Adaptive Detection and Policy Transformation for Insider Threats

Nicholas B. Harrell, Alexander Master, and J. Eric Dietz
Center for Education and Research in Information Assurance and Security
Purdue Homeland Security Institute
Purdue University
nharrel@purdue.edu, amaster@purdue.edu, jedietz@purdue.edu

Abstract Insider threats are among the most costly and prevalent cybersecurity incidents. Modern organizations lack an effective way to detect and deter insider threat events; traditional mitigation approaches that focus on recruitment processes and workplace behavior have proven insufficient. Current analytic detection tools do not map technical indicators to organizational policies. This limitation results in poor risk calculations, rendering inaccurate risk mitigation decisions regarding insider threats. This paper proposes a pragmatic, data-driven approach that uses policy-mapped technical indicators to assess insider threat risk. Our approach provides a quantitative insider threat risk score to facilitate informed decision-making by policymakers. Using computer simulation modeling and synthetic data to iterate common threat scenarios, we increase the probability of detecting an insider threat event. This novel approach provides quantitative analysis with distinct advantages over qualitative risk matrices commonly used in industry to forecast and assess organizational risk.

INTRODUCTION

According to Verizon's most recent data breach investigation report, internal employees caused 19 percent of data breaches (Bassett et al., 2023). The report also noted that external actors likely took advantage of internal errors—more than 40 percent of breaches involved stolen credentials from a legitimate user. Verizon ranked ransomware and phishing among the top four actions contributing to breaches (Bassett et al., 2023). Note that most of these attacks required some action by an internal employee. The rise of remote work, accelerated by the global COVID-19 pandemic (Manokha, 2020), has further complicated detection of insider threats; data from the IT industry in the same year indicated a concerning trend of increased insider disruptions. Insider threats are defined as individuals, either current or former employees, who possess particular access to an organization's internal resources. Their actions, whether unintentional or intentional, cause harm or increase the risk of harm within the organization (Collins, 2016). Harm can be monetary loss from service downtime, loss of intellectual property, liability for disclosure of personally identifiable information, or reputational damage.

To study information environments in organizations, various researchers have emphasized the significance of the recruitment process, workplace behavior, and interpersonal interactions among colleagues to determine how psychometric traits can predict insider behavior. Most organizations employ internal mechanisms to flag behavioral indicators via background checks, which include past employment records, credit reports, credential verification, criminal convictions, and insights from previous coworkers and supervisors who may reveal more intrinsic details about the candidate (Collins, 2016). Based on current industry statistics, we can see that despite these approaches, insider threats are still prevalent and on the rise (Bassett et al., 2023).

This paper proposes a pragmatic approach based on the premise that insider threats are inevitable due to human error. Rather than attempting to discern indicators from human behavior, we focus on the relationship between an organization's policy and its monitoring and auditing capabilities. Our primary contribution is to offer a tool and an approach that utilizes quantitatively measured technical indicators to provide policymakers with an insider threat risk score. Rather than relying solely on subject matter expert opinion, our model provides policymakers with a composite score that supports informed decision-making. We adopt a process known as dynamic adaptive management or dynamic adaptive policy pathways (DAPP), which has been employed in various industries (Haasnoot et al., 2013). We demonstrate how looking at insider threats as an event of uncertainty can assist decision-makers in making better risk management decisions regarding insider threat activity. We call this approach adaptive detection and policy transformation for insider threats (ADAPT-IT).

PRIOR WORK

Analytic and Monitoring Capabilities

Legg et al. (2015) demonstrated the mapping of technical indicators to user behavior in insider threat detection using log data, such as login attempts, removable media,

email, web, and file logs. This approach allowed for monitoring activities and incidents indicative of insider behavior, supporting the identification of insider threat activity through profiling.

Analytical strategies in insider threat detection have encompassed anomaly-based and heuristic-based approaches (Yamin et al., 2020; Collins, 2016; Eldardiry et al., 2013; Caputo et al., 2009). However, the mapping of features to policy violations remains underdeveloped.

Intrusion detection systems are commonly used to detect network threats. They can be either signature-based, relying on known attacks, or anomaly-based, relying on deviations from normal behavior. Anomaly-based systems require more time for setup and rely on establishing a baseline of normal behavior, making them susceptible to mimicry attacks.

Detection of insider threats depends on capturing logs documenting specific activities and associated features (Legg et al., 2015). Each feature contributes to the user's profile, providing insights into their behavior. Organizations can more effectively detect insider threats by classifying activities into risk categories based on user roles. They utilize various methods, including user activity monitoring, data loss prevention, security information and event management, analytics, and digital forensics (Spooner et al., 2018). Organizations can detect deviations from normal behavior and identify potential insider threats by leveraging technical indicators such as file transfers, database queries, and login activities. Monitoring specific actions (e.g., file transfers and logins) can effectively narrow the focus and increase the likelihood of insider threat detection.

Effective Decision-Making

Insider threats involve uncertainty, posing challenges for decision-makers who struggle with intangibles (Hubbard and Seiersen, 2023). Researchers have explored models addressing security policy compliance and noncompliance. The cause of uncertainty lies in the behaviors of employees regarding adherence to security policies (Warkentin and Willison, 2009).

Traditionally, policymakers in many industries assumed they could predict the future; they created a static "optimal" plan based on a single "most likely" future (Haasnoot et al., 2013). This strategy solved the short-term problems; however, when different results happen in the assumed future, a new "optimal" plan must be created. Collingridge (1980) suggested that when there is limited knowledge regarding the potential side effects of emerging technologies, it is crucial to prioritize decision correctability, thorough monitoring of effects, and adaptability. In the context of uncertainty, Rosenhead (1990) and Rosenhead et al. (1972) proposed that evaluating the robustness of strategies can be done by considering the degree of flexibility, specifically by measuring the number of available options.

Adaptive policymaking, proposed by Haasnoot et al. (2013), presented a structured approach for designing dynamic and robust plans. It emphasized the importance of decision correctability, extensive monitoring, and flexibility in the face of limited knowledge about the potential side effects of emerging technologies. The approach consisted of five steps: (1) analyze existing system conditions and set objectives, (2) formulate a basic plan, (3) enhance plan resilience through mitigating, hedging, seizing, and shaping actions, (4) continuously monitor plan performance, and (5) implement triggered actions based on signpost information. Adaptive policymaking enabled data-driven decision-making and reduced uncertainty by establishing a bidirectional relationship between policy and monitored technical indicators.

The adaptation pathways approach, summarized by Haasnoot et al. (2011, 2012), offered a different perspective on planning for adaptation. This concept considered adaptation tipping points; these signify the conditions under which an action no longer aligns with specified objectives. After reaching a tipping point, the approach presented a sequence of possible actions through adaptation trees, similar to decision trees or road maps. It utilized computational scenario approaches to assess the timing of tipping points across different scenarios. The adaptation pathways map provided an overview of alternative routes to achieve desired future outcomes, considering different actions and their potential performance. By incorporating stakeholder perspectives, cultural mapping, and cost-benefit analyses, decision-makers could make informed choices about the pathways to follow. The adaptation pathways approach provided a framework to adapt to changing conditions and support decision-making in uncertain and dynamic environments.

Integrating adaptive policymaking and adaptation pathways into DAPP aligns with Hubbard's strategy of measuring data points based on observable events. As Hubbard emphasized, quantitative measurement reduces uncertainty by focusing on observable technical indicators rather than subjective human behavior. The integrated approach of DAPP incorporated a monitoring system that tracks signpost information, which represents observable events and triggers related to the plan's success (Hassnoot et al., 2013). This data-driven approach enabled decision-making based on real-time information and facilitated continuous adjustment of actions and strategies to ensure alignment with preferred pathways (Haasnoot et al., 2011, 2012).

Objectives and Indicators

Haasnoot et al. (2013) posited that understanding indicators within asset pathways is crucial. The authors explored adversarial modeling, identifying potential actions that deviate from standard pathways and pose threats to organizations. Looking at their approach from a cyber perspective, we can see pathways of data extraction within an organization's cyber infrastructure. For instance, an email could lead to the dissemination of sensitive information to unauthorized entities.

Haasnoot et al. (2013) also offered that it is vital to monitor system actions to collect signpost information related to triggers. Monitoring enhances quantitative assessment of risk posture, facilitating the mapping of actions to policies and their categorization based on selected methods. Our contribution suggests incorporating a prioritization mechanism using weights. Haasnoot et al. (2013) described an adaptive system that measures the actions' effectiveness after each iteration. The system allowed policymakers to adjust weights based on indicators' effectiveness in detecting threat-like behavior, enabling quick adaptation of risk mitigation techniques. Similar challenges exist in ecology and other fields, where adaptive strategies require streamlined decision-making to address knowledge gaps (Scarlett, 2013).

Weights and Composite Indicators

Composite indicators are used in many fields of study regarding human development sustainability, perceived corruption competitiveness, or other complex phenomena (Becker et al., 2017). Studies that perform uncertainty and sensitivity analysis on composite indicator assumptions rely on subjective choices (Saisana et al., 2005). It is important to realize the bias in these assumptions and choices made, which will deviate from the importance factor placed on the overall aggregated score. Optimization is generally effective in defining the impact and importance of weights on a composite indicator (Becker et al., 2017). Due to corollary relationships between features, weights can often have negative scores (Becker et al., 2017). Further investigation using correlation analysis is required to ensure that the weights do not contribute to the same variation in the outcome, effectively canceling each other out.

Policy and Training

Kweon et al. (2021) conducted a study on the impact of security training on organizations. Many factors were considered, including managerial knowledge, employee knowledge, security policies, time spent on training, firm size, and budget for training. An interesting finding in this study was that security policy programs were an indicator of having many security incidents. Kweon et al. concluded that organizations with many security policies most likely have the policies in place due to numerous security incidents. The author found that the more the population is aware of security concerns, the fewer incidents occur. However, many security policies in an organization are due to previous security incidents. The author inferred that the security policy is usually a consequence of a prior incident (Kweon et al., 2021).

AnyLogic Modeling

AnyLogic is simulation software that provides mechanisms to simulate real-life scenarios. It allows policymakers to make informed decisions without having to allocate immense resources. This study will use the optimizer in AnyLogic to calibrate the weights we use to determine the appropriate importance for each weight. Several research groups have used simulations to inform risk management decisions (Master et al., 2022; Tzvetanov et al., 2022; Lerums et al., 2018).

THE ADAPT-IT MODEL

Methodology

This paper offers an iterative approach to reduce policy infringements by increasing detection rates of malicious actors. Through exploratory data analysis, our model calculates a composite score for each user and ranks them based on deviations from their exponential moving averages (EMA) over short and long periods. We evaluate model success by measuring all malicious actors that rank within the top 40 of composite scores.

Validity

We used an open-source dataset from the United States Computer Emergency Readiness Team (US-CERT) to test and validate the model's effectiveness in identifying insider threat behavior.[1] Using publicly available data promotes transparency and ease of reproducibility of our work. We used the version 5.2 release from the US-CERT data repository in this study to offer an assortment of scenarios encapsulating various facets of network interaction. These included email reception events, the structure of directories on removable media, properties of email attachments such as size, user login attempts, and web content.

The data derived from these scenarios were systematically cataloged across four distinct csv files: file, http, email, and device. Each of these files comprised more than 800,000 entries and encompassed many features that facilitate the mapping and understanding of user behavior within the network context. Our work underscores the value of linking technical indicators with policy frameworks to evaluate insider threat risk.

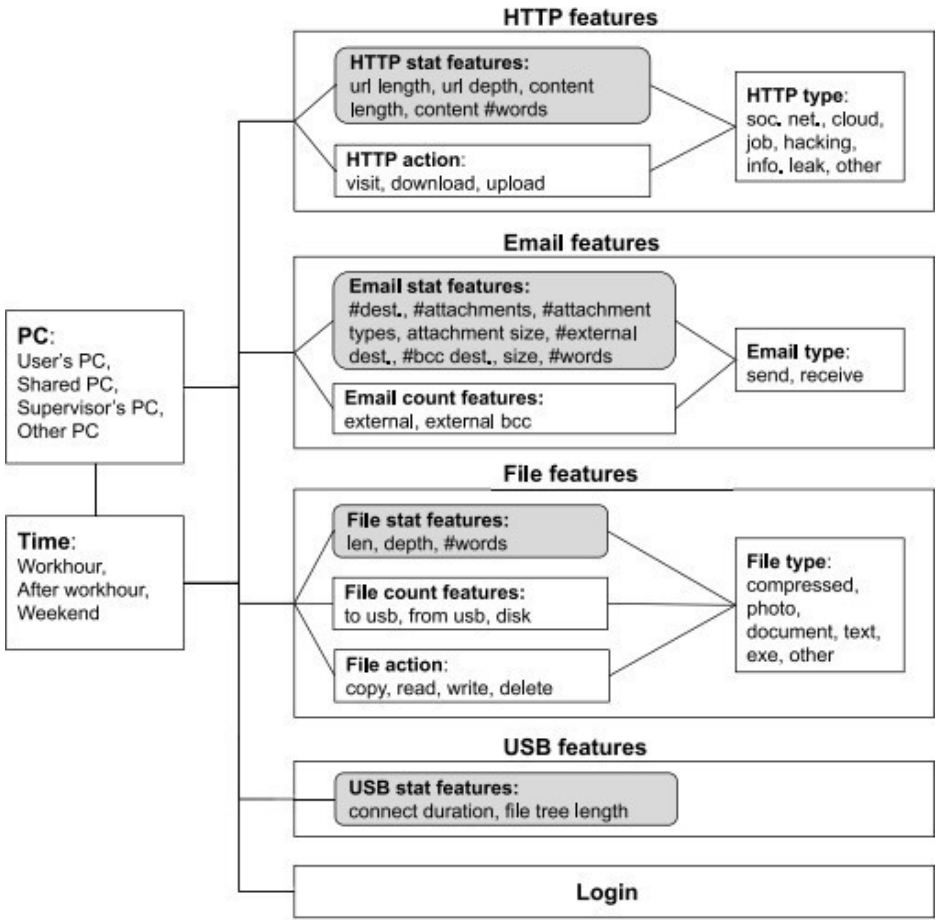

Figure 1. Example of feature extraction (Le et al., 2020)

Le et al. (2020) demonstrated the different elements within a cyber infrastructure that can be used to model a user's day-to-day work behavior, as seen in Figure 1. By relating the features on the right of Figure 1 with temporal and spatial elements, we extracted sequential features that indicated a higher probability of a malicious actor event. Our work did not focus on content filtering to simplify the results. However, we did use file names and URLs to flag suspicious behavior.

Our features reflect realistic scenarios that may enable us to use the model in production environments in future work. The dataset we used in this study is synthetic data, which alleviates privacy and ethical concerns in the conduct of our research (Glasser and Lindauer, 2013; US Department of Homeland Security, 2012).

Limitations and Delineations

Given the synthetic nature of the dataset we used in this study, our findings support the simulation model's feasibility and internal consistency. Due to hardware resource limitations, we evaluate only the users in scenarios 3 and 4. Interested readers may refer to the US-CERT dataset for these scenario definitions. We defined our general policies using scenarios 3 and 4.

The implemented model provides preliminary support for the detection of other events. To enhance the robustness of our findings and minimize sources of variation, we deliberately narrowed the focus of our experiment to two specific scenarios and a fixed group of forty identified malicious actors. Due to limited research showing the effectiveness of corrective measures on an employee population, we used a 5% reduction factor to simulate an organization implementing a corrective measure on its population (e.g., cyber awareness training).

We do not directly compare our model outputs to detection model effectiveness rates in the literature. Our experiment aims to simulate how an organization could leverage user event monitoring tools to formulate a composite score and reduce policy infringements. The model's success is demonstrated by its ability to increase the probability of flagging a malicious actor using an iterative approach.

Table 1. Policy to feature mapping

Policy	Feature/event	Trigger
No files, large emails, or attachments sent to external organizations without internal organization awareness	Monitor sent email	Employee sends a large email with no internal employee on the email to an organization outside of domain
No files, large emails, or attachments sent to self or personal email without internal organization awareness	Monitor sent email	Employee sends a large email with no internal employee on the email to own email or a personal email
Internet will be used for work-related activity; no browsing unauthorized content	Monitor web traffic URL	User browses unauthorized or non-work-related content
Users are prohibited from accessing files that are not related to their projects	Monitor file tree access	User access file not within normal file tree
Users are prohibited from copying proprietary information on to personal removable media or copying foreign files that can be executable on to organization devices	Monitor connect/disconnect and reading/writing of files to personal devices	User copies file outside of file tree or writes unauthorized file to file tree
Users are only authorized to perform work-related activities on the device assigned to them	Monitor logon attempts and PC names	User logs onto more than one device in set period
User event sequences will be monitored for suspicious activity	Monitor unusual sequences of activity	User performs activities outside of normal behavior

By concentrating on controlled parameters based on features derived from policy, we mitigated potential confounding factors, minimized the standard deviation, and attenuated the influence of outliers within the model's detection capability.

Preprocessing

We imported each csv file and processed them using Python parsing techniques. We applied filtering techniques based on the appropriate policy. All files were aggregated into a main csv file with a heading of (date, user, pc, event, score, neg_event). The score was binary, signifying whether the event was a policy infringement. The neg_event was also binary and signified whether the event was part of a scenario that represented an incident inside the organization. We used the event and timestamp columns together to extract sequential features that rely on temporal data. Because the model does not directly align with the dataset's intent, many more scores than neg_events exist.

Model Dynamics

We collected a priority queue for each user, timestamp, and associated score. We applied a modified exponential moving average (MEMA) that calculated a score based on a set period and a second score based on a third of the previous set period. We applied a smoothing factor (SF) to the formula that gave more weight and influence on the nearer-term period. Using the changes in the user's EMA and distance from the standard deviation, we calculated each user's z-score, which became larger when there was more disruption in the user's composite score. Consider the following representation:

let SPS = short period score; LPS = long period score; SF = $\frac{2}{(1 + (SPS))}$

$$MEMA = (SPS - LPS)(SF) + SPS$$

We created a distribution based on the user population z-scores. Users who demonstrated the largest z-scores (a larger deviation from their normal behavior compared to the population) within a certain period were moved to the top of an array. This distribution was consistently updated. A representation of the z-score is as follows:

Let z represent the z-score; x represents the MEMA; μ = mean average of MEMA for a user; σ is the standard deviation of the population or sample.

$$z = \frac{x - \mu}{\sigma}$$

We measured where the forty malicious events occurred throughout the time window in which the incidents occurred and derived the average rank based on z-scores to determine how well the model successfully classifies the users associated with bad events. We sorted the z-scores from highest to lowest and ranked them in order. To

Table 2. Weights used in model

	Weights
personalP	User email's attachment or large email to self or personal email
emailP	User email's attachment or large email to external organization with no
webP	User browses unauthorized websites
logonP	Logon attempts outside of normal work hours
deviceP	Unauthorized copying of files or transferring files to a work device
fileP	Unauthorized file tree access
multiP	Accessing multiple computers within a set period of time
afP	Multiplier that is used to weigh events that happen after hours
evP	Tracks suspicious sequences of events (e.g., multiple unauthorized websites followed by unauthorized device)

minimize the error, we set checks throughout the model to ensure that the events were classified appropriately and aligned with the timestamps within the answer dataset. We also ensured that only malicious actors with ground truth negative events were assessed during the evaluation period.

We used AnyLogic's optimization feature to calibrate the weights associated with each event within the dataset based on the collected features. Two types of features were considered: sequential features and frequency features. We scored frequency features according to their corresponding weight. We based weights in our study on email transactions, file tree access, web behavior, removable media behavior, and logons.

The sequential features relied on users being logged in to multiple PCs during a set period, or were judged based on access patterns. We purposely left the features generic to reflect simple, monitorable organizational policies.

Let weights be represented as **W**, where W $\in \{\mathbb{R}\}$ and $\{w_i, \ldots, w_n\} \setminus \{0\}$.

We collected results by assessing how many malicious actors (n = 40) out of the total population (N = ~2000) were correctly classified by the model. We assessed all individuals that were ranked based on the absolute value of their z-scores, from highest to lowest. The short period was set to 15 days and the long period to

user's z-score $\leftarrow F(username, timestamp, weights, event, score, pc, w_i, ..., w_n)$

Initialize User Priority Queue to PQ
Initialize ZScore HashMap to HM
for each entry in event log **do**
  if username exists in PQ **then**
    Load all username entries for set time-period in PQ into a temporary set
  end if
  if score = 1 **then**
    score *= (appropriate weight + 1)
  end if
  if user has been on another PC **then**
    score *= (multP + 1)
  end if
  if user was active after hours **then**
    score *= (afP + 1)
  end if
  if user falls in a suspicious sequence **then**
    score += (evP + 1)
  end if
  Calculate MEMA = (SPS - LPS) * SF + SPS
  if MEMA $\neq$ 0 **then**
    Update score with the $MEMA_{i-1} - MEMA_i$
  else
    Update score with 0
  end if
  Update PQ with timestamp, username, event, score, pc
  Calculate z-score using the formula: z-score = $\frac{score-\mu}{\sigma}$
  if HM has username and z score < current z score **then**
    Update HM with current username and z score
    Sort z score from highest to lowest
  end if
  return 100 Users with highest z scores
end for

Figure 2. Algorithm for composite scoring process

45 days. The MEMA used these two time frames throughout the experiment to calculate the user's scores. We chose these time periods because they allowed enough time for the model to correctly classify malicious actors over the variation of detection time windows throughout the experiment.

To demonstrate corrective measures, we performed the same experiment with a reduction factor applied. The reduction factor removes benign actor policy infringements to assess if the change improves the model's ability to detect insider threats. The reduction factor was set to .05, which is an estimation of an effective reduction factor when considering similarly sized organizations with adequate training versus those with a high amount of security incidents (Kweon et al., 2019). The experiment's purpose was exploratory analysis to determine if there was adequate support for using composite scores to classify potential insider threats. Our study also explored whether corrective measures could help improve detection rates with the ADAPT-IT framework, which implied an adaptive and iterative noise reduction approach.

RESULTS AND ANALYSIS

The initial experiment without a reduction factor successfully classified 38 out of the 40 malicious actors, resulting in an error rate of 5 percent. The model flagged many of the bad actors with a rank of three or lower during their detection window. Three malicious actors were classified early, but their ranks remained below 40 during their detection window. As the events were classified, the ranks of the malicious actors shifted. At the end of the experiment, the 38 correctly classified malicious actors had a mean rank of 36.05, with a maximum rank of 106 and a minimum rank of 1. The range of rankings spanned 105 positions.

In the second experiment, the model rankings significantly changed with the reduction factor applied. The composite scores for the bad actors in the first experiment ($M = 43.30$, $SD = 12.65$) remained identical; however, the ranks changed noticeably. Two bad actors were not classified correctly in the second experiment, resulting in an error rate of 5 percent. Only two bad actors were classified earlier than their detection window. The range of ranks slightly improved to 100. To assess the significance of the results, we compared the composite scores of each correctly classified malicious actor divided by their respective rank at the end of the experiment using a two-tailed pairwise *t*-test. The second experiment with the reduction factor demonstrated a significantly higher ranking of malicious actors compared to the first experiment without the reduction factor, $t(38) = 3.28$, $p < .02$, $r = .22$.

DISCUSSION AND FUTURE WORK

This paper demonstrates support for an alternative approach to mitigating cybersecurity risk built around adaptive policymaking informed by network monitoring. Following the iterative cycle illustrated in Figure 3, organizations will have an adaptive, quantifiable score to facilitate better-informed decision-making. As demonstrated by applying appropriate corrective measures and adjusting policy, malicious actors will become more detectable as the organization reduces overall policy infringements within its cyber infrastructure. This approach will assist organizations in improving the security of their information systems by giving them risk reduction measures to improve their security posture iteratively.

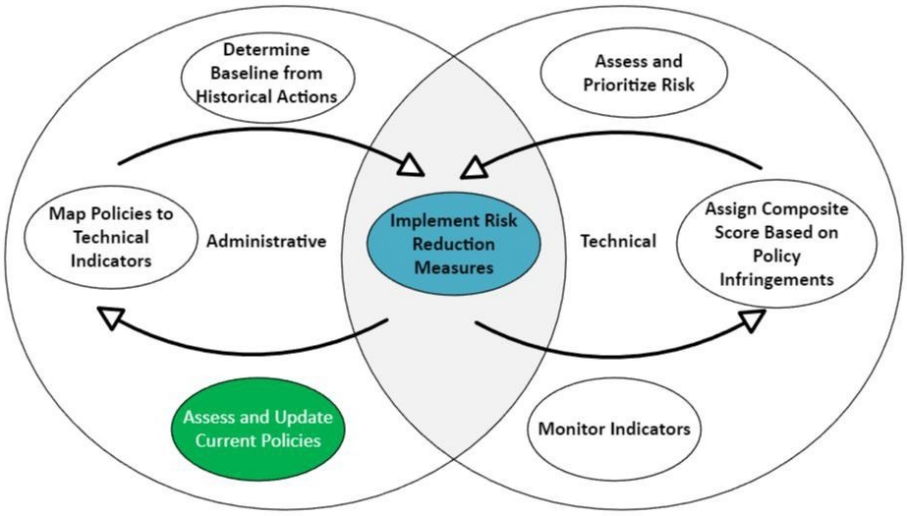

Figure 3. ADAPT-IT approach

For future work, researchers should assess more scenarios of the US-CERT dataset. These experiments would require widening the scope of the policies to capture policy infringements that relate to other malicious actor scenarios. Our model successfully identifies users that changed their behavior abruptly; however, persistent threats that lasted over two months were consistently classified at a higher rank, signifying that the malicious actor events were deviating very little from their normal behavior. This phenomenon is a known issue, described above as related to mimicry attacks. Our work illustrates how more analysis of convolution techniques that use windowing to create a more precise flagging signal is needed.

NOTES

1. The US-CERT dataset is available at https://kilthub.cmu.edu/articles/dataset/Insider_Threat_Test_Dataset/12841247/1

REFERENCES

Balozian, P., & Leidner, D. (2017). Review of IS security policy compliance: Toward the building blocks of an IS security theory. *ACM SIGMIS Database*, 48(3), 11–43. doi:10.1145/3130515.3130518.

Bassett, G., Hylender, C. D., Langlois, P., Pinto, A., & Widup, S. (2023). *Verizon 2022 data breach investigations report*.

Becker, W., Saisana, M., Paruolo, P., & Vandecasteele, I. (2017). Weights and importance in composite indicators: Closing the gap. *Ecological Indicators, 80*, 12–22. doi:10.1016/j.ecolind.2017.03.056.

Caputo, D., Maloof, M., & Stephens, G. (2009). Detecting insider theft of trade secrets. *IEEE Security & Privacy Magazine, 7*(6), 14–21. doi:10.1109/MSP.2009.110.

Chatterjee, S., Sarker, S., & Valacich, J. S. (2015). The behavioral roots of information systems security: Exploring key factors related to unethical IT use. *Journal of Management Information Systems, 31*(4), 49–87. doi:10.1080/07421222.2014.1001257.

Collingridge, D. (1980), *The social control of technology.* St. Martin's Press, New York.

Collins, M., Theis, M., Trzeciak, R., Strozer, J., Clark, J., Costa, D., & Moore, A. (2016). *Common sense guide to mitigating insider threats* (Technical report CMU/SEI-2015-TR-010). CERT Insider Threat Center.

Couretas, J. M. (2018). *An introduction to cyber modeling and simulation.* John Wiley & Sons, Hoboken, NJ.

Eldardiry, H., Bart, E., Liu, J. Hanley, J., Price, B., & Brdiczka, O. (2013). Multi-domain information fusion for insider threat detection. In *Proceedings of the Security and Privacy Workshops* (pp. 45–51). IEEE, San Francisco, CA. doi:10.1109/SPW.2013.14

Glasser, J., & Lindauer, B. (2013). Bridging the gap: A pragmatic approach to generating insider threat data. In *Proceedings of the Security and Privacy Workshops* (pp. 98–104). IEEE, San Francisco, CA. doi:10.1109/SPW.2013.37

Haasnoot, M., Middelkoop, H., Van Beek, E., & Van Deursen, W. P. A. (2011). A method to develop sustainable water management strategies for an uncertain future. *Sustainable Development*, *19*(6), 369–381.

Haasnoot, M., Van Deursen, W., Middelkoop, H., Beek, E. V., & Wijermans, N. (2012). *An integrated assessment metamodel for developing adaptation pathways for sustainable water management in the lower Rhine Delta* [Conference presentation]. Sixth International Congress on Environmental Modelling and Software.

Haasnoot, M., Kwakkel, J. H., Walker, W. E., & Ter Maat, J. (2013). Dynamic adaptive policy pathways: A method for crafting robust decisions for a deeply uncertain world. *Global Environmental Change, 23*(2) 485–498. doi:10.1016/j.gloenvcha.2012.12.006.

Hadlington, L. (2020). The 'human factor' in cybersecurity: Exploring the accidental insider." In *Research anthology on artificial intelligence applications in security* (pp. 1960–1977). IGI Global. doi:10.4018/978-1- 7998-7705-9.ch087.

Hubbard, D. W., & Seiersen, R. (2023). *How to measure anything in cybersecurity risk* (2nd ed.). John Wiley & Sons, Hoboken, NJ.

Khan, M. I., Foley, S. N., & O'Sullivan, B. (2022). Database intrusion detection systems (DIDs): Insider threat detection via behaviour-based anomaly detection systems—a brief survey of concepts and approaches. In W. Meng & S. K. Katsikas (Eds.), *Emerging information security and applications* (pp. 178–197). Springer, Cham. doi:10.1007/978-3-030-93956-4_11

Kim, J., Park, M., Kim, H., Cho, S., & Kang, P. (2019). Insider threat detection based on user behavior modeling and anomaly detection algorithms. *Applied Sciences, 9*(19), 4018. doi:10.3390/app9194018

Kweon, E., Lee, H., Chai, S., & Yoo, K. (2021). The utility of information security training and education on cybersecurity incidents: An empirical evidence. *Information Systems Frontiers, 23*(2), 361–373. doi:10.1007/s10796-019-09977-z

Le, D. C., Zincir-Heywood, N., & Heywood, M. I. (2020). Analyzing data granularity levels for insider threat detection using machine learning. *IEEE Transactions on Network and Service Management*, *17*(1), 30–44.

Legg, P. A. (2015). Visualizing the insider threat: Challenges and tools for identifying malicious user activity. In *Proceedings of the IEEE Symposium on Visualization for Cyber Security (VizSec)* (pp. 1–7). Chicago. doi:10.1109/VIZSEC.2015.7312772

Lerums, J. E., Poe, L. D., & Dietz, J. E. (2018). *Simulation modeling cyber threats, risks, and prevention costs* [Paper presentation]. IEEE International Conference on Electro/Information Technology (EIT), Rochester, MI. doi:10.1109/EIT.2018.8500240

Manokha, I. (2020). The implications of digital employee monitoring and people analytics for power relations in the workplace. *Surveillance and Society*, *18*(4).

Master, A., Hamilton, G., & Dietz, J. E. (2022). *Optimizing cybersecurity budgets with AttackSimulation* [Paper presentation]. IEEE International Symposium on

Technologies for Homeland Security (HST), Boston, MA. doi:10.1109/HST56032.2022.10024984

Mingers, J., & Rosenhead, J. (Eds.). (2001). *Rational analysis for a problematic world revisited: Problem structuring methods for complexity, uncertainty and conflict* (2nd ed.). John Wiley & Sons.

Rosenhead, J., Elton, M., & Gupta, S. K. (1972). Robustness and optimality as criteria for strategic decisions. *Journal of the Operational Research Society, 23*(4). doi:10.1057/jors .1972.72.

Rosenhead, J. (1990). Rational analysis: Keeping your options open. In *Rational analysis for a problematic world: Problem structuring methods for complexity, uncertainty and conflict*. John Wiley & Sons.

Saisana, M., Saltelli, A., & Tarantola, S. (2005). Uncertainty and sensitivity analysis techniques as tools for the quality assessment of composite indicators. *Journal of the Royal Statistical Society Series A: Statistics in Society, 168*(2), 307–323. doi:10.1111/j.1467-985X.2005.00350.x

Schultz, E. E. (2002). A framework for understanding and predicting insider attacks", *Computers & Security, 21*(6), 526–531. doi:10.1016/S0167- 4048(02)01009-X

Senator, T. E., Goldberg, H. G., Memory, A., Young, W. T., Rees, B., Pierce, R., Huang, D., et al. (2013). Detecting insider threats in a real corporate database of computer usage activity (pp. 1393–1401). In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. doi:10.1145/2487575.2488213.

Shaw, E., Ruby, K., & Post, J. (1998). *The insider threat to information systems: The psychology of the dangerous insider.* Security Awareness Bulletin 2–98. Defense Security Service.

Spooner, D., Silowash, G., Costa, D., & Albrethsen, M. (2018). Navigating the insider threat tool landscape: Low cost technical solutions to jump start an insider threat program (pp. 247–257). In Proceedings of the IEEE Security and Privacy Workshops (SPW), San Francisco. doi:10.1109/SPW .2018.00040.

Tzvetanov, K., Riegsecker, A., Frantz, B., Xiong, C., Bott, R., Cline, T., & Dietz, J. E. (2022). Agent-based modeling for theme park evacuation. *Journal of Emergency Management, 20*(2), 157–173.

US Department of Homeland Security. (2012). *The Menlo report: Ethical principles guiding information and communication technology research*.

Warkentin, M., & Willison, R. (2009). Behavioral and policy issues in information systems security: The insider threat. *European Journal of Information Systems, 18*(2), 101–105. doi:10.1057/ejis.2009.12

Yamin, M. M., Katt, B., Sattar, K., & Ahmad, M. B. (2020). Implementation of insider threat detection system using honeypot based sensors and threat analytics (pp. 801–829). In *Advances in information and communication: Proceedings of the 2019 future of information and communication conference (FICC), volume 2*. Springer International.