University of Kent

Kent Academic Repository

Nurse, Jason R. C., Dobrontei, Suzie, Webster, Jonathan and Alashe, Oz (2026) ***Toward a behavioral ontology for cybersecurity: Introducing SebDB*****. In: HCI for Cybersecurity, Privacy and Trust: 8th International Conference, HCI-CPT 2026, Held as Part of the 28th HCI International Conference, HCII 2026, Montreal, QC, Canada, July 26–31, 2026, Proceedings. . pp. 492-512. Springer ISBN 978-3-032-297 E-ISBN 978-3-032-29723-5.**

Downloaded from

https://kar.kent.ac.uk/113131/ The University of Kent's Academic Repository KAR

The version of record is available from

https://doi.org/10.1007/978-3-032-29723-5_28

This document version

Publisher pdf

DOI for this version

Licence for this version

CC BY (Attribution)

Additional information

Versions of research works

Versions of Record

If this version is the version of record, it is the same as the published version available on the publisher's web site. Cite as the published version.

Author Accepted Manuscripts

If this document is identified as the Author Accepted Manuscript it is the version after peer review but before type setting, copy editing or publisher branding. Cite as Surname, Initial. (Year) 'Title of article'. To be published in ***Title of Journal*** , Volume and issue numbers [peer-reviewed accepted version]. Available at: DOI or URL (Accessed: date).

Enquiries

If you have questions about this document contact ResearchSupport@kent.ac.uk. Please include the URL of the record in KAR. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from https://www.kent.ac.uk/guides/kar-the-kent-academic-repository#policies).



Toward a Behavioral Ontology for Cybersecurity: Introducing SebDB

Jason R. C. Nurse[1,2(✉)], Suzie Dobrontei[2], Jonathan Webster[2], and Oz Alashe[2]

[1] Institute of Cyber Security for Society (iCSS), University of Kent, Canterbury, UK
j.r.c.nurse@kent.ac.uk
[2] CybSafe, London, UK
{jason.nurse,oz.alashe}@cybsafe.com

Abstract. Cyberattacks that target individuals remain a significant threat to an organization's security, placing human behavior at the center of effective defense. This paper focuses on this topic and makes two novel contributions to the field. First, we introduce the Security Behaviors Database (SebDB), an open-source resource that documents security behaviors, links them to organizational risks, and maps them to established security standards. SebDB currently includes more than 100 core behaviors, associated impacts when behaviors are not followed, and alignments with widely used frameworks, including the NIST Cybersecurity Framework (CSF) and MITRE ATT&CK. Developed by a consortium of practitioners and researchers, SebDB is intended to improve the understanding, modeling, and measurement of human cyber risk and to support behavior-change initiatives that can measurably improve security posture. In this article, we describe the motivation, design, and structure of SebDB, including behaviors and their risk and control mappings. The second novel contribution of this research is the proposal of a behavioral ontology for cybersecurity. This ontology, which realizes the overarching vision for SebDB, presents one of the first comprehensive conceptualizations of knowledge in the behavioral security and human aspects of the cybersecurity domains. Through this ontology, we model how behaviors can be linked to risks, threat actor tactics, vulnerabilities, and security controls, while also covering how behavioral interventions can be leveraged as part of an organization's security strategy.

Keywords: security behaviors · behavior change · human cyber risk · awareness and training · database · behavioral ontology · knowledge base · human risk management · behavioral security · psychology

1 Introduction

The human aspect is one of the most actively discussed topics in cybersecurity today. This is largely due to the persistent stream of attacks that exploit people, the increased prevalence of human error when interacting with complex

© The Author(s) 2026
A. Moallem (Ed.): HCII 2026, LNCS 16738, pp. 492–512, 2026.
https://doi.org/10.1007/978-3-032-29723-5_28

digital technologies, and the difficulty in adequately addressing such a multidisciplinary challenge [32]. These realities have compelled security practitioners and researchers to move beyond the purely technical aspects of the cyber domain to better understand people. This covers aspects such as how we perceive security, how we reason in security contexts, and the heuristics and biases that lead us to succumbing to attacks and making predictable mistakes [21]. As a part of this progression, a steady stream of work focused on solutions oriented around humans and their behavior has emerged.

Focusing first on the industry and policy communities, a host of standards, frameworks, and knowledge bases have been proposed that aim to provide a robust foundation for research and practice. Prominent examples include NIST SP 800-50: Building a Cybersecurity and Privacy Learning Program, the Cyber Security Body of Knowledge (CyBOK) Human Factors Knowledge Area, and the SANS Security Awareness & Culture Maturity Model™. In the academic community, proposals are considerably wider but also significantly more disparate. These range from frameworks for assessing organizational security culture [10,32,33] to models focusing on human behavior in security [26], and standardized artifacts for studying and measuring security practices [13,27,36]. Each proposal targets a part of the problem and presents novel steps toward its resolution.

Although the contributions to date have merit, critical gaps remain in how we as a community approach the human-centered cybersecurity and build effective initiatives, systems, and programs [4,23,28]. This paper addresses one such gap by advancing knowledge of how human behavior relates to cybersecurity. Specifically, we make two novel contributions.

First, we introduce a comprehensive database oriented around security behaviors, examining how human actions shape risk and can reinforce or undermine core security standards. This database, SebDB, is recognized in practice (SebDB.com) but is being presented here for the first time to the research community. SebDB is an open-source resource and the foundation of a rapidly evolving behavioral ontology for cybersecurity. It was created with input from a consortium of professionals in the practitioner and academic security awareness and training (SAT), cybersecurity culture, and human risk management (HRM) communities, to improve understanding, modeling, and measurement of human cyber risk.

In this article, we introduce SebDB, including its motivation and the robust process through which it was developed. We further present the core of SebDB, including examples from the more than 100 documented behaviors within it, the associated impacts when behaviors are not followed, and mappings of behaviors to pertinent security standards (e.g., the NIST Cybersecurity Framework (CSF) and the MITRE ATT&CK framework).

The second novel contribution is the proposal of a behavioral ontology for cybersecurity. This ontology represents an early glimpse into the future of SebDB, and one of the first comprehensive conceptualizations of the people and cybersecurity domain. Through this ontology, we define key domain concepts and

relationships, and explain how they fit together in a domain model. Examples of concepts include security behaviors, behavioral interventions and behavior change techniques, threats, vulnerabilities, risks, incidents, risk outcomes, and security controls. To our knowledge, this is the first academic article to present a behavioral ontology for cybersecurity, which marries the domains of behavioral science and cybersecurity in this way.

The remainder of this paper is structured as follows. Section 2 reviews related literature on the topics of security culture, awareness, human risk management, and behavior change in cybersecurity. In Sect. 3, SebDB is introduced and detailed. This covers the methodology supporting its creation, its core concepts, and how these can be used. Section 4 then presents our initial ontology complete with a visual depiction as well as definitions of its concepts and relationships. We then conclude and briefly outline future work in Sect. 5.

2 Related Work

Research into how humans and cybersecurity aspects interact has progressed significantly over the last two decades. Early research in the domain focused largely on the usability of security systems and functionality, considering how complex interfaces were and the need to redesign them with the human (i.e., their context, perceptions, decision-making nuances) in mind. Noteworthy examples span the usability of authentication systems, where password setup and recovery were arduous and overly complicated [6], to the appropriate communication of cyber risk to aid understanding, given a particular target audience [22].

Since then, the variety and amount of research on the human aspect has grown substantially. For instance, one stream of work has centered on the psychology of security, i.e., how psychology and behavioral science can help to understand an individual's security attitudes and practices [34]. This research covers a clear gap in traditional literature and combines two topical disciplines: cybersecurity and psychology. Another popular stream of work focuses on security awareness and training and the broader subject of security culture [2,23]. In this domain, research has examined the impact of awareness on knowledge, attitude, and behavior [2,36], the challenges of changing behavior and the need for a more comprehensive approach [4,28], and how behavior and culture are linked [32,33]. A primary conclusion of this work is that awareness itself does not lead to behavior change, and instead, a more targeted approach is required.

While the human aspect of security domain has developed significantly in the past two decades, one topic that has yet to receive more formal attention is how human actions shape risk and can either reinforce or undermine core security standards. This goes beyond the general reality that security behaviors are important and should be promoted [5,30], to defining specific behaviors that are critical, understanding their influence on specific risks, and how those behaviors contribute to wider organizational security postures. Such an understanding could allow more customized approaches toward researching and addressing human cyber risk, as well as providing a pathway for potential standardization of knowledge in the domain.

3 SebDB: The Security Behavior's Database

3.1 Motivation

Three factors motivated the design and development of SebDB. The first factor was grounded in the premise that individual security behaviors are central to the conceptualization and mitigation of human cyber risk. While the importance of secure online practices is well known [14], there is limited work cataloging the range of human behaviors that affect an organization's or individual's security. Such a characterization would help define not only the actions to monitor (e.g., clicking on phishing emails), but also the positive actions to actively promote (e.g., checking messages for signs of legitimacy). Specifying behaviors also creates a transparent path for targeted, and ultimately measurable, security initiatives [28].

The second factor acknowledged that people's behaviors are tightly connected to other security constructs. Prior research documents relationships among security attitudes, culture, and behavior [2,36], yet structured analyses of how specific behaviors shape impacts of attacks and risk outcomes, or how they support or undermine compliance with security standards, remain nascent. A better appreciation of these topics would enable more rigorous risk reasoning (e.g., prioritizing behaviors by their influence on risk) and provide clearer evidence of how behaviors and behavioral interventions (i.e., targeted activities designed to change specific behaviors [19]) contribute to an organization's security posture.

As with other more technical domains in security, there are significant benefits to be gained from standardization; this formed the third factor. Technical standards and frameworks, such as MITRE's CVE, CWE, and ATT&CK, FIRST's CVSS, and the OWASP Top 10, give practitioners a shared foundation and vocabulary and help improve responses to threats through common characterizations and learning. In the human aspects domain, comparable standardization is largely absent, particularly in the form of well-structured, typed frameworks or comprehensive data and knowledge bases. Guided by these three factors, the next section details how we designed SebDB.

3.2 Design and Development

SebDB was created and refined over five years through a series of development stages. In what follows, we focus on the two most significant stages: stage one, which produced the initial version of the database, and stage four, which led to SebDB's most substantial and recent update. The intermediate stages involved smaller, incremental revisions informed by feedback from security professionals and by advances in the SAT/HRM, technology, and business domains.

Stage one focused on establishing the foundational structure and content of SebDB. The process began with a systematic review of academic literature (e.g., [7,29]), complemented by a comprehensive review of industry guidance and standards from authoritative sources such as the UK's NCSC, ISO (e.g., ISO 27000 series), NIST (e.g., NIST SP 800-50, NIST CSF), and SANS. The primary

objective was to identify and define all relevant security-related behaviors and articulate how their adoption or lack of use could introduce risks to individuals and their employers. We also leveraged existing research to provide an initial estimate of each behavior's influence on risk, which was defined in a four-tier structure: Tier 1 behaviors were deemed to have the greatest impact on risk, and Tier 4 the least.

Following the initial data capture, we undertook a mixed-method validation process. First, a series of expert interviews was conducted with ten security professionals from academia, industry, and government, all specializing in human-centered security. These sessions were designed to validate the initial set of behaviors, risks, and tiers, and critically, to co-design the structure of the database, ensuring the final resource directly addressed practitioner needs. Subsequently, a survey was administered to 599 members of the general public to assess the feasibility and current practice rates of the identified behaviors among average users. The survey results were then used to finalize the behavior set, particularly by refining their names, structure, and contexts. The culmination of stage one was the release of SebDB version 1 (v1), which contained 76 behaviors, their mapping to eight risk outcomes (e.g., account compromise, malware infection), and an assessment of the influence they had on risk (i.e., their tier).

Stage four was driven by significant shifts in the risk landscape and aimed to comprehensively update SebDB to address these changes. This involved a thorough review and reflection process, incorporating direct feedback gathered from security professionals actively using the database. Every behavior, risk outcome, and tier was re-evaluated for its value, utility, and suitability. To guide this process and ensure consistency, a set of SebDB principles was established to specify how core components should be interpreted and applied. The review was further informed by new advances in the field, including updated industry publications (e.g., NIST SP 800-50 Revision 1), recent academic research [3,4, 28], and insights pertaining to the emergence of new security behaviors (e.g., associated with new technologies (AI) or practices (remote working)).

Through review, reflection, and refinement, we produced an updated database, which was subsequently reviewed by a group of engaged practitioners and revised to incorporate their input. Similar to SebDB v1, this iteration specified behaviors (now 102 in total), their mappings to risks (through seven impacts), and their influence on risk (via four tiers). Critically, a more refined and robust assessment methodology was adopted for determining each behavior's tier. This involved evaluating how significantly performing, or failing to perform, the behavior affects the plausibility of any of the undesirable outcomes (i.e., impacts) occurring. Plausibility is regarded as the logical and practical likelihood that a specific impact may occur as a direct or indirect consequence of an individual failing to perform a specific behavior. This assessment considers three dimensions: directness, i.e., whether the behavior acts as a primary safeguard or one of multiple defenses; conditionality which considers whether the impact requires additional system failures or attacker actions to materialize; and sever-

ity of consequence if the behavior is not performed. These judgments represent expert consensus based on established security knowledge, real-world incident analysis, and industry frameworks, rather than empirically derived probabilities. Through combining plausibility scores for each impact for a behavior, we were able to determine the behavior's wider influence and then define its respective tier (i.e., how significant an influence it would have in contributing to risk).

The updated database also incorporated a comprehensive mapping of behaviors to core security frameworks (NIST CSF and MITRE ATT&CK), enabling cybersecurity professionals and researchers to examine how behaviors tangibly influence threat actions (e.g., MITRE's initial access or privilege escalation) and support security functions (e.g., NIST's Protect or Recover). To aid this mapping assessment as well as the analysis of behavior tiers, an AI model specifically trained in cybersecurity and human risk management was used.

The model served as an analytical tool to systematically synthesize research, industry benchmarks, and contextual data relevant to our specific scope. We leveraged this capability to evaluate three specific areas: first, the behavior's cumulative influence on risk by determining how its performance (or omission) shifts the plausibility of adverse outcomes; second, the extent to which a behavior supports specific NIST CSF core functions; and third, the potential for a user's actions to disrupt or facilitate known MITRE ATT&CK tactics. By using the model to process these diverse data sources, we were able to generate a data-driven assessment of how individual behaviors impact the broader organizational security posture. All output from the model was manually assessed for accuracy and quality by the research team, and then was shared with a group of security professionals for further review.

The outcome of this process, and stage four more generally, was SebDB version 4 (v4); this is the latest version.

3.3 The Database

As discussed above, a primary goal of SebDB is to improve the understanding, modeling, and measurement of human cyber risk and to support behavior-change initiatives that can measurably improve an organization's security posture. To achieve this aim, the database is structured around security behaviors, their associated impacts, and their alignment with recognized security frameworks.

Security Behaviors. Each security behavior in SebDB represents a discrete, observable action that individuals can take to protect information, systems, or assets. What constitutes a security behavior is governed by thirteen core principles to ensure consistency and utility. These principles establish that security behaviors must focus on activities of everyday users of technology rather than technical specialists (Principle 1), represent single specific actions as opposed to composite activities (Principle 2), and focus on cybersecurity and physical security as they relate to protecting information (Principle 3). Building on these points, behaviors also represent actions that could result in a negative impact

on an organization if not performed (Principle 6). This is an important facet as SebDB is built to help reduce the risk faced by organizations.

Behaviors are deliberately framed in positive, affirmative language (Principle 5) to emphasize constructive security practice. For instance, "SB010: Keeps a password (or passphrase) private" uses positive language to explain what to do, rather than "SB010: Does not share passwords", which uses negative language to convey what not to do. This positive framing is critical because it clearly defines the desired "good" state and expected actions, and moves practitioners away from relying on lists of prohibitions toward guiding people proactively. Positively phrased and framed messages are more persuasive than negatively framed ones, particularly when there is minimal emphasis on detailed processing [22]. This approach aligns with calls to shift from a "human-as-problem" to a "human-as-solution" mindset in cybersecurity, where individuals are viewed as active contributors to organizational security rather than the primary security risk [35].

SebDB's security behaviors are created to be technology-agnostic (Principle 7), applying broadly across hardware and software contexts rather than being narrowly scoped to particular platforms or tools. This universality prevents fragmentation of the database into numerous overlapping sub-behaviors and ensures that they remain relevant as technology evolves. Additionally, a security behavior must be a unique action in the SebDB context, which avoids overlap with other SebDB behaviors (Principle 10). Behaviors must also be measurable (Principle 8 [8]) in order for targeted behavior interventions to be evaluated for their effectiveness [17,28]. Measurement approaches vary by context and may include direct observation, behavioral telemetry from security tools, self-report assessments, or controlled simulations, depending on organizational capabilities and privacy considerations. Finally, behaviors must be realistic and actionable within everyday contexts [4], acknowledging the practical constraints individuals face (Principle 9). This is critical given the reality that if behaviors are atypical, cumbersome, or irrelevant, decades of usable security research have shown that they are much less likely to be performed [22,25].

To maintain relevance and accuracy over time, each behavior is assigned a permanent unique identifier (UID) that persists throughout its lifecycle (Principle 12). When behaviors are removed from the database, their UIDs are retired rather than reassigned, ensuring consistency and maintaining a clear historical record. These practices are aligned with standard guidance [9]. Recognizing that security threats, technologies, and organizational practices continually evolve, behaviors are subject to change over time (Principle 11). Each behavior is documented in a changelog that tracks its status: whether adopted (newly introduced), unchanged (reviewed and still valid), renamed (clarified or updated), split (divided for specificity), merged (combined due to redundancy), or deprecated (retired due to irrelevance).

The database includes 102 behaviors that were identified and validated through collaboration among security professionals and academics, informed by extensive academic research as mentioned earlier. These behaviors align with established research on actions that place the organization and the individual

at increased threat from malicious attacks [13,14]. For example, risky behaviors documented in prior research, sharing passwords with colleagues ("SB010: Keeps a password (or passphrase) private" in SebDB), downloading copyrighted material from illegal websites ("SB019: Downloads content only from an approved or official website"), ignoring warnings to update apps and computer software ("SB024: Keeps software up-to-date"), use of the same password for multiple websites ("SB016: Authenticates with a unique password (or passphrase)"), and clicking on potentially malicious links in emails ("SB159: Clicks only a legitimate link") [14], are all addressed within SebDB via its positive framing approach. Table 1 lists several other SebDB behaviors; the full set can be found at SebDB.com.

Table 1. Examples of security behaviors

UID	Security behavior name
SB001	Authenticates with MFA
SB003	Authenticates with a strong password (or passphrase)
SB013	Reports a security incident
SB058	Checks a website for signs of deception
SB061	Backs up data

Each behavior has a "Description" that articulates the specific action in a clear and concise way, alongside a "Why is it important?" section connecting the behavior to the security outcomes it influences. To facilitate ease of search and contextual relevance, behaviors are further organized into ten distinct Categories, including "Authentication & access control", "Device security", and "Threat detection & prevention". A category defines a grouping of related security actions; therefore, SebDB allows for a many-to-many relationship where a single behavior can be associated with multiple categories. For instance, "SB153: Opens a file only from a legitimate source" is mapped to both "Threat detection & prevention" and "Safe online practices & communication".

Impacts. Another key component of SebDB is impact. An impact is focused on the harmful or undesirable result of a security risk (Principle 1). It represents a significant concern for an organization (Principle 5) and is a result that could arise due to a risky behavior (e.g., not backing up data), or that could be prevented by a security behavior (e.g., backing up systems) (Principle 3). This conceptualization of impact draws on existing practice and research, such as NIST SP 800-30 Rev 1. Most notably, here we interpret impact as a type of outcome that could arise, rather than the magnitude of harm that might result. This, we acknowledge, is another valid interpretation.

SebDB categorizes impacts into seven distinct types: "IMP001: System compromise", "IMP002: Data compromise", "IMP003: Physical (asset) compromise",

"IMP005: Identity theft and fraud", "IMP007: Account compromise", "IMP009: Business interruption", and "IMP010: Financial loss". Examples for impact definitions are provided in Table 2. These definitions were derived from a critical synthesis of industry standards, including NIST SP 800-30 Rev. 1, the NIST CSF, ISO 31073:2022, and NCSC's Risk Management guidance, as well as academic research [20]. This methodology ensures that the proposed framework aligns with established practices and remains directly applicable to the functional requirements of practitioners and researchers.

Table 2. Examples of impacts and their definitions

Impact	Definition
IMP001: System compromise	System compromise occurs when there is the unauthorized disclosure, modification, substitution, deletion or use of systems. This can lead to operational instability, unauthorized control over critical processes, and serve as a gateway for further attacks.
IMP010: Financial loss	Financial loss occurs when there has been an unplanned decrease in monetary value or reduction in financial assets. Such losses often follow cyber incidents, regulatory fines, or fraud, and can undermine business continuity and investor confidence.

Impacts are recognized as being able to cause, or lead to, other impacts (Principle 4). For example, "IMP001: System compromise" may lead to "IMP007: Account compromise", which in turn could bring about "IMP002: Data compromise". These relationships allow organizations to better understand cascading risks and the full scope of potential consequences when certain security behaviors are not adopted. Several principles for impact mirror those established for behaviors to ensure the integrity of SebDB. These include the requirement for impacts to be a unique result in the SebDB context to avoid overlap (Principle 6), documented in a changelog over time to track their evolution (Principle 7), and assigned a permanent unique identifier (UID) that is retired rather than reassigned if an impact is removed (Principle 8).

Each security behavior is mapped to every impact through a plausibility metric, with plausibility interpreted as defined in Sect. 3.2. This metric is represented by a four-point scale: 0 (None - no credible scenario exists where failure to perform the behavior leads to this impact); 1 (Low - the link is possible but indirect and requires other system failures or specific attacker actions); 2 (Medium - there is a notable chance of impact, but the impact is not guaranteed or immediate); and 3 (High - the link is clear, direct, and significant because the behavior acts as a primary safeguard against this impact).

Examples of these scores for each impact if "SB013: Reports a security incident" is not performed are presented in Table 3. In such a case, we view "IMP001: System compromise" as high because failure to report a security incident allows an attacker to persist and move laterally. An initial foothold on one machine can escalate to the compromise of entire networks and systems. Whereas, we regard

"IMP003: Physical (asset) compromise" as low due to the reality that while it is conceivable that an unreported system compromise could eventually affect systems controlling physical access, this represents a remote and indirect potential impact of the reporting failure. These justifications are also part of SebDB, given their value in building the evidence base for these decisions.

Table 3. Plausibility scores for each impact if SB013 is not performed

Security behavior	Impact	Plausibility
SB013: Reports a security incident	IMP001: System compromise	3 (High)
	IMP002: Data compromise	3 (High)
	IMP003: Physical (asset) compromise	1 (Low)
	IMP005: Identity theft and fraud	2 (Medium)
	IMP007: Account compromise	3 (High)
	IMP009: Business interruption	3 (High)
	IMP010: Financial loss	3 (High)

More broadly, such scores provide insight that allows professionals and researchers to assess and meaningfully link behaviors to impacts.

Tiers. A primary goal of SebDB is to provide a tool that closely links behaviors to risk. Each behavior is therefore assigned a tier based on its influence on risk (Principle 13). This is defined in a four-tier structure: Tier 1 captures the most critical behaviors where failure creates multiple high-impact pathways, while Tier 4 includes behaviors with more limited or indirect influence. As introduced above in Sect. 3.2, a behavior's tier is determined by assessing how significantly performing (or not performing) the behavior affects the plausibility of undesirable outcomes (i.e., impacts) from related security risks. These plausibility judgments were combined into a score, which determined the behavior's tier: Tier 1 (17–21), Tier 2 (15–16), Tier 3 (13–14), and Tier 4 (0–12). These ranges were defined through a combination of expert consensus and analysis of score distribution patterns. We examined the distribution of cumulative plausibility scores across all 102 behaviors and identified natural clustering points that aligned with practitioner intuitions about relative risk significance. We acknowledge that these thresholds represent expert judgment based on current understanding and may be refined as the field evolves and additional empirical evidence becomes available. In total, there are 22 behaviors in Tier 1, 35 in Tier 2, 25 in Tier 3, and 20 in Tier 4.

To illustrate, "SB013: Reports a security incident" is a Tier 1 behavior. As shown in Table 3, this is largely due to the high plausibility of multiple critical

impacts occurring if the behavior is not performed: "IMP001: System compromise", "IMP002: Data compromise", "IMP007: Account compromise", "IMP009: Business interruption", and "IMP010: Financial loss". The resulting total score of 18 (a tally of the scores in Table 3) places the behavior within the Tier 1 range (17–21).

Framework Mappings. Security behaviors in SebDB are mapped to well-established security frameworks to enable SAT and HRM professionals to demonstrate how behaviors tangibly inform or support security functions, and reduce the success likelihood of threat actions. Below we present our mappings to NIST CSF and MITRE ATT&CK; however, we expect this to be expanded to other frameworks (e.g., ISO 27000 series, CIS Critical Security Controls, Cyber Essentials) in the future.

NIST Cybersecurity Framework (CSF): The NIST CSF 2.0 is a globally recognized framework for managing cybersecurity risks. The tool provides a taxonomy of cybersecurity outcomes that can be used to better understand, assess, prioritize, and communicate cybersecurity efforts. The CSF Core consists of six Functions: Govern, Identify, Protect, Detect, Respond, and Recover. These Core Functions offer a structural guide for organizations, industry, and government to assess security risks, implement protective measures, identify and respond to security events, and recover from incidents [15].

Each security behavior in SebDB is mapped to the specific CSF Core Functions it supports. More specifically, this mapping captures whether a user performing a behavior contributes to the organization achieving that Function. For example, "SB001: Authenticates with MFA" is associated with the Govern and Protect functions. The Protect function focuses on utilizing safeguards, such as identity management and access control, to secure assets and reduce the likelihood of adverse events. "SB001: Authenticates with MFA" serves as a critical protective control by requiring multiple forms of authentication. Another significant benefit of associating behaviors to Functions is that it provides a clear way to demonstrate that behavioral interventions (e.g., security nudge campaigns, webinars) are supporting real security functions and controls, and not just awareness or communications.

MITRE ATT&CK: MITRE ATT&CK is a comprehensive knowledge base of adversarial tactics and techniques based on real-world observations. It has gained extensive recognition across both academia and industry, leading to its implementation in various practical applications, including adversary emulation or behavioral analysis development [10].

The framework is a model that systematically categorizes adversary behavior into two main components: tactics and techniques. Tactics represent the "why": the adversary's tactical goal or the reason for performing an action. There are 14 tactics in the framework, such as "TA0043: Reconnaissance", where the adversary gathers information to plan future operations. Conversely, techniques represent "how" adversaries achieve these tactical goals by performing specific actions. These can also be broken into subtechniques.

In SebDB, each security behavior is currently mapped to the MITRE ATT&CK tactics that it relates to, thereby contextualizing human actions within the technical phases of an adversary's lifecycle. This mapping assesses whether a user performing a behavior can potentially influence an adversary in achieving the goal of the tactic. For example, "SB001: Authenticates with MFA" is linked to "TA0042: Resource development", "TA0001: Initial access", "TA0003: Persistence", "TA0004: Privilege escalation", "TA0006: Credential access", "TA0008: Lateral movement", and "TA0040: Impact". In the specific case of "TA0004: Privilege Escalation", the use of multi-factor authentication on privileged accounts increases the effort required for an adversary to compromise them or escalate permissions, whereas the absence of such controls facilitates easier target acquisition once an initial, lower-privileged foothold is gained.

By aligning behaviors with these tactical goals, SebDB enables organizations, academics, and policymakers to identify where the adoption of a security behavior serves as a point of disruption across the kill chain. Another noteworthy benefit of this mapping is its provision of a shared language between human risk professionals and security operations teams, helping practitioners align behaviors to real-world threat scenarios, prioritize what to influence first, and target interventions where they are most likely to reduce adversary success.

4 Toward a Behavioral Ontology for Cybersecurity

4.1 The Move Toward an Ontology

As a database, SebDB offers a useful lens for viewing a comprehensive set of security behaviors and how they relate to risks, support protective actions, and enable or disrupt threat activity. While this is valuable as exemplified in Sect. 3, additional benefit could be realized by extending SebDB toward a richer knowledge representation structure, such as an ontology. An ontology can be regarded as a shared vocabulary that defines the concepts in a specific domain and the relationships among them [12,31]. Ontologies are increasingly common in cybersecurity research and practice because they support shared, extensible representations that can be made machine readable. They also provide a semantic foundation that enables reasoning over domain concepts. For instance, an ontology focused on security behaviors and cyber risks could support inferences such as identifying that, if a particular behavior is not performed, then an attack tactic is more likely to succeed, thereby leading to increased probability of a certain risk materializing.

There have been several ontologies proposed for the cybersecurity domain [1]. One under-researched area, however, concerns ontologies that incorporate human aspects, particularly those informed by behavioral science and psychology. Existing work includes an information security ontology that incorporates human behavioral factors and is intended to support security management processes [26]. Another example is a human factors ontology for cybersecurity, which describes how individual and situational characteristics, and the relationships between them, influence the trust placed in an individual within security contexts [24].

A further contribution is an insider risk ontology that models sociotechnical factors across individuals and organizations to support understanding of insider threats [11].

In our work, we build on these ontologies and on relevant ontologies from the behavioral sciences [18] to develop an ontology that connects cybersecurity with behavioral science. This is particularly important given the central role of people's behaviors in holistic security risk management strategies [4,23,28]. Accordingly, our ontology's goal is to provide an evidence-based conceptualization that enables the systematic application of behavioral science to cybersecurity. The ontology defines a shared vocabulary and explicit relationships between behavioral constructs, including interventions, determinants, and behavior change techniques, and the security behaviors, risks, and security controls they interact with. This structure supports its users in designing, describing, and evaluating behaviorally informed approaches to managing human cyber risk, which align with wider security control environments. It also provides a common knowledge base that facilitates communication between security awareness and training and human risk management professionals and more technically focused security teams in relation to security risks. These aims, particularly our focus on connecting behavioral science and cybersecurity, considering the full behavior intervention lifecycle, and supporting security initiatives, set our ontology apart from those explored in the related works mentioned above.

4.2 SebDB as an Ontology

The ontology is structured specifically to fulfill its goal of bridging cybersecurity and behavioral science, and was built following a standard ontology creation methodology [16]. There are over 50 concepts in the ontology spread across six thematic domains. Connecting these concepts are more than 70 relationships, which describe how concepts relate to each other. The thematic domains are: Asset and vulnerability, Threat, Behavior, Behavioral intervention, Risk and outcome, and Security control. The ontology can be visualized in Fig. 1.

To start, the Asset and vulnerability domain covers organizational assets and the weaknesses that they may possess. It contains two concepts: asset and vulnerability, which are associated as follows: a vulnerability *exposes* an asset. For instance, high susceptibility to business email compromise in certain finance departments *exposes* the accounts payable system.

In the Threat domain there are six concepts: threat actor, threat driver, tactic, technique, or procedure (TTP), threat vector, threat, and threat scenario. This domain models the source of an attack or incident, their motives, and approach. It contains several relationships, including: a threat driver *drives* a threat actor; a threat actor *employs* a TTP; a TTP *uses* a threat vector; a threat vector *enables* a threat; a threat actor *initiates* a threat; and a threat *is contextualized by* a threat scenario. An example of how these interact is as follows. The potential for significant financial gain from selling stolen healthcare data *drives* a financially-motivated cybercriminal group, who *employs* spear phishing with malicious attachments *using* email to deliver the attack. The email

channel *enables* a credential harvesting attack, which *is contextualized by* a scenario describing how the group targets accounting staff to obtain credentials and initiate fraudulent transfers.

The Behavior domain is central to the ontology and emphasizes the importance of human actions in influencing cyber risk. The concepts span user, behavior, security behavior, risky behavior, behavior category, and behavior tier. Supporting these concepts are a number of relationships, namely: a user *performs* a behavior; a security behavior and a risky behavior *are types of* behavior; a behavior *is grouped in* a behavior category; and a behavior *is assigned to* a behavior tier. These can be instantiated as follows. An accounts payable clerk *performs* a range of behaviors daily, including reporting suspicious emails to the security team and, on occasion, clicking links without verifying sender authenticity. Reporting suspicious emails to the security team *is grouped in* the Incident response & reporting category and *is assigned to* Tier 3 due to its lower direct influence on reducing risk. Here we also see three concepts that originated in earlier versions of SebDB, i.e., security behavior, behavior category, and tier.

The Behavioral intervention domain features alongside the Behavior domain and encompasses the contexts and activities pertaining to changing people's security-related behaviors. In this domain, the core concepts are: behavioral intervention, behavioral determinant, behavior change technique (BCT), intervention catalyst, directive catalyst, elective catalyst, and corrective catalyst. We also have example/sub-concepts for behavioral intervention (e.g., persuasion, education), behavioral determinant (e.g., psychological capability, reflective motivation), and behavior change technique (e.g., prompt, social comparison) [19]. The relationships for the core concepts are: an intervention catalyst *triggers* a behavioral intervention; directive catalyst, elective catalyst and corrective catalyst *are types of* intervention catalyst; a behavioral intervention *leverages* a behavior change technique; a behavior change technique *targets* a behavioral determinant; and a behavioral determinant *causes changes in* a behavior.

Practically, new regulations mandating security awareness training, a proactive leadership decision to improve security culture, and a mandated remediation program following a phishing incident *are types of* intervention catalyst that can lead to organizational responses. The regulation can also *trigger* the launching of a phishing resilience program, which *leverages* showing employees their click rates compared to departmental averages to *target* employees' overconfidence in their ability to spot phishing emails, thus *causing changes in* their behavior.

Next is the Risk and outcome domain. This domain's concepts include: risk, likelihood, potential impact, incident, risk outcome, and actual impact. Its relationships are as follows: likelihood *contributes to* risk; potential impact *contributes to* risk; risk *materializes* as an incident; an incident *results in* a risk outcome; and a risk outcome *is quantified* by actual impact. When applied in practice: the high probability of a successful phishing attack given current employee susceptibility and the potential for significant regulatory fines if customer data is compromised both *contribute to* the risk facing the organization. This risk can *materialize as* the exposure of a large number of customer records

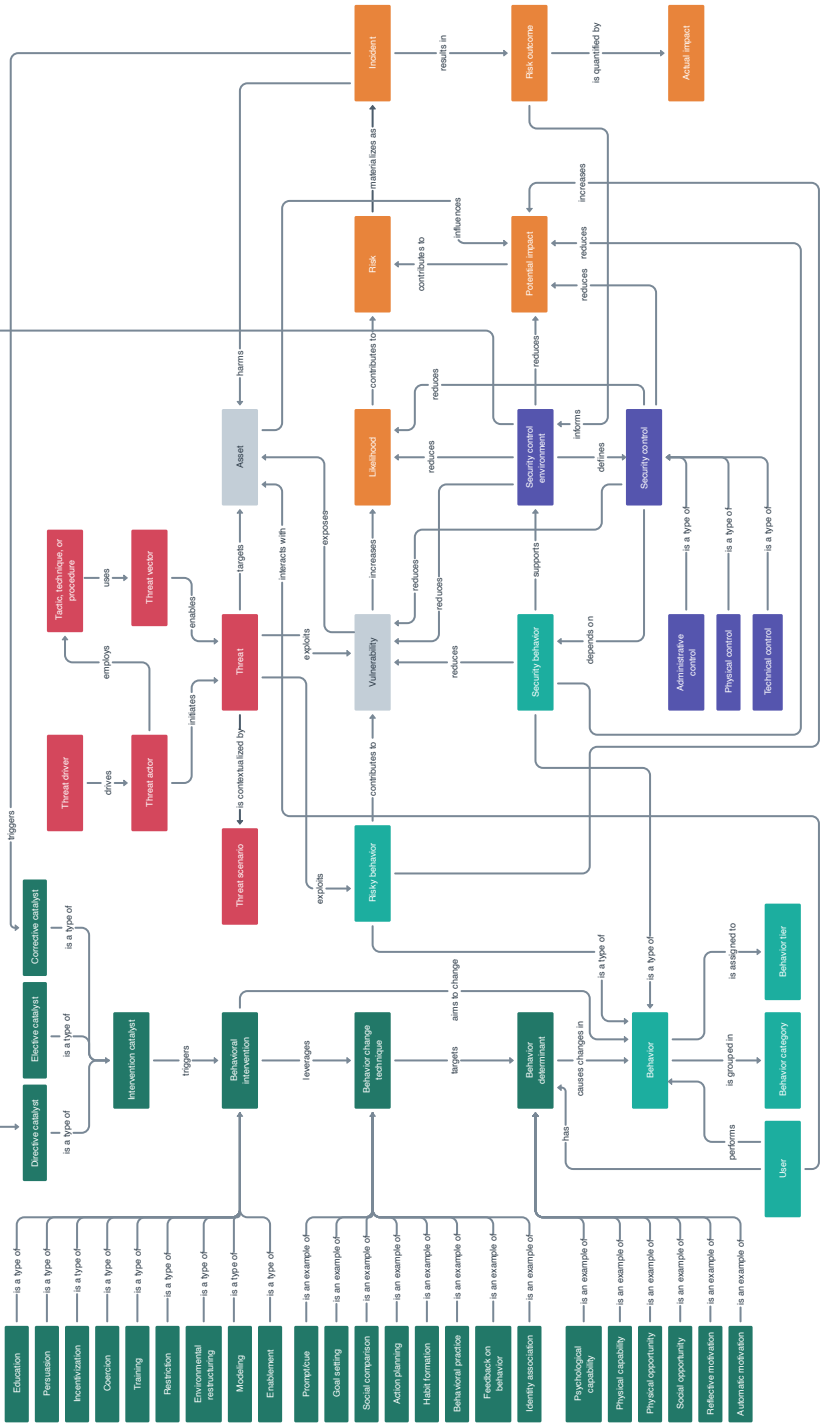


Fig. 1. A behavioral ontology for cybersecurity

via a compromised employee account, which *results in* reputational damage, regulatory penalties, and operational disruption. These outcomes are quantified by a total of £2 million in fines and lost business.

Finally, there is the Security control domain, which defines the measures focused on risk reduction. The concepts are: security control environment, security control, administrative control, physical control, and technical control. Between these concepts are the following relationships: security control environment *defines* a security control; and administrative controls, physical controls, and technical controls *are types of* security controls. A real-world example of these is: an organization's PCI DSS compliance program *defines* the requirement to protect cardholder data. This is implemented by several *types of* controls, e.g., through a combination of the data handling policy requiring encryption for all sensitive data, badge-controlled access to the server room housing encryption key management systems, and encryption applied automatically by the storage platform.

For a full list of the concepts and their definitions, see Table 4.

Table 4. Concept definitions

Concept	Definition
Asset	Any resource of value to the organization that requires protection.
Vulnerability	A weakness that can be exploited by a threat, creating the conditions for harm to occur.
Threat actor	A source of malicious pressure against organizational assets, including external adversaries and malicious insiders.
Threat driver	A condition or motive that increases the likelihood, or amplifies the impact, of threat actor activity.
Tactic, technique, or procedure	The methods employed by threat actors to achieve their objectives, as cataloged in frameworks such as MITRE ATT&CK.
Threat vector	The pathway or channel through which a threat reaches its target.
Threat	Any source of pressure capable of exploiting a vulnerability to cause an incident.
Threat scenario	A structured narrative describing how a specific threat could plausibly unfold against a defined target.
Behavior	An observable action or omission by an individual or other entity that influences security risk.
Security behavior	An action by an individual or other entity that reduces exposure of assets to unauthorized access, damage, or compromise.
Risky behavior	An action by an individual or other entity that increases exposure of assets to unauthorized access, damage, or compromise.
User	A person (e.g., employee, contractor) or entity (e.g., AI) who performs behaviors and interacts with organizational assets.
Behavior category	A thematic classification that groups related behaviors based on shared characteristics.
Behavior tier	A prioritization classification that ranks behaviors according to their influence on risk.

(*continued*)

Table 4. (*continued*)

Concept	Definition
Behavioral intervention	A deliberate effort to modify behavior through mechanisms such as environmental design, process change, or persuasion (e.g., nudges).
Behavior change technique	A discrete, evidence-based method for influencing behavior, representing the active component within an intervention.
Behavior determinant	A modifiable factor that influences whether and how a behavior occurs, corresponding to constructs within the Capability, Opportunity, Motivation, and Behavior (COM-B) model or the Theoretical Domains Framework (TDF).
Intervention catalyst	The triggering condition that initiates a behavioral intervention, shaping its design and urgency.
Elective catalyst	A catalyst arising from proactive organizational choice, motivated by improvement rather than external pressure.
Directive catalyst	A catalyst arising from imposed requirements such as regulatory mandates, policy changes, or leadership directives.
Corrective catalyst	A catalyst arising from an observed failure, incident, or identified gap requiring remediation.
Likelihood	The probability of an incident occurring.
Potential impact	The estimated severity of harm if an incident were to occur.
Risk	The possibility that an incident will occur and cause harm. It is the combination of likelihood and potential impact.
Incident	A security event that occurred when risk materialized, causing harm or requiring response. For instance, a data breach, malware infection, account compromise, system compromise or policy violation.
Risk outcome	The organizational harm resulting from an incident. Examples are regulatory penalties, reputational damage, operational disruption.
Actual impact	The measured severity of harm after an incident occurred.
Security control environment	The integrated architecture of controls that collectively reduces vulnerability, likelihood, and impact.
Security control	A safeguard that reduces the vulnerability, likelihood or impact of a security risk through prevention, detection, or response.
Administrative control	A control implemented through governance mechanisms such as policies, procedures, and standards.
Physical control	A control implemented through tangible measures that restrict or monitor physical access.
Technical control	A control implemented through technological mechanisms that enforce security requirements.

Thus far we have covered the core concepts of the ontology and a sample of the relationships present. As can be seen in Fig. 1 and understood given the wide remit of the ontology, there are several more relationships that exist. In particular, these act to connect the different domains and thereby allow the ontology to fulfill its interdisciplinary and cross-enterprise goals. Due to space limitations, we are not able to discuss all of these relationships, but there are three relationship chains worth highlighting given their significance.

The first chain links the Threat, Behavior, Asset and vulnerability, and Risk and outcome domains. This chain demonstrates that threats exploit human behavior, not only technical weaknesses. Traditional security frameworks often treat threat exploitation as primarily technical (e.g., exploiting a software vulnerability). These relationships make explicit that human actions are themselves exploitable attack surfaces, and traces the full causal pathway from threat through behavior to organizational risk. The chain is: a business email compromise attack [threat] *exploits* employees' tendency to comply with urgent payment requests without verification [risky behavior], which *contributes to* weak financial transaction controls [vulnerability] *increasing* the probability of successful fraud [likelihood], thereby *contributing to* the overall financial crime risk [risk].

The second chain comprises the Behavior, Security control, and Risk outcome domains. Here we reveal how security behavior influences risk through two distinct pathways: by supporting the organization's defensive architecture, and by directly reducing the potential severity of harm should an incident occur. This dual pathway demonstrates that behavior is not merely a component of security controls but an independent factor in risk reduction. This is a critical insight often overlooked when behavioral security is treated as subordinate to technical measures. The chain is: the practice of employees maintaining verified offline backups of critical data [security behavior] *supports* the organization's business continuity program [security control environment], while also *reducing* the potential operational disruption from ransomware [potential impact] that *contributes to* the overall ransomware risk [risk].

The third chain captures the organizational learning loop, i.e., how security failures drive behavioral improvement. It spans the Risk and outcome, Behavioral intervention, and Behavior domains. This chain demonstrates that the ontology is not purely descriptive but actionable: incidents are not endpoints but triggers for systematic behavioral intervention design. This relationship connects reactive incident response to behavior change initiatives, demonstrating how the ontology supports continuous improvement. The chain is: the discovery that employees pasted confidential client data into a public generative AI tool, resulting in data exposure [incident] *triggers* a mandatory review of AI usage practices [corrective catalyst], which then *triggers* the implementation of an AI security program [behavioral intervention]. The program *leverages* just-in-time nudges that remind employees to check data sensitivity before pasting content into AI tools [behavior change technique], *targeting* employees' automatic habits when interacting with AI tools [behavior determinant]. This *causes changes in* (ideally increasing) the practice of verifying data classification before sharing with AI [security behavior].

These examples demonstrate the utility of the ontology in supporting organizations, security professionals, and researchers to reason about and address behavioral security and human-centered cybersecurity. As with previous iterations of SebDB, we expect this ontology to evolve as the field matures.

5 Conclusion

This paper has made two contributions to advance existing knowledge of how human behavior relates to cybersecurity. First, we introduced SebDB, an open-source resource designed to improve the understanding, modeling, and measurement of human behaviors and risks in cybersecurity. Second, we presented our vision for SebDB as a behavioral ontology for cybersecurity. Building on these foundations, our next step is to formalize this ontology into a model that supports not only explanation but also formal reasoning across the behavioral security domain.

Acknowledgements. The authors would like to thank Inka Karppinen, John Blythe, Joe Giddens, Veronika Bondareva, Ben Donaldson, and their respective teams, who contributed significantly to the development of SebDB. We also thank all those whose valuable feedback helped to refine this important resource.

References

1. Adach, M., Hänninen, K., Lundqvist, K.: Security ontologies: a systematic literature review. In: International Conference on Enterprise Design, Operations, and Computing, pp. 36–53. Springer, Cham (2022)
2. Alkhazi, B., Alshaikh, M., Alkhezi, S., Labbaci, H.: Assessment of the impact of information security awareness training methods on knowledge, attitude, and behavior. IEEE Access **10**, 132132–132143 (2022)
3. Baltuttis, D., Teubner, T., Adam, M.T.: A typology of cybersecurity behavior among knowledge workers. Comput. Secur. **140**, 103741 (2024)
4. Chaudhary, S.: Driving behaviour change with cybersecurity awareness. Comput. Secur. 103858 (2024)
5. Chaudhary, S., Gkioulos, V., Katsikas, S.: Developing metrics to assess the effectiveness of cybersecurity awareness program. J. Cybersecurity **8**(1) (2022)
6. Chiasson, S., van Oorschot, P.C., Biddle, R.: A usability study and critique of two password managers. In: USENIX Security Symposium, vol. 15, pp. 1–16 (2006)
7. Egelman, S., Peer, E.: Scaling the security wall: Developing a security behavior intentions scale (SEBIS). In: ACM CHI, pp. 2873–2882 (2015)
8. Fertig, T., Schütz, A.E., Weber, K.: Current issues of metrics for information security awareness. In: ECIS, vol. 11, pp. 19–20 (2020)
9. Främling, K., Harrison, M., Brusey, J., Petrow, J.: Requirements on unique identifiers for managing product lifecycle information: comparison of alternative approaches. IJCIM **20**(7), 715–726 (2007)
10. Georgiadou, A., Mouzakitis, S., Askounis, D.: Assessing mitre att&ck risk using a cyber-security culture framework. Sensors **21**(9), 3267 (2021)
11. Greitzer, F., Purl, J., Leong, Y.M., Becker, D.S.: Sofit: Sociotechnical and organizational factors for insider threat. In: IEEE SPW, pp. 197–206. IEEE (2018)
12. Gruber, T.R.: A translation approach to portable ontology specifications. Knowl. Acquis. **5**(2), 199–220 (1993)
13. Hadlington, L.: Human factors in cybersecurity; examining the link between internet addiction, impulsivity, attitudes towards cybersecurity, and risky cybersecurity behaviours. Heliyon **3**(7) (2017)

14. Hadlington, L.: Employees attitude towards cyber security and risky online behaviours: an empirical assessment in the united kingdom. Int. J. Cyber Criminol. **12**(1), 269–281 (2018)
15. Jangampeta, S., Makani, S.T.: Optimizing security controls: aligning threat intelligence with nist csf for effective decision-making. J. ID **9471**, 1297 (2022)
16. King, B., Reinold, K.: Finding the concept, not just the word: a librarian's guide to ontologies and semantics. Elsevier (2008)
17. Kirlappos, I., Parkin, S., Sasse, M.A.: "shadow security" as a tool for the learning organization. ACM SIGCAS **45**(1), 29–37 (2015)
18. Marques, M.M., et al.: The behaviour change technique ontology: transforming the behaviour change technique taxonomy v1. Wellcome Open Res. **8**, 308 (2024)
19. Michie, S., Van Stralen, M.M., West, R.: The behaviour change wheel: a new method for characterising and designing behaviour change interventions. Implement. Sci. **6**(1), 42 (2011)
20. Mott, G., et al.: 'There was a bit of PTSD every time I walked through the office door': Ransomware harms and the factors that influence the victim organization's experience. J. Cybersecurity **10**(1) (2024)
21. Nurse, J.R.C.: Cybercrime and you: How criminals attack and the human factors that they seek to exploit. In: The Oxford Handbook of Cyberpsychology, pp. 663–690. Oxford University Press (2019)
22. Nurse, J.R.C., Creese, S., Goldsmith, M., Lamberts, K.: Guidelines for usable cybersecurity: Past and present. In: Cyberspace Safety and Security Workshop (2011)
23. Nurse, J.R.C., Milward, J., Alashe, O.: From security awareness and training to human risk management in cybersecurity. In: International Conference on Human-Computer Interaction, pp. 86–104. Springer, Cham (2025)
24. Oltramari, A., Henshel, D.S., Cains, M., Hoffman, B.: Towards a human factors ontology for cyber security. Stids **2015**, 26–33 (2015)
25. Ortloff, A.M., et al.: Replication: "no one can hack my mind" - 10 years later: An update and outlook on experts' and non-experts' security practices and advice. In: SOUPS, pp. 435–454 (2025)
26. Parkin, S.E., van Moorsel, A., Coles, R.: An information security ontology incorporating human-behavioural implications. In: Proceedings of the 2nd International Conference on Security of Information and Networks, pp. 46–55 (2009)
27. Parsons, K., Calic, D., Pattinson, M., Butavicius, M., McCormac, A., Zwaans, T.: The human aspects of information security questionnaire (HAIS-Q): Two further validation studies. Comput. Secur. **66**, 40–51 (2017)
28. Prümmer, J., van Steen, T., van den Berg, B.: Assessing the effect of cybersecurity training on end-users: a meta-analysis. Comput. Secur. **150**, 104206 (2025)
29. Redmiles, E.M., et al.: A comprehensive quality evaluation of security and privacy advice on the web. In: USENIX Security, vol. 20, pp. 89–108 (2020)
30. Stanton, J.M., Stam, K.R., Mastrangelo, P., Jolton, J.: Analysis of end user security behaviors. Comput. Secur. **24**(2), 124–133 (2005)
31. Studer, R., Benjamins, V.R., Fensel, D.: Knowledge engineering: principles and methods. Data Knowl. Eng. **25**(1–2), 161–197 (1998)
32. Sutton, A., Tompson, L.: Towards a cybersecurity culture-behaviour framework: a rapid evidence review. Comput. Secur. **148**, 104110 (2025)
33. Uchendu, B., Nurse, J.R.C., Bada, M., Furnell, S.: Developing a cyber security culture: Current practices and future needs. Comput. Secur. **109** (2021)
34. West, R.: The psychology of security. Comms. ACM **51**(4), 34–40 (2008)

35. Zimmermann, V., Renaud, K.: Moving from a 'human-as-problem' to a 'human-as-solution' cybersecurity mindset. IJHCS **131**, 169–187 (2019)
36. Zwilling, M., Klien, G., Lesjak, D., Wiechetek, Ł, Cetin, F., Basim, H.N.: Cyber security awareness, knowledge and behavior: a comparative study. J. Comput. Inf. Syst. **62**(1), 82–97 (2022)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

