

OPERATIONAL TECHNOLOGY IN CONSTRUCTION: A CYBERSECURITY TAXONOMY

SUBMITTED: May 2026

PUBLISHED: October 2026

EDITOR: Farzad Rahimian

DOI: [10.36680/j.itcon.2026.052](https://doi.org/10.36680/j.itcon.2026.052)

Muammer Semih Sonkor, Postdoctoral Associate

Division of Engineering, New York University Abu Dhabi, United Arab Emirates

<https://orcid.org/0000-0002-6851-2462>

semih.sonkor@nyu.edu

Borja García de Soto, Associate Professor

Division of Engineering, New York University Abu Dhabi, United Arab Emirates

<https://orcid.org/0000-0002-9613-8105>

garcia.de.soto@nyu.edu

SUMMARY: Construction projects are increasingly adopting operational technologies (OT), such as autonomous machinery, robotic manipulators, drones, and 3D concrete printers. While these technologies improve productivity, quality, and safety, they also expand the cyberattack surface of construction sites. Although many cybersecurity standards, frameworks, and guidelines have been developed, they have been used primarily in more technologically advanced sectors, such as manufacturing and energy. Since OT integration is relatively recent in the construction industry, identifying the relevant documents for a construction OT scenario requires understanding the construction process, the OT family, and the system components involved. To address this need and given the increasing number of available documents, this paper proposes a novel taxonomy that maps cybersecurity documents to OT used in construction. The taxonomy is structured as a two-dimensional table, with columns representing construction OT families and components, and rows representing construction processes involving OT. Each cell in the table contains the cybersecurity references for the OT–process combination. To populate the taxonomy, this paper proposes a human-in-the-loop workflow combining large language model (LLM)-assisted retrieval, quality control, deduplication, scoring, and manual review. The workflow is demonstrated for a single taxonomy cell (the communication component of robotic manipulators used in masonry and bricklaying), together with a parallel workflow for construction-specific cybersecurity references. By supporting secure OT-enabled construction, resilient digital infrastructure, and safer construction work environments, the proposed approach aligns with the United Nations Sustainable Development Goals (SDGs), particularly SDG 9 (Industry, Innovation and Infrastructure), SDG 11 (Sustainable Cities and Communities), and SDG 8 (Decent Work and Economic Growth).

KEYWORDS: construction 4.0, cybersecurity, cybersecurity frameworks, operational technology, taxonomy.

REFERENCE: Sonkor, M. S., & García de Soto, B. (2026). Operational technology in construction: A cybersecurity taxonomy. *Journal of Information Technology in Construction (ITcon)*, Special issue: 'Digital Twins for Smart and Sustainable Built Environments: Perspectives from SASBE 2025', 31, 1245-1269. <https://doi.org/10.36680/j.itcon.2026.052>

COPYRIGHT: © 2026 The author(s). This is an open access article distributed under the terms of the Creative Commons Attribution 4.0 International (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

The built environment lifecycle is increasingly being transformed by the rapid development of technologies used in design, construction, operations, and facility management. These technologies are integrated into daily tasks through digital and connected systems and tools such as building information modeling (BIM), common data environments (CDEs) for information exchange (Brelh and Klinc, 2025), and building management systems (BMSs) to monitor and control operational built environments. During the construction phase, this transformation is reflected in the digitalization of information and the automation of physical tasks, where robots and autonomous machinery are increasingly proposed or used to perform activities conventionally carried out by humans (Sawhney *et al.*, 2020). These digitalized and automated construction operations rely on both information technologies (IT) and operational technologies (OT), which have historically been treated as separate domains. However, the increasing need for improved control, interconnectivity, remote monitoring, and better operational performance has made IT/OT convergence unavoidable (Harp and Gregory-Brown, 2015). The convergence of these two domains increases technological heterogeneity and expands the cyber threat surface, making robust cybersecurity more difficult to achieve (Givehchi *et al.*, 2017). Therefore, construction sites that integrate OT and automation face similar challenges (García de Soto *et al.*, 2021).

Almost every industry is affected by increasing digitalization and, as a result, faces heightened cyber risk. To address such challenges, governmental institutions and private organizations have been developing cybersecurity documents varying from standards to frameworks and guidelines. These documents often target a specific technology or sector. An example of a technology-specific standard is the ISA/IEC 62443 series (ISA/IEC, 2026), developed by the International Society of Automation (ISA) and the International Electrotechnical Commission (IEC), which focuses on cybersecurity for industrial automation and control systems. On the other hand, ISO/SAE 21434 (ISO/SAE, 2021), developed by the International Organization for Standardization (ISO) and the Society of Automotive Engineers (SAE), presents cybersecurity requirements for road vehicles, positioning it as a sector-specific document (i.e., the automotive sector).

In recent years, several standards and frameworks have been published to address cybersecurity challenges in construction, and, more generally, in the built environment sector. For example, ISO 19650-5 (ISO, 2020) presents a security-minded approach to information management in BIM-based projects. Similarly, the code of practice, “Cyber Security in the Built Environment”, by the Institution of Engineering and Technology (IET) (2021) provides management teams with guidance on protecting built assets across lifecycle stages. These documents are useful, as they establish sector-relevant security principles; however, their scope is often broad and does not directly target individual construction phases or specific OT types. In particular, these documents primarily emphasize organizational governance and information security rather than the technical and operational controls required for cyber-physical systems (CPSs) that interact with construction equipment, site processes, and workers. They lack the specialized guidance required to safeguard complex construction OT with safety-critical functions and stringent availability requirements. Therefore, construction projects that introduce autonomous machinery, robotics, or other OT systems require the use of cybersecurity documents that may not be specifically targeted to the construction sector. This is crucial for securing digitalized and interconnected construction sites against the evolving cyber threat landscape.

As OT becomes more common in construction applications, including autonomous machinery, drones, robotic manipulators, and 3D concrete printers, selecting appropriate cybersecurity references requires more than a general understanding of cybersecurity. It also requires knowledge of the supported construction activity, the OT family being used, and the specific system components involved. To address this need, and given the large number of available guidelines, control catalogs, standards, and frameworks, this paper proposes a taxonomy that maps cybersecurity documents to relevant construction processes and OT types. This taxonomy provides construction teams with a structured approach to identifying the cybersecurity references applicable to their specific cases. As a result, the paper aims to support the broader adoption of cybersecurity documents within the construction sector. This paper is an extended version of the previously published conference paper by Sonkor and García de Soto (2026). Compared to the conference paper, this paper revises and expands the methodology by introducing a human-in-the-loop LLM-assisted web search workflow for document retrieval, quality control, deduplication, and scoring, and demonstrates the workflow through a taxonomy cell-specific implementation and a parallel construction-specific workflow. Therefore, the methodology, implementation, and discussion sections were significantly revised.

2. BACKGROUND

2.1 Operational technology in the construction sector

According to the National Institute of Standards and Technology (NIST), OT includes programmable devices and systems that monitor, control, or directly affect the physical environment (Stouffer *et al.*, 2023). This definition is relevant to construction sites, where OT is increasingly used to automate repetitive physical tasks, monitor activities or worker behavior, support quality control, and produce structures using systems such as 3D concrete printers. These construction technologies either monitor the physical environment or cause a change in it, which aligns with NIST's definition of OT. Their cyber counterparts, IT components, store, exchange, and process data to support the same activities. The integration of IT and OT layers creates CPSs, which are systems of systems, in which sensing, computation, communication, and actuation connect digital processes with the physical environment in real time (Griffor, 2017).

Examples of repetitive construction tasks that can be automated by OT include earthwork tasks such as excavation (Seo *et al.*, 2011), ground leveling (Jeon *et al.*, 2025), and compaction (Schöberl *et al.*, 2024). Manufacturers such as Hitachi, Komatsu, and Caterpillar, as well as start-ups like Teleo and Lumina, have been developing autonomous earthmoving equipment. The deployment of such machinery has primarily been limited to mining environments, which are relatively controlled, involve fewer humans, and require less dynamic tasks. Another increasingly popular construction task that utilizes OT is 3D concrete printing. Its popularity has been increasing exponentially in recent years due to its advantages, such as its ability to create complex geometries and reduced time and material requirements (Tay *et al.*, 2017). Other use cases for OT in construction projects include automated safety inspections (Johansen *et al.*, 2024), progress monitoring using autonomous robots (Prieto *et al.*, 2020), quality inspection utilizing drones (Obradović *et al.*, 2019), and wall frame installation (Oyediran *et al.*, 2025).

The common objective across these applications is to make construction operations more efficient, accurate, and safe, particularly by reducing the need for workers to perform repetitive or hazardous tasks directly. At the same time, the use of connected OT introduces new cybersecurity concerns because construction equipment and robotic systems depend on sensors, controllers, actuators, communication links, and data-processing pipelines. Cyber incidents reported in the construction sector have mainly affected digital information and business processes rather than site-level OT systems. For example, the engineering firm Arup recently became the victim of a deepfake attack in which attackers convinced an employee to transfer \$25 million (Elliott, 2025). At the beginning of 2025, DragonForce, a ransomware-as-a-service (RaaS) group, claimed that it had stolen 6 TB of data from the Saudi contractor Al Bawani (Lemos, 2025). These incidents suggest that the sector is becoming a more visible target for malicious actors. As OT deployment increases, similar attention may extend to construction technologies that interact with the physical site environment. This is particularly important because compromised OT can affect nearby workers, equipment, and ongoing site operations. Despite this concern, research on OT cybersecurity in construction remains limited (Salami Pargoo and Ilbeigi, 2023).

2.2 Cybersecurity documents

The development of cybersecurity guidance has been shaped by major cyber incidents and the increasing dependence of organizations on connected systems. One early example was the Morris Worm, which spread across the internet in 1988 and brought wider attention to cybersecurity (Orman, 2003). As cyber threats became more frequent and complex, public institutions, standards bodies, and industry organizations began developing documents to guide the protection of digital assets. In the US, National Security Decision Directive Number 145 (NSDD-145), issued under the Reagan administration in 1984, gave the National Security Agency (NSA) a central role in preparing cybersecurity-related standards and guidance and supporting research in the field (Warner, 2012). In the UK, British Standard 7799, titled “Code of Practice for Information Security Management,” was released in 1995 and provided structured guidance for managing information security (Humphreys, 2008).

Cybersecurity documents can be legally binding or voluntary, and they are developed for different purposes and target groups. Some examples of mandated ones are laws and regulations, such as the General Data Protection Regulation (GDPR) (EU, 2016) and the NIS2 Directive (EU, 2022), with which organizations operating within certain jurisdictions must comply. Documents, such as cybersecurity frameworks, control catalogs, guidelines, and codes of practice, are voluntarily used by organizations to improve their security levels. The Cybersecurity Framework 2.0 (NIST, 2024) and SP 800-82, Guide to OT Security (Stouffer *et al.*, 2023), both published by NIST,

are among the most commonly used in the US. Standards are generally voluntary, but they can become legally binding when mandated by governments or incorporated into sector-specific regulations. For instance, the Federal Energy Regulatory Commission (FERC) mandates the North American Electric Reliability Corporation Critical Infrastructure Protection (NERC-CIP) series of standards for all bulk electric system operators in the US (NERC, 2025). Some other cybersecurity document types are maturity models, baselines, and playbooks, which serve different roles in safeguarding cyber-connected assets.

2.3 Cybersecurity taxonomies

Taxonomies hierarchically organize concepts within a domain into structured categories using relationships to show how broader and narrower concepts are connected (IEEE, 1987). In cybersecurity, most taxonomies have focused on organizing terminology or categorizing attacks, risks, vulnerabilities, and threats. Some examples from academic publications are the ransomware taxonomy proposed by Al-Rimy et al. (2018), the cybercrime taxonomy by Brar and Kumar (2018), the IoT cyberattack taxonomy by Xenofontos et al. (2021), and the cyber risk treatment taxonomy proposed by Sánchez-García et al. (2025). Kalogeraki and Polemi (2024) proposed a taxonomy of cybersecurity standards with a motivation similar to that of this paper: to support stakeholders in determining which documents to use under different conditions. In addition to academic taxonomies, public and private entities have developed taxonomies widely used by cybersecurity experts. For example, the MITRE ATT&CK framework (MITRE, 2020) classifies adversary behavior based on procedures, tactics, and techniques across technology platforms and attack stages. Similarly, Common Attack Pattern Enumeration and Classification (CAPEC) (MITRE, 2025), also developed by MITRE, hierarchically organizes known cyberattack patterns.

Another crucial aspect of taxonomies is the way they are generated. Based on a systematic literature review, Spyros et al. (2025) categorized cybersecurity taxonomies into two groups depending on how they are generated: manual and dynamic. Manual taxonomies rely on human domain expertise, such as that of researchers and professionals, while dynamic taxonomies use automated or semi-automated approaches. Most cybersecurity taxonomies have traditionally been developed manually, prior to the broader adoption of natural language processing (NLP) methods. Examples include the advanced persistent threat actors taxonomy proposed by Bahrami et al. (2019) and the cyber threat taxonomy developed by Ferdinand and Benham (2017). For dynamic taxonomy development, Kougoumtzidou et al. (2024) introduced an AI-based framework that applies Named Entity Recognition (NER) and pattern matching to generate and maintain cybersecurity taxonomies.

The taxonomy proposed by Kalogeraki and Polemi (2024) is the most closely related work; however, it differs from this paper's approach in three important respects. First, it is a cross-sector taxonomy for navigating cybersecurity documents, whereas the taxonomy in this paper is specific to construction and particularly related to OT. Second, their semantic ontology organizes documents through four attributes: stakeholders, sectors, security domains, and technologies. On the other hand, the taxonomy presented in this study maps references to the intersections of construction OT families, OT components, and OT-enabled construction processes. Finally, Kalogeraki and Polemi's taxonomy does not employ a systematic methodology to identify relevant standards, frameworks, and best practices. In contrast, this paper proposes a dynamic, semi-automated workflow for retrieving, prioritizing, and manually validating cybersecurity references using LLMs.

3. METHODOLOGY

The proposed methodology (see Figure 1) is designed as a human-in-the-loop process in which LLMs support document retrieval, organization, and scoring, while human review is retained to reduce unreliable or misleading outputs. This is similar to the approach suggested by Mahaini et al. (2019), which involves human-machine teaming to build taxonomies. Also, since the proposed methodology is semi-automated, it can be considered dynamic, as described by Spyros et al. (2025), using automation to improve efficiency, adaptability, and scalability while still relying on manual human review. These characteristics are crucial because the number of cybersecurity documents continues to grow, and LLMs are advancing rapidly. This paper presents the proposed taxonomy as a table with two dimensions: (1) columns - OT families and their components, and (2) rows - construction processes. The columns include two tiers: Tier 1 for broad OT families used in construction and Tier 2 for specific components within each family. Finally, each cell in the table contains relevant cybersecurity documents. Table 1 demonstrates the proposed taxonomy table structure. The following subsections describe the criteria to define the columns and rows and to identify the cybersecurity references to be included.




Figure 1: Human-in-the-loop workflow for populating the proposed taxonomy.

Table 1: Schematic representation of the proposed taxonomy structure.

Construction process	OT Family 1			OT Family 2			...
	OT Component 1	OT Component 2	...	OT Component 1	OT Component 2	...	
Process 1	Documents	Documents	...	Documents	Documents	...	...
Process 2	Documents	Documents	...	Documents	Documents	...	...
...	...	...	...	...	...	...	...

3.1 Identifying OT families and components (columns)

Tier 1 – Construction OT Families: Although automation and robotics have been proposed in construction research for many years, not all proposed systems have been widely deployed on real-world construction sites. Therefore, academic resources are more suitable than industry examples for generating a comprehensive list of OT types used in construction. Accordingly, relevant publications can be identified through a literature search using keyword combinations that capture both construction and OT concepts.

Since many academics have already published review papers on construction automation and robotics and identified relevant categories, this paper searched for such publications and used them to derive the OT families. Scopus was selected as the citation database for this search because of its broad coverage of peer-reviewed literature and reliable indexing of high-quality journal articles, conference papers, and book chapters. This search used the proximity operator “Within” (i.e., “W”) to combine keyword groups and limit results to relevant publications. The literature-search criteria are summarized below:

- Search field: article titles, abstracts, and keywords;
- Search string: ("construction" W/3 ("robot*" OR "automation" OR "operational technolog*")) W/10 ("trends" OR "advances" OR "challenges" OR "review" OR "survey" OR "research areas");
- Publication period: From 2021 to the end of 2025;
- Eligible publication types: conference papers, journal articles, and book chapters;
- Publication language: English.

The keywords "trends," "challenges," "advances," "survey," "review," and "research areas," included in the search string, were used to identify publications that summarize the developments and advances in construction OT. The initial search identified 305 publications. These publications were first screened based on their titles and abstracts to identify the ones that provide a general overview of OT research related to construction rather than a specific type of OT. After this preliminary screening, 32 publications remained. Then, these publications were reviewed in full to determine whether they were eligible for inclusion. Only publications that categorized OT based on

operational characteristics and hardware, rather than only based on the tasks performed, were included in the final publication list. Three publications, namely Bogue (2025), Gao et al. (2025), and Wu et al. (2024), met this criterion. Wu et al. (2024) analyzed 498 publications. They identified several categories of construction robots based on their applications and technologies, including field robots (e.g., in-pipe inspection robots), mobile robots (e.g., 3D printers, indoor mapping robots), bolting robots, climbing robots, and swarm robots. Gao et al. (2025) reviewed 64 publications and classified construction robots into six main categories, namely autonomous mobile robots, quadruped robots, robotic manipulators, unmanned aerial vehicles, autonomous construction machinery, and automated guided vehicles. Unlike the other two publications, Bogue (2025) focused on robotics solutions developed by commercial organizations for the construction industry. He discussed these industry applications across six categories: robotic ground vehicles, drones, robotic 3D printers, welding robots, bricklaying robots, and others. In the “others” category, he presented robotic micro-factories, powered exoskeletons, and glazing robots.

In this paper, cybersecurity documents are mapped to construction OT types (columns) and construction processes (rows). To reduce duplicate documents across columns, the OT categories should represent distinct groups based on hardware configuration, mobility, and mode of operation. Keeping this in mind, the categories from the three reviewed publications were assessed and consolidated into the eight groups shown in Table 2.

Table 2: Construction OT families used in the proposed taxonomy, adapted and revised from Sonkor and García de Soto (2026).

Construction OT Family	Definition	Connection to Reviewed Papers
Autonomous heavy construction machinery	Large construction equipment used for earthmoving or material handling, including excavators, dozers, loaders, and similar machines.	Corresponds to “field robots” by Wu <i>et al.</i> (2024) and “autonomous construction machinery” by Gao <i>et al.</i> (2025).
Ground mobile robots	Small or medium-sized wheeled or tracked robotic platforms, commonly ROS-based, that use sensors for navigation, mapping, and site-level mobility.	Combines “mobile robots” by Wu <i>et al.</i> (2024), “autonomous mobile robots” and “automated guided vehicles” by Gao <i>et al.</i> (2025), and “robotic ground vehicles” and “swarm robots” by Bogue (2025).
Legged robots	Mobile robotic systems that use electromechanical legs to move across the construction environment.	Aligned with “quadruped robots” by Gao <i>et al.</i> (2025).
Aerial robots (UAVs)	Unmanned aircraft systems that can be remotely operated or fly autonomously and are available in forms such as multirotor or fixed-wing.	Covers “unmanned aerial vehicles” by Gao <i>et al.</i> (2025) and “drones” by Bogue (2025).
Climbing robots	Robotic systems designed to move on vertical, inclined, or overhead surfaces using adhesion mechanisms such as suction, magnets, or other attachment methods.	Covers “climbing robots” by Wu <i>et al.</i> (2024) and partially covers “bolting robots” by Wu <i>et al.</i> (2024) that have climbing capabilities.
Robotic manipulators	Fixed or mobile robotic arm systems that perform construction tasks through programmable manipulation, positioning, or tool operation.	Partially covers “bolting robots” by Wu <i>et al.</i> (2024) that use robotic arms, covers “robotic manipulators” by Gao <i>et al.</i> (2025), and “bricklaying robots” and “welding robots” discussed by Bogue (2025).
Construction 3D printers	Large-scale additive manufacturing systems that use gantry-based, cable-suspended, or arm-mounted extrusion mechanisms to print structural elements or full-scale components.	Partially covers “mobile robots” by Wu <i>et al.</i> (2024), which includes 3D printers, and covers the “robotic 3D printers” category by Bogue (2025).
Wearable/assistive robotics	Robotic systems worn by workers, such as partial or full exoskeletons, that provide powered assistance during physical tasks.	Partially covers the “others” category by Bogue (2025), which includes “powered exoskeletons”.

Tier 2 – OT Components: ISA-95 is a standards series that provides a structure for implementing control systems in manufacturing environments. Since a comparable reference model for organizing OT in construction is not available, this paper uses selected parts of the ISA-95 hierarchy (ISA, 2025) as a basis and modifies them to align with the proposed taxonomy. This modification was necessary because relevant cybersecurity references may vary depending on the OT component under consideration. In ISA-95, Level 1 refers to “sensing and manipulating the production process.” For the purpose of this paper, this level was separated into two components: sensing and actuation. Level 2 of ISA-95 is “monitoring and supervising control,” which was also separated into two components: human-machine interface (HMI), and control and compute. In addition, “communication” was included as a separate component to cover networking, connectivity, and data exchange among the OT system layers. These components, along with some examples of each, are presented in Table 3.

3.2 Identifying construction processes using OT (rows)

The identification of construction processes followed the same literature search approach used to define the OT families in the Tier 1 taxonomy columns. However, the screening focus was different at this stage. Instead of looking for publications that categorize OT systems by their hardware and operational characteristics, the review focused on publications that describe the construction tasks and processes in which OT is used. Since review papers on construction OT often enumerate such construction tasks, the 32 publications retained after the preliminary title and abstract screening were reviewed again for this purpose. Publications focused on enabling technologies or on a single OT type were not included. Following the full-text review, five publications were selected: Attalla et al. (2023), Gao et al. (2025), Liu et al. (2024), Wang et al. (2024), and Zeng et al. (2024). The construction processes reported in these studies were then consolidated into the list presented in Table 4.

Table 3: OT components and examples.

OT component	Examples
Sensing	Cameras, Light Detection and Ranging (LiDAR), inertial measurement units (IMUs), sensors
Actuation	Actuators, hydraulic or electric drives, extruder pumps for 3D printers
Control and compute	programmable logic controllers (PLCs), controllers
HMI	HMI panels, web user interface (UIs)
Communication	Fieldbus, Wi-Fi 6, 4G/5G

Table 4: Construction processes utilizing OT.

Construction Process	References
Site preparation and earthworks	(Attalla et al., 2023; Gao et al., 2025; Liu et al., 2024; Wang et al., 2024)
Surveying	(Gao et al., 2025)
Material handling	(Attalla et al., 2023; Gao et al., 2025; Wang et al., 2024)
Concrete placement and finishing	(Gao et al., 2025; Liu et al., 2024; Wang et al., 2024)
Reinforcement and structural steel assembly	(Gao et al., 2025; Liu et al., 2024; Wang et al., 2024)
Prefabrication and modular assembly	(Gao et al., 2025; Zeng et al., 2024)
Additive manufacturing (3D printing)	(Gao et al., 2025; Wang et al., 2024; Zeng et al., 2024)
Masonry and bricklaying	(Wang et al., 2024; Zeng et al., 2024)
Finishing works	(Attalla et al., 2023; Gao et al., 2025; Wang et al., 2024; Zeng et al., 2024)
Progress, quality, and safety monitoring	(Gao et al., 2025; Liu et al., 2024; Wang et al., 2024)
Recycling and sorting construction waste	(Gao et al., 2025)

3.3 Retrieving candidate documents

The earlier steps of the methodology defined the OT families and components used as taxonomy columns, as well as the construction processes used as taxonomy rows (see Table 1). The purpose of this step is to identify candidate cybersecurity documents for each cell in the final taxonomy table, which consists of these rows and columns. Because standards, frameworks, guidelines, and control catalogs are continually published and updated, manually finding the relevant documents for each OT type and construction process is inefficient and impractical. For this reason, this paper uses the advanced reasoning, retrieval, and web search capabilities of LLMs to identify candidate documents for the taxonomy cells in a semi-automated manner.

This paper employed an LLM-assisted, iterative web search workflow to generate a candidate reference pool for each target cell. An LLM prompt was designed for this purpose. Primary references were defined as the following categories: (1) regulation/directive, (2) standard/code of practice, (3) guidance/framework, and (4) control catalog/benchmark. Sources that do not fall under one of these categories, such as news pages, announcements, blog posts, or newsletters, were excluded. In the output, document type was accompanied by structured metadata, including publication year, publisher organization, scope level (i.e., international, regional), official URL, official description, access type, and cell applicability (i.e., a short explanation for the document's applicability to the target cell). The official description is intended to be based solely on information extracted from the official source page, without any additional interpretation or assumptions. In contrast, cell applicability involved model-based inference. The reason for this distinction will be explained in the "Scoring and Prioritizing the Documents" step.

Web search was constrained to a predefined list of domains (rather than any domain on the internet) and loaded from a separate comma-separated values (CSV) file for each OT component under consideration. This restriction was introduced to limit results to authoritative sources, avoid unreliable documents, and increase the output's determinism. The search results of each run were parsed into a structured table and converted into a cumulative list of candidate references. To reduce duplication, URLs were canonicalized, and titles were normalized; documents were treated as identical when either their canonical URL or their normalized title matched an item from a previous run.

Since the LLM-assisted web search was performed iteratively, a stopping point had to be chosen to determine when retrieval had become sufficiently exhaustive. This paper adopted the saturation logic proposed by Guest et al. (2020). Even though Guest et al.'s approach was originally developed to assess thematic saturation for interviews in qualitative research, it was used in this paper as a transparent, reproducible stopping rule for iterative document retrieval. Guest et al. (2020) define three metrics for determining saturation: base size (B), run length (L), and new information threshold (T). In this paper, base size refers to the number of initial web search iterations. The number of unique documents identified across these B searches is then used as the denominator when calculating the proportion of new information obtained in subsequent runs. Run length is the window of consecutive web searches for which new information is looked. Finally, new information threshold is the metric that defines how small new information should appear in the window of L to determine sufficient saturation. In this paper, B was set to 6, L to 2, and T to 5%. After the first six runs, the cumulative number of unique documents identified constituted the denominator D . Retrieval then continued in increments of two runs, and saturation was considered achieved when the proportion of newly identified unique documents in the most recent two runs, relative to D , fell to 5% or below. These metrics could be adjusted (e.g., increasing L , decreasing T) to increase saturation, which would increase the number of web searches required.

In parallel, a separate web search workflow was conducted to identify cybersecurity references explicitly written for the construction sector. In this parallel workflow, the same iterative retrieval and saturation logic was utilized while removing the domain restriction since the number of cybersecurity documents directly addressing the construction sector was expected to be limited. Instead, the prompt explicitly asked the model to provide a short list of up to five construction-sector references per run, which were then accumulated and checked for saturation in the same manner as the cell-specific searches. Moreover, in the output metadata, "construction-sector applicability" was included instead of "cell applicability" since this search was independent of the taxonomy cells.

Prompt design is an important part of LLM-assisted workflows since the instructions given to the model influence the scope, format, and consistency of the generated output. Giray (2023) describes prompts as containing four main elements: instruction, context, input data, and output indicator. In this paper, these elements were used to structure the prompts for document retrieval so that the model could search within the intended scope, filter

references according to the defined inclusion criteria, and return outputs in a consistent format. Prompt 1 (see the Appendix) was developed for taxonomy cell-specific candidate documents, while Prompt 2 was developed for construction-specific references.

3.4 Quality control and deduplication

The candidate references identified during the retrieval stage were then subjected to quality control (QC) and deduplication to avoid potential errors in the output and minimize manual intervention in the “manual review and final selection” step. An LLM-assisted script was used to review the raw candidate references and correct or standardize key metadata fields, including the document title, publication year, publisher organization, document type, scope level, official URL, and access type. One of the main objectives of this step was to merge duplicate items. The model was instructed to merge rows only when they clearly referred to the same document, such as when they shared the same canonical URL, the same official identifier or version, or a near-identical title under the same publisher. At the same time, the workflow explicitly prohibited merges across different parts of the same standards family. For example, different parts of the IEC 62443 series were not collapsed into a single item, since each is considered an individual standard with distinct purposes.

Following the LLM-assisted QC and deduplication, an additional deterministic post-processing layer was applied. This layer canonicalized URLs, normalized document-type labels to the four predefined categories, standardized the scope and access type fields, and implemented a final deduplication step (if still necessary after LLM-assisted deduplication). It also removed older versions when multiple versions of the same document remained. As a result, the output of this step was a cleaner and more internally consistent candidate reference set suitable for subsequent scoring and prioritization. Prompt 3 in the Appendix shows the LLM prompt used in this step. Please note that the “Output JSON schema” included in Prompt 3 is specific to the taxonomy cell-specific references. In the construction-specific workflow, this part of the prompt was revised to reflect the relevant columns (e.g., construction-sector applicability). Moreover, the “payload” included in the “Input Data” section of Prompt 3 refers to the list of candidate references identified in the previous step.

3.5 Scoring and prioritizing the documents

After QC and deduplication, the resulting candidate documents were scored to help prioritize the most relevant references. This scoring step was intended to support selection rather than to replace expert judgment. For the taxonomy cell-specific workflow, each document was assessed using six criteria: implementation usefulness, OT family relevance, OT component relevance, construction process relevance, currency, and authoritativeness. The first four criteria were assigned by the LLM using anchored rubrics, based solely on the structured metadata available after the QC and deduplication stage, particularly the title, official description (which serves as evidence), and cell applicability (which involves LLM inference). This stage did not include any web lookup to increase determinism.

The scoring rubric was designed to distinguish between documents that were merely related to cybersecurity in a broad sense and documents that were genuinely useful for the target taxonomy cell. Implementation usefulness was scored on a 0–10 scale, ranging from purely descriptive references to implementation-ready documents that include explicit requirements, controls, baselines, checklists, or assessment-oriented content. OT family, OT component, and construction process relevance were each scored on separate 0–10 scales using anchored definitions that distinguished indirect applicability from direct targeting. To improve scoring transparency, the information used for evaluation was separated into two fields in the “retrieving candidate documents” step: “Official Description (Evidence)”, only based on information extracted from the official source page, and “Cell Applicability (Inference)”, which involves the interpretation of the LLM explaining how the document could apply to the specific taxonomy cell. To prevent overinterpretation, relevance scores above 6 were assigned only when the title or official description (i.e., evidence) explicitly supported the claim. The scoring rubric is presented in Table 5. Intermediate integers (i.e., 1, 3, 5, 7, and 9) were allowed when the documents did not clearly meet any of the criteria.

Scores for the two additional criteria, currency and authoritativeness, were assigned deterministically. The currency score, ranging from 1 to 5, was assigned based on publication year using a rule-based scale, with the most recent documents receiving the highest score. On the other hand, documents with unknown publication years received a neutral mid-range value of 3. The scoring rubric for currency is shown in Table 6. The authoritativeness

score ranged from 1 to 4. It was derived from document type, with regulations and directives assigned the highest value (score of 4), followed by standards and codes of practice (score of 3), control catalogs and benchmarks (score of 2), and finally guidance and framework documents (score of 1). The overall score for each document was obtained by summing the six component scores, enabling the prioritization of the candidate reference list in preparation for the manual review and selection.

For the construction-specific workflow, the same scoring logic was utilized. However, the OT family, OT component, and construction process relevance were combined into a single construction sector relevance score, ranging from 0 to 10 (i.e., 0, 2, 4, 6, 8, and 10) and again allowing intermediate scores (i.e., 1, 3, 5, 7, and 9) when necessary. The scoring rubric for construction sector relevance is shown in Table 7. This rubric considered the relevance of the scored document to construction project delivery. Therefore, documents addressing the overall lifecycle of built environments received lower scores than those directly targeting the construction (i.e., delivery) phase. Prompts 4 and 5 in the Appendix show the prompts used for LLM-based scoring.

Table 5: Scoring rubric for cell-specific candidate references.

Score	Implementation usefulness	OT Family relevance	OT Component relevance	Construction process relevance
0	Not actionable (purely descriptive, background, concept definitions only)	No clear link to the OT family	No clear link to the OT component	No plausible applicability to the construction process
2	High-level principles only (general guidance with few concrete steps)	Indirect relevance (general OT cybersecurity, OT family not implied)	Indirect relevance (general cybersecurity governance; could apply but OT component not mentioned)	Indirect relevance (generic construction or OT context only)
4	Some actionable guidance (recommended practices, partial procedures, but not structured as controls/requirements)	Some relevance (OT family-related cues appear, but not central)	Some relevance (OT component-related theme appears, but not central)	Some relevance (construction process-related cues appear, but weakly)
6	Actionable governance guidance (clear process steps, roles, lifecycle activities, or program elements that can be implemented)	Clear relevance (OT family-specific concerns are explicitly discussed)	Clear relevance (OT component-related issues are explicitly discussed)	Clear relevance (site deployment/operations relevance, even if the process is not explicitly named)
8	Control/requirements oriented (explicit security requirements, control statements, checklists, baselines)	Strong relevance (multiple explicit OT family-specific cues)	Strong relevance (explicitly addresses the OT component area with multiple cues)	Strong relevance (multiple explicit construction process-specific cues)
10	Implementation-ready (structured requirements/controls + operational detail such as mappings, checklists, assessment/audit approach, or measurable security levels)	Directly targeted (title/rationale clearly focuses on this OT family)	Directly targeted (title/rationale clearly focuses on that OT component's security requirements/controls)	Directly targeted (title/evidence clearly focuses on this construction process context)

Table 6: Currency score rubric for candidate references.

Currency Criterion	Score
Current or previous year	5
2-3 years old	4
4-6 years old	3
7-10 years old	2
More than 10 years old	1
For unknown publication year	3



Table 7: Scoring rubric for construction sector relevance.

Construction Sector Relevance	Score
Not relevant to construction project delivery (clearly unrelated sector/context)	0
Indirect relevance (generic cybersecurity; construction not implied)	2
Some relevance (plausibly usable by construction companies, but not central)	4
Clear relevance (construction project delivery context is explicitly discussed)	6
Strong relevance (explicit construction sector cues in evidence)	8
Directly targeted (title/evidence clearly focuses on construction sector cybersecurity)	10

3.6 Manual review and final selection

The final stage of the workflow involved manual review and selection. Even though the LLM-assisted retrieval, QC, deduplication, and scoring steps identified the list of most relevant cybersecurity references for the targeted taxonomy cell, this step is necessary to avoid misleading output, as generative AI models may exhibit reliability and accuracy issues (e.g., hallucinations) (Augenstein *et al.*, 2024). Moreover, the purpose of LLM-assisted scoring is to support final selection and accelerate the overall process, rather than to replace expert judgment. The manual review included checking whether each record genuinely represented a primary cybersecurity reference (i.e., one of the four identified categories), whether its metadata (e.g., publisher, title) was accurate, whether any residual duplicates or version conflicts remained, and whether the stated applicability to the cell or the construction sector was defensible. This human review stage was especially important for borderline cases in which a document included relevant cybersecurity content but was not written exclusively for the target technology or construction context. Therefore, the purpose of the preceding LLM-assisted steps was to reduce the screening work, improve consistency, support final selection, and accelerate the overall process, not to automate the final decision entirely.

This review stage also included human scoring of the candidate references remaining after QC and deduplication, using the same rubric as in the LLM-based scoring step. These human-assigned scores were used to determine the final prioritization of the documents and to identify the most suitable references for inclusion in the taxonomy. Where LLM-based and human judgments were not aligned, final inclusion decisions were based on human assessment. This way, the resulting taxonomy remained human-validated while still benefiting from the efficiency of LLM-assisted retrieval and preprocessing. In the demonstrative implementation presented in this paper, this human scoring stage was further strengthened through a consensus-based review involving three experts. However, since such a panel-based evaluation would not be practical for every taxonomy cell, it is included in this paper to enhance the credibility of the human scoring rather than as a mandatory step in the proposed workflow.

This step resulted in a set of ranked cybersecurity references for each taxonomy cell, highlighting the most relevant and actionable ones, along with a separate list of construction-specific cybersecurity documents applicable across cells. Finally, to keep the resulting reference sets practical and interpretable, the final inclusion threshold was defined as the top 10 manually prioritized documents for each taxonomy cell and the top 5 manually prioritized construction-sector references applicable across all cells. For the cell-specific documents, where two documents had the same overall manual score, priority was given to the document with the higher component relevance score, followed by implementation usefulness, family relevance, construction process relevance, currency, and authority. For construction-specific documents, priority was given to the document with the higher construction-sector relevance score, followed by implementation usefulness, currency, and authoritativeness.

4. IMPLEMENTATION

To demonstrate the proposed methodology, the taxonomy cell-specific and construction-specific workflows were implemented through three Python scripts each, corresponding to the retrieval, QC and deduplication, and scoring steps, followed by manual review and selection. In this implementation, the human scoring component of the final review stage was strengthened through a consensus-based evaluation involving three experts. Characteristics of the experts are presented in Table 8. The implementation used OpenAI's LLMs via its application programming interface (API) due to their advanced capabilities, and the web search tool (required for the retrieval step) that is

enabled in certain models (via the Responses API (OpenAI, 2026a)). For the cell-specific workflow, the following cell was selected:

- OT Family (Level 1 column): Robotic manipulators
- OT Component (Level 2 column): Communication
- Construction Process (row): Masonry and bricklaying

In addition to the taxonomy cell-specific workflow, a parallel workflow was implemented to identify construction-specific cybersecurity references. This second workflow was executed independently of the taxonomy cell, since the construction-specific references apply to the entire taxonomy table. The following subsections present the details of implementation and results for both workflows.

Table 8: Characteristics of the experts involved in the consensus-based evaluation.

Expert #	Cybersecurity Experience	OT/ICS cybersecurity experience	Role
Expert 1	4-6 years	4-6 years	Academic (Professor or Researcher)
Expert 2	1-3 years	1-3 years	Academic (Professor or Researcher)
Expert 3	4-6 years	4-6 years	Academic (Professor or Researcher)

4.1 Taxonomy cell-specific workflow

Retrieving Candidate Documents: The first Python script used OpenAI’s gpt-5.1-2025-11-13 model with the web search tool enabled via the Responses API to retrieve candidate references. The model was called with a reasoning setting of medium. Since enabling web search makes the model’s output inherently nondeterministic, the Responses API does not accept the temperature parameter. However, the web search tool in the Responses API allows restricting the search to certain domains, which was used to reduce the variability of the results. For the selected “communication” component of “robotic manipulators”, web search was limited to the following domains:

- International standards and professional bodies: iec.ch, isa.org, iso.org, etsi.org, itu.int, and ieee.org;
- Governmental and regulatory cybersecurity sources: nist.gov, cisa.gov, ncsc.gov.uk, npsa.gov.uk, enisa.europa.eu, and eur-lex.europa.eu;
- Industry and technology associations: ciscure.org, gsma.com, profibus.com, and odva.org.

These domains were chosen to limit the retrieval to authoritative standards bodies, governmental organizations, and widely recognized cybersecurity or communications-related institutions. These domains include international, US-, UK-, and EU-based organizations to limit the identified documents to the English language. However, the list of domains can be revised in different cases, depending on the country and the specific technology. For example, “bsi.bund.de” could be added to include the Federal Office for Information Security (BSI) of Germany if the construction site was in Germany, or “faa.gov” could be added to include the Federal Aviation Administration (FAA) of the US if the OT family was “aerial robots (UAVs)”.

Each retrieval run produced a structured table containing the identified candidate references and the associated metadata. The outputs of individual runs were stored separately for tracing purposes, while a cumulative unique document list was generated by comparing canonicalized URLs and normalized titles across runs and eliminating duplicates. Duplicate handling during this stage was intentionally conservative, considering only documents with the exact same canonical URLs or normalized titles identical, since more detailed deduplication and cleaning were performed in the following step.

LLM-based retrieval via web search was iteratively executed until the adapted saturation criterion described in the Methodology section was met. The implementation used base size $B = 6$, run length $L = 2$, and new information threshold $T = 5\%$. The first six runs yielded 26 unique documents, which formed the base denominator D . Retrieval then continued in increments of two runs. At run 9, the most recent two runs had produced only 1 new unique document, corresponding to 3.85% of the base denominator, which satisfied the stopping criterion. Therefore, the retrieval process stopped after 9 runs, resulting in a final cumulative set of 29 unique candidate references. The number of unique documents identified at each run and the percent change over D in the last L runs are shown in Table 9.

Table 9: Saturation statistics for the taxonomy cell-specific workflow.

Metric	Runs 1–6 (base)	Run 7	Run 8	Run 9
New unique documents per run	9, 7, 5, 3, 1, 1	2	1	0
Cumulative unique documents	26	28	29	29
New documents in the last L = 2 runs	—	—	3	1
% change over base D = 26	—	—	11.54%	3.85%

Quality Control and Deduplication: The second Python script implemented the QC and deduplication step. In this script, OpenAI’s gpt-5-chat-latest model was used since it accepts the temperature parameter. The temperature parameter was set to 0 to make the output as deterministic as possible. However, while setting the temperature parameter to 0 significantly increases determinism, it still does not guarantee the same output across runs. Moreover, although the name gpt-5-chat-latest appears to indicate a changing model because of the term “latest,” it actually refers to a snapshot model, as shown by OpenAI (2026b), which allows users to lock in a particular version.

QC and deduplication were applied to the 29 unique candidate references identified during the retrieval stage. After this step, 8 documents were removed, leaving 21 candidate references for scoring. The removed items were excluded because they were either duplicates of other records or older versions of documents that already had a more recent version in the list. The remaining 21 candidate references after QC and deduplication are presented in Table 10.

Table 10: Candidate references remaining after QC and deduplication for the taxonomy cell-specific workflow.

No.	Title	Publication Year	Publisher Organization	Document Type	Scope Level
1	CIS Critical Security Controls Version 8.1	2024	CIS	Control catalog / Benchmark	International
2	Cross-Sector Cybersecurity Performance Goals (CPGs)	2023	CISA	Control catalog / Benchmark	Regional
3	Cybersecurity Best Practices for Industrial Control Systems	2020	CISA	Guidance / Framework	Regional
4	Recommended Practice: Configuring and Managing Remote Access for Industrial Control Systems	2023*	CISA**	Guidance / Framework	Regional
5	Recommended Practice: Improving Industrial Control System Cybersecurity with Defense-in-Depth Strategies	2016	CISA	Guidance / Framework	Regional
6	Baseline Security Recommendations for Internet of Things in the Context of Critical Information Infrastructures	2017	ENISA	Guidance / Framework	Regional
7	Good Practices for Security of Internet of Things in the Context of Smart Manufacturing	2018	ENISA	Guidance / Framework	Regional
8	Industry 4.0 – Cybersecurity Challenges and Recommendations	2019	ENISA	Guidance / Framework	Regional
9	Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union (NIS 2 Directive)	2022	European Union	Regulation / Directive	Regional
10	IEC 62443-2-1:2024 Security for Industrial Automation and Control Systems – Part 2-1: Security Program Requirements for IACS Asset Owners	2024	IEC	Standard / Code of practice	International

No.	Title	Publication Year	Publisher Organization	Document Type	Scope Level
11	IEC 62443-3-3:2013 Industrial Communication Networks – Network and System Security – Part 3-3: System Security Requirements and Security Levels	2013	IEC	Standard / Code of practice	International
12	IEC PAS 62443-2-2:2025 Security for Industrial Automation and Control Systems – Part 2-2: IACS Security Protection Scheme	2025	IEC	Standard / Code of practice	International
13	ISO/IEC 27001:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Management Systems — Requirements	2022	ISO/IEC	Standard / Code of practice	International
14	ISO/IEC 27002:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Controls	2022	ISO/IEC	Standard / Code of practice	International
15	Cyber Assessment Framework (CAF) 4.0	2025	NCSC UK	Guidance / Framework	Regional
16	Secure Connectivity Principles for Operational Technology (OT)	Unknown	NCSC UK	Guidance / Framework	Regional
17	Cybersecurity Framework Version 1.1 Manufacturing Profile (NISTIR 8183 Rev. 1)	2020	NIST	Guidance / Framework	Regional
18	Guide to Operational Technology (OT) Security (NIST SP 800-82 Rev. 3)	2023	NIST	Guidance / Framework	Regional
19	Security and Privacy Controls for Information Systems and Organizations (NIST SP 800-53 Rev. 5, Release 5.2.0)	2025	NIST	Guidance / Framework	Regional
20	The NIST Cybersecurity Framework (CSF) 2.0	2024	NIST	Guidance / Framework	Regional
21	PROFINET Security Guideline	2013*	PROFIBUS & PROFINET International	Guidance / Framework	International

**These publication years were assigned incorrectly by the LLM. To reflect the LLM output as it is, these incorrect publication years were not corrected in this table.*

***While the publisher organization of this document is the US Department of Homeland Security (DHS), the LLM incorrectly assigned the Cybersecurity and Infrastructure Security Agency (CISA), likely because the document is hosted on the cisa.gov domain. To reflect the LLM output as it is, this incorrect publisher organization was not corrected in this table.*

Scoring and Prioritizing the Documents: The third Python script implemented the scoring and prioritization step for the 21 candidate references that remained after QC and deduplication. Similar to the QC and deduplication script, OpenAI's gpt-5-chat-latest model was used with temperature set to 0 to reduce response variability as much as possible. The LLM assigned scores for implementation usefulness, OT family relevance, OT component relevance, and construction process relevance using the rubric defined in the methodology section and provided in the prompt. The remaining two criteria, currency and authoritativeness, were assigned deterministically based on publication year and document type, respectively. The overall score for each document was then calculated as the sum of these six criteria.

Manual Review and Final Selection: After LLM-based scoring, the same 21 candidate references were manually scored using the same rubric. To increase the credibility of the final prioritization, this review was conducted through a consensus-based evaluation involving three experts with experience in cybersecurity, OT systems, and the built environment. The experts jointly examined the candidate references and agreed on a single final score for

each criterion. This consensus-based review was used to determine the final prioritization of the documents and to assess the extent to which the LLM-based ranking aligned with expert judgment. As shown in Table 11, the expert review did not change the authoritativeness scores for any reference, as the LLM accurately assigned all document types. However, currency scores changed after the review for three references due to the LLM assigning inaccurate publication years or not being able to identify the year and assigning “unknown”. For example, while the publication year of “Recommended Practice: Configuring and Managing Remote Access for Industrial Control Systems” is 2010 and the publisher organization is DHS, the LLM incorrectly assigned the year 2023 and the publisher organization CISA. This error was likely caused by the fact that the document is hosted on the cisa.gov domain and its URL includes “2023-01”, which might have misled the LLM. Another observation based on the scores in Table 11 is that while LLM tended to rely on the anchored rubric (i.e., 0, 2, 4, 6, 8, and 10), even though intermediate integers were allowed, the expert consensus frequently involved intermediate scores (i.e., 1, 3, 5, 7, and 9) to show finer distinctions. To keep Table 11 compact, the scores are reported as “LLM/Expert Consensus” for each criterion. The references in Table 11 are ordered from highest to lowest overall consensus-based expert scores.

As explained in the Methodology section, the top 10 documents from the manually scored list (expert consensus in the case of this implementation) were included in the final selection for the target taxonomy cell. These 10 documents are shown in Table 12. When compared to the top 10 documents from the LLM-scored list, 8 out of 10 documents were the same. This indicates substantial overlap between the LLM-based and human judgments. The documents that were included in the LLM-scored top 10 list but not included in the manually scored list are “Security and Privacy Controls for Information Systems and Organizations (NIST SP 800-53 Rev. 5, Release 5.2.0)” by NIST and “Cross-Sector Cybersecurity Performance Goals (CPGs)” by CISA.

4.2 Construction-specific workflow

For the construction-specific workflow, the same models from OpenAI, the same temperature parameters (where applicable) for the LLMs, and the same saturation criteria were used. Therefore, these aspects are not explained again in this subsection.

Retrieving Candidate Documents: The construction-specific workflow also used an iterative web search until the predefined saturation criteria were met. The first six runs yielded 10 unique documents, which formed the base denominator D . Retrieval then continued in increments of two runs. At run 10, the most recent two runs had produced 0 new unique documents, corresponding to 0% of the base denominator, which satisfied the stopping criterion. Therefore, the retrieval process stopped after 10 runs, resulting in a final cumulative set of 11 unique candidate references. The number of unique documents identified at each run and the percent change over D in the last $L = 2$ runs are shown in Table 13.

Quality Control and Deduplication: The QC and deduplication process was applied to the 11 unique construction-specific candidate references identified during the retrieval step. This step did not identify any duplicate documents or older versions of the same reference. Therefore, the number of documents after this step stayed the same (i.e., 11). The final list of the candidate references is presented in Table 14.

Scoring and Prioritizing the Documents: At this step, the LLM-based script was used to score 11 construction-specific candidate references based on the rubric presented in the Methodology section.

Manual Review and Final Selection: After LLM-based scoring, the 11 candidate references were scored using the same rubric in a consensus-based evaluation involving three experts, as presented in Table 15. In this case, neither the authoritativeness nor the currency scores changed after the manual review for any reference. The last two references in Table 15 were found ineligible and thus were not scored by the experts. “ISO 23234:2021 – Buildings and Civil Engineering Works — Security — Planning of Security Measures in the Built Environment” by ISO was excluded as a false positive under the study’s inclusion criteria, as its primary focus is physical security rather than cybersecurity. Moreover, “Cyber Security for Construction Businesses” by the National Cyber Security Centre (NCSC) was excluded from scoring as it is not available anymore, and its link redirects to another page titled “Cyber security advice for small to medium sized organisations” by NCSC. Prompt 2 included the instruction “If the URL redirects, follow the redirect and only include the final destination if it is still the same document and still fits the construction-sector scope; otherwise exclude it.”; however, this outcome shows that the LLM did not fully comply with that instruction. Similar to the taxonomy cell-specific workflow, intermediate scores (i.e., 1, 3, 5, 7, and 9) were allowed and were more frequently used by the experts compared to the LLM.

Table 11: LLM-based and manual scores for the 21 taxonomy cell-specific candidate references remaining after QC and deduplication.

No.	Title	Implementation usefulness (LLM/ Expert Consensus)	Family relevance (LLM/Expert Consensus)	Component relevance (LLM/ Expert Consensus)	Construction process relevance (LLM/ Expert Consensus)	Currency (LLM/ Expert Consensus)	Authoritativeness (LLM/ Expert Consensus)	Overall score (LLM/Expert Consensus)
1	IEC 62443-2-1:2024 Security for Industrial Automation and Control Systems – Part 2-1: Security Program Requirements for IACS Asset Owners	8/10	6/5	6/8	4/4	4/4	3/3	31/34
2	IEC 62443-3-3:2013 Industrial Communication Networks – Network and System Security – Part 3-3: System Security Requirements and Security Levels	10/10	6/5	8/10	4/4	1/1	3/3	32/33
3	IEC PAS 62443-2-2:2025 Security for Industrial Automation and Control Systems – Part 2-2: IACS Security Protection Scheme	8/10	6/5	8/6	4/4	5/5	3/3	34/33
4	PROFINET Security Guideline	8/7	6/6	10/10	4/3	1/5	1/1	30/32
5	Secure Connectivity Principles for Operational Technology (OT)	6/6	6/5	8/10	4/5	3/4	1/1	28/31
6	CIS Critical Security Controls Version 8.1	8/10	4/4	6/7	4/4	4/4	2/2	28/31
7	Recommended Practice: Configuring and Managing Remote Access for Industrial Control Systems	8/8	6/5	8/10	4/4	4/1	1/1	31/29
8	Guide to Operational Technology (OT) Security (NIST SP 800-82 Rev. 3)	8/6	6/5	8/7	4/5	4/4	1/1	31/28
9	ISO/IEC 27002:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Controls	10/10	2/2	6/5	4/4	3/3	3/3	28/27
10	ISO/IEC 27001:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Management Systems — Requirements	8/10	2/2	4/5	4/4	3/3	3/3	24/27
11	Security and Privacy Controls for Information Systems and Organizations (NIST SP 800-53 Rev. 5, Release 5.2.0)	10/9	4/2	6/5	4/4	5/5	1/1	30/26

No.	Title	Implementation usefulness (LLM/ Expert Consensus)	Family relevance (LLM/Expert Consensus)	Component relevance (LLM/ Expert Consensus)	Construction process relevance (LLM/ Expert Consensus)	Currency (LLM/ Expert Consensus)	Authoritativeness (LLM/ Expert Consensus)	Overall score (LLM/Expert Consensus)
12	Cybersecurity Framework Version 1.1 Manufacturing Profile (NISTIR 8183 Rev. 1)	6/7	4/4	6/6	4/4	3/3	1/1	24/25
13	Good Practices for Security of Internet of Things in the Context of Smart Manufacturing	6/6	6/5	6/6	4/4	2/2	1/1	25/24
14	Baseline Security Recommendations for Internet of Things in the Context of Critical Information Infrastructures	6/7	4/4	6/5	4/4	2/2	1/1	23/23
15	Recommended Practice: Improving Industrial Control System Cybersecurity with Defense-in-Depth Strategies	6/6	6/5	6/5	4/4	2/2	1/1	25/23
16	The NIST Cybersecurity Framework (CSF) 2.0	4/6	2/3	4/4	4/4	4/4	1/1	19/22
17	Cross-Sector Cybersecurity Performance Goals (CPGs)	8/5	4/3	6/4	4/4	4/4	2/2	28/22
18	Cyber Assessment Framework (CAF) 4.0	6/6	4/3	4/3	4/4	5/5	1/1	24/22
19	Directive (EU) 2022/2555 on Measures for a High Common Level of Cybersecurity Across the Union (NIS 2 Directive)	6/6	2/2	2/3	4/4	3/3	4/4	21/22
20	Industry 4.0 – Cybersecurity Challenges and Recommendations	4/3	4/5	4/4	4/4	2/2	1/1	19/19
21	Cybersecurity Best Practices for Industrial Control Systems	6/2	6/4	6/4	4/4	3/3	1/1	26/18

Finally, the top 5 documents from the manually scored (i.e., expert consensus) construction-specific list were included in the final selection to be applicable across all cells of the taxonomy. These documents are shown in Table 16. When compared to the top 5 documents from the LLM-scored list (excluding the three disqualified documents), 3 out of 5 matched. The two documents that were included in the LLM-scored top 5 list but not included in the manually scored list were “Security Considerations Assessment (SCA) – including SCA – New Built Assets and Related Guides” and “Major Infrastructure Projects – Security Considerations for Clients” by the National Protective Security Authority (NPSA). The experts agreed that these two documents did not provide sufficient detail for implementation and did not place sufficient emphasis on construction.

Table 12: Final selected cybersecurity references for the target taxonomy cell.

No.	Title
1	IEC 62443-2-1:2024 Security for Industrial Automation and Control Systems – Part 2-1: Security Program Requirements for IACS Asset Owners
2	IEC 62443-3-3:2013 Industrial Communication Networks – Network and System Security – Part 3-3: System Security Requirements and Security Levels
3	IEC PAS 62443-2-2:2025 Security for Industrial Automation and Control Systems – Part 2-2: IACS Security Protection Scheme
4	PROFINET Security Guideline
5	Secure Connectivity Principles for Operational Technology (OT)
6	CIS Critical Security Controls Version 8.1
7	Recommended Practice: Configuring and Managing Remote Access for Industrial Control Systems
8	Guide to Operational Technology (OT) Security (NIST SP 800-82 Rev. 3)
9	ISO/IEC 27002:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Controls
10	ISO/IEC 27001:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Management Systems — Requirements

Table 13: Saturation statistics for the construction-specific workflow.

Metric	Runs 1–6 (base)	Run 7	Run 8	Run 9	Run 10
New unique documents per run	5, 1, 1, 2, 1, 0	0	1	0	0
Cumulative unique documents	10	10	11	11	11
New documents in the last L = 2 runs	—	—	1	1	0
% change over base D = 10	—	—	10.00%	10.00%	0.00%

Table 14: Candidate references remaining after QC and deduplication for the construction-specific workflow.

No.	Title	Publication Year	Publisher Organization	Document Type	Scope Level
1	CIBSE DE6: Security Requirements	2017	CIBSE	Guidance / Framework	regional
2	The Role of Security in the Construction Industry	2018	CIOB	Guidance / Framework	regional
3	Code of Practice: Cyber Security in the Built Environment (Revised Second Edition)	2022	IET	Standard / Code of practice	regional
4	ISO 19650-5:2020 – Organization and Digitization of Information about Buildings and Civil Engineering Works, including Building Information Modelling (BIM) — Information Management using Building Information Modelling — Part 5: Security-Minded Approach to Information Management	2020	ISO	Standard / Code of practice	international
5	ISO 23234:2021 – Buildings and Civil Engineering Works — Security — Planning of Security Measures in the Built Environment	2021	ISO	Standard / Code of practice	international
6	Cyber Security for Construction Businesses	2022	NCSC	Guidance / Framework	regional
7	Joint Ventures in the Construction Sector: Information Security Best Practice Guidance	2022	NCSC	Guidance / Framework	regional
8	Information Management – Document Release Guidance	2021	NPSA	Guidance / Framework	regional
9	Major Infrastructure Projects – Security Considerations for Clients	2026	NPSA	Guidance / Framework	regional
10	Security Considerations Assessment (SCA) – including SCA – New Built Assets and Related Guides	2022	NPSA	Guidance / Framework	regional
11	Security-Minded Approach to Digital Engineering	2025	NPSA	Guidance / Framework	regional

Table 15: LLM-based and manual scores for the 11 construction-specific candidate references remaining after QC and deduplication.

No.	Title	Implementation usefulness (LLM/Expert Consensus)	Construction sector relevance (LLM/ Expert Consensus)	Currency (LLM/ Expert Consensus)	Authoritativeness (LLM/ Expert Consensus)	Overall score (LLM/ Expert Consensus)
1	ISO 19650-5:2020 – Organization and Digitization of Information about Buildings and Civil Engineering Works, including Building Information Modelling (BIM) — Information Management using Building Information Modelling — Part 5: Security-Minded Approach to Information Management	8/10	8/10	3/3	3/3	22/26
2	Code of Practice: Cyber Security in the Built Environment (Revised Second Edition)	6/6	6/8	3/3	3/3	18/20
3	Joint Ventures in the Construction Sector: Information Security Best Practice Guidance	6/5	10/10	3/3	1/1	20/19
4	The Role of Security in the Construction Industry	6/5	10/10	2/2	1/1	19/18
5	Security-Minded Approach to Digital Engineering	4/3	6/6	5/5	1/1	16/15
6	Security Considerations Assessment (SCA) – including SCA – New Built Assets and Related Guides	8/3	8/6	3/3	1/1	20/13
7	Major Infrastructure Projects – Security Considerations for Clients	6/2	8/4	5/5	1/1	20/12
8	CIBSE DE6: Security Requirements	8/3	8/4	2/2	1/1	19/10
9	Information Management – Document Release Guidance	6/3	4/2	3/3	1/1	14/9
10	Cyber Security for Construction Businesses	6/NA	10/NA	3/NA	1/NA	20/NA
11	ISO 23234:2021 – Buildings and Civil Engineering Works — Security — Planning of Security Measures in the Built Environment	8/NA	6/NA	3/NA	3/NA	20/NA

Table 16: Final selected construction-specific cybersecurity references.

No.	Title
1	ISO 19650-5:2020 – Organization and Digitization of Information about Buildings and Civil Engineering Works, including Building Information Modelling (BIM) — Information Management using Building Information Modelling — Part 5: Security-Minded Approach to Information Management
2	Code of Practice: Cyber Security in the Built Environment (Revised Second Edition)
3	Joint Ventures in the Construction Sector: Information Security Best Practice Guidance
4	The Role of Security in the Construction Industry
5	Security-Minded Approach to Digital Engineering

5. DISCUSSION AND LIMITATIONS

The results demonstrate that the proposed semi-automated workflow can generate a structured, prioritized set of cybersecurity references for construction OT scenarios. In the demonstrated taxonomy cell-specific workflow, iterative LLM-assisted web search produced 29 unique candidate references before saturation was reached. After QC and deduplication, 21 references remained for scoring, and the final manual scoring and selection resulted in the top 10 documents. In the parallel construction-specific workflow, iterative retrieval yielded 11 unique candidate references, all of which remained after QC and deduplication, and the final manual scoring and selection involved 5 documents. This human-in-the-loop approach uses LLMs to accelerate the process while keeping human review to reduce hallucinations and inaccuracies. Therefore, the generated taxonomy aims to help connect cybersecurity references with the practical needs of the construction industry.

In this paper, the taxonomy cell-specific workflow captures documents that are more directly relevant to the selected OT family, OT component, and construction process, while the parallel construction-specific workflow identifies references that concern the construction industry as a whole. Therefore, the methodology does not treat construction cybersecurity as a purely technology-specific issue; rather, it combines technology-focused references with the sector-oriented ones. This structure is especially important in construction, where project teams may need to consider both component-level cybersecurity controls and broader security-minded management practices.

Several limitations of the proposed approach should be noted. One of the limitations to address is the false positives and unavailable documents identified by the construction-specific workflow. Further adjustments need to be made to the prompt and the scripts used in the workflow to ensure that documents that address only physical security or lack functional links are excluded from the final list of documents. Also, the proposed taxonomy can help construction teams find cybersecurity references that fit their specific scenarios; however, interpreting these references and applying their requirements in practice still requires cybersecurity expertise. Another limitation is that this paper did not compare the performance of different LLMs for document retrieval, review, and scoring. However, this paper does not focus on selecting the best-performing language model. Moreover, the dynamic nature of the methodology enables efficient updating of the output as future LLMs emerge. Additionally, the final prioritization relied on the consensus of three subject matter experts. More expert reviewers could increase the credibility and strength of the results. Finally, the implementation presented in this paper uses OpenAI models whose future availability cannot be guaranteed. Therefore, the reproducibility of this specific implementation is limited, while the overall workflow is reproducible as it is model-agnostic.

6. CONCLUSIONS AND FUTURE WORK

This paper addressed the challenge of identifying cybersecurity references for OT-enabled construction scenarios. The proposed taxonomy organizes the problem around three dimensions: construction OT families, OT components, and construction processes. Eight OT families and eleven OT-enabled construction processes were identified from the literature, while five OT components were defined by adapting selected levels of the ISA-95 hierarchy. These elements form a two-dimensional taxonomy structure that links OT–process combinations with relevant cybersecurity standards, frameworks, regulations, guidelines, and best practices. To populate the cells of this taxonomy, this paper proposed an iterative semi-automated workflow combining LLM-assisted retrieval, QC and deduplication, scoring, and manual review and selection. This hybrid approach demonstrated that a

comprehensive, up-to-date set of references can be generated with much less manual effort than traditional approaches. In doing so, the research provides construction management teams a starting point for embedding security controls into project lifecycles. It also aligns with the United Nations Sustainable Development Goals (SDGs), particularly SDG 9 (Industry, Innovation and Infrastructure), SDG 11 (Sustainable Cities and Communities), and SDG 8 (Decent Work and Economic Growth), by supporting secure technology integration and innovation, more resilient built environment lifecycles, and safer OT-enabled construction sites.

To demonstrate the usefulness of the proposed approach, the implementation was executed through a taxonomy cell-specific workflow for the “communication” component of “robotic manipulators” performing “masonry and bricklaying”. This process was supported by a parallel workflow for identifying construction-specific cybersecurity references applicable across cells. In the demonstrated cell-specific workflow, iterative document retrieval identified 29 unique candidate references, which were reduced to 21 after QC and deduplication, and finally to a shortlist of 10 manually prioritized references. In the construction-specific workflow, 11 unique candidate references were retrieved, of which 5 were finally selected after manual review and scoring. The manual review followed a consensus-based approach involving three experts.

The paper also showed that LLM-assisted prioritization can support, but not replace, expert judgment. The substantial overlap between the LLM-supported and manually prioritized shortlists (especially for the taxonomy cell-specific workflow) suggests that the automated stages can meaningfully reduce the effort required to identify relevant cybersecurity references. Therefore, the main contribution of this paper is the proposed taxonomy table structure and the human-in-the-loop workflow for populating each table cell, rather than the identified documents in the implementation section.

Future work is explicitly aligned with the limitations identified above. First, to address the limitation of the continuing need for cybersecurity expertise when interpreting the selected references, a retrieval-augmented generation (RAG) approach will be developed to process the identified documents and extract actionable items for construction teams. Second, to address the model dependence of the implementation and the absence of cross-model evaluation, experiments with multiple language models will be conducted, comparing their performance. Third, to reduce false positives and unavailable links, the LLM prompts and workflow scripts will be refined to exclude ineligible documents before final selection. Finally, to address the limited expert panel included in this implementation, future validation will involve a broader group of subject matter experts.

DATA AVAILABILITY STATEMENT

The Python scripts developed and used for document retrieval, quality control, deduplication, and scoring are available from the corresponding author upon reasonable request.

ACKNOWLEDGMENTS

This paper is an extended version of our previous work, which was presented at the Smart and Sustainable Built Environment Conference, Lille – France | 3rd to 5th November 2025. The authors acknowledge the support and feedback from the conference chairs, Prof. Farzad Rahimian, Prof. Abiola Akanmu, and Dr. M. Reza Hosseini, along with their team, throughout the peer review process of SASBE 2025 and during the conference, which helped improve our submissions. This research was partially supported by the Center for Cyber Security at New York University Abu Dhabi (CCS-AD), funded by Tamkeen under the NYUAD Research Institute Award G1104. This research was partially supported by the Center for Sand Hazards and Opportunities for Resilience, Energy, and Sustainability (SHORES), funded by Tamkeen under the NYUAD Research Institute Award CG013.

REFERENCES

- Al-rimy, B. A. S., Maarof, M. A. and Shaid, S. Z. M. (2018). Ransomware threat success factors, taxonomy, and countermeasures: A survey and research directions, *Computers and Security* 74: 144–166. <https://doi.org/10.1016/j.cose.2018.01.001>
- Attalla, A., Attalla, O., Moussa, A., Shafique, D., Raean, S. B. and Hegazy, T. (2023). Construction robotics: review of intelligent features, *International Journal of Intelligent Robotics and Applications* 7(3): 535–555. <https://doi.org/10.1007/s41315-023-00275-1>



- Augenstein, I., Baldwin, T., Cha, M., Chakraborty, T., Ciampaglia, G. L., Corney, D., DiResta, R., Ferrara, E., Hale, S., Halevy, A., Hovy, E., Ji, H., Menczer, F., Miguez, R., Nakov, P., Scheufele, D., Sharma, S. and Zagni, G. (2024). Factuality challenges in the era of large language models and opportunities for fact-checking, *Nature Machine Intelligence* 6(8): 852–863. <https://doi.org/10.1038/s42256-024-00881-z>
- Bahrami, P. N., Dehghantanha, A., Dargahi, T., Parizi, R., Choo, K.-K. R. and Javadi, H. (2019). Cyber Kill Chain-Based Taxonomy of Advanced Persistent Threat Actors: Analogy of Tactics, Techniques, and Procedures, *J. Inf. Process. Syst.* <https://www.semanticscholar.org/paper/cyber-kill-chain-based-taxonomy-of-advanced-threat-bahrami-dehghantanha/7867503f1fe19b548c524a5398df5a5e6075845e>
- Bogue, R. (2025). The role of robots in the construction industry, *Industrial Robot: the international journal of robotics research and application* 52(2): 157–164. <https://doi.org/10.1108/IR-01-2025-0027>
- Brar, H. and Kumar, G. (2018). Cybercrimes: A Proposed Taxonomy and Challenges, *J. Comput. Networks Commun.* 2018: 17986591–179865911. <https://doi.org/10.1155/2018/1798659>
- Brelhi, A. and Klinc, R. (2025). Building digital trust in CDE-based BIM workflows: Key strategies, *Journal of Information Technology in Construction (ITcon)* 30(22): 524–543. <https://doi.org/10.36680/j.itcon.2025.022>
- Elliott, D. (2025, February 4). Cybercrime: Lessons learned from a \$25m deepfake attack, *World Economic Forum*. <https://www.weforum.org/stories/2025/02/deepfake-ai-cybercrime-arup/>
- EU. (2016). *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)* (No. Regulation (EU) 2016/679). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- EU. (2022). *Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union*, EU. <https://eur-lex.europa.eu/eli/dir/2022/2555>
- Ferdinand, J. and Benham, R. (2017, October 2). The Cyber Security Ecosystem: Defining a Taxonomy of Existing, Emerging and Future Cyber Threats. <https://papers.ssrn.com/abstract=3047753>
- Gao, Y., Shu, J., Xia, Z. and Luo, Y. (2025). From muscular to dexterous: A systematic review to understand the robotic taxonomy in construction and effectiveness, *Journal of Field Robotics* 42(1): 180–205. <https://doi.org/10.1002/rob.22409>
- García de Soto, B., Georgescu, A., Mantha, B. R. K., Turk, Ž., Maciel, A. and Sonkor, M. S. (2021). Construction Cybersecurity and Critical Infrastructure Protection: New Horizons for Construction 4.0, *Journal of Information Technology in Construction*.
- Giray, L. (2023). Prompt Engineering with ChatGPT: A Guide for Academic Writers, *Annals of Biomedical Engineering* 51(12): 2629–2633. <https://doi.org/10.1007/s10439-023-03272-4>
- Givchchi, O., Landsdorf, K., Simoens, P. and Colombo, A. W. (2017). Interoperability for Industrial Cyber-Physical Systems: An Approach for Legacy Systems, *IEEE Transactions on Industrial Informatics* 13(6): 3370–3378. <https://doi.org/10.1109/TII.2017.2740434>
- Griffor, E. (2017). *Framework for Cyber-Physical Systems: Volume 1, Overview* (No. NIST Special Publication 1500-201), NIST. <https://doi.org/10.6028/nist.sp.1500-201>
- Guest, G., Namey, E. and Chen, M. (2020). A simple method to assess and report thematic saturation in qualitative research, *PLOS ONE* 15(5): e0232076. <https://doi.org/10.1371/journal.pone.0232076>
- Harp, D. R. and Gregory-Brown, B. (2015). *IT / OT Convergence Bridging the Divide*, NexDefense.
- Humphreys, E. (2008). Information security management standards: Compliance, governance and risk management, *Information Security Technical Report* 13(4): 247–255. <https://doi.org/10.1016/j.istr.2008.10.010>
- IEEE. (1987). IEEE Standard Taxonomy for Software Engineering Standards.

- IET. (2021). *Code of Practice: Cyber Security in the Built Environment - 2nd Edition*, IET. <https://electrical.theiet.org/guidance-codes-of-practice/publications-by-category/cyber-security/code-of-practice-cyber-security-in-the-built-environment-revised-second-edition/>
- ISA. (2025). ISA-95 Standard: Enterprise-Control System Integration, *Isa.Org*. <https://www.isa.org/standards-and-publications/isa-standards/isa-95-standard>
- ISA/IEC. (2026). ISA/IEC 62443 Series of Standards, *ISA/IEC 62443 Series of Standards*. <https://www.isa.org/standards-and-publications/isa-standards/isa-iec-62443-series-of-standards>
- ISO. (2020). *ISO 19650-5:2020 Organization and digitization of information about buildings and civil engineering works, including building information modelling (BIM) — Information management using building information modelling — Part 5*, ISO. <https://www.iso.org/standard/74206.html>
- ISO/SAE. (2021). ISO/SAE 21434:2021. <https://www.iso.org/standard/70918.html>
- Jeon, J., Bae, J., Kwon, O., Ko, Y., Kim, C., Won, S., Shin, W. and Hong, D. (2025). Automating motor grader leveling operations: Kinematic analysis for blade pose control, *Automation in Construction* 175. <https://doi.org/10.1016/j.autcon.2025.106199>
- Johansen, K. W., Teizer, J. and Schultz, C. (2024). Automated rule-based safety inspection and compliance checking of temporary guardrail systems in construction, *Automation in Construction* 168. <https://doi.org/10.1016/j.autcon.2024.105849>
- Kalogeraki, E.-M. and Polemi, N. (2024). A taxonomy for cybersecurity standards, *Journal of Surveillance, Security and Safety*. <https://doi.org/10.20517/jsss.2023.50>
- Kougioumtzidou, A., Papoutsis, A., Kavallieros, D., Mavropoulos, T., Tsikrika, T., Vrochidis, S. and Kompatsiaris, I. (2024). An End-to-End Framework for Cybersecurity Taxonomy and Ontology Generation and Updating, In *Proceedings of CSR 2024*, pp. 247–254. <https://doi.org/10.1109/CSR61664.2024.10679346>
- Lemos, R. (2025). Ransomware Builds Against Saudi Construction Firms, *Dark Reading*. <https://www.darkreading.com/cyberattacks-data-breaches/ransomware-attacks-saudi-construction-firms>
- Liu, Y., A.H., A., Haron, N. A., N.A., B. and Wang, H. (2024). Robotics in the Construction Sector: Trends, Advances, and Challenges, *Journal of Intelligent & Robotic Systems* 110(2): 72. <https://doi.org/10.1007/s10846-024-02104-4>
- Mahaini, M. I., Li, S. and Sağlam, R. B. (2019). Building Taxonomies based on Human-Machine Teaming: Cyber Security as an Example, In *Proceedings of ARES '19*, New York, NY, USA: Association for Computing Machinery: 978-1-4503-7164-3, pp. 1–9. <https://doi.org/10.1145/3339252.3339282>
- MITRE. (2020). *Mitre Att&Ck*, MITRE.
- MITRE. (2025). CAPEC - Common Attack Pattern Enumeration and Classification. <https://capec.mitre.org/index.html>
- NERC. (2025). US Reliability Standards, *NERC*. <https://www.nerc.com/pa/Stand/Pages/USRelStand.aspx>
- NIST. (2024). *The NIST Cybersecurity Framework (CSF) 2.0*, NIST. <https://nvlpubs.nist.gov/nistpubs/cswp/nist.cswp.29.pdf>
- Obradović, R., Vasiljević, I., Kovačević, D., Marinković, Z. and Farkas, R. (2019). Drone Aided Inspection during Bridge Construction, In *ZINC 2019*, pp. 1–4. <https://doi.org/10.1109/ZINC.2019.8769345>
- OpenAI. (2026a). Responses, *OpenAI API Reference*. <https://developers.openai.com/api/reference/resources/responses/>
- OpenAI. (2026b). GPT-5 Chat Model | OpenAI API, *GPT-5 Chat Model | OpenAI API*. <https://developers.openai.com/api/docs/models/gpt-5-chat-latest>
- Orman, H. (2003). The Morris worm: a fifteen-year perspective, *IEEE Security & Privacy* 1(5): 35–43. <https://doi.org/10.1109/MSECP.2003.1236233>

- Oyediran, H., Turner, W., Kim, K. and Barrows, M. (2025). Integration of 4D BIM and Robot Task Planning: Creation and Flow of Construction-Related Information for Action-Level Simulation of Indoor Wall Frame Installation, *Journal of Information Technology in Construction (ITcon)* 30(15): 352–374. <https://doi.org/10.36680/j.itcon.2025.015>
- Prieto, S., García de Soto, B. and Adan, A. (2020). A Methodology to Monitor Construction Progress Using Autonomous Robots, In *Proceedings of ISARC 2020*, Kitakyshu, Japan. <https://doi.org/10.22260/ISARC2020/0210>
- Salami Pargoo, N. and Ilbeigi, M. (2023). A Scoping Review for Cybersecurity in the Construction Industry, *Journal of Management in Engineering* 39(2): 3122003. <https://doi.org/10.1061/JMENEA.MEENG-5034>
- Sánchez-García, I. D., Feliu, T. S. and Calvo-Manzano, J. A. (2025). Building a cyber risk treatment taxonomy, *Cluster Computing* 28(3): 205. <https://doi.org/10.1007/s10586-024-04899-1>
- Sawhney, A., Riley, M. and Irizarry, J. (2020). *Construction 4.0: An Innovation Platform for the Built Environment*, (A. Sawhney, M. Riley, & J. Irizarry, Eds), Routledge. <https://doi.org/10.1201/9780429398100>
- Schöberl, M., Schneider, M., Lechner, C. and Fottner, J. (2024). Virtual Construction Equipment Sensor for Determining Soil Stiffness during Compaction, Presented at the Proceedings of ISARC 2024, pp. 1182–1189. <https://doi.org/10.22260/ISARC2024/0153>
- Seo, J., Lee, S., Kim, J. and Kim, S.-K. (2011). Task planner design for an automated excavation system, *Automation in Construction* 20(7): 954–966. <https://doi.org/10.1016/j.autcon.2011.03.013>
- Sonkor, M. S. and García de Soto, B. (2026). A Cybersecurity Taxonomy for Operational Technology in Construction, In F. Rahimian, M. R. Hosseini, A. Akanmu, Z. Lafhaj, & L. Ducoulombier (Eds), *Proceedings of the International Conference on Smart and Sustainable Built Environment*, Singapore: Springer Nature: 9789819584895, pp. 593–607. https://doi.org/10.1007/978-981-95-8489-5_45
- Spyros, A., Kougioumtzidou, A., Papoutsis, A., Darra, E., Kavallieros, D., Tziouvaras, A., Tsikrika, T., Vrochidis, S. and Kompatsiaris, I. (2025). A comprehensive survey of manual and dynamic approaches for cybersecurity taxonomy generation, *Knowledge and Information Systems* 67(6): 4785–4811. <https://doi.org/10.1007/s10115-025-02382-w>
- Stouffer, K., Pease, M., Tang, C., Zimmerman, T., Pillitteri, V., Lightman, S., Hahn, A., Saravia, S., Sherule, A. and Thompson, M. (2023). *NIST Special Publication (SP) NIST SP 800-82r3, NIST Special Publication 800-82 Rev 2*, Gaithersburg, MD: NIST. <https://doi.org/10.6028/nist.sp.800-82r3>
- Tay, Y. W. D., Panda, B., Paul, S. C., Noor Mohamed, N. A., Tan, M. J. and Leong, K. F. (2017). 3D printing trends in building and construction industry: a review, *Virtual and Physical Prototyping* 12(3): 261–276. <https://doi.org/10.1080/17452759.2017.1326724>
- Wang, T., Mao, C., Sun, B. and Li, Z. (2024). Genealogy of construction robotics, *Automation in Construction* 166: 105607. <https://doi.org/10.1016/j.autcon.2024.105607>
- Warner, M. (2012). Cybersecurity: A pre-history, *Intelligence and National Security* 27(5): 781–799.
- Wu, H., Liu, Y., Chang, R. and Wu, L. (2024). Research Status Quo and Trends of Construction Robotics: A Bibliometric Analysis, *Journal of Computing in Civil Engineering* 38(1): 03123001. <https://doi.org/10.1061/JCCEE5.CPENG-5274>
- Xenofontos, C., Zografopoulos, I., Konstantinou, C., Jolfaei, A., Khan, M. and Choo, K. (2021). Consumer, Commercial, and Industrial IoT (In)Security: Attack Taxonomy and Case Studies, *IEEE Internet of Things Journal* 9: 199–221. <https://doi.org/10.1109/JIOT.2021.3079916>
- Zeng, L., Guo, S., Wu, J. and Markert, B. (2024). Autonomous mobile construction robots in built environment: A comprehensive review, *Developments in the Built Environment* 19: 100484. <https://doi.org/10.1016/j.dibe.2024.100484>

APPENDIX

This appendix presents the large language model (LLM) prompts used in this paper.

Prompt 1. The prompt used for retrieving OT- and process-specific candidate documents for each cell of the taxonomy:

“Instruction:

Using the Context to interpret taxonomy terms and the Input Data to define the active taxonomy cell, generate a list of key cybersecurity references that a construction management team can utilize for the specified cell. Prioritize references that provide implementable security requirements or controls.

Include only references whose primary purpose is cybersecurity; exclude documents that are mainly safety, machinery, functional safety, or general product compliance, even if they mention cybersecurity. Exclude items written mainly for OT manufacturers or equipment vendors.

Document Type rules (MANDATORY):

Each reference MUST be classified into EXACTLY ONE of the following document categories:

- 1) Regulation / Directive (e.g. GDPR, NIS 2 Directive)*
- 2) Standard / Code of practice (e.g. ISA/IEC 62443, ISO/IEC 27001)*
- 3) Guidance / Framework (e.g., NIST SP 800-53, NIST Cybersecurity Framework (CSF))*
- 4) Control catalog / Benchmark (e.g., CIS Critical Security Controls, CIS Benchmarks)*

Hard exclusions (MANDATORY):

- Exclude news pages, announcements, press releases, blog posts, event pages, webinar pages, podcasts, newsletters, and other editorial updates.*
- Include only documents that are primary references (regulations/directives, standards/codes of practice, guidance/frameworks, control catalogs/benchmarks).*

Also provide:

- Access Type: one of [open, paywalled, unknown]*
- Publication Year: 4-digit year when available; otherwise use unknown*
- If multiple years are visible (e.g., original release and later revisions), use the MOST RECENT publication/revision year from the official source page (prefer "latest version", "updated", or "version" metadata over older legacy dates).*

Context:

Taxonomy Framework:

· Columns (two levels)

— Level 1 — OT Family: Autonomous heavy construction machinery; Ground mobile robots; Legged robots; Aerial robots (UAVs); Climbing robots; Robotic manipulators; Construction 3D printers; Wearable/assistive robotics.

— Level 2 — OT Component (for each OT family): Sensing (e.g., cameras, LiDAR, IMUs, proximity sensors); Actuation (e.g., hydraulic or electric drives, extruder pumps for 3D printers); Control and compute (e.g., PLCs, embedded computers, microcontrollers); HMI (e.g., HMI panels, web UIs); Communication (e.g., fieldbus, Wi-Fi 6, 4G/5G).

· Rows — Construction Process Utilizing OT: Site preparation and earthworks; Surveying; Material handling; Concrete placement and finishing; Reinforcement and structural steel assembly; Prefabrication and modular assembly; Additive manufacturing (3D printing); Masonry and bricklaying; Finishing works; Progress, quality, and safety monitoring; Recycling and sorting construction waste.

Input Data:



- OT Family (Level 1 column): {OT_FAMILY}
- OT Component (Level 2 column): {OT_COMPONENT}
- Construction Process (row): {CONSTRUCTION_PROCESS}

Output Indicator:

Required Output Format (single markdown table only):

| Title | Publication Year | Publisher Organization | Document Type | Scope Level (international / regional) | Official URL | Official Description (Evidence) | Cell Applicability (Inference) | Access Type (open/paywalled/unknown) |

Formatting rules:

- Put links *ONLY* in the "Official URL" column. Do not include markdown links inside the Title cell.
- "Title" *MUST* be the official document title (not a website/domain name like nist.gov or iso.org).
- Do *NOT* wrap URLs in quotes or backticks. Use plain URLs only.
- Do *NOT* include URLs/citations/markdown links in any column except "Official URL".
- Official Description (Evidence): 1–2 sentences, paraphrase *ONLY* the official source page's scope/purpose. Do not tailor it to construction; do not add assumptions.
- Cell Applicability (Inference): 1–2 sentences explaining how the document can be used for *THIS* cell (OT family + OT component + construction process). You *MUST* label it as inference by starting the first sentence with exactly: "Inference:"
- Do *NOT* claim the document explicitly targets the OT family or construction process unless the Title or Official Description explicitly says so.
- Keep it practical (how it would be used), not promotional."

Prompt 2. The prompt used for retrieving construction-specific candidate documents:

"Instruction:

Provide a short list (max 5) of the most important cybersecurity references written specifically for the construction sector as a whole.

Document Type rules (MANDATORY):

Each reference *MUST* be classified into *EXACTLY ONE* of the following document categories:

- 1) Regulation / Directive (e.g. GDPR, NIS 2 Directive)
- 2) Standard / Code of practice (e.g. ISA/IEC 62443, ISO/IEC 27001)
- 3) Guidance / Framework (e.g., NIST SP 800-53, NIST Cybersecurity Framework (CSF))
- 4) Control catalog / Benchmark (e.g., CIS Critical Security Controls, CIS Benchmarks)

Hard exclusions (MANDATORY):

- Exclude news pages, announcements, press releases, blog posts, event pages, webinar pages, podcasts, newsletters, and other editorial updates.
- Include only documents that are primary references (regulations/directives, standards/codes of practice, guidance/frameworks, control catalogs/benchmarks).
- Exclude withdrawn, archived, superseded, or discontinued documents (e.g., pages that explicitly say "Withdrawn", "This guidance was withdrawn", "no longer updated", "archived", "superseded", "replaced").
- Exclude inaccessible pages (e.g., 404 Page not found, 403 Access denied, login/authorization required).
- If the URL redirects, follow the redirect and only include the final destination if it is still the same document and still fits the construction-sector scope; otherwise exclude it.



Also provide:

- *Publication Year*: 4-digit year when available; otherwise use *unknown*
- *Access Type*: one of [open, paywalled, unknown]

Context:

Definition of “construction sector as a whole”: References written for construction organizations and project delivery (e.g., contractors, project teams, construction site operations, construction supply chain), not for building operations/facilities management only.

Output Indicator:

Required Output Format (single markdown table only):

| Title | Publication Year | Publisher Organization | Document Type | Scope Level (international / regional) | Official URL | Official Description (Evidence) | Construction-sector Applicability (Inference) | Access Type (open/paywalled/unknown) |

Formatting rules:

- Put links *ONLY* in the “Official URL” column. Do not include markdown links inside the Title cell.
- “Title” *MUST* be the official document title (not a website/domain name like *nist.gov* or *iso.org*).
- Do *NOT* wrap URLs in quotes or backticks. Use plain URLs only.
- Do *NOT* include URLs/citations/markdown links in any column except “Official URL”.
- “Official URL” must be the final, currently accessible landing page (no withdrawn/archived warnings; no 403/404; no “no longer updated” banners; no redirect-only URLs).
- Official Description (Evidence): 1–2 sentences, paraphrase *ONLY* the official source page’s scope/purpose (do not add assumptions).
- Construction-sector Applicability (Inference): 1–2 sentences explaining how the document can be used by construction organizations/project delivery teams.
- Do *NOT* claim the document is “construction-specific” unless the Title or Official Description explicitly says so.
- Keep it practical (how it would be used), not promotional.”

Prompt 3. The prompt used for QC and deduplication of candidate documents:

“Instruction:

You are doing a quality-control pass on a CSV of cybersecurity references written for the construction sector as a whole.

Return *ONLY* valid JSON (no markdown, no code fences, no commentary).

Tasks (MANDATORY):

1) Title QC:

- Detect wrong entries in Title (e.g., a domain like “*cisa.gov*” or “*iso.org*”).
- Correct them using the URL, publisher, Official Description (Evidence), and Construction-sector Applicability (Inference).
- If you cannot confidently infer a proper title, set Title to “UNKNOWN TITLE (manual review)”.

2) Deduplication:

- Remove duplicate references by comparing titles and identifiers.
- Example duplicates: “CIS Critical Security Controls Version 8.1” == “CIS Critical Security Controls v8.1”.
- When merging duplicates, keep *ONE* row and choose:
- the most informative Title (include version/year if known),



- the clearest Publisher Organization,
- the best Official URL (prefer stable canonical landing pages; remove tracking params if obvious).
- NEVER merge different parts/documents from the same standards family/series.

Example: do NOT merge "IEC 62443-3-2" with "IEC 62443-2-1" or "IEC PAS 62443-2-2".

3) Document Type QC:

- Correct Document Type so it is EXACTLY ONE of:

A) Regulation / Directive (e.g. GDPR, NIS 2 Directive)

B) Standard / Code of practice (e.g. ISA/IEC 62443, ISO/IEC 27001)

C) Guidance / Framework (e.g., NIST SP 800-53, NIST Cybersecurity Framework (CSF))

D) Control catalog / Benchmark (e.g., CIS Critical Security Controls, CIS Benchmarks)

- IMPORTANT: ISO/IEC/IEC/ISA/ANSI standard documents MUST be "Standard / Code of practice" (not Regulation / Directive).

4) Publisher Organization unification:

- Normalize publisher names (e.g., "ENISA (EU Agency for Cybersecurity)" -> "ENISA").
- Keep well-known official names short and consistent.

5) Scope Level QC:

- Scope Level must be exactly one of: "international" or "regional".
- Decide primarily based on issuer/publisher mandate:
- International examples: ISO, IEC, ISA, CIS, ITU, ETSI, IEEE, GSMA (global/standards bodies)
- Regional examples: EU bodies (e.g., ENISA) or national agencies (e.g., NIST, CISA, NCSC UK)
- If unclear, choose the most defensible option based on publisher.

6) Publication Year QC:

- Must be a 4-digit year or "unknown".
- If multiple years are visible for the same document (original release vs later revision), keep the MOST RECENT publication/revision year.
- If the same document appears in multiple versions, keep ONLY the latest version and remove older versions.

Coverage and merge constraints (MANDATORY):

1) Input row coverage is REQUIRED.

- There are exactly {total_rows} input rows, with row_id values 1..{total_rows}.
- Every input row_id MUST appear in source_row_ids of exactly ONE output object.
- Do not omit any row_id.

2) Merge conservatively. Merge only if rows clearly refer to the SAME document.

- Acceptable merge evidence (at least one):

a) same canonical URL (ignoring tracking query params and trivial URL formatting differences), OR

b) same official identifier/version in title (e.g., "IEC 62443-3-3", "SP 800-82 Rev. 3", "ISO/IEC 27001:2022"), OR

c) near-identical title meaning AND same publisher.

- If uncertain, DO NOT merge.

3) Do not discard rows due to low confidence.



- If title is uncertain, keep the row and set Title to "UNKNOWN TITLE (manual review)".
- Still include its row_id in source_row_ids.

Input Data:

Input rows: `{json.dumps(payload, ensure_ascii=False)}`

Output Indicator:

Output JSON schema (MANDATORY):

```
{{
  "cleaned_documents": [
    {{
      "source_row_ids": [1, 7, 12],
      "Title": "...",
      "Publication Year": "2024",
      "Publisher Organization": "...",
      "Document Type": "Regulation / Directive | Standard / Code of practice | Guidance / Framework | Control
catalog / Benchmark",
      "Scope Level (international / regional)": "international|regional",
      "Official URL": "...",
      "Official Description (Evidence)": "...",
      "Construction-sector Applicability (Inference)": "...",
      "Access Type (open/paywalled/unknown)": "open|paywalled|unknown"
    }}
  ]
}}
```

Prompt 4. The prompt used for scoring taxonomy cell-specific candidate documents:

“Instruction:

You are scoring a list of cybersecurity references for this OT taxonomy cell:

- OT Family: `{ot_family}`
- OT Component: `{ot_component}`
- Construction Process: `{construction_process}`

Return ONLY valid JSON (no markdown, no code fences, no commentary).

Scoring criteria definitions (MANDATORY to apply):

<SCORING RUBRIC*>

Scoring ranges (MANDATORY):

- Implementation usefulness: integer 0..10
- Family relevance: integer 0..10
- Component relevance: integer 0..10
- Construction process relevance: integer 0..10

Rules:



- 1) Score *EVERY* input row_id exactly once.
- 2) Use only the provided row content (no web lookup).
- 3) Use the rubric anchors above for consistency across rows.
- 4) Intermediate integers (1,3,5,7,9) are allowed only when clearly between adjacent anchor definitions.
- 5) Do not award scores >6 for Family/Component/Process relevance unless the Title or Official Description (Evidence) contains strong support (not only Cell Applicability (Inference)).

Input rows: {json.dumps(payload, ensure_ascii=False)}

Output JSON schema (MANDATORY):

```
{
  "scores": [
    {
      "row_id": 1,
      "Implementation usefulness": 0,
      "Family relevance": 0,
      "Component relevance": 0,
      "Construction process relevance": 0
    }
  ]
}
```

*The full prompt used in the Python script includes the scoring rubric presented in Table 5.

Prompt 5. The prompt used for scoring construction-specific candidate documents:

“Instruction:

You are scoring a list of cybersecurity references for the construction sector as a whole (construction organizations and project delivery).

Return ONLY valid JSON (no markdown, no code fences, no commentary).

Scoring criteria definitions (MANDATORY to apply):

<SCORING RUBRIC>*

Scoring ranges (MANDATORY):

- Implementation usefulness: integer 0..10

- Construction-sector relevance: integer 0..10

Rules:

- 1) Score *EVERY* input row_id exactly once.
- 2) Use only the provided row content (no web lookup).
- 3) Use the rubric anchors above for consistency across rows.
- 4) Intermediate integers (1,3,5,7,9) are allowed only when clearly between adjacent anchor definitions.
- 5) Do not award Construction-sector relevance >6 unless the Title or Official Description (Evidence) contains strong support (not only Construction-sector Applicability (Inference)).

Input rows: {json.dumps(payload, ensure_ascii=False)}

Output JSON schema (MANDATORY):

```
{{
  "scores": [
    {{
      "row_id": 1,
      "Implementation usefulness": 0,
      "Construction-sector relevance": 0
    }}
  ]
}}
```

*The full prompt used in the Python script includes the scoring rubric presented in Table 7 and the implementation usefulness rubric from Table 5.