method for detecting masqueraders using session folding and hidden markov models. In Mexican Int. Conference on Artificial Intelligence. Springer, 622–631.
- C. Posey, R. J. Bennett, and T. L. Roberts. 2011. Understanding the mindset of the abusive insider: An examination of insiders' causal reasoning following internal security changes. Computers & Security 30, 6 (2011), 486–497.
- S. Pramanik, V. Sankaranarayanan, and S. Upadhyaya. 2004. Security policies to mitigate insider threat in the document control domain. In Computer Security Applications Conference, 2004. 20th Annual. IEEE, 304–313.
- J. Predd, S. L. Pfleeger, J. Hunker, and C. Bulford. 2008. Insiders behaving badly. IEEE Security & Privacy 6, 4 (2008), 0066–70.
- C. W. Probst, R. R. Hansen, and F. Nielson. 2006. Where can an insider attack?. In Int. Workshop on Formal Aspects in Security and Trust. Springer, 127–142.
- C. W. Probst and J. Hunker. 2010. The risk of risk analysis and its relation to the economics of insider threats. In Economics of information security and privacy. Springer, 279–299.
- C. W. Probst, J. Hunker, M. Bishop, and D. Gollmann. 2008. Summary - Countering Insider Threats. In Countering Insider Threats (Dagstuhl Seminar). Leibniz-Zentrum fuer Informatik, Germany.
- C. W. Probst, J. Hunker, D. Gollmann, and M. Bishop. 2010. Aspects of insider threats. In Insider Threats in Cyber Security. Springer, 1–15.
- PWC. 2017. Global Economic Crime Survey 2016: US Results. (2017). <https://www.pwc.com/us/en/forensic-services/economic-crime-survey-us-supplement.html> Accessed on August/2017.
- M. Raissi-Dehkordi and D. Carr. 2011. A multi-perspective approach to insider threat detection. In 2011-MILCOM 2011 Military Communications Conference. IEEE, 1164–1169.
- M. R. Randazzo, M. Keeney, E. Kowalski, D. Cappelli, and A. Moore. 2005. Insider threat study: Illicit cyber activity in the banking and finance sector. Technical Report. CERT Coordination Center, Software Engineering Institute, Carnegie Mellon University (PA, USA).
- V. Raskin, J. M. Taylor, and C. F. Hempelmann. 2010. Ontological semantic technology for detecting insider threat and social engineering. In Workshop on New security paradigms. ACM, 115–128.
- I. Ray and N. Poolsapassit. 2005. Using attack trees to identify malicious attacks from authorized insiders. In European Symposium on Research in Computer Security. Springer, 231–246.
- Raytheon. 2015. SureView Insider Threat. (2015). http://www.raytheon.com/capabilities/rtnwcm/groups/cyber/documents/content/rtn_244832.pdf Accessed on August/2017.
- J. Reason. 1990. Human error. Cambridge university press.
- E. Rich, I. J. Martinez-Moyano, S. Conrad, D. M. Cappelli, and others. 2005. Simulating insider cyber-threat risks: a model-based case and a case-based model. In Int. Conference of the System Dynamics Society. 17–21.
- G. P. Richardson. 2001. System dynamics. Springer US, Boston, MA, 807–810.
- S. C. Roberts, J. T. Holodnak, T. Nguyen, S. Yuditskaya, M. Milosavljevic, and W. W. Streilein. 2016. A Model-Based Approach to Predicting the Performance of Insider Threat Detection Systems. In Security and Privacy Workshops. IEEE, 314–323.
- P. R. Sackett. 2002. The structure of counterproductive work behaviors: Dimensionality and relationships with facets of job performance. Int. Journal of Selection and Assessment 10, 1-2 (2002), 5–11.
- M. B. Salem, S. Hershkop, and S. J. Stolfo. 2008. A Survey of Insider Attack Detection Research. In Insider Attack and Cyber Security. Springer US, 69–90.
- M. B. Salem and S. J. Stolfo. 2009. Masquerade attack detection using a search-behavior modeling approach. Columbia University, Computer Science Department, Technical Report CUCS-027-09 (2009).
- M. B. Salem and S. J. Stolfo. 2011a. Decoy Document Deployment for Effective Masquerade Attack Detection. In Conference on Detection of Intrusions and Malware & Vulnerability Assessment, Springer, 35–54.
- M. B. Salem and S. J. Stolfo. 2011b. Modeling user search behavior for masquerade detection. In Int. Workshop on Recent Advances in Intrusion Detection, Springer, 181–200.
- V. Sankaranarayanan, S. Pramanik, and S. Upadhyaya. 2006. Detecting masquerading users in a document management system. In Int. Conference on Communications, Vol. 5. IEEE, 2296–2301.
- E. Santos, H. Nguyen, F. Yu, K. Kim, D. Li, J. T. Wilkinson, A. Olson, and R. Jacob. 2008. Intent-driven insider threat detection in intelligence analyses. In Int. Conference on Web Intelligence and Intelligent Agent Technology. IEEE, 345–349.
- E. Santos, H. Nguyen, F. Yu, K. J. Kim, D. Li, J. T. Wilkinson, A. Olson, J. Russell, and B. Clark. 2012. Intelligence analyses and the insider threat. IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans 42, 2 (2012), 331–347.

- A. Sanzgiri and D. Dasgupta. 2016. Classification of Insider Threat Detection Techniques. In Annual Cyber and Information Security Research Conference. ACM, 25.
- K. R. Sarkar. 2010. Assessing insider threats to information security using technical, behavioural and organisational measures. information security technical report 15, 3 (2010), 112–133.
- T. Sasaki. 2012. A Framework for Detecting Insider Threats using Psychological Triggers. JoWUA 3, 1/2 (2012), 99–119.
- M. Schonlau, W. DuMouchel, W. Ju, A. F. Karr, M. Theus, and Y. Vardi. 2001. Computer intrusion: Detecting masquerades. Statistical science (2001), 58–74.
- E. Schultz. 2002. A framework for understanding and predicting insider attacks. Computers & Security 21, 6 (2002), 526–531.
- E. Schultz and R. Shumway. 2001. Incident response: a strategic guide to handling system and network security breaches. SAMS.
- Securonix Inc. 2017. Securonix. <http://www.securonix.com/>. (May 2017). Accessed on May/2017.
- S. Sen. 2014. Using instance-weighted naive Bayes for adapting concept drift in masquerade detection. Int. Journal of Information Security 13, 6 (2014), 583–590.
- T. E. Senator, H. G. Goldberg, A. Memory, W. T. Young, B. Rees, R. Pierce, D. Huang, M. Reardon, D. A. Bader, E. Chow, and others. 2013. Detecting insider threats in a real corporate database of computer usage activity. In Int. Conference on Knowledge Discovery and Data Mining. ACM, 1393–1401.
- Sentinel One. 2017. Sentinel One. <https://sentinelone.com>. (May 2017). Accessed on May/2017.
- D. Servos and S. L. Osborn. 2017. Current Research and Open Problems in Attribute-Based Access Control. ACM Computing Surveys (CSUR) 49, 4 (2017), 65.
- A. Shabtai, Y. Elovici, and L. Rokach. 2012. A survey of data leakage detection and prevention solutions. Springer Science & Business Media.
- N. Shalev, I. Keidar, Y. Moatti, and Y. Weinsberg. 2016. WatchIT: Who Watches Your IT Guy?. In Int. Workshop on Managing Insider Security Threats. ACM, 93–96.
- J. Shavlik and M. Shavlik. 2004. Selection, combination, and evaluation of effective software sensors for detecting abnormal computer usage. In Int. conference on Knowledge Discovery and Data Mining. ACM, 276–285.
- E. Shaw, K. Ruby, and J. Post. 1998. The insider threat to information systems: The psychology of the dangerous insider. Security Awareness Bulletin 2, 98 (1998), 1–10.
- E. D. Shaw. 2006. The role of behavioral research and profiling in malicious cyber insider investigations. Digital investigation 3, 1 (2006), 20–31.
- E. D. Shaw and L. F. Fischer. 2005. Ten tales of betrayal: The threat to corporate infrastructure by information technology insiders analysis and observations. Technical Report. DTIC Document.
- F. M. Sibai and D. A. Menascé. 2011. Defeating the insider threat via autonomic network capabilities. In Int. Conference on Communication Systems and Networks. IEEE, 1–10.
- S. Sinclair and S. W. Smith. 2008. Preventative directions for insider threat mitigation via access control. In Insider Attack and Cyber Security. Springer, 165–194.
- Y. Song, M. B. Salem, S. Hershkop, and S. J. Stolfo. 2013. System level user behavior biometrics using Fisher features and Gaussian mixture models. In Security and Privacy Workshops. IEEE, 52–59.
- L. Spitzner. 2003a. Honeypots: Catching the insider threat. In Annual Computer Security Applications Conference. IEEE, 170–179.
- L. Spitzner. 2003b. Honeypots: tracking hackers. Vol. 1. Addison-Wesley Reading.
- S. Steele and C. Wargo. 2007. An introduction to insider threat management. Information Systems Security 16, 1 (2007), 23–33.
- S. J. Stolfo, M. B. Salem, and A. D. Keromytis. 2012. Fog computing: Mitigating insider data theft attacks in the cloud. In Security and Privacy Workshops. IEEE, 125–128.
- D. W. Straub and R. J. Welke. 1998. Coping with systems risk: security planning models for management decision making. MIS quarterly (1998), 441–469.
- B. K. Szymanski and Y. Zhang. 2004. Recursive data mining for masquerade detection and author identification. In Information Assurance Workshop. IEEE, 424–431.
- K. Tang, M. Zhao, and M. Zhou. 2011. Cyber insider threats situation awareness using game theory and information fusion-based user behavior predicting algorithm. Journal of Information & Computational Science 8, 3 (2011), 529–545.
- P. J. Taylor, C. J. Dando, T. C. Ormerod, L. J. Ball, M. C. Jenkins, A. Sandham, and T. Menacere. 2013. Detecting insider threats through language change. Law and human behavior 37, 4 (2013), 267.
- M. Theoharidou, S. Kokolakis, M. Karyda, and E. Kiountouzis. 2005. The insider threat to information systems and the effectiveness of ISO17799. Computers & Security 24, 6 (2005), 472–484.
- F. Toffalini, I. Homoliak, A. Harilal, A. Binder, and M. Ochoa. 2018. Detection of Masqueraders Based on Graph Partitioning of File System Access Events. In Security and Privacy Workshops. IEEE, 217–227.

- R. F. Trzeciak. 2017. SEI Cyber Minute: Insider Threats. (2017). <http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=496626> Accessed on August/2017.
- A. Vance, B. Molyneux, and P. B. Lowry. 2012. Reducing unauthorized access by insiders through user interface design: Making end users accountable. In *Hawaii Int. Conference on System Science*. IEEE, 4623–4632.
- Veriato. 2017a. Veriato 360. (2017). <http://www.veriato.com/products/veriato-360> Accessed on August/2017.
- Veriato. 2017b. Veriato Recon. (May 2017). <http://www.veriato.com/products/veriato-recon> Accessed on May/2017.
- V. Viduto, C. Maple, and W. Huang. 2010. An analytical evaluation of network security modelling techniques applied to manage threats. In *Int. Conference on Broadband, Wireless Computing, Communication and Applications*. IEEE, 117–123.
- N. Virvilis, B. Vanautgaerden, and O. Serrano. 2014. Changing the game: The art of deceiving sophisticated attackers. In *Int. Conference on Cyber Conflict*. IEEE, 87–97.
- D. S. Wall. 2013. Enemies within: Redefining the insider threat in organizational security policy. *Security journal* 26, 2 (2013), 107–124.
- K. Wang and S. J. Stolfo. 2003. One-class training for masquerade detection. In *Workshop on Data Mining for Computer Security*, Melbourne, Florida, 10–19.
- X. Wang, Y. Sun, and Y. Wang. 2014. An abnormal file access behavior detection approach based on file path diversity. In *Int. Conference on Information and Communications Technologies*. 1–5.
- R. Willison and M. Siponen. 2009. Overcoming the insider: reducing employee computer crime through Situational Crime Prevention. *Commun. ACM* 52, 9 (2009), 133–137.
- R. Willison and M. Warkentin. 2009. Motivations for employee computer crime: understanding and addressing workplace disgruntlement through the application of organisational justice. In *Int. Workshop on Information Systems Security Research*. IFIP, 127–144.
- R. Willison and M. Warkentin. 2013. Beyond Deterrence: An Expanded View of Employee Computer Abuse. *MIS quarterly* 37, 1 (2013), 1–20.
- J. F. Wolfswinkel, E. Furtmueller, and C. P. Wilderom. 2013. Using grounded theory as a method for rigorously reviewing literature. *European Journal of Information Systems* 22, 1 (2013), 45–55.
- B. Wood. 2000. An insider threat model for adversary simulation. *Research on Mitigating the Insider Threat to Information Systems*, *SRI Int.* 2 (2000), 1–3.
- G. Z. Wu, S. L. Osborn, and X. Jin. 2009. Database intrusion detection using role profiling with role hierarchy. In *Workshop on Secure Data Management*. Springer, 33–48.
- J. Wu, J. Zhou, J. Ma, S. Mei, and J. Ren. 2011. An active data leakage prevention model for insider threat. In *Int. Symposium on Intelligence Information Processing and Trusted Computing*. IEEE, 39–42.
- M. Wurzenberger, F. Skopik, R. Fiedler, and W. Kastner. 2016. Discovering Insider Threats from Log Data with High-Performance Bioinformatics Tools. In *Int. Workshop on Managing Insider Security Threats*. ACM, 109–112.
- L. Yang, Z. Hu, J. Long, and T. Guo. 2011. 5W1H-based conceptual modeling framework for domain ontology and its application on STPO. In *Int. Conference on Semantics Knowledge and Grid*. IEEE, 203–206.
- Q. Yaseen and B. Panda. 2011. Enhanced insider threat detection model that increases data availability. In *Int. Conference on Distributed Computing and Internet Technology*. Springer, 267–277.
- W. T. Young, A. Memory, H. G. Goldberg, and T. E. Senator. 2014. Detecting unknown insider threat scenarios. In *Security and Privacy Workshops*. IEEE, 277–288.
- Y. Yu and J. H. Graham. 2006. Anomaly Instruction Detection of Masqueraders and Threat Evaluation Using Fuzzy Logic. In *Int. Conference on Systems, Man and Cybernetics*, Vol. 3. IEEE, 2309–2314.
- K. H. Yung. 2004. Using self-consistent naive-bayes to detect masquerades. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 329–340.
- N. Zhang, W. Yu, X. Fu, and S. K. Das. 2010. Maintaining defender’s reputation in anomaly detection against insider attacks. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 40, 3 (2010), 597–611.
- Zoho Corporation. 2017. Manage Engine. (May 2017). <https://www.manageengine.com/> Accessed on May/2017.

APPENDIX

A INSIDER THREAT FROM DIFFERENT PERSPECTIVES

The borders of the insider threat problem are difficult to determine, as the insider threat can be part of various definitions pertaining to several research topics that have emerged in different eras of the history of computer systems. In this section, we demonstrate this issue on four definitions

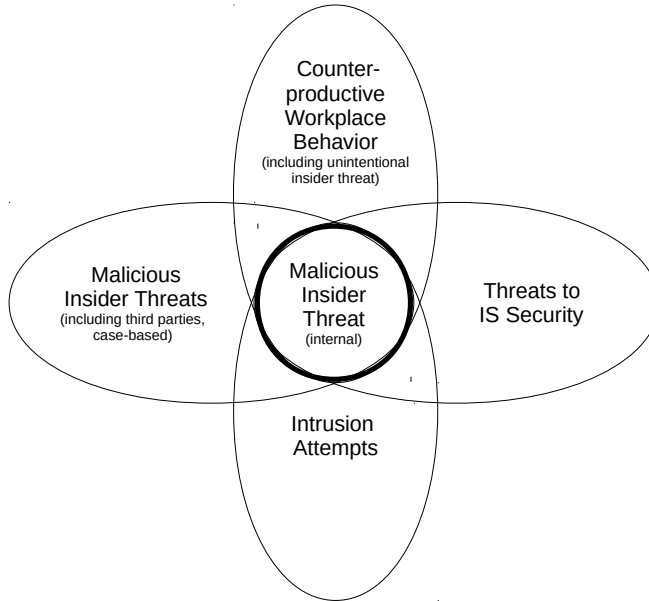


Fig. 9. Malicious insider threat from different perspectives

from four disparate areas: 1) *counterproductive workplace behavior*, 2) *intrusion attempts*, 3) *threats to IS security*, and 4) *malicious insider threat involving third parties*.

Counterproductive Workplace Behavior. Counterproductive workplace behavior (CWB) is defined as “intentional behaviors that are contrary to legitimate organizational interests” [Sackett 2002], and thus involve insider threat. Unlike insider threat, CWB includes, for example, misuse of IT resources and work time for personal business, misuse of sick leave, purposely slow work, engagement in P2P sharing, etc. Note that CWB inherently only encompasses the internal employees of an organization.

Intrusion Attempts. Anderson [1980] defined threats against computer systems (also denoted as intrusion attempts) as the possibility of an intentional unauthorized attempt to: 1) *access information*, 2) *modify information*, or 3) *render a system unavailable for other legitimate users*. In terms of information security, Anderson referred to attempts to violate *confidentiality*, *integrity*, and *availability* of information (i.e., CIA triad). According to Anderson, threats are further divided into internal and external threats, where internal threats refer to insider threats and contain masqueraders, misfeasors, and clandestine users (see Section 3.3), each of which is considered an internal employee of an organization. Thus, insider threats can be conceptualized as a subset of intrusion threats [Myers et al. 2009].

Threats to IS Security. Loch et al. [1992] proposed four-dimensional categorization of threats to information system security, also based on the CIA concept, which, in contrast to Anderson [1980], also covers unintentional insider threats, natural disasters, and mechanical/electrical failures. The dimensions of this categorization, along with their top-down order is as follows: 1) *internal and external*, 2) *human and non-human*, 3) *accidental and intentional*, 4) *disclosure, modification, destruction, and denial of use*; yielding 16 subcategories.

Malicious Insiders Involving Third Parties. Similar to Anderson’s CIA-based definition of threat against computer systems, Cappelli et al. [2012] defined malicious insider threat as “a

current or former employee, contractor or business partner who has or had authorized access to an organization's network, system or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity or availability of the organization's information or information systems." This definition is based on real case studies from the CERT database. In contrast to previous definitions, this definition also explicitly mentions external employees such as contractors or business partners.

If we consider the CIA-based definition of malicious insider threat by Cappelli et al. [2012], the CIA-based definition of intrusion attempts by Anderson [1980], the CIA-based categorization of threats to information systems from Loch et al. [1992], as well as the definition of CWB by Sackett [2002], we can see that there is an intersection of these definitions that is shared by all of them and refers to the internal malicious insider threat for information systems. The situation is depicted abstractly in Figure 9. Note that unintentional insider threat can be an aspect of CWB (e.g., not respecting policies), external intrusion attempts (e.g., thwarting by social engineering), and threats to IS security (i.e., inherently covered by the dimensions).

B THE APPLICATION OF PROPOSED STRUCTURAL TAXONOMY OF INSIDER INCIDENTS: EXAMPLES

We present two examples of the application of our proposed structural taxonomy of insider incidents utilizing 5W1H questions. The first example is aimed at a case study of a malicious insider incident and the second example contains a case study of an unintentional insider incident.

Example of Application for a Malicious Insider Incident. Consider a case study involving a malicious insider employed as a security guard (a contractor) in a hospital facility [Collins et al. 2016]. In this case, the insider was engaged in the Internet underground, and moreover was the leader of a hacking group. He worked for the targeted organization at night without any supervision. The insider's unauthorized activities involved an interaction with the heating, ventilation, and air conditioning (HVAC) computer. "The HVAC computer was placed in a locked room, but the insider used his security key to obtain physical access to the computer. The insider remotely accessed the HVAC computer five times over a two-day period, using password cracking programs to attack the organization, and installed a botnet with the intention of conducting a DDoS attack against an unknown external target. The insider's malicious activities caused the HVAC system to become unstable, which eventually led to a one-hour system outage. In addition, the insider accessed a nurses station computer, which was connected to all of the victim organization's computers, stored medical records, and patient billing information."

In this case, the answers to the 5W1H questions of the proposed categorization are as follows: **What** – The outcome of the incident is *miscellaneous* as it includes the outage of the HVAC computer and also has some aspects of *fraud*, because the insider accessed private patient information that might be sold on the dark web, and it also involves the unauthorized use of resources for the purpose of a DDoS attack. **How** – The insider obtained *unauthorized access* to the HVAC computer by using his authorized access to enter the locked room. **Why** – We consider the insider as *planted*, because it is likely that the reason the leader of a hacking group worked as a security guard was so that he could exploit the physical access provided by this line of work. The motivation of the insider is not clear from the description, but it seems to be financial, as providing the botnet may yield money, and there is money to be gained in selling stolen private records. Moreover, the motivation could also be considered as *personal* if the insider perpetrated the incident solely because of curiosity. **Who** – Despite the nature of the insider's position which is not typical for *high-end* insiders, he has some characteristics of this profile, and we can consider his job as a mission. The incident belongs to the *inside associate* category, *hacker* subtype, with an *external* job contract and a system role of

none, since the insider was not granted legitimate access to any system. **Where** – The inside attack might be detected at the *operating system*, *network*, and *physical* levels (unusual time and activities on the HVAC computer; a password cracker that attacked other machines; and entering the locked room), while the DDoS attack that followed might be detected at the *operating system* and *network* levels (suspicious processes of DDoS application; execution of the DDoS attack). **When** – The two-day duration of the malicious activities makes it *short-term*, and the incident was committed *before* the insider’s job was terminated.

Example of Application for an Unintentional Insider Incident. Consider a case study involving a social engineering attack from [CERT 2014]. In this case, due to previous incidents of social engineering attacks, employees received training in this area, however some time later some of them were deceived by a phishing email about HR benefits. The clicking on a vulnerable link caused exploitation of a zero-day vulnerability, followed by download of malicious code. “Only a few megabytes of encrypted data were stolen, but the organization failed to recognize additional dormant malicious code. The organization was forced to disconnect Internet access after administrators discovered data being externally siphoned from a server.”

In this case, the answers to the 5W1H questions of the proposed categorization are as follows: **What** – *Data leakage* was the outcome of this incident. **How** – *Installed malicious code* through phishing was the type of the attack, while the origin of the attack was an *outsider* who *duped* an insider. An unintentional violation of policy was inflicted on by using a legitimate access. **Why** – Despite the fact that the employees had previous training in order to resist phishing, they succumbed to it. This could have been the result of *insufficient training* that did not cover highly obfuscated incidents of phishing, as well as *insufficient technical controls* that enabled the attacker to deliver email with a spoofed internal sender address from the outside. **Who** – The insiders belong to the *social engineered* type, and they had an *internal* job contract with no system role mentioned in the description of the incident. **Where** – The phishing part of the attack might be detected at the *application* level (by advanced mechanisms that monitor and profile memory accesses of an application that contains a vulnerability), and the following DDoS attack might be detected at the *operating system* and *network* levels (suspicious processes of the DDoS application; execution of the DDoS attack). **When** – The incident *recurred* in the past.

C TAXONOMY OF DETECTION FEATURES

Gheyas and Abdallah [2016] proposed a taxonomy of insider threat detection features based on the MOC model: 1) *motive*, 2) *opportunity*, and 3) *capability*, which is further divided hierarchically by feature domain. Below we present some specific feature domains from these three categories, along with a few examples of features:

Motive. The motive refers to “the reason why an insider or group of insiders perpetrate a crime.” The examples include: **predisposition to malicious behavior** (e.g., trap-based decoys such as honeypots or honeytokens, information inferred from social networks such as delinquent behavior in the past, swearing or vulgar comments toward law enforcement and authorities); **mental disorders** (e.g., depression, schizophrenia, paranoia, bipolarity); **personality factors** (e.g., neuroticism, narcissism, conscientiousness); **current emotional state** (e.g., anger, fear, stress, surprise, disgust).

Opportunity. The opportunity features refer to various roles of insiders and the activities they may perform. The examples include: **roles** (e.g., system and project roles assigned to a user); **login** (e.g., logins on a user’s PC, logins on other PCs, session duration, login frequency); **file** (e.g., accessed directories, files which were copied, created, deleted, moved, modified); **database** (e.g., accessed

and modified database items, the number of tables/schemes accessed per time interval); **HTTP** (e.g., accessed URLs and domains information, encryption of websites, browser information); **removable Devices** (e.g., device name and type, file operations per time); **email** (e.g., source and destination, destinations outside of the organization, communication patterns, names and types of attachment); **mobile calls** (e.g., source and destination, date and time of calls, duration, communication patterns); **printing** (e.g., name of documents, the number of copies, time-based copy patterns); **network flows** (e.g., source and destination IP addresses, amount of data sent over the network, duration of connections, time-based communication patterns); **specific applications** (e.g., features derived from Microsoft Office Suite or other applications typical for an insider).

Capability. The capability features represent “the demonstrated skill level of an insider monitored by the IT system.” Examples include: **consumption of system resources** (e.g., consumption of CPU and RAM by the user across different sessions – the higher the RAM and CPU usage, the higher the insider’s level of sophistication may be); **application usage** (e.g., unique applications run by an insider in different sessions – the higher the number of unique applications, the higher the insider’s sophistication may be; multiple applications simultaneously run by the insider across different sessions – the higher is the number of applications simultaneously run per session, the more sophisticated the insider may be).

D BEST PRACTICES AND GUIDELINES

Best practices and guidelines represent procedural defense countermeasures against insider threat which may, however, involve recommendations about the usage of technical controls. In this section we summarize different works in this area developed by academia, industry, and government.

General. The major contribution to general best practices was made by CERT and the US Secret Service, both of which released several versions of *common sense guides* (e.g., [Collins et al. 2016]) and in addition addressed specific areas such as government [Kowalski et al. 2008], the financial sector [Cummings et al. 2012], the critical infrastructure sector [Keeney et al. 2005], and the cloud environment [Claycomb and Nicoll 2012]. Mitigation of malicious insiders from the cloud computing perspective is also described in [Kandias et al. 2011]. Technologies that can be used to control insider threat are addressed by Cole and Ring in chapter 9 of their book [2005], while other countermeasures aimed at survivability are present in chapter 10. In addition to technical controls, a wide spectrum of procedural best practices and countermeasures are discussed in [Hayden 1999], [Anderson 1999], and [Sarkar 2010]. Based on experiences with different organizations, Guido and Brooks [2013] described the components needed for an insider threat mitigation and auditing program and discussed several best practices. A wide range of factors influencing unintentional insider threat, along with best practices for its mitigation were discussed by Greitzer et al. [2014]. Similarly aimed at unintentional data breaches, Liginlal et al. [2009] proposed three strategies for managing human error: *avoidance*, *interception*, and *correction of error*, which were projected onto various data representations. McCormick [2008] proposed a number of ways for improving both administrative and technical controls for data leakage in a typical enterprise, while highlighting enterprise DLP programs. Data leakage was also addressed by Miller and Maxim [2015] who highlighted identity and access management along with DLP.

Management and High Level View. This subsection includes papers dealing with insider threat from the managerial perspective, often presenting a high level view of addressing this problem. Steele and Wargo [2007] provided guidelines for the prevention and mitigation of insider threats. Their guidelines include cyclic stages of assessment, prioritization & review, and remediation; all of these are governed by administrative methods (policies, procedures, human resources, awareness

training, and education). Blackwell [2009] discussed aspects of a security model for the investigation and evaluation of organizational security in three layers: the *social/organizational layer*, *logical layer* (i.e., intangible entities), and *physical layer* (i.e., tangible entities). In a similar vein, Viduto et al. [2010] proposed a generic onion skin model for hardening the defense against malicious insiders, which consists of four layers (physical, technical, logical, and assets), each of which contains specific countermeasures. Gritzalis et al. [2014] proposed mitigating strategies at the business process level by: 1) designing secure business processes, 2) performing risk assessment, and 3) monitoring each business process, while inferring conclusions. Coles-Kemp and Theoharidou [2010] examined information security management practices with regard to the insider threat and further elaborated on crime theories that provide methods for information security management design. Colwill [2009] examined many human factors (including technical, social, business, economic, and cultural factors) that can be used for assessing insider threat.

Monitoring and Analysis of Incidents. Murphy et al. [2012] proposed best practices and recommendations for the design of detection methods, as well as for the analysts of insider threat incidents working with them. In a similar vein, Gheyas and Abdallah [2016] identified several best practices related to the design of detection methods. Inspired by the overwhelming amount of monitoring alerts, Colombe and Stephens [2004] proposed unsupervised anomaly-based visualization of RealSecure alerts, which enables the analyst to distinguish between non-automated and automated insider activity and pay attention to more interesting alerts. Doss and Tejay [2009] developed a model describing how analysts can use existing security tools, such as anti-virus, intrusion detection systems, and log analysis tools, to detect insider attacks.

Standards. Insider threat has also been addressed by several standards. For instance, ISO 17799 provides a set of recommendations for information security management, and more specifically, the *personnel security* category mentions controls aimed at the protection of an IS from accidental and deliberate insider threats [Theoharidou et al. 2005]. Other related standards are ISO 27001 and ISO 27002, in which *control domains* and *control areas* address this topic [Coles-Kemp and Theoharidou 2010; Humphreys 2008]. According to Coles-Kemp and Theoharidou [2010], three distinct categories of controls can be identified in ISO 27002: “controls identifying insiders from outsiders, controls used to identify unexpected insider behavior, and controls used to influence the development of an organization’s security culture.”

E DECOY-BASED SOLUTIONS

According to Spitzner [2003b], “a honeypot is a security resource whose value lies in being probed, attacked, or compromised.” Therefore, a honeypot has no production value, and any activity performed with it is suspect by nature. Although the primary purpose of this concept is the detection of external threats, the honeypot can also be utilized in the area of insider threat detection. A honeypot is an example of a decoy-based concept that provides fake information delivered by an internal legitimate system (e.g., record in database, files in a repository).

One of the first works in this area is [Spitzner 2003a], where the author discussed a few options for the detection of spies by placing false emails in mailboxes, planting honeypots within search engine results, and “*enriching*” network traffic with spurious data. Honeypots as websites decoying insiders were used by Maybury et al. [2005]. Decoys for malicious insiders were also addressed by Bowen et al. [2009] who proposed using trap-based documents with bogus credentials, as well as stealthily embedded beacons that signal an alert when the document is opened. Detection of masqueraders perpetrating data exfiltration in the Windows XP environment was addressed by Salem and Stolfo [2011a] who designed a decoy document access sensor that compares a keyed-hash

message authentication code (HMAC) embedded in each decoy document with HMAC computed over a document loaded into memory. Virvilis et al. [2014] proposed several deceiving techniques for insiders and APTs, which were divided into two categories: 1) *attack preparation* (e.g., fake entries in robots.txt, passwords in HTML comments, invisible HTML links, fake DNS records, and fake social network avatars), and 2) *exploitation & exfiltration* (e.g., honeytokens in databases, decoy files, and accounts). As part of their reconnaissance deception system, Achleitner et al. [2016] utilize honeypots to detect insider adversaries such as APTs performing reconnaissance.

F OTHER PRACTICAL SOLUTIONS

This section provides examples of other practical tools that can be applicable for the detection or forensic investigation of insider threat incidents. In the majority of the cases, the authors of these tools do not provide detailed implementation descriptions, as these tools are usually commercial products. The examples included here are grouped into three categories: *detection*, *SIEM*, and *audit tools*. The detection category represents products that use several data sources to raise alarms for suspicious or anomalous activities. Security information and event management (SIEM) systems serve as log aggregation engines that provide consolidated views of the disparate events and alerts they analyze. Audit tools can serve forensic and audit purposes, which may potentially lead to the detection of suspicious activities attributed to insider threat. Note that this section does not contain references from our input literature database used for categorization, but rather it provides an overview of practical solutions that can be used for defense against insider threat.

Detection. The examples of tools that we include in this category can be divided into two groups: *misuse-based* tools, and *machine learning-based* tools. Misuse-based tools are based on rule matching and include Raytheon's SureView [2015], the Ekran System [2017], and SentinelOne [2017]. ObserveIT [2017] utilizes a different solution that involves a record of screen activity that may help in providing evidence in the case of an insider threat incident. In addition, there have been efforts made and incentives have been used to deploy machine learning techniques in commercial environments. Darktrace [2017] is based on advanced unsupervised machine learning techniques that utilize Bayesian probabilistic approaches to assess users' behavior, an entity, and software. Veriato 360 [2017a] and Veriato Recon [2017b] (from Veriato, formerly known as SpectorSoft) combine various machine learning techniques in order to detect anomalies in users' behavior. Moreover, this tool integrates psychological profiling of users, which is obtained through the users' language. Securonix [2017] provides a solution based on machine learning techniques for the detection of anomalous behavior in a group of users, referred to as a control group. This tool is particularly aimed at the detection of data exfiltration.

SIEM. Security information and event management (SIEM) [Miller et al. 2010] can provide statistical summaries of events over time, which can be aggregated and reported using various dimensions. We describe some examples of commercial and open-source tools, starting with a commercial tool, called ArcSight Express [Hawlett Packard 2017] that provides a user tracking system based on events that can be correlated in real-time, and in the case of suspicious activity, an alert can be raised. Carbon Black's endpoint security platform [2017] is capable of performing a full spectrum analysis that includes behavioral, reputation, machine learning, and SIEM analysis that could also be used for behavioral alerts regarding insider threat. Kibana [2017] is an open-source SIEM based on the Elastic search engine that has several features, such as geographical activity representation and time series of alarms. In addition, thanks to the nature of this tool and the availability of its source code, it is easy to build new custom components.

Audit Tools. LUARM, an audit engine for insider IT misuse detection, was proposed by Magklaras et al. [2011], and it logs actions associated with file system access and process execution, as well as actions at the network endpoint level. In the same vein, other commercial audit solutions propose similar features, such as MKinsight [2017] and ADAudit Plus [Zoho Corporation 2017], both of which are fully compatible with Windows enterprise deployments.

G DETECTION AND THREAT ASSESSMENT APPROACHES: ADDITIONAL WORKS

In addition to the detection and threat assessment approaches discussed in Section 8.2, in this section we briefly describe the remaining papers of our input literature database which fit this subcategory, maintaining the categorization structure used in the main text of this survey.

G.1 Conceptual Works

Anomaly-Based. Dealing with **role-based monitoring**, Phyo et al. [2004] focused on insiders that misuse their privileges and proposed a detection concept for separation of duties violations; this concept involves matching input to a database of role-assigned actions. Incorporating attribute-based access control (ABAC), Bishop et al. [2009b] proposed an insider threat assessment concept for prioritizing users according to their access to resources, while considering **psychological indicators** gleaned from logs, HR records, and language affectation used in emails, IMs, and blogs. Greitzer and Hohimer [2011] proposed CHAMPION, a model-based belief propagation conceptual framework for reasoning about insider threat, which utilizes auto-associative memories and integrates psychological and cyber observables. In later work, Greitzer et al. [2012] proposed a risk assessment concept based on 12 behavioral and psychosocial indicators that serve as input to a Bayesian network, linear regression, and an artificial neural network (ANN). Sasaki [2012] proposed a concept that creates a stimulating event that is used to prompt malicious insiders to behave anomalously, and consequently monitors the users' reactions. Berk et al. [2012] proposed BANDIT, a system that assesses the three components of the MOC model with respect to the user's baseline as well as to group behavior, producing two scores by calculating the Euclidean distance. In terms of **assessing the environment**, Chinchani et al. [2010] presented a modeling approach for insider threat assessment based on their previously proposed key challenge graph (KCG), in order to address the problem of finding an attack with minimal cost. The KCG concept was also applied by Ha et al. [2007] in their insider threat assessment graphical tool that displays possible attack trails. Naghmouchi et al. [2016] designed risk assessment graphs capturing the topological information of the network, including the assets and vulnerabilities associated with each asset, as well as the way these elements vary over time.

G.2 Operational Works

Misuse-Based. Roberts et al. [2016] designed a BN aimed at data exfiltration, consisting of four binary detectors: the presence of connected devices, after-hours work, visiting file sharing websites, and information about employees' dismissal. Graph-based misuse detection of malicious user interactions within *Web applications* was proposed by Jaballah and Kheir [2016] who designed a gray box approach based on the subgraph isomorphic matching.

Anomaly-Based. The bulk of the operational studies dealing with anomaly detection of insiders were conducted on **Unix/Linux command histories** and aimed at the masquerader attacker's model. Coull et al. [2003] proposed a masquerader detection approach that uses pair-wise sequence

alignment to represent similarity between any two aligned sequences of commands in SEA. Wang and Stolfo [2003] utilized one-class Naïve Bayes and one-class SVM for masquerader detection in SEA, comparing multivariate Bernoulli (binary) and multinomial (bag-of-words) approaches for feature representation. Yu and Graham [2006] employed a concept of finite state machine and a fuzzy inference system for modeling users; evaluation was performed on the PU dataset, as well as their custom dataset collected from 44 users. Szymanski and Zhang [2004] used one-class SVM and recursive mining that recursively searches for frequent patterns in a string of commands in SEA, and then they encoded such patterns with unique symbols and rewrote the string using this new coding. Oka et al. [2004] employed Eigen co-occurrence matrices and layered networks for the detection of masqueraders in SEA. Latendresse [2005] employed a customized version of the Sequitur algorithm to generate context-free grammar from sequences of commands executed by a user; the author evaluated the approach on SEA. Similarly, dealing with the compression of commands in SEA, Posadas et al. [2006] proposed a local hybrid method that combines HMM and session folding accomplished by the extraction of the user's context-free grammar. In order to deal with concept drift in the masquerader detection problem, Sen [2014] proposed applying instance weighting as an updating scheme for the Naïve Bayes model; she evaluated the approach on SEA. Considering **file system logs** as a data source, Camiña et al. [2014] proposed detection systems for masqueraders utilizing Markov chains and Naïve Bayes as one-class techniques that were evaluated on the WUIL dataset. Gates et al. [2014] proposed detection of information theft by comparing input with a history of the user's file access (and that of his/her peers) by similarity functions working over hierarchical file system structure; evaluation was performed on logs from a source code repository containing injected attackers. Aimed at per process file path diversity assessment, Wang et al. [2014] proposed a system for detection of masqueraders searching for sensitive files in Windows machines. Dealing with **system call sequences**, Nguyen et al. [2003] proposed a system for the detection of insiders in Unix systems based on the profiling of: 1) file system access executed by system users, and 2) human and system processes having either a fixed number of children or accessed files. The evaluation was performed on detection of insiders exploiting local buffer overflow vulnerabilities. The detection of insiders in system calls was also addressed by Parveen et al. [2011b] who employed an ensemble of one-class SVMs as a stream-based approach that coped with the concept drift problem. In contrast to the previous papers, we now present several examples dealing with **general cyber observables**. Shavlik and Shavlik [2004] proposed a masquerader detection method that creates statistical profiles of Windows 2000 users and intakes over 200 features measured in seven historical variants, weighted using a custom algorithm called Winnow. Raissi-Dehkordi and Carr [2011] utilized one-class SVM to analyze features derived from file/database sever access and network monitoring in order to address information theft involving colluding insiders. User identification was the focus of Song et al. [2013] who compared a Gaussian mixture model, Parzen method, and one-class SVM (in a one-vs-all classification setting); these were further optimized by the Fisher criterion for feature selection and evaluated on benign data of the RUU dataset. Park and Giordano [2006] proposed a combination of role-based and individual-based anomaly detection of information theft in the intelligence community. Considering **psychosocial observables**, Greitzer et al. [2013] implemented a tool based on their previous work on personality traits [2012], which performs word use analysis by LIWC with the Mahalanobis distance for the identification of outliers; they evaluated the approach on an email corpus containing injected outliers, and again later [Brown et al. 2013] on the Enron dataset. Considering MOC indicators at an abstract level, Axelrad et al. [2013] developed a list of 83 indicators potentially related to insider threat (e.g., psychosocial and CWB indicators), which were further ranked and combined into a single score using a BN. Leveraging **NLP** techniques, Santos et al. [2008] presented intent-driven

detection of malicious insiders spreading disinformation in intelligence reports by utilizing a BN that captures interests of analyst, context of knowledge, and changes of preferences over time. Later, these authors extended their approach by a normalization procedure that divided an analyst's discrepancy value by his/her global correlation value [Santos et al. 2012].

Classification-Based. Wu et al. [2009] proposed the detection of anomalous **database** transactions by two Naïve Bayes classifiers modeling role profiles and user profiles, respectively, while utilizing six input features: user ID, role ID, time, IP address, access type, and SQL statement. Kim and Cha [2005] applied SVM with radial basis function for the detection of masqueraders in blocks of **Unix commands**, while they applied a simple voting scheme to assess a super-block of these commands. Detection of masquerades in Unix commands was also addressed by Maxion [2003], who applied a Naïve Bayes classifier with an updating scheme, which considered frequency of particular commands and their parameters for each user. Yung [2004] extended a Naïve Bayes classifier by self-consistency that kept both of the classes updated consistently with new data by the expectation-maximization algorithm. Killourhy and Maxion [2008] proposed the detection of super-masqueraders – those issuing never-before-seen-commands (NBSC) – in Unix command sequences using an enhanced Naïve Bayes classifier with a simple NBSC detector. An approach to modeling user interaction with a **GUI-based application** was proposed by El Masri et al. [2014] who focused on masquerader detection and used ensemble-based classifiers, such as Random Forest and Ada-Boost with Random Forest, and evaluated the results on Microsoft Word commands of the MITRE OWL dataset. Considering **sentiment analysis** and **NLP**, Taylor et al. [2013] proposed detection of insiders perpetrating data exfiltration by binary logistic regression using several LIWC categories, together with linguistic style matching in email messages.

Unsupervised Detection of Outliers. In the general **cyber observables** vein, Young et al. [2014] evaluated previously designed methods [Senator et al. 2013] as an ensemble against individual methods and scenario-based outlier detectors. Investigating various activity domains, Eldardiry et al. [2013] proposed detection of anomalies representing insider threat by *k*-means clustering and Markov chain algorithms. Dealing with the concept drift problem, Parveen et al. [2011a] combined a stream mining approach with an ensemble of GBADs for the detection of insiders in **system call** activities. Wurzenberger and Kastner [2016] proposed a scalable approach for the detection of data exfiltration in **databases** by applying high performance bioinformatics clustering tools. Considering aspects of **mutual interactions of users**, data exfiltration in collaborative information systems (CIS) was addressed by Chen et al. [2012] who utilized nearest neighbor networks applied onto access logs of a healthcare CIS, observing that normal users tend to form communities unlike illicit insiders.

H ADDITIONAL CATEGORIZATIONS OF OPERATIONAL APPROACHES OF THE DETECTION AND THREAT ASSESSMENT SUBCATEGORY

In addition to the main categorization that is based on intrusion detection and machine learning, we propose the use of two other categorizations that can also be applied on operational works dealing with the detection and assessment of insider threat: 1) categorization based on **the dataset setting used for evaluation**, and 2) categorization based on **the feature domains**. The former distinguishes among six types of selected dataset settings used in experiments, which include detection of *masquerader attacks*, *traitor attacks*, *miscellaneous malicious attacks*, *substituted masquerader attacks*, *unintentional insider incidents*, and *identification of users*. The majority of these settings are in accordance with the dataset categorization introduced in Section 5.2, however we

emphasize that various types of datasets can be utilized in different types of settings; for example, the masquerader-based datasets can be utilized in the user identification task (e.g., [Song et al. 2013]), and the identification/authentication-based datasets can be utilized in the detection of substituted masqueraders (e.g., [El Masri et al. 2014; Maxion 2003; Yu and Graham 2006]). The latter type of categorization considers a taxonomy of detection feature domains [Gheyas and Abdallah 2016] based on the MOC model. For details about this taxonomy, we refer the reader to Appendix C. The three types of categorization (the main one and two additional ones) are applied to the operational works in Table 1 and Table 2. In comparison to the original categorization of feature domains [Gheyas and Abdallah 2016], several new feature domains were added to the tables in order to reflect all of the operational works of our literature database. Note that some of the approaches listed in the tables contain a particular source of data in the description, however the feature domain referring to that data source is not marked (or vice versa). This is because some approaches apply non-data-inherent feature domains for data analysis. For example, [Brown et al. 2013; Greitzer et al. 2013] use emails as a data source, however their approach works only with the text information of those emails and considers motive-based feature domains (i.e., personality factors, current emotional states); this is in contrast with [Legg et al. 2015], whose approach also analyzes recipients and senders of emails.

Approach		Data		Motive		Feature Domains										Capability																		
Reference	Category	Description	Dataset Reference	Setting of Evaluation		Predisposition to Malicious Behavior	Mental Disorders	Personality Factors	Current Emotional State	Logon	File	Database	HTTP	Removable Devices	Email	Mobile Calls	Printing	Roles	Instant Messaging	Processes	Social Networks	HID	Business Processes	Physical World	System Calls	Shell commands	Security Logs	Network Flows	Specific Applications	Consumption of System Resources	Applications Usage			
	Anomaly (A), Misuse (M), Classification (C), Outliers (O)			Masquerades (M), Trojans (T), Denial of Service (DoS), Identification of Users (IDEN), Unintentional Misuse (UMI)																														
[Lane and Brody 1998]	O	Unix commands	PU	SM																														
[Schonlau et al. 2001]	A	Unix commands	SEA	SM																														
[Maxion and Townsend 2002]	C	Unix commands	SEA	SM																														
[Coudill et al. 2003]	A	Unix commands	SEA	SM																														
[Wang and Stolfo 2003]	A	Unix commands	SEA	SM																														
[Maxion 2003]	C	Unix commands	Greenberg	SM																														
[Chia et al. 2004]	A	Unix commands	SEA	SM																														
[Yung 2004]	C	Unix commands	SEA	SM																														
[Latenfresse 2005]	C	Unix commands	SEA	SM																														
[Kin and Chra 2005]	C	Unix commands	SEA, Greenberg	SM																														
[Szymanski and Zhang 2006]	A	Unix commands	SEA	SM																														
[Posadas et al. 2006]	A	Unix commands	SEA	SM																														
[Yu and Graham 2006]	A	Unix commands	PU, Custom	SM, M																														
[Kilbury and Maxion 2008]	C	Unix commands	Greenberg	SM																														
[Sen 2014]	A	Unix commands	SEA	SM																														
[Gama et al. 2014a]	A	File access logs	WUL	M																														
[Gama et al. 2014b]	A	File access logs	WUL	M																														
[Tofalini et al. 2018]	A	File access logs	WUL, TMOS	M																														
[Wang et al. 2014]	A	File access logs	Custom	M																														
[Gates et al. 2014]	A	File access logs	Custom	T																														
[Gama et al. 2016]	C	File access logs	WUL	M																														
[Nguyen et al. 2003]	A	System calls	Custom	T																														
[Lu et al. 2005]	A	System calls	Custom	T																														
[Parveen et al. 2011a]	A	System calls	DARPA, 1998	M																														
[Parveen et al. 2011b]	A	System calls	DARPA, 1998	M																														
[Wu et al. 2009]	C	Database transactions	Custom	IDEN																														
[Mathew et al. 2010]	C, O	Database transactions	Custom	SM																														
[Prangrati et al. 2013]	A	Database transactions	Custom	SM																														
[Wurzenberger and Kasner 2016]	A	Database transactions	Custom	M																														
[Santharajayyan et al. 2006]	O	MS Word commands	Custom	MM																														
[El Masri et al. 2014]	A	MS Word commands	MITRE OWL	SM																														
[Malhot and Stephens 2007]	H	Network traces	Custom	T																														
[Siba and Menase 2011]	M	Network traces	Custom	T																														
[Raisi-Dehkordi and Carr 2011]	A	Network traces	Custom	T																														
[Iokica et al. 2008]	A	Emails	Enron	T																														
[Geitner et al. 2013]	O	Emails	Custom	T																														
[Trayor et al. 2013]	A	Emails	Custom	T																														
[Brown et al. 2013]	C	Emails	Enron	T																														
[Gang et al. 2006]	A	Mouse actions	Custom	SM																														
[Kandas et al. 2013a]	C	Youtube data	Custom	T																														
[Kandas et al. 2013b]	C	Twitter data	Custom	T																														
[Kandas et al. 2013a]	O	Virtual directory logs	Custom	T																														
[Chaycomb and Shim 2010]	M	Access logs of IS	Custom	T																														
[Chen et al. 2012]	O	Access logs of IS	Custom	M																														
[Jaballah and Kheir 2016]	M	HTTP requests in syslog format	Custom	M																														

Table 1. Categorization of the operational approaches (part 1/2)

I ALL RESEARCH CONTRIBUTIONS IN INSIDER THREAT FIELD

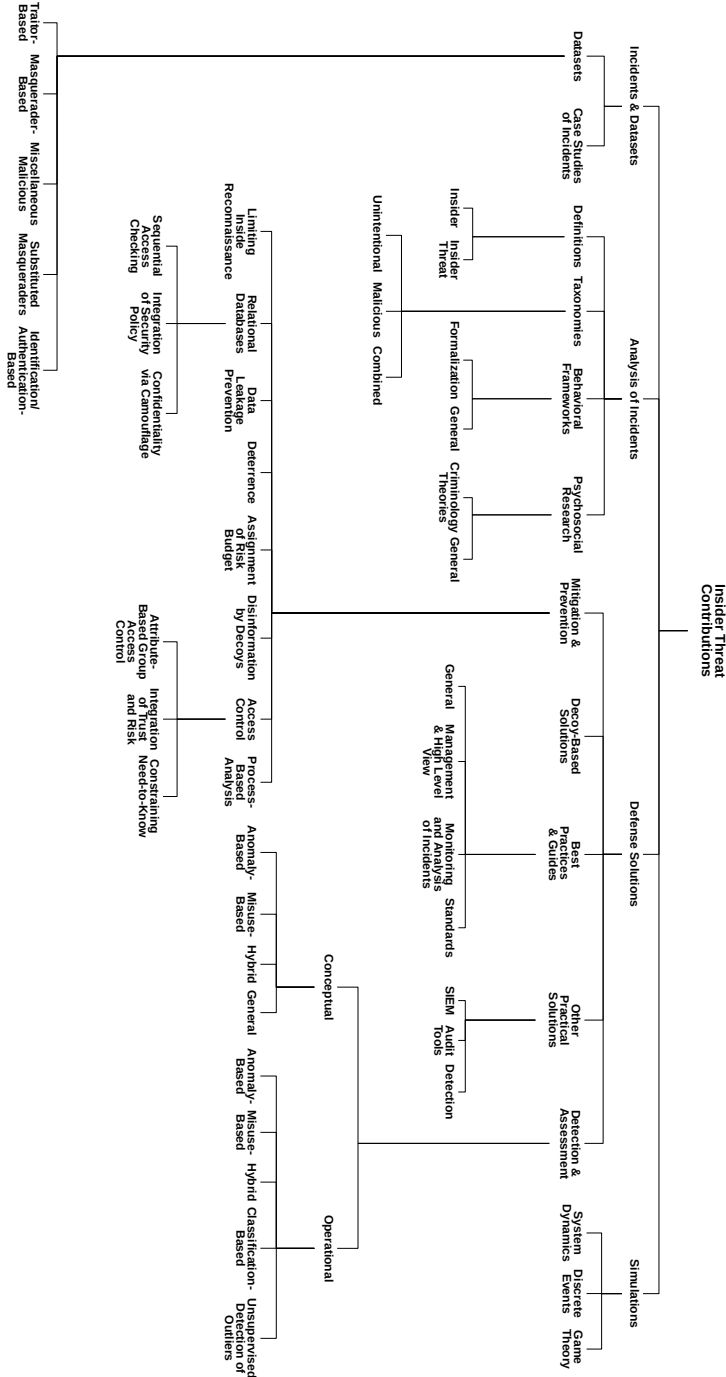


Fig. 10. Categorization of research contributions in insider threat field

Received September 2017; revised XXXX XXXX; accepted XXXX XXXX