Building an Open AIBOM Standard in the Wild: An Experience Report on Extending the SPDX SBOM (ISO/IEC 5962:2021) for AI Supply Chains

Gopi Krishnan Rajbahadur
Queen's University
Kingston, Canada
grajbahadur@acm.org

Keheliya Gallaba
Queen's University
Kingston, Canada
gallabak@sigsoft.org

Elyas Rashno
Queen's University
Kingston, Canada
elyas.rashno@queensu.ca

Arthit Suriyawongkul
ADAPT Centre, Trinity College Dublin
Dublin, Ireland
suriyawa@tcd.ie

Karen Bennet
IEEE
Toronto, Canada
karen.bennet@gmail.com

Kate Stewart
Linux Foundation
Austin, USA
kstewart@linuxfoundation.org

Ahmed E. Hassan
Queen's University
Kingston, Canada
ahmed@cs.queensu.ca

Abstract

Modern software engineering increasingly relies on open, community-driven standards, yet how such standards are created in fast-evolving domains like AI-powered systems remains underexplored. This paper presents a detailed experience report on the development of the AI Bill of Materials (AIBOM) specification, an extension of the ISO/IEC 5962:2021 Software Package Data Exchange (SPDX) software bill of materials (SBOM) standard, which captures AI components such as datasets and iterative training artefacts. Framed through the lens of Action Research (AR), we document a global, multi-stakeholder effort involving over 90 contributors and structured AR cycles. The resulting specification was validated through four complementary approaches: alignment with major regulations and ethical standards (e.g., EU AI Act and IEEE 7000 standards), systematic mapping to six industry use cases, semi-structured practitioner interviews, and an industrial case study. Beyond delivering a validated artefact, our paper documents the process of building the AIBOM specification "in the wild," and reflects on how it aligns with the AR cycle, and distills lessons that can inform future standardisation efforts in the software engineering community.

CCS Concepts

• **Computing methodologies** → **Machine learning**; • **Software and its engineering** → **Software development techniques**.

Keywords

standardisation, open source community, Action Research, AIBOM, software bill of materials, SBOM, AI governance, trustworthy AI, software supply chain security, Software Package Data Exchange, SPDX

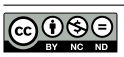

ICSE-SEIP '26, Rio de Janeiro, Brazil

ACM ISBN 979-8-4007-2426-8/26/04
https://doi.org/10.1145/3786583.3786921

ACM Reference Format:
Gopi Krishnan Rajbahadur, Keheliya Gallaba, Elyas Rashno, Arthit Suriyawongkul, Karen Bennet, Kate Stewart, and Ahmed E. Hassan. 2026. Building an Open AIBOM Standard in the Wild: An Experience Report on Extending the SPDX SBOM (ISO/IEC 5962:2021) for AI Supply Chains. In *2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE-SEIP '26), April 12–18, 2026, Rio de Janeiro, Brazil.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3786583.3786921

1 Introduction

Standards play a critical role in Software Engineering (SE) by ensuring consistency, interoperability, reliability, and trustworthiness across systems and end users [26]. They underpin much of the discipline. For instance, ISO/IEC 5962:2021, Software Package Data Exchange (SPDX), offers a shared vocabulary for documenting software supply chains (SSCs) that has been widely adopted to assess vulnerabilities and compliance in received software.

Regulators also rely on standards as mechanisms for compliance. In the EU, conformity with the Medical Device Regulation can be shown through harmonised standards such as ISO 13485 (quality management) and IEC 62304 (software lifecycle processes). In the U.S., Software Bill of Material (SBOM) standards like SPDX and CycloneDX are explicitly referenced in federal cybersecurity guidance [13, 60]. These SBOM standards have spurred a plethora of recent work that has examined SBOM adoption, proposed taxonomies, and developed extensions for vulnerability management [86, 98], which underscores their growing importance.

These SBOM standards target traditional software and cannot meaningfully capture AI-specific artefacts such as datasets, trained models, fine-tuned checkpoints, and data pipelines that dominate modern systems [96]. To address this limitation, the AI community has introduced transparency frameworks such as model cards [58], datasheets [33], and factsheets [21], as well as machine-readable formats like Croissant [22]. Yet, these approaches often isolate AI-specific details from the broader software engineering context in which practitioners build and deploy systems. Sculley et al. [78] showed in their seminal work on hidden technical debt that AI components typically form only a small fraction of much larger systems, underscoring the need to document them *in situ* within

the broader software context. Practitioners thus face a critical gap: they lack a unified, standard mechanism to describe AI components as first-class citizens within the software supply chain [37, 71].

Our paper presents an experience report on tackling this challenge through the creation of the AI Bill of Materials (AIBOM) specification, an extension to the SPDX 3.0 standard. Our extension comprises 36 new fields that treat datasets, models, and their provenance as first-class supply-chain elements to address the trustworthiness challenges. Our work makes a distinct contribution by showing how AR can be adapted from its traditional use within a single organisation [5] to guide a global, multi-stakeholder standardisation effort "in the wild." Through this lens, our AIBOM specification itself emerges as a validated artefact of a disciplined, iterative process. Over 90 contributors participated in *diagnose-design-evaluate-reflect* cycles, providing a replicable blueprint for developing standards in fast-moving domains in SE, while delivering a practical, community-driven specification.

AIBOM specification has since landed in the SPDX standard, been piloted in industrial settings, and informed regulatory discussions. While the technical specification and schema live in our whitepaper [7], this paper reports our *experience* in open standardisation, detailing our governance, cadence, and pitfalls. We also demonstrate AIBOM specification's practical utility using four complementary validation methods: (1) **regulatory/standards alignment** (EU AI Act, US and EU medical-device guidance, IEEE 7000 series) to support compliance practice; (2) **systematic mapping to industry use cases** (compliance, risk management, supply-chain governance) with coverage analysis; (3) **practitioner interviews** (n=10) across roles confirming field relevance and surfacing gaps; and (4) a **multinational industrial case study** showing AIBOM covers ethics/legal checklists, generates portions of model cards, and enables partial automation for third-party AI assets. Together, these results position AIBOM specification as part of an actionable standard and our process as a blueprint for creating standards in fast-moving, multi-stakeholder SE domains.

Contributions. We report how we created *AIBOM* specification through an open, global AR process and why this approach worked in practice. Our aim is not to claim that SPDX AIBOM specification is the definitive AIBOM solution as credible alternatives such as CycloneDX exist and continue to evolve. Specifically:

- **AR at scale.** We demonstrate how AR, traditionally used within single organisations [5], can be adapted to guide a global, multi-stakeholder standardisation effort "in the wild." Our account documents governance design, release cadence, decision traceability, and iterative cycles that extended SPDX SBOM standard with 36 AI and Dataset specific fields.
- **Validated artefact.** We present AIBOM and demonstrate its practicality via four complementary validations.
- **Lessons and blueprint.** We distill actionable patterns and lessons learned for creating standards in fast-moving, multi-stakeholder domains.

2 Background

This section contextualises our work by examining the formal processes for creating software engineering standards, the technical foundation of the SBOM that requires extension.

Table 1: Representative categories of software engineering standards and their scope.

Category	Examples	Purpose / Scope
Lifecycle process	ISO/IEC/IEEE 12207, 15288	Define processes for development, maintenance, and complete system lifecycles.
Quality	ISO/IEC 25010	Provide models for evaluating core software quality attributes.
Testing & verification	ISO/IEC/IEEE 29119	Standardise testing methods, documentation, and reporting practices.
Supply chain & compliance	ISO 28000, GS1, CIS Controls	Secure supply chains and ensure regulatory and ethical conformance.
Emerging AI-related	ISO/IEC JTC 1/SC 42	Address risk, transparency, and trust in AI systems.

2.1 Development and Extension of Software Standards

Standards in SE. Standards are vital in SE, providing common processes and terminology that minimise ambiguity, enhance interoperability, and ensure reliable high-quality system delivery [94]. Current standards span lifecycle processes, quality assessment, testing, supply-chain security, and emerging AI considerations (Table 1). These frameworks underpin the structured development and governance of modern enterprise software, fostering consistent engineering practices across diverse domains and organisations.

Standards development and extension pathways. Standards in SE typically evolve through two complementary pathways. The first is the **formal route**, followed by standards development organisations (SDOs) such as IEEE and ISO/IEC, which emphasise stability, consensus, and global legitimacy. This approach is characterised by well-defined, gated stages including project authorisation, working group (WG) formation, draft development, public review, and final balloting [43]. These stages typically span 18–48 months for IEEE standards and 24–36 months for ISO/IEC standards [42, 43, 94]. The first version of IEEE POSIX standard, for example, progressed through this process via multiple iterations and ballots before becoming ISO/IEC 9945 [45].

The second pathway is an **upstream, community-led extension model**, in which specifications evolve collaboratively in open WGs, consortia, or technical bodies before entering the formal process. This "implementation-first" or "parallel" model [8] underpins the evolution of major standards: POSIX is maintained by the Austin Group prior to IEEE and ISO ratification [92]; amendments to IEEE 802.11 (Wi-Fi) are incubated in dedicated task groups [40]; and SPDX itself matured within the Linux Foundation community before being standardised as ISO/IEC 5962:2021 [44]. Such upstream-first approaches offer advantages such as faster iteration, broader participation, and validation against real-world practice. These factors are critical in fast-evolving domains such as AI. Our work deliberately adopts this second model.

2.2 Software Bill of Materials (SBOMs)

From BOM to SBOM. The concept of a Bill of Materials (BOM) originated in manufacturing as a structured inventory of components and sub-assemblies within a product [49]. Applying this principle to software, a Software Bill of Materials (SBOM) provides a machine-readable inventory of software components, enabling organisations to understand what software they run, where it originates, and

how it can be trusted [97]. SBOMs have become central to software supply chain security as modern systems increasingly depend on third-party and open-source components [86].

SBOM standards. Three SBOM standards dominate the landscape: SPDX [32], developed under the Linux Foundation and standardised as ISO/IEC 5962:2021 [44]; CycloneDX [64], introduced by OWASP and standardised as EMCA-424 [19]; and SWID Tags [46], maintained by the U.S. National Institute of Standards and Technology (NIST) to provide transparent software identification. While SPDX and CycloneDX are both widely adopted and offer comparable core capabilities, their selection in practice is typically guided by organisational priorities.

Limitations for AI systems. Despite their value, current SBOMs standards do not capture the full complexity of AI systems. New artefact types, such as datasets and models, introduce deeper provenance, explainability, and compliance requirements. Dataset quality and lineage directly influence performance and fairness [54, 71]; model opacity necessitates metadata on interpretability, hyperparameters, and operational constraints [3]; and iterative lifecycles, including pre-training, fine-tuning, and redeployment, create intricate provenance chains and evolving compliance obligations [37]. These limitations have led both researchers and practitioners to call for extensions to SBOM standards that address the specific needs of AI systems, laying the foundation for the concept of an AIBOM.

3 Methodology

We followed a participatory and iterative approach to develop the AIBOM specification in the open, engaging multiple, diverse stakeholders. Figure 1 illustrates the process we followed.

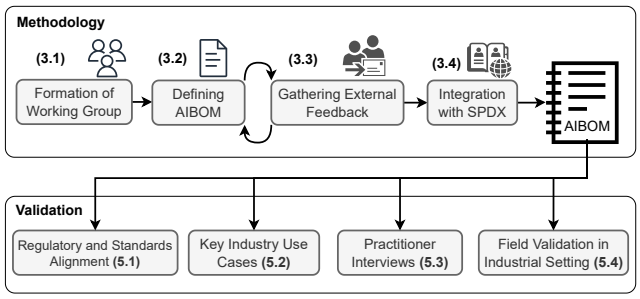

Figure 1: Overview of AIBOM standard creation process and multi-faceted validation.

3.1 Working Group Formation

Motivation and scope. The AI Working Group (WG) [36] was formed in 2021 under the SPDX community to explore how existing SBOM standards could be extended to represent AI systems. Its primary goal was to determine whether SBOMs captured the artefacts and metadata needed for regulatory compliance, license obligations, auditability, and transparency in AI systems.

Recruitment and participation. From the outset, the WG held weekly one-hour virtual meetings open to anyone interested. Early participation (2021 to mid-2022) relied on direct invitations and word-of-mouth recruitment through professional networks. To broaden engagement, the group launched a public mailing list on 7 June 2022, allowing contributors to sign up independently and join ongoing discussions.

Structure and operation. The WG ran open agendas, made decisions by consensus, and used mailing lists and GitHub for asynchronous discussion and review. Participants included scientists, researchers, professors, CTOs, product managers, AI and software developers, and legal and licensing experts.

> **Outcome.** By May 2024, shortly after the official release of SPDX 3.0, the WG had held 82 meetings with 92 unique participants, 20 of whom attended more than six sessions. All meetings since April 2022 were recorded and archived publicly for transparency and traceability [35].

3.2 Defining the AIBOM Specification

Initial field definition. The WG's first major task was to identify the core information an AIBOM specification should capture. Drawing on their diverse expertise and relevant research, participants proposed an initial set of fields during weekly meetings in early 2022. The objective was to ensure that the AIBOM could represent essential aspects of AI systems including datasets, models, and their associated metadata. At this stage, the focus was on brainstorming and cataloging all potentially relevant fields.

Incorporating existing practices. From July 2022, once a comprehensive list had been drafted, members systematically analysed established documentation artefacts such as model cards [58], datasheets [33], and factsheets [21] to refine and expand the field set. The goal was not exhaustive mapping, but to identify the most relevant and widely applicable fields recognised as useful and appropriate in the AIBOM context. A detailed mapping of the fields from these sources that were incorporated, excluded and the reasons we did so, is provided in our whitepaper [7].

Balancing adoption and completeness. Throughout this process, the WG prioritised adoption over comprehensiveness by defining a minimal set of *required fields*—those most likely to exist in real-world projects and satisfy regulatory or auditing needs. Each field had to meet strict inclusion criteria: (1) relevance to AIBOM goals, (2) availability of a representation method, and (3) consensus on its necessity. More ambitious metadata elements were intentionally deferred to future releases. In line with NTIA [59] recommendations and emerging best practices, the WG focused on a small, readily adoptable set of required fields. These fields were refined continuously until May 2024 through iterative review and discussion.

> **Outcome.** The WG evaluated 103 candidate fields derived from existing tools and community proposals. The final specification defined 36 fields: 20 for the AI profile (five required) and 18 for the Dataset profile (six required).

3.3 External Feedback and Public Consultation

Presenting drafts and gathering early feedback. WG members frequently presented draft versions of the AI and Dataset profiles (that form the AIBOM specification along with the rest of the SPDX profiles, please see Section 4 for more details) at several formal venues, including Open Source Summit Europe (2022, 2023), Open Source Summit Japan (2022), and Open Source Summit North America (2023, 2024, 2025) and semi-formal WGs including the OpenChain Licensing WG [12], OpenSSF AI WG [62], and IEEE P7014.1 WG [41]. These events brought together open-source experts, software practitioners, and legal professionals, whose feedback was then discussed in WG meetings and incorporated into subsequent revisions of the specification.

Public release candidates and community review. Starting on 8 May 2023, the WG initiated a year-long public consultation process, releasing two versions of the AIBOM specification: two release candidates (RC1 and RC2) [31]. The SBOM tooling community was invited to review the specification, model, and profiles and to submit proposed changes as pull requests (PRs) to the public repository. In total, 20 PRs specific to AIBOM were submitted, 16 of which were accepted after detailed discussion in WG meetings.

Integration with SPDX. In parallel with the release candidate process, the WG collaborated with the broader SPDX project to ensure that the proposed extensions were incorporated and could work seamlessly with other elements of the SPDX specification. Since several new fields required changes to the underlying data model, the group regularly presented proposed modifications to the SPDX Technical Team (responsible for the specification's architecture and publication). These discussions led to coordinated updates across all SPDX WGs, enabling seamless integration of the AIBOM related profiles and fields into SPDX 3.0 prior to its final release.

Outcome. After multiple rounds of refinement, SPDX 3.0, including the AI and Dataset profiles (please see Section 4 for more details) that encapsulate the AIBOM specification was officially released on 16 April 2024. The release was accompanied by a Linux Foundation blog post [30] and a detailed whitepaper [7] to support adoption and provide implementation guidance to the broader community.

4 AIBOM

Building on the methodology described in Section 3, the working group integrated two new profiles: **Dataset** and **AI**, into the SPDX 3.0 specification. These extensions expand the SBOM standard beyond traditional software components to capture the artefacts underpinning AI systems. Figure 2 shows how the various SPDX profiles together represent an AIBOM specification (thereby extending the SBOM standard to represent AI systems. In our paper, we refer to this as the AIBOM standard or AIBOM specification.

Designed to support regulatory compliance, licence attribution, provenance tracking, and responsible development, these extensions remain fully compatible with existing SPDX tooling. A detailed overview of the AIBOM specification is available in the

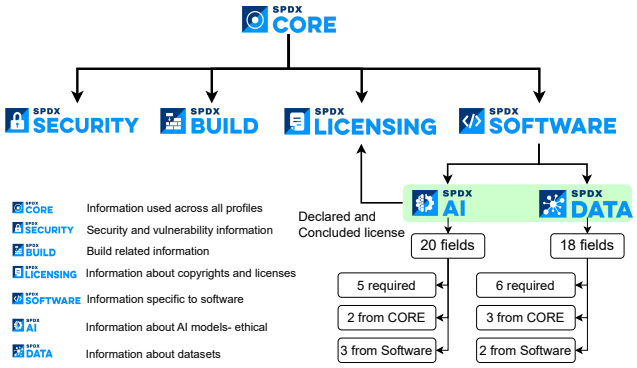

Figure 2: Overview of SPDX profiles and their extensions for AI and Dataset.

SPDX 3.0 white paper [7] and the public repository.[1] As this paper focuses on the specification's creation process over its technical details, we provide only a high-level summary below.

Profiles and extensibility in SPDX. SPDX 3.0 adopts a profile-based modular architecture that enables domain-specific extensions while preserving interoperability. A *profile* defines a coherent subset of the model: the *Core* profile establishes foundational classes and vocabularies, while additional profiles extend these capabilities to represent licencing, build metadata, security, and software composition.

Within this framework, the AIBOM specification for an AI system is realised through the *AI* and *Dataset* profiles working alongside existing components (though they can also be used independently to document models or datasets). Rather than exist as standalone artefacts, these profiles extend SPDX in a modular way without fragmenting the broader ecosystem (Figure 2). They build on existing classes and relationships from the *Core*, *Licensing*, and *Software* profiles while introducing domain-specific metadata essential for describing AI artefacts. Together, they allow datasets and models to be represented as first-class supply-chain elements within the same SBOM framework.

AI profile. The AI profile captures AI model-specific components, capturing provenance (identifiers, versions, and licences), architecture (type, domain, and evaluation metrics), training details (data sources, configurations, and resource usage), and model risks (limitations, safety considerations, and compliance evidence). It provides a comprehensive view of how models are developed, deployed, and maintained.

Dataset profile. The Dataset profile standardises how datasets are described, capturing provenance (origins, sources, suppliers, and collection processes), descriptive details (size, modality, structure, and preprocessing), intended use (applications and purposes), and risk information (biases, sensitive data, and anonymisation methods). This enables traceability, accountability, and compliance throughout the AI lifecycle.

[1] https://github.com/spdx/spdx-3-model

Table 2: Overview of key standards and regulations considered for AIBOM alignment

Studied Standard/Regulation	Objective of the Standard/Regulation
EU Artificial Intelligence Act (EU AI Act) [24]	Ensure safe deployment, fundamental rights protection, and risk-based regulation of AI.
EU Medical Device Regulation (EU MDR) [23]	Demonstrate conformity and document risk management for medical devices placed on the EU market.
US Food and Drug Administration (FDA) Guidance [9]	Support pre-market evaluation of software-based medical devices and promote cybersecurity risk management.
IEEE 7000 Series [85]	Promote ethical design, transparency, bias mitigation, and data governance in autonomous and intelligent systems.

5 Validation

To demonstrate that our proposed AIBOM specification is comprehensive, practical, and fit for purpose, we designed a multi-faceted validation strategy. This process evaluated our specification from four distinct but complementary angles which we detail in the following subsections.

5.1 Regulatory and Standards Alignment

Motivation. It is essential that our AIBOM specification captures all metadata required to verify compliance with key regulations and standards governing AI systems. Practical adoption depends on whether a supply chain standard enables Open Source Program Offices (OSPOs) and compliance auditors to fulfill their legal and governance obligations. In this section, we validate whether the fields defined in AIBOM specification can represent the information demanded by widely adopted and emerging standards. We focus on the EU AI Act, U.S. and EU medical device regulations, and the IEEE 7000 series, as they collectively span regulatory, safety-critical, and ethical requirements across diverse domains. Table 2 summarises the studied standards and regulations. Due to their extensive scope, we focused our analysis on the most relevant clauses and information elements that directly impact software supply chain compliance.

Methodology. We analysed the alignment between fields captured in the AIBOM specification and the regulatory and ethical requirements defined by the standards and regulations shown in Table 2. The analysis was performed by the fourth and fifth author, who have extensive experience in standards development and regulatory alignment, and was further independently reviewed by first and the sixth authors for additional validation.

EU AI Act., we examined the obligations defined in Articles 49 and 60, which govern mandatory registration in the EU high-risk AI database prior to market entry or testing in sandbox. Required information was grouped into four categories: *Identification*, *System Details*, *Verification Details*, and *Application Details*, comprising a total of 14 subcategories (e.g., provider contact, intended purpose, classification, testing plan, and involved parties).

US and EU medical device regulations, we analysed our AIBOM specification's alignment with Article 10 of Regulation (EU) 2017/745 and FDA cybersecurity guidance. Information elements were organised into three categories: *package details*, *model details*, and *data details*.

IEEE 7000 series, we evaluated over 40 subclauses from eight standards (7000, 7001, 7002, 7005, 7007, 7009, 7010, and 7014), covering key areas such as transparency, privacy, data governance, and algorithmic bias.

Results. Our analysis shows that the SPDX 3.0 AIBOM specification provides strong coverage of regulatory, safety, and ethical requirements. We found that we could successfully represent **13 of 14** information obligations under Articles 49 and 60 of the EU AI Act, capturing **all required elements** from US and EU medical device regulations. Our validation shows that AIBOM specification aligns with **over 40 subclauses** across eight IEEE 7000 series standards i.e., more than **90%** overall coverage indicating practical readiness for compliance adoption.

However, the analysis also revealed limitations that highlight opportunities for future refinement. While AIBOM captures most metadata required by the EU AI Act, it cannot yet explicitly represent the "parties involved in testing" relationship mandated by Annex IX (2), even though contact details can be documented in existing fields. Similarly, the strong alignment with medical device regulations is partly due to their reliance on free-text "summary" and "description" fields, which map to generic SPDX constructs such as `informationAboutApplication` and `comment` but may limit traceability and automation. Finally, although more than 40 IEEE 7000 subclauses on transparency, privacy, data governance, and algorithmic bias are covered, areas such as child and student data governance, robotic nudging, and environmental and social governance remain outside the scope of the current AIBOM specification. A detailed field-level mapping of these standards and their alignment with SPDX 3.0 is available in the official white paper [7].

5.2 Validation against Key Industry Use Cases

Motivation. While regulatory alignment ensures that our AIBOM specification meets legal and compliance obligations, its practical value ultimately depends on how well it addresses real-world needs identified by industry stakeholders. In June 2025, the *SBOM for AI Tiger Team* [75], convened by the U.S. Cybersecurity and Infrastructure Security Agency (CISA), articulated six foundational use cases that any AI-focused SBOM standard must support. These use cases span the full lifecycle of AI system deployment: *regulatory compliance*, *Vulnerability and Incident Management*, *legal and intellectual property assurance*, *third-Party AI Risk Management*, *Open Source Model Risk Assessment* and *Model Lifecycle and Asset Management*. A detailed overview of these usecases can be found here [28].

Validating our AIBOM specification against these use cases is therefore essential to demonstrate that the specification is not only standards-compliant but also operationally relevant, usable by organisations in practice, and capable of supporting the core security, legal, and governance workflows envisioned by the broader AIBOM consumer ecosystem.

Methodology. The CISA "SBOM for AI" use cases are intentionally high-level; however, assessing whether our AIBOM specification is capable of capturing the metadata outlined in them requires concrete, testable requirements. We therefore operationalised these use cases by consolidating them into broader themes, assembling a focused evidence base from recent studies, extracting atomic requirements, and mapping them to SPDX 3.0 AIBOM fields before performing external validation.

Step 1: Consolidate use cases into themes (with rationale). Because several of the six use cases overlap conceptually and are instances of broader goals, we collapsed them into three themes to avoid double-counting and align with the structure of available evidence: (i) **Compliance** (Compliance; Legal & IP Protections), (ii) **Open Source Risk Management** (Vulnerability/Incident Management; Open Source Model Risk Assessment), and (iii) **Supply Chain Management** (Third-Party AI Risk Management; Model Lifecycle & Asset Management).

Step 2: Build an evidence base per theme. For each theme, we identified relevant research papers and technical reports that analyse concrete instances of the use cases. For example, Rajbahadur et al. [71], which outlines methods for conducting dataset license compliance analysis, was mapped to the *Compliance* category. All selected studies are listed in Table 3.

Step 3: Extract atomic requirements. The third author systematically reviewed each source and extracted concrete, verifiable requirements specific to its theme (template: *requirement statement*, *evidence snippet/page*, *rationale*). For example, Rajbahadur et al. [71] examine licensing and provenance risks in using public datasets for commercial AI; from this, we derived the requirement to *capture original data sources and licensing/redistribution constraints for each dataset and derivative build.* This process was repeated for all sources under each theme.

Step 4: Map requirements to AIBOM. We created a traceability matrix linking each extracted requirement to specific SPDX 3.0 AIBOM fields, indicating whether and how the requirement could be addressed by existing fields.

Step 5: Internal and external validation. The first author independently verified the extractions and mappings. We then presented the traceability matrix and representative examples to the *AI SBOM Tiger Team* (29 Sept 2025, attended by 12 participants) and the *SPDX AI and Dataset WG* (attended by four participants). Authors of this paper who were part of the working group did not participate in the validation exercise. Feedback was solicited on whether AIBOM could be effectively operationalised to address the six use cases. Meeting minutes from both working group sessions are publicly available.[2]

Results. Our analysis demonstrated that the AIBOM specification could represent all **46 distinct requirements** extracted from the eight studies reviewed across the three consolidated use-case categories. Table 3 summarises the number of AIBOM fields across the AI, Dataset, and other SPDX profiles that satisfied the requirements identified in each study. In several cases, a single requirement could be addressed by multiple fields, underscoring the flexibility and composability of our AIBOM specification. A detailed traceability matrix, including the full set of extracted requirements and their field-level mappings, is available in our extended technical report [73].

External validation. Feedback from participants in both the *SPDX AI & Datasets* and *AI SBOM Tiger Team* working groups was broadly positive. Reviewers agreed that the methodology and resulting mappings were sound and well-aligned with practical needs, while noting that additional time would be required for an exhaustive, line-by-line review. As our primary objective was to identify any major conceptual gaps or misinterpretations rather than obtain final endorsement, this feedback was encouraging. Both groups have committed to conducting deeper evaluations in future iterations, and we intend to incorporate any outcomes available before the camera-ready version of this paper.

[2]SBOM for AI (AIBOM) Tiger Team working group meeting minutes https://docs.google.com/document/d/1IpXG7XBOJnPl_hwFf3JZkDaFb0k2CnI0/edit?usp=sharing&ouid=110194678381965933391&rtpof=true&sd=true

Table 3: Alignment of consolidated AI SBOM use cases and the number of mapped fields across the AI, Dataset, and other SPDX profiles.

Categories	References	AI profile	Dataset profile	Other profiles
Compliance	[29]	18	7	8
	[71]	5	11	3
	[53]	4	3	3
Open Source Risk Management	[93]	6	5	4
	[48]	3	8	4
	[87]	5	4	2
Supply Chain Management	[51]	7	3	3
	[76]	5	6	3

5.3 Validation based on Practitioner Interviews

Motivation. While the design of our AIBOM specification was driven by a large, multi-stakeholder working group, it was essential to validate the resulting fields through an independent lens. To ensure that our specification was both understandable and practical beyond the community that developed it, we conducted semi-structured interviews with practitioners who were *not* involved in the AIBOM working groups. Their perspectives served as an external "sanity check", allowing us to assess whether the proposed fields align with real-world needs, uncover overlooked requirements, and identify potential barriers to adoption in industry practice.

Methodology. This validation occurred prior to the official AIBOM release, ensuring participants identified essential fields based on practical needs rather than reacting to the limitations of a finalised specification. Our methodology is detailed below.

Participant recruitment. We recruited ten practitioners via professional networks and targeted emails between 2022 and mid-2023. Participants were selected based on their expertise in SBOM development or AI software engineering to ensure a multi-disciplinary evaluation of the AIBOM framework.

Table 4 details the ten participants, who represent diverse work domains, experience levels, and potential future roles within the AIBOM ecosystem.

Interview protocol. After refining an initial questionnaire based on feedback from two academic and industry experts, we finalised an 18-question protocol. Interviews explored four themes: (1) participants' experience in AI and software development; (2) working definitions of AI software/system traceability; (3) current traceability practices and tools; and (4) expectations for the new AIBOM specification.

We conducted semi-structured video interviews (30–60 minutes) to allow for improvisation and exploration [95]. To avoid bias, the

Table 4: Study Participant Demographics and Experience

Part. ID	Software/AI Dev./SBOM Experience (yrs)	Domain	AIBOM Role
P1[a]	11 / 9 / 0	Entertainment Software	Consumer
P2[b]	25 / 0 / 4	Healthcare Software	Consumer
P3[c]	4 / 2 / 1	ICT	Consumer
P4[c]	14 / 2 / 2	ICT	Both
P5[d]	15 / 6 / 0.5	Academia	Consumer
P6[e]	20 / 0 / 16	Open Source Foundation	Both
P7[f]	4 / 3 / 0	Logistics Software	Creator
P8[g]	12 / 12 / 0	IT Consultancy	Both
P9[d]	9 / 5 / 0	Software Research	Both
P10[h]	15 / 0 / 5	IT and Networking	Both

[a] Lead Data Scientist; [b] Chief Operating Officer; [c] SE Researcher; [d] Research Scientist; [e] General Manager; [f] ML Engineer; [g] Consultant; [h] SBOM Advocate.

proposed AIBOM fields were not disclosed during the sessions. With participant consent, interviews were recorded, transcribed via software, and manually verified for accuracy.

Transcript coding and analysis. Two authors initially co-coded two transcripts to induce a baseline coding scheme. The remaining transcripts were then coded independently. Following this, the authors met to filter irrelevant codes and reach a consensus on categories. Finally, we computed the overlap between these participant-derived concepts and our proposed AIBOM specification fields.

Results. Excluding *Configuration* and *Deployment* fields—reserved for future versions due to capture complexity—all participant-proposed fields map to the SPDX AIBOM 3.0 specification. Table 5 details this mapping between participant suggestions and the specification.

The primary theme emerging from the interviews was the widespread inadequacy of current traceability practices for AI systems, which participants described as ad hoc and insufficient for ensuring proper governance. This sentiment was powerfully articulated by an executive director (P2), who stated, *"I think it's a general mystery across the industry... I don't think that [traceability] adequately exists within the AI community around algorithms and training sets to be able to demonstrate that traceability from an auditing standpoint."* This lack of a standardised, auditable record is exacerbated by real-world development pressures. A machine learning engineer (P7) from a startup described this vividly: *"It's something that we lack and it's because... the general nature of a startup, we usually deprioritise documentation and it usually also ends with a lot of pain for us... because we don't have the documentation, I usually spend so much time trying to explain anyway."*

Beyond identifying these gaps, participants also provided clear insights into effective traceability requirements. When asked what information they need when consuming a pre-trained model, a lead data scientist (P1) offered a detailed wishlist: *"What do I look for? I look for license... support... When, which training data was used, what demographic was used, and what biases do they have? ...And these are stuff that you want captured in the document, because that's what I'm looking for [as a model consumer]... what's the [reported] accuracy? How did you test it?"* P1's requirements—spanning licensing, provenance, and evaluation—directly align with the AIBOM framework's core categories, confirming that its design addresses practitioner needs and will guide the specification's evolution.

Table 5: AIBOM fields validated by interviews.

Category	Proposed Field	SPDX 3.0 AI/Dataset Fields
Data	Lineage	originatedBy, dataCollectionProcess, datasetUpdateMechanism, downloadLocation
	Provenance	downloadLocation, originatedBy
	Metadata	spdxId, name, packageVersion, buildTime, releaseTime, primaryPurpose, datasetSize
	Assumptions	datasetNoise
	Preprocessing	dataPreprocessing, anonymizationMethodUsed
	Licenses	hasConcludedLicense, hasDeclaredLicense
	Intended use	intendedUse
	Personal info	hasSensitivePersonalInformation
Models	Behaviour	metric, metricDecisionThreshold, safetyRiskAssessment, modelExplainability
	Parameters	hyperparameter
	Algorithms	informationAboutTraining, typeOfModel
	Intended use	primaryPurpose, informationAboutApplication
	Assumptions	informationAboutApplication, limitation
	Biases	knownBias, limitation
	Licenses	hasConcludedLicense, hasDeclaredLicense
	Features	modelExplainability, domain, typeOfModel
Code	License	hasConcludedLicense, hasDeclaredLicense
Environment	Hardware	sensor, energyConsumption, inferenceEnergyConsumption, trainingEnergyConsumption
	Configuration	proposed in SPDX 3.1
	Deployment	proposed in SPDX 3.1
Process	Data Collection	dataCollectionProcess, sensor
	Training	informationAboutTraining, trainingEnergyConsumption, finetuningEnergyConsumption
Governance	Dependencies	contains, downloadLocation
	Attribution	originatedBy, suppliedBy, name, spdxId
	Ethical considerations	safetyRiskAssessment, knownBias, useSensitivePersonalInformation, hasSensitivePersonalInformation, confidentialityLevel, anonymizationMethodUsed, standardCompliance, intendedUse
	Design decisions & limitations	limitation, modelDataPreprocessing, dataPreprocessing, metricDecisionThreshold, typeOfModel, hyperparameter

5.4 Field Validation in Industrial Setting

We evaluated AIBOM's practical suitability through an external study at a large multinational software organisation. From June to September 2024, a multi-disciplinary team of security architects, data scientists, DevOps experts, and AI developers mapped the overlap between our specification fields and those required during the organisation's AI system development.

Methodology. The organisation already had a checklist for developers to evaluate how their AI software development affects the internal ethics policy. Moreover, the legal department had a checklist to assess the legal impact of each AI-related software release. The researchers investigated how many of the fields in these checklists could be populated if the organisations already had AIBOMs constructed as part of the AI software development process. Furthermore, they investigated whether our AIBOM specification can be used to automatically generate model cards. Then, they investigated to what extent AIBOM generation for third-party models can be automated using web scraping or external APIs, such as Hugging Face [39].

Results. Validation revealed that AIBOM fields cover all internal checklist requirements for developers, proving its practical comprehensiveness. They observed that 60% of existing model card fields

are extractable from the fields in the AIBOM specification. Furthermore, external APIs and web scraping successfully populated 40% of fields for third-party models, demonstrating that AIBOM generation can be partially automated to reduce manual effort.

6 Standards Development as Action Research

A primary contribution of this work is demonstrating how the development of a complex, multi-stakeholder standard can be systematically understood through the lens of **Action Research (AR)**. By retrospectively framing our effort to extend SPDX and create an AIBOM specification as an AR process, we illustrate how community-driven standards development can maintain a disciplined, research-grounded approach within rapidly evolving domains.

While not initially conceived as AR, this project's conduct and outcomes align closely with its core principles. AR focuses on changing practice through real-world intervention [5, 20, 27]; similarly, our objective was to develop a functional, community-driven standard for AI transparency and compliance. We therefore adopt the canonical AR cycle [88] to retrospectively structure and reflect on this process.

Our AR framing is particularly relevant in software engineering (SE), where controlled experiments, case studies, and surveys dominate, while AR remains comparatively rare [65]. Sjøberg et al. [79] attribute this rarity partly to a limited understanding of AR's role in SE. By analysing the AIBOM specification development via an AR lens, we contribute a concrete example of its application in the wild to the software engineering body of knowledge.

Table 6 maps the main steps of our methodology to the canonical AR stages. This perspective highlights the key enablers of our success: continuous stakeholder engagement, iterative adaptation, and feedback-driven evolution. Further, it demonstrates how standards development can operate simultaneously as a scientific method and a practical intervention.

Table 6: Mapping AIBOM Development to the AR Cycle

AR Phase	AIBOM Development Stage
Diagnosing *Identifying a practical problem*	**WG formation and scoping:** Addressed the lack of a standardised way to describe AI components for traceability, compliance, and transparency.
Action Planning *Designing an intervention*	**Field definition and framework design:** Extended SPDX to include AI artefacts by analysing model cards, datasheets, and existing practices.
Action Taking *Implementing the intervention*	**Drafting and presenting profiles:** Created draft AI and Dataset profiles and presented them at major forums to gather feedback.
Observation *Evaluating impact*	**Feedback collection:** Collected input via release candidates, pull requests, and discussions to assess utility and design.
Reflection *Learning and refining*	**Profile refinement:** Iteratively updated profiles based on feedback, aligning them with real-world adoption needs.
Iterative Cycles *Repeating and improving*	**Release iterations:** Integrated feedback into successive release candidates, refining the specification before final release.
Knowledge Dissemination *Sharing outcomes*	**Publishing and outreach:** Released SPDX 3.0 with AI and Dataset profiles, supported by a whitepaper, blog posts, and documentation.

7 Lessons Learned and the Road Ahead

This section distills lessons from developing the AIBOM standard in a multi-stakeholder setting and discusses how these insights inform future extensions for foundation model-powered software.

7.1 Lessons Learned

Present work-in-progress early and solicit feedback often. Public presentations of early drafts, in venues such as Open Source Summit (OSS) sessions (see Section 3.3) and other WGs provided critical feedback between mid-2022 and 2024. This early engagement not only strengthened the evolving standard but also built practitioner trust and buy-in. Feedback from outside our WG often prompted significant design changes. For instance, input from the OpenChain WG led to the introduction of a `standardsCompliance` field to capture standards an AI model or dataset already adheres to, which is a key requirement for audits. A question during an OSS North America 2023 panel about the queryability of fields such as `energyConsumption` and `metrics` prompted us to replace free-form text with structured types. Similarly, discussions at OSS Japan 2022 revealed that tracking AI and data lineage separately would complicate adoption, leading us to merge them into a unified `dataCollectionProcess` field. Based on our experience, in line with previous studies [10], we suggest that continuous feedback-driven development is not only advantageous but essential for open standards in fast-moving domains like AI.

Prioritise adoption over comprehensiveness. Early drafts that attempted to capture every conceivable detail of an AI system consistently faced pushback from practitioners. Most organisations simply do not maintain information at that level of granularity, and *a standard that demands it becomes impractical* [45, 47]. A similar pattern is evident with model cards, which prescribe numerous fields; over 60% of models and 70% of datasets on Hugging Face lack a complete model card [63, 87]. Several prior studies document similar resistance to heavyweight standards [66, 77]. Another responsible AI field study reports an organisation expressing "a real fear that they may adopt principles that they cannot uphold in practice" [74].

We therefore optimised our AIBOM specification for adoption by defining a small set of readily recordable *required fields* and enforcing strict entry criteria. In some cases, we intentionally excluded ambitious goals to improve practicality. For example, rather than attempting to enumerate every potential form of bias, we introduced a single `knownBias` field to capture only documented biases. Additional fields were deferred to future releases, enabling incremental evolution as practices mature.

We also found that two factors: metadata scale and availability, strongly influence adoption. For instance, a comprehensive AIBOM for a self-driving car could span several gigabytes, making generation, storage, and maintenance impractical. At the same time, much of the desired metadata simply does not exist in accessible form.

Even foundational datasets such as ImageNet and CIFAR-10 do not fully disclose their data sources [71]. Supporting fields that can be automatically collected or derived are therefore essential. Recognising this, we are now collaborating with the SPDX Implementers WG [84] and CISA SBOM-o-RAMA [14] to improve automation support in future iterations. These experiences highlight a broader reality: premature attempts at exhaustive coverage can undermine adoption. While similar lessons have been echoed several decades ago during POSIX standardisation [45, 47], they ring true even today. Prioritising a minimal, usable core builds early momentum and a foundation for richer requirements as the ecosystem matures.

Plan for churn and conflict from day one. In long-running, community-driven standardisation, contributor turnover and conflicting priorities are inevitable; planning for them early is essential. Several participants disengaged as organisational priorities shifted, while others joined late to contribute to specific phases. This irregular participation often left critical tasks unfinished, hindering progress. Consequently, we established a rotating *core group* with decision-making authority to ensure continuity alongside broader community input.

Disagreements were a constant reality in this multi-stakeholder setting. Participants frequently proposed modifications with vague or narrow motivations, risking developmental stalls. To mitigate this, we applied the evidence-based acceptance criteria defined in Section 7.1, significantly reducing friction and maintaining focus on shared priorities. As the AI ecosystem diversifies, sustaining continuity amid organisational churn and aligning competing interests will grow more challenging. Embedding these governance mechanisms early ensures open standards remain stable yet adaptable.

7.2 Road Ahead

Evolving AIBOM for FM-powered software. The 2023 surge in foundation models (FMs) and FM-powered software (FMware) necessitated an evolution of the AIBOM specification. Unlike static ML artefacts, FMware is continuous, adaptive, and open-ended [37], introducing new assets like prompts, fine-tuned variants, and agents—the latter acting as first-class entities in multi-model pipelines [37, 70]. Operational layers such as *grounding* (decision source verification) and *guarding* (safety/policy enforcement) generate further dependencies requiring tracking [70]. Driven by regulatory and risk-management needs, the FMwareBOM iteration expands AIBOM to capture agent behaviours, orchestration context, and lifecycle transformations. While the current specification handles traditional AI (see Section 5), adapting it for FMware is now urgent.

To meet FMware's requirements, we are updating the AI and Dataset profiles for the SPDX 3.1 release candidate [80]. This update introduces fields for agent identity and capabilities [82], enhanced provenance linking datasets to prompts and retrieval contexts [83], and mechanisms to track adaptation events and fine-tuning lineage [81]. These enhancements ensure FMwareBOM captures the full AI lifecycle while remaining automatable, auditable, and compliant with evolving standards.

Making the future AIBOM tool-native and automatable. As Section 7.1 notes, effective tooling is vital for adoption, particularly as FMware increases tracking scale and complexity. To address this, we are designing the evolved version of AIBOM to be tooling-native from the outset. Our goal is to ensure that key fields can be extracted directly from development pipelines or inferred through lightweight analysis, while still meeting regulatory and risk-management requirements. Our hope is that this approach will enable greater automation throughout the FMware lifecycle—from generating BOMs and validating schema conformance during CI stages to assessing compliance and policy obligations automatically. These design choices aim to make AIBOM not only richer in representation but also a practical, automatable, and auditable part of the FMware development process.

Enabling interoperability across the AI supply-chain ecosystem. Our roadmap positions the next-generation SPDX AIBOM as a connective layer across a fragmented ecosystem. Rather than evolving in isolation, SPDX aims to provide a unified integration point for complementary standards—including SLSA (build provenance) [68], OpenChain (organisational compliance) [67], Croissant (data provenance) [22], and Coalition for Secure AI [11]—ensuring interoperability across the AI supply chain. To address this, we are actively working to align and interoperate with these initiatives. For example, we are exploring how FMwareBOM can serve as the canonical machine-readable artefact for AI compliance within the OpenChain AI WG and collaborating with OpenSSF and CISA to integrate provenance, verification, and security metadata. To this end, our Open Source Summit North America 2025 panel on harmonising SLSA and SPDX [56] gathered community input on reducing metadata duplication and ensuring standard reinforcement. Concurrently, the formal SPDX 3.0 specification is being prepared to update the ISO/IEC 5962 standard. This standardisation process, involving Technical Committee reviews, is expected to accelerate global adoption and ecosystem maturity.

8 Related Work

To contextualise our contribution, we present prior work on emerging approaches for trustworthy AI, the challenges of multi-stakeholder collaboration in software engineering, and the application of AR as an intervention methodology.

Emerging approaches for trustworthy AI. Various initiatives address AI trustworthiness through auditability and provenance, yet gaps remain. DataBOM [52] focuses on dataset provenance but lacks model-level detail, while CycloneDX ML-BOM [15] extends SBOMs to ML without capturing full lifecycle complexities. Reporting frameworks like factsheets [21], model cards [25], and datasheets [33, 58] enhance documentation but lack standardisation and system-level coverage. AI Cards [34] offer rich risk data but lack tooling. Broader governance frameworks (NIST AI RMF [61], Microsoft RAI [57], and Croissant [22]) promote accountability but remain decoupled from software supply chain realities. Recently, TAIBOM [91] aims to provide formal component and dependency descriptions. Despite their contributions, these solutions remain fragmented, treating datasets or models in isolation or offering high-level frameworks detached from implementation. This piecemeal

landscape lacks a machine-readable artifact providing an end-to-end view of AI system composition. AIBOM addresses this gap by extending the dependency-aware structure of SBOMs to unify data, models, and lifecycle processes into a cohesive, auditable standard.

SBOM research in SE. SBOMs are widely used for vulnerability management, transparency, risk assessment, and supply chain integrity [55, 69]. Empirical work shows that they reduce remediation times by enabling rapid dependency analysis [6] and improve auditability through provenance and licensing metadata [55]. However, their impact depends heavily on tooling completeness and metadata quality [86, 96]. Researchers and practitioners alike have proposed extensions, such as blockchain-enabled SBOMs for tamper resistance [98] and DataBOMs for capturing provenance in data pipelines [4, 52]. However, our study is the first document the process of creating AIBOM specification in a global, multi-stake holder setting.

Standardisation experience reports. SE standardisation reports often provide historical narratives rather than reproducible extension methodologies. POSIX retrospectives trace evolution from user groups to ISO ratification but lack detail on how extensions were conceived [45, 47]. Similarly, the Austin Group's SD/6 details governance but lacks practitioner perspectives [50], while SystemVerilog reports document IEEE transitions without codifying development methods [90]. We extend this literature by detailing the end-to-end development of an AI extension—capturing decision-making and validation—and mapping standards development as a replicable Action Research (AR) cycle.

Multi-stakeholder studies in SE. Software engineering's socio-technical nature requires balancing the conflicting goals of developers, managers, customers, and regulators. Research shows that analysing these multi-stakeholder dependencies improves requirements elicitation [16] and risk management [1], particularly within large-scale agile environments [38]. Challenges like interest balancing and cross-organisational communication [72] are magnified in decentralised standards development. Lacking a single organisational authority, these volunteer-driven efforts face unique consensus-building and governance hurdles. We therefore use open standard development as a case study to analyse high-complexity multi-stakeholder collaboration.

Action Research in SE. AR has emerged as a key methodology for industry–academia collaboration, enabling iterative refinement of tools in real-world settings [2, 65]. While recent guidelines [89] and hybrid models like Design Science Action Research [17] have matured the field, AR remains largely limited to internal process improvement within single organisations [18]. We extend AR to open, community-governed initiatives, using it as a framework to coordinate global stakeholders around a shared standardisation goal.

9 Conclusion

Our experience establishes Action Research as a robust framework for standardisation in dynamic, multi-stakeholder domains. By structuring the development of AIBOM specification around iterative diagnose-design-evaluate-reflect cycles, we showed how open collaboration, continuous feedback, and evidence-based iteration can transform fragmented community efforts into a coherent, widely applicable standard. The lessons distilled from this process provide a blueprint for future initiatives like our evolution of AIBOM specification in SPDX 3.1. Looking ahead, we envision AR-driven approaches playing a central role in shaping agile standards that evolve alongside the technologies they govern.

Acknowledgements. We thank Helen Oakley and Rhea Michael Anthony for the industrial validation in Section 5.4; Prof. Daniel M. German and Prof. Zhen Ming (Jack) Jiang for their mentoring; and the members of the SPDX AI and Dataset Working Group for their contributions and feedback. Arthit Suriyawongkul's contribution to the WG is made possible through the financial support of Taighde Éireann – Research Ireland under Grant number 18/CRT/6224 (Research Ireland Centre for Research Training in Digitally-Enhanced Reality (d-real)).

Disclaimer. Generative AI tools were used for copy-editing and table formatting. All experiments, analysis, and writing were performed by the authors, who thoroughly reviewed the content. This complies with IEEE and ACM policies on AI use in publications.

References

[1] Abdullah Alqahtani, Shadi Banitaan, and Sawsan Banitaan. 2017. A systematic literature review of software risk management for global software development. In *2017 IEEE/ACS 14th International Conference on Computer Systems and Applications (AICCSA)*. IEEE, 1310–1317. https://doi.org/10.1109/AICCSA.2017.71

[2] David Avison, Richard Baskerville, and Michael Myers. 1999. Action research. *Commun. ACM* 42, 1 (1999), 94–97.

[3] Iain Barclay, Alun Preece, Ian Taylor, Swapna Krishnakumar Radha, and Jarek Nabrzyski. 2023. Providing assurance and scrutability on shared data and machine learning models with verifiable credentials. *Concurrency and Computation: Practice and Experience* 35, 18 (2023), e6997.

[4] Iain Barclay, Alun Preece, Ian Taylor, and Dinesh Verma. 2019. Towards traceability in data ecosystems using a bill of materials model. *arXiv preprint arXiv:1904.04253* (2019).

[5] Richard L. Baskerville. 1999. Investigating Information Systems with Action Research. *Communications of the Association for Information Systems* 2 (1999). https://doi.org/10.17705/1cais.00219

[6] Carlo Benedetti, Francesco Cofano, et al. 2024. The Impact of SBOM Generators on Vulnerability Assessment in Python: A Comparison and a Novel Approach. In *International Conference on Software Engineering and Security*. Springer. https://link.springer.com/chapter/10.1007/978-3-031-95764-2_19

[7] Karen Bennet, Gopi Krishnan Rajbahadur, Arthit Suriyawongkul, and Kate Stewart. 2024. *Implementing AI Bill of Materials (AI BOM) with SPDX 3.0.* Technical Report. The Linux Foundation. https://doi.org/10.70828/RNED4427

[8] Knut Blind and Margarete Böhm. 2019. *The Relationship Between Open Source Software and Standard Setting.* Technical Report JRC116106. Joint Research Centre (JRC), European Commission. https://publications.jrc.ec.europa.eu/repository/handle/JRC116106

[9] Center for Devices and Radiological Health. 2025. Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions.

[10] European Commission. Joint Research Centre. 2019. *The relationship between open source software and standard setting.* Publications Office, LU. https://doi.org/10.2760/163594

[11] Coalition for Secure AI. 2024. SAIF Risk Assessment: A new tool to help secure AI systems across industry. https://www.coalitionforsecureai.org/google-saif-risk-assessment/

[12] Shane Coughlan. 2022. OpenChain Work Groups – New and Improved Structure. https://openchainproject.org/news/2022/10/12/wg-structure

[13] Cybersecurity and Infrastructure Security Agency. 2025. 2025 Minimum Elements for a Software Bill of Materials (SBOM) Public Comment Draft. https://www.cisa.gov/resources-tools/resources/2025-minimum-elements-software-bill-materials-sbom.

[14] Cybersecurity and Infrastructure Security Agency. 2025. CISA SBOM-a-rama. https://www.cisa.gov/resources-tools/resources/cisa-sbom-rama

[15] CycloneDX. 2024. Machine Learning Bill of Materials (ML-BOM). https://cyclonedx.org/capabilities/mlbom/.

[16] Alan M. Davis. 1993. The software requirements specification: A roadmap. *Journal of Systems and Software* 21, 2 (1993), 179–188. https://doi.org/10.1016/0164-1212(93)90040-T

[17] Valeria de Castro, María Luz Martín-Peña, Esperanza Marcos Martínez, and Maricela Salgado. 2025. Combining Action Research With Design Science as

a Qualitative Research Methodology. An Application to Service (Operations) Management Research. *International Journal of Qualitative Methods* 24 (2025), 15 pages. https://doi.org/10.1177/16094069241312018

[18] Yvonne Dittrich, Johan Bolmsten, and Cathrine Seidelin. 2024. *Action Research with Industrial Software Engineering: An Educational Perspective.* Springer, 413–461. https://doi.org/10.1007/978-3-031-71769-7_15

[19] Ecma International. 2024. CycloneDX Bill of Materials Specification. https://ecma-international.org/publications-and-standards/standards/ecma-424/

[20] Max Elden and Rupert F. Chisholm. 1993. Emerging Varieties of Action Research: Introduction to the Special Issue. *Human Relations* 46, 2 (Feb. 1993), 121–142. https://doi.org/10.1177/001872679304600201

[21] Arnold et al. 2019. FactSheets: Increasing trust in AI services through supplier's declarations of conformity. *IBM Journal of Research and Development* 63, 4/5 (July 2019), 6:1–6:13. https://doi.org/10.1147/jrd.2019.2942288

[22] Akhtar et al. 2024. Croissant: A Metadata Format for ML-Ready Datasets. In *Proc. 8th Workshop on Data Management for End-to-End ML (DEEM).*

[23] European Parliament and Council of the European Union. 2017. Regulation (EU) 2017/745 of the European Parliament and of the Council of 5 April 2017 on Medical Devices, Amending Directive 2001/83/EC, Regulation (EC) No 178/2002 and Regulation (EC) No 1223/2009 and Repealing Council Directives 90/385/EEC and 93/42/EEC. https://data.europa.eu/eli/reg/2017/745/oj

[24] European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828. https://doi.org/10.5040/9781782258674

[25] Hugging Face. 2022. Model Cards. https://huggingface.co/docs/hub/en/model-cards.

[26] Norman E. Fenton and Martin Neil. 1998. A Strategy for Improving Safety Related Software Engineering Standards. *IEEE Trans. Softw. Eng.* 24, 11 (1998), 1002–1013. https://doi.org/10.1109/32.730547

[27] Maria Angela Ferrario, Will Simm, Peter Newman, Stephen Forshaw, and Jon Whittle. 2014. Software engineering for "social good": integrating action research, participatory design, and agile development. In *Companion Proc. of the Int'l Conf. on Software Engineering (ICSE '14).* ACM, 520–523. https://doi.org/10.1145/2591062.2591121

[28] SBOM for AI Tiger Team. 2025. AIBOM Squad. https://github.com/aibom-squad.

[29] Stanford Center for Research on Foundation Models. 2023. Do Foundation Model Providers Comply with the Draft EU AI Act?: Requirements. https://github.com/stanford-crfm/EUAIActJune15/blob/main/requirements.md.

[30] The Linux Foundation. 2024. SPDX 3.0 Revolutionizes Software Management in Systems with Enhanced Functionality and Streamlined Use Cases — linuxfoundation.org. https://www.linuxfoundation.org/press/spdx-3-revolutionizes-software-management-in-systems-with-enhanced-functionality-and-streamlined-use-cases.

[31] The Linux Foundation, SPDX contributors, OMG, and OMG contributors. 2025. SPDX Specification Releases. https://github.com/spdx/spdx-spec/releases

[32] The Linux Foundation, SPDX contributors, OMG, and OMG contributors. 2025. SPDX Specifications. https://spdx.dev/specifications/ Accessed: 2025-09-04.

[33] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé Iii, and Kate Crawford. 2021. Datasheets for datasets. *Commun. ACM* 64, 12 (2021), 86–92.

[34] Delaram Golpayegani, Isabelle Hupont, Cecilia Panigutti, Harshvardhan J. Pandit, Sven Schade, Declan O'Sullivan, and Dave Lewis. 2024. AI Cards: Towards an Applied Framework for Machine-Readable AI and Risk Documentation Inspired by the EU AI Act. In *Privacy Technologies and Policy*, Meiko Jensen, Cédric Lauradoux, and Kai Rannenberg (Eds.). Springer Nature Switzerland, Cham, 48–72. https://doi.org/10.1007/978-3-031-68024-3_3

[35] SPDX AI Working Group. 2024. Meeting Minutes. https://github.com/spdx/meetings/tree/main/ai Accessed: 2024-08-09.

[36] SPDX AI Working Group. 2025. SPDX AI Working Group. https://spdxai.github.io/

[37] Ahmed E Hassan, Dayi Lin, Gopi Krishnan Rajbahadur, Keheliya Gallaba, Filipe Roseiro Cogo, Boyuan Chen, Haoxiang Zhang, Kishanthan Thangarajah, Gustavo Oliva, Jiahuei Lin, et al. 2024. Rethinking software engineering in the era of foundation models: A curated catalogue of challenges in the development of trustworthy fmware. In *Companion Proceedings of the 32nd Int'l Conf. on the Foundations of Software Engineering.* 294–305.

[38] Rashina Hoda, James Noble, and Stuart Marshall. 2017. The impact of inadequate customer collaboration on self-organizing Agile teams. *Information and Software Technology* 88 (2017), 20–30. https://doi.org/10.1016/j.infsof.2017.03.004

[39] Hugging Face. 2025. Hugging Face Hub API Endpoints. https://huggingface.co/docs/hub/en/api Accessed: 2025-01-12.

[40] IEEE 802.11 Working Group. 2023. IEEE 802.11 Standards Development Process. https://www.ieee802.org/11/.

[41] IEEE P7014.1 Working Group (EEEPG). 2024. *IEEE Draft Recommended Practice for Ethical Considerations of Emulated Empathy in Partner-based General-Purpose Artificial Intelligence Systems.* Technical Report. IEEE Standards Association. https://standards.ieee.org/ieee/7014.1/11609/

[42] IEEE Standards Association. 2023. IEEE-SA Standards Development Lifecycle. https://standards.ieee.org/develop/.

[43] International Organization for Standardization. 2023. ISO/IEC Directives, Part 1: Procedures for the technical work. https://www.iso.org/sites/directives/current/consolidated/index.html

[44] International Organization for Standardization and International Electrotechnical Commission. 2021. ISO/IEC 5962:2021 Information Technology — SPDX Specification V2.2.1.

[45] Jim Isaak. 2005. POSIX – Inside: A Case Study. In *Proceedings of the IEEE International Symposium on Technology and Society (ISTAS).* 1–6. https://doi.org/10.1109/ISTAS.2005.1451989

[46] ISO/IEC. 2015. ISO/IEC 19770-2:2015 Information technology — Software asset management — Software identification tag (SWID tag). https://www.iso.org/standard/65666.html

[47] Hal Jespersen. 1995. POSIX Retrospective. *StandardView* 3, 1 (March 1995), 2–10. https://doi.org/10.1145/210308.210313

[48] Wenxin Jiang, Jerin Yasmin, Jason Jones, Nicholas Synovic, Jiashen Kuo, Nathaniel Bielanski, Yuan Tian, George K Thiruvathukal, and James C Davis. 2024. Peatmoss: A dataset and initial analysis of pre-trained models in open-source software. In *Proceedings of the 21st International Conference on Mining Software Repositories.* 431–443.

[49] Jianxin Jiao, Mitchell M Tseng, Qinhai Ma, and Yi Zou. 2000. Generic bill-of-materials-and-operations for high-variety production management. *Concurrent Engineering* 8, 4 (2000), 297–321.

[50] Andrew Josey. 2012. Committee Maintenance Procedures for the Approved Standard. https://www.opengroup.org/austin/docreg.html

[51] Omer F Keskin, Kevin Matthe Caramancion, Irem Tatar, Owais Raza, and Unal Tatar. 2021. Cyber third-party risk management: A comparison of non-intrusive risk scoring reports. *Electronics* 10, 10 (2021), 1168.

[52] Yue Liu, Dawen Zhang, Boming Xia, Julia Anticev, Tunde Adebayo, Zhenchang Xing, and Moses Machao. 2024. Blockchain-Enabled Accountability in Data Supply Chain: A Data Bill of Materials Approach. In *Int'l Conf. on Blockchain (Blockchain).* IEEE, 557–562.

[53] Shayne Longpre, Robert Mahari, Anthony Chen, Naana Obeng-Marnu, Damien Sileo, William Brannon, Niklas Muennighoff, Nathan Khazam, Jad Kabbara, Kartik Perisetla, et al. 2023. The Data Provenance Initiative: A large scale audit of dataset licensing & attribution in AI. (2023).

[54] Qinghua Lu, Liming Zhu, Xiwei Xu, Jon Whittle, and Zhenchang Xing. 2022. Towards a roadmap on software engineering for responsible AI. In *Proceedings of the 1st Int'l Conf. on AI Engineering: software Engineering for AI.* 101–112.

[55] Daniele Martini. 2023. *Improving Transparency, Trust, and Automation in the Software Supply Chain.* Ph. D. Dissertation. University of Camerino. https://tesidottorato.depositolegale.it/handle/20.500.14242/193708

[56] Mihai Maruseac, Kate Stewart, Hasan Yasar, and David A. Wheeler. 2025. Panel Discussion: Strengthening Software Supply Chains: Harmonizing SLSA Provenance and SPDX SBOM for Better Adoption. In *Open Source Summit North America 2025.* https://ossna2025.linuxfoundation.org

[57] Microsoft. 2022. Microsoft Responsible AI Standard, v2 General Requirements.

[58] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT* '19).* ACM. https://doi.org/10.1145/3287560.3287596

[59] National Telecommunications and Information Administration (NTIA). 2021. *The Minimum Elements for a Software Bill of Materials (SBOM).* Technical Report. U.S. Department of Commerce. https://www.ntia.gov/report/2021/minimum-elements-software-bill-materials-sbom

[60] The United States Department of Commerce. 2021. The Minimum Elements for a Software Bill of Materials (SBOM). https://www.ntia.gov/report/2021/minimum-elements-software-bill-materials-sbom.

[61] National Institute of Standards and Technology. 2023. *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* Technical Report. https://www.nist.gov/itl/ai-risk-management-framework

[62] Open Source Security Foundation. 2025. OpenSSF AI/ML Security Working Group. https://openssf.org/groups/ai-ml-security/

[63] Ernesto Lang Oreamuno, Rohan Faiyaz Khan, Abdul Ali Bangash, Catherine Stinson, and Bram Adams. 2024. The state of documentation practices of third-party machine learning models and datasets. *IEEE Software* 41, 5 (2024), 52–59.

[64] OWASP. 2025. CycloneDX Specification. https://github.com/CycloneDX/specification/. Accessed: 2025-09-04.

[65] Kai Petersen, Cigdem Gencel, Hamed Asghari, Dejan Baca, and Stefanie Betz. 2014. Action research as a model for industry-academia collaboration in the software engineering context. In *Proc. of the Int'l Conf. on Software Engineering (ICSE) Companion.* ACM, 43–46.

[66] Shari Lawrence Pfleeger, Norman Fenton, and Stella Page. 2002. Evaluating software engineering standards. *Computer* 27, 9 (2002), 71–79.

[67] OpenChain Project. 2025. OpenChain. https://openchainproject.org/

[68] SLSA Project. 2025. Supply-chain Levels for Software Artifacts. https://slsa.dev/

[69] Erik Nordström Qvarfordt. 2024. Exploring the Dynamics of Software Bill of Materials (SBOMs) and Security Integration in Open Source Projects. https://www.diva-portal.org/smash/record.jsf?pid=diva2%3A1844757

[70] Gopi Krishnan Rajbahadur, Gustavo A Oliva, Dayi Lin, and Ahmed E Hassan. 2024. From Cool Demos to Production-Ready FMware: Core Challenges and a Technology Roadmap. *arXiv preprint arXiv:2410.20791* (2024).

[71] Gopi Krishnan Rajbahadur, Erika Tuck, Li Zi, Dayi Lin, Boyuan Chen, Zhen Ming, Daniel M German, et al. 2021. Can I use this publicly available dataset to build commercial AI software?–A Case Study on Publicly Available Image Datasets. *arXiv preprint arXiv:2111.02374* (2021).

[72] Paul Ralph. 2016. The role of power in requirements engineering. *Requirements Engineering* 21, 3 (2016), 297–313. https://doi.org/10.1007/s00766-014-0207-5

[73] Elyas Rashno. 2025. AIBOM_SPDX_Mapping_To_UseCases. https://doi.org/10.6084/m9.figshare.30234535.v1

[74] Lorenn P Ruster and Jenny L Davis. 2025. The Gaps That Never Were: Reconsidering Responsible AI's Principle-Practice Problem. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25)*. Association for Computing Machinery, New York, NY, USA, 350–360. https://doi.org/10.1145/3715275.3732024

[75] SBOM for AI Tiger Team. 2025. SBOM for AI (AIBOM) Tiger Team Working Group Document. https://docs.google.com/document/d/1IpXG7XBOJnPl_hwFf3JZkDaFb0k2CnI0/edit?rtpof=true#heading=h.gjdgxs

[76] Marius Schlegel and Kai-Uwe Sattler. 2023. Management of machine learning lifecycle artifacts: A survey. *ACM SIGMOD Record* 51, 4 (2023), 18–35.

[77] Michael E Schmidt. 2000. *Implementing the IEEE software engineering standards*. Sams.

[78] David Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-Francois Crespo, and Dan Dennison. 2015. Hidden technical debt in machine learning systems. *Advances in neural information processing systems* 28 (2015).

[79] Dag I. K. Sjøberg, Tore Dyba, and Magne Jorgensen. 2007. The Future of Empirical Methods in Software Engineering Research. In *Future of Software Engineering (FOSE '07)*. IEEE, 358–378. https://doi.org/10.1109/fose.2007.30

[80] SPDX contributors. 2025. SPDX 3 Model - Milestone 3.1. https://github.com/spdx/spdx-3-model/milestone/3

[81] SPDX contributors. 2025. SPDX 3.0 Proposal: Adaptation Events, Fine-tuning Lineage, and Downstream Modifications. https://github.com/spdx/spdx-3-model/pull/1100. Accessed: 2025-09-28.

[82] SPDX contributors. 2025. SPDX 3.0 Proposal: Agent Identity and Capabilities Fields. https://github.com/spdx/spdx-3-model/pull/1091 Accessed: 2025-09-28.

[83] SPDX contributors. 2025. SPDX 3.0 Proposal: Provenance Links for Dataset Generation Prompts and Retrieval Contexts. https://github.com/spdx/spdx-3-model/pull/892 Accessed: 2025-09-28.

[84] SPDX Project / Linux Foundation. 2025. SPDX Implementers Mailing List. https://lists.spdx.org/g/spdx-implementers

[85] Sarah Spiekermann. 2017. IEEE P7000—The First Global Standard Process for Addressing Ethical Concerns in System Design. *Proceedings* 1, 3 (2017), 159. https://doi.org/10.3390/IS4SI-2017-04084

[86] Trevor Stalnaker, Nathan Wintersgill, Oscar Chaparro, Massimiliano Di Penta, Daniel M German, and Denys Poshyvanyk. 2024. BOMs Away! Inside the Minds of Stakeholders: A Comprehensive Study of Bills of Materials for Software Systems. In *Int'l Conf. on Software Engineering (ICSE '24)*. ACM, 1–13. https://doi.org/10.1145/3597503.3623347

[87] Trevor Stalnaker, Nathan Wintersgill, Oscar Chaparro, Laura A Heymann, Massimiliano Di Penta, Daniel M German, and Denys Poshyvanyk. 2025. The ML supply chain in the era of software 2.0: Lessons learned from Hugging Face. *arXiv preprint arXiv:2502.04484* (2025).

[88] Miroslaw Staron. 2020. *Action Research in Software Engineering: Theory and Applications*. Springer International Publishing. https://doi.org/10.1007/978-3-030-32610-4

[89] Miroslaw Staron. 2025. Guidelines for Conducting Action Research Studies in Software Engineering. *e-Informatica Software Engineering Journal* 19, 1 (2025), 250105. https://doi.org/10.37190/e-Inf250105

[90] Stuart Sutherland. 2002. Verilog, the Next Generation: Accellera's SystemVerilog. In *Proceedings of the HDL Conference*.

[91] TAIBOM. 2025. TAIBOM. https://taibom.org/

[92] The Open Group. 2022. The Austin Group: POSIX Standardization. https://www.opengroup.org/austin

[93] David Thiel. 2023. Identifying and Eliminating CSAM in Generative ML Training Data and Models. https://purl.stanford.edu/kh752sm9123?ref=404media.co.

[94] Frank Tsui, Orlando Karam, and Barbara Bernal. 2022. *Essentials of software engineering*. Jones & Bartlett Learning.

[95] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in Software Engineering*. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-29044-2

[96] Boming Xia, Tingting Bi, Zhenchang Xing, Qinghua Lu, and Liming Zhu. 2023. An Empirical Study on Software Bill of Materials: Where We Stand and the Road Ahead. In *Int'l Conf. on Software Engineering (ICSE)*. IEEE. https://doi.org/10.1109/icse48619.2023.00219

[97] Boming Xia, Tingting Bi, Zhenchang Xing, Qinghua Lu, and Liming Zhu. 2023. An empirical study on software bill of materials: Where we stand and the road ahead. In *Int'l Conf. on Software Engineering (ICSE)*. IEEE, 2630–2642.

[98] Boming Xia, Dawen Zhang, Yue Liu, Qinghua Lu, Zhenchang Xing, and Liming Zhu. 2024. Trust in software supply chains: Blockchain-enabled SBOM and the AIBOM future. In *Int'l Workshop on Engineering and Cybersecurity of Critical Systems (EnCyCriS) and 2nd Int'l Workshop on Software Vulnerability*. 12–19.