

# Explainability as a Requirement for Hardware: Introducing Explainable Hardware (XHW)

Timo Speith\*, Julian Speith†, Steffen Becker‡†, Yixin Zou†, Asia Biega†, and Christof Paar†

\*University of Bayreuth, Bayreuth, Germany

†Max Planck Institute for Security and Privacy, Bochum, Germany

‡Ruhr-University Bochum, Bochum, Germany

Email: timo.speith@uni-bayreuth.de, {julian.speith, yixin.zou, asia.biega, christof.paar}@mpi-sp.org, steffen.becker@rub.de

**Abstract**—In today’s age of digital technology, ethical concerns regarding computing systems are increasing. While the focus of such concerns currently is on requirements for *software*, this article spotlights the *hardware* domain, specifically *microchips*. For example, the opaqueness of modern microchips raises security issues, as malicious actors can manipulate them, jeopardizing system integrity. As a consequence, governments invest substantially to facilitate a secure microchip supply chain. To combat the opaqueness of hardware, this article introduces the concept of *Explainable Hardware (XHW)*. Inspired by and building on previous work on Explainable AI (XAI) and explainable software systems, we develop a framework for achieving XHW comprising relevant stakeholders, requirements they might have concerning hardware, and possible explainability approaches to meet these requirements. Through an exploratory survey among 18 hardware experts, we showcase applications of the framework and discover potential research gaps. Our work lays the foundation for future work and structured debates on XHW.

**Index Terms**—hardware requirements, explainable hardware, explainability, explainable systems, non-functional requirements, explainable artificial intelligence, XAI, trustworthiness

## I. INTRODUCTION

In recent years, researchers have started examining the ethical implications of digital technologies. While most work focuses on quality aspects or non-functional requirements (NFRs) of the *software* that runs on these systems [2], [3], the focus of this article is the *hardware*—viz., the microchips—that make these systems run in the first place. Consequential initiatives about microchip manufacturing capabilities are currently taking shape around the world, e. g., the European Chips Act [4] and the US CHIPS and Science Act [5]. Through these multi-billion dollar investments, the governments behind them seek to gain more control over the supply of microchips and become less dependent on (untrusted) foreign manufacturers.

A prominent concern is that adversarial (foreign) actors can modify hardware undetected in such a way that the software running on top of it can be manipulated at will. Indeed, malicious hardware manipulations can have catastrophic consequences, potentially leading to a complete loss of security or incorrect algorithmic decisions. Such manipulations include the insertion of a kill switch to render military hardware inoperable under specifiable conditions [6], manipulation of Machine Learning (ML) accelerators [7]–[9], or the compromising of hardware security primitives [10].

The primary problem is that modern microchips are *opaque*. Microchips have become increasingly complex, culminating, e. g., in the Apple M1 Ultra with 114 *billion* transistors [11]. Furthermore, non-deterministic design tools automate vast parts of the hardware design process, alienating the designers from the final microchip schematics. Meanwhile, microchip supply chains are globally distributed and subject to increasing geopolitical tensions [12], leading to opaque manufacturing processes. Overall, the resulting opaqueness affects not only downstream stakeholders, such as the end users (i. e., consumers or operators) who interact with the systems but also the experts who design them. However, solutions on how to address this opaqueness have not yet been established.

Towards a possible solution, we adapt a prominent concept from discussions on the ethics of other digital technologies: *explainability*. In Requirements Engineering (RE), the concept of explainability has received a lot of attention recently, and it is rapidly establishing itself as a vital NFR [3], [13]–[15].

As explainability promises to ease concerns about the security and safety of software and Artificial Intelligence (AI) systems by making them more comprehensible to various stakeholders [3], [16], we argue that adopting the concept of explainability to hardware requirements—thus designing *Explainable Hardware (XHW)*—has the potential to achieve a similar goal: making hardware more comprehensible to various stakeholders and addressing concerns about security and more.

Based on these considerations, we develop a comprehensive XHW framework. Our specific contributions include:

- **Motivating and Defining XHW.** We argue that hardware is opaque and justify the need for more hardware transparency through, among other things, legislation for trustworthy hardware. As a solution, we propose to transfer the concept of explainability to the hardware domain. We argue that XHW is essential to achieve explainability at the system level, building on definitions and models for Explainable AI (XAI). (Section II and Section III)
- **A Framework for XHW.** We conceive a *framework* for XHW encompassing different *stakeholders*, their explainability needs (i. e., requirements and quality aspects related to hardware they are interested in, called *desiderata*), and *approaches* for enhancing the explainability of hardware, drawing on literature from hardware design, manufacturing, and analysis. (Section IV)

- **An Exploratory Study with Hardware Experts.** To demonstrate the applicability of our framework, we conduct an exploratory survey among 18 hardware experts. Our survey findings hint at distinct needs across stakeholders (e.g., only manufacturers may not care about explainability) and potential limitations of existing approaches (e.g., no approaches are thought to work well for end users). (Section V)
- **Potential Applications of the XHW Framework.** Finally, we demonstrate how our framework, informed by the results of our study, can help identify directions for future research in RE. (Section VI)

## II. BACKGROUND

In this work, we refer to a system as something that is usually built and operated by humans and consists of different hardware and software components to serve a specific purpose. In other words, a system can be anything from a smartphone, a car, to a PC. A system may also incorporate smaller subsystems, thereby introducing a hierarchy of systems. We define hardware as physical electronic components, focusing on microchips. Software is executed on hardware and includes operating systems, user applications, and algorithms like AI.

### A. Reasons for Hardware Opacity

Modern microchips are specified and designed using high-level Hardware Description Languages (HDLs). The resulting schematics are mapped to a technology library that describes all circuit elements available for realizing the design. This process is known as *synthesis*. The technology libraries used in this process are typically provided by large manufacturers, also called *foundries*. The implemented design, stripped of all high-level information like hierarchy, labels, and comments, is passed on to the manufacturer. They produce the actual chip using the chosen manufacturing technology in one of their production facilities, i. e., a semiconductor manufacturing plant also referred to as *fab*. In the last step, the fabricated chip can be integrated into various types of systems.

With this background in place, we synthesize the reasons why hardware is opaque by drawing on the three types of opacity in Burrell’s work about ML algorithms [17] (see also Mann et al. [18]). First, hardware is opaque to most people due to *technical illiteracy*, even to some experts. Their ever-growing complexity makes physical inspection and verification of microchips a challenging task, mastered by only a few highly specialized companies or government agencies worldwide. Second, modern synthesizers are based on efficient heuristic or even AI-based algorithms, which (for complex circuits) leads to different results for every synthesis run. This correlates to what Burrell calls *opacity characteristic of ML*. Third, developers and manufacturers use obfuscation to protect their Intellectual Property (IP) and thus their investments, which relates to opacity as *intentional corporate secrecy*. The situation is exacerbated by conflicting interests of nation states: they demand openness of foreign hardware while striving to protect domestic systems from external access [4], [5].

### B. Towards Hardware Transparency

The opaqueness of hardware, regardless of its origin, is increasingly recognized as a problem by various stakeholders, as we will outline below. In particular, recent research, as well as industry and government initiatives point to the need for a better understanding of hardware to facilitate its transparency.

While the *hardware industry* has historically been seclusive, an open hardware movement has formed in recent years [19], primarily driven by the RISC-V initiative [20], Google’s OpenTitan project [21], and the CHIPS Alliance [22]. However, these approaches are rarely designed explicitly to foster the understanding of hardware; rather, they tend to cater to vastly different objectives across various stakeholders.

The few studies available to date—although not directly about microchips, but rather about hardware devices in general—indicate that *end users* have limited understanding of hardware [23]–[25]. For instance, end users rarely notice hardware-based webcam LED indicators attached to laptops and smartphones [23]; in theory these should serve as privacy safeguards, but in reality they fail to improve users’ risk awareness. Likewise, end users exhibit limited technical knowledge about smart home devices [24], [25], despite existing research on the information leakage and security vulnerabilities among these devices [26]. Across these examples, a lack of understanding applies to both the software and hardware aspects, jeopardizing end users’ privacy and security.

The need for hardware transparency is also increasingly recognized by *governments and policy makers*, making the design and manufacturing of high-end semiconductors a political issue. For instance, the USA [5] and the EU [4] have pledged 52 billion USD and 43 billion EUR, respectively, to invest in domestic semiconductor manufacturing [27]. One primary concern with reliance on foreign-made semiconductors is implanted backdoors especially in secure hardware components used by the military or in space [4], [28], [29]. To this end, the European Chips Act demands “a solid understanding of a chip’s architecture” [4, p. 44].

Altogether, understanding hardware requires efforts from multiple stakeholders as well as dialogues among them. What is needed is a comprehensive view of possible directions to achieve hardware understanding for relevant stakeholders. This view should not only address technical aspects, but also highlight how and what information must be conveyed to which of the stakeholders involved in the hardware ecosystem.

## III. HARDWARE EXPLAINABILITY AS A SOLUTION FOR HARDWARE OPACITY

The importance of devising solutions to make hardware understandable to various stakeholders—across diverse contexts and irrespective of the reasons for its opacity—grows. The objective of explainability aligns perfectly with this goal, aiming to render certain aspects of (AI) systems understandable to different stakeholders [16], [30], [31], irrespective of context or source of opacity [18]. This alignment raises the question: How can explainability be transferred to the hardware domain to make hardware understandable to various stakeholders?

### A. From System Explainability to a Definition of Hardware Explainability

As hardware is part of a system, the transfer of explainability to the hardware domain can take place quite straightforwardly. Recent work has extended explainability from AI to software systems *holistically* [3], [14], [15], [32], as these can also be highly opaque [18]. While the opacity of AI may be particularly drastic, other kinds of opacity can already evoke legitimate desires for explainability. Chazette et al. [3] offer a general definition of explainability for software systems:

**Definition 1** (System Explainability). *A system  $S$  is explainable with respect to an aspect  $X$  of  $S$  relative to an addressee  $A$  in context  $C$  if and only if there is an entity  $E$  (the explainer) who, by giving a corpus of information  $I$  (the explanation of  $X$ ), enables  $A$  to understand  $X$  of  $S$  in  $C$ .* [3, p. 200]

In this definition, the aspect  $X$  to be explained may well be the hardware, since it is a constituent—and, thus, certainly also an aspect—of a software system  $S$ . With this finding, a gap in prior research becomes apparent: To make a system *truly* explainable, its hardware components must be made explainable as well. Merely making an AI—or the software that implements it—explainable is insufficient.

To the best of our knowledge, there is no prior research on the explainability of hardware, despite hardware being an essential constituent of systems. Furthermore, stakeholders’ interests cannot be fully met if we only focus on explainability at the algorithmic level. To see why this is the case, let us consider *trustworthiness* as one of the goals of explainability. Kästner et al. propose the following operationalization of trustworthiness for software systems [33]:

**Definition 2** (Trustworthiness). *A system  $S$  is trustworthy to a stakeholder  $H$  in a context  $C$  if and only if*

- (a)  $S$  works properly in  $C$ , and
- (b)  $H$  would be justified to believe that (a) if  $H$  came to believe that (a).

[33, p. 171]

For a system to work properly and thus be trustworthy, we must factor in all its constituents—software *and* hardware. After all, a system whose hardware does not work properly is hardly worth our trust. To find out whether the hardware works properly—and be justified in believing that it does so—we need information on it. That is where explainability comes into play [33]: by explaining a certain aspect of the system (e. g., its hardware), the understanding of a stakeholder is facilitated, letting them judge whether this aspect works properly [34].

### B. A Model of Explainability

Explainability not only promises to help facilitate trustworthiness; it also supports other desirable quality aspects of systems, such as debuggability, safety, and security [3]. These desirable properties are often referred to as *desiderata* [16], [35]. Coming back to the definition of system explainability

(see Definition 1), all desiderata are *downstream* goals. In other words, gaining an understanding of a system through explainability facilitates the satisfaction of a desideratum for a stakeholder. Langer et al. [16] and Hoffman et al. [36] propose similar models of this process for XAI (see Figure 1).




Fig. 1. A simplified version of the explainability models that Langer et al. [16] and Hofmann et al. [36] have proposed.

The models propose that explainability approaches (i. e., ways of achieving explainability) provide explanatory information to facilitate a stakeholder’s understanding. This understanding, in turn, affects the satisfaction of certain desiderata. Founded in these models, we develop our XHW framework.

## IV. A FRAMEWORK FOR HARDWARE EXPLAINABILITY

Through expert knowledge and literature review, we develop our framework for XHW in three components: relevant stakeholders (Section IV-A), desiderata (Section IV-B), and—in absence of established approaches to XHW—existing methods and techniques from hardware design, manufacturing, and analysis that could enhance explainability (Section IV-C). Figure 2 shows a high-level illustration of the framework.




Fig. 2. Our XHW Framework: Stakeholders want ① desiderata and can use ② explainability approaches that contribute to ③ satisfying these desiderata.

### A. Stakeholder Ecosystem

Various stakeholders interact with a hardware-based system throughout its lifetime, each of them having individual desiderata concerning the explainability of hardware. To address these individual desiderata, we introduce stakeholder categories to classify entities that interact with the hardware by developing, manufacturing, integrating, regulating, or using it. These categories are adapted from existing work by Langer et al. [16] and Tomlinson et al. [37]: (i) designers, (ii) manufacturers, (iii) system integrators, (iv) policymakers and watchdogs, and (v) end users.<sup>1</sup> Below, we outline each stakeholder category (see Table I) as well as the interactions between the stakeholders and the system (see Figure 3).

<sup>1</sup>An entity (a person, company, organization, or government agency) may belong to multiple stakeholder categories at the same time.

TABLE I  
THE FIVE STAKEHOLDER CATEGORIES CONSIDERED IN OUR FRAMEWORK AS WELL AS A SHORT DESCRIPTION FOR EACH OF THEM.

Stakeholder	Description
Designers	Describe the desired hardware functionality in a high-level language.
Manufacturers	Produce the hardware using highly specialized tools and equipment.
System Integrators	Integrate pieces of hardware into a larger system.
Policymakers & Watchdogs	Set the legal framework in which a system operates and attest adherence thereof.
End Users	Directly use the system or offer a service that is dependent on the system.




Fig. 3. Stakeholder interactions among each other, with the hardware, and with the system as a whole.

### B. Desiderata

The goals and purposes of explainability—also referred to as desiderata—depend on the perspectives and requirements of the stakeholders involved. Based on these factors, some aspects of explainability may be more relevant than others.

In Table II, we give a concise overview of the desiderata most relevant in the context of hardware explainability by adapting existing work on XAI and explainable software by Chazette et al. [3], Langer et al. [16], and Speith [38].

### C. Leveraging Approaches from Hardware Design, Analysis, and Manufacturing for Explainability

The final component of our framework is the approaches that can be leveraged to reach XHW. As we are the first to introduce the notion of XHW, there are no established approaches in the literature yet. Nevertheless, we can adopt approaches from the domains of hardware design, analysis, and manufacturing that contribute to the understanding of hardware and, thereby, improve its explainability according to Definition 1. We have compiled five such techniques in Table III.

TABLE II  
THE EIGHT DESIDERATA CONSIDERED IN OUR FRAMEWORK AS WELL AS A SHORT DESCRIPTION FOR EACH OF THEM.

Desideratum	Description
Safety	Avoid physical harm to people and the surrounding system.
Accountability	Identify who is responsible in case of failure.
Debuggability	Identify, trace, and correct bugs in order to prevent malfunction.
Legal Compliance	Adhere to the legal framework in which the system operates.
Security	Ensure confidentiality, integrity, and availability of data.
Verifiability	Test correct operations and rule out targeted manipulations.
Trustworthiness	Have correct functionality and be able to demonstrate it.
Trust	Calibrate dependence of a user or component on another component, fitting to its trustworthiness.

## V. AN EXPLORATORY SURVEY WITH HARDWARE EXPERTS

To showcase the applicability of our framework, we conducted an exploratory online survey with 18 experts from the hardware domain.

### A. Survey Method

The primary goal of the survey (see our additional material [1] for the full survey) was to demonstrate that our framework is suitable for identifying hardware explainability gaps (i. e., desiderata that are relevant to a stakeholder but may not be met with the approaches we distilled). Following the framework proposed in Figure 2, we asked our participants to ❶ determine the relevance of the desiderata for the various stakeholders, ❷ identify whether these desiderata can already be achieved using the presented approaches, and ❸ how applicable these approaches are to the different stakeholders.

1) *Study Procedures*: We conducted our survey as an online questionnaire using LimeSurvey. The median completion time was 19 minutes (min: 9; max: 60; average: 23). At the beginning, we asked participants for their consent to use the collected data for scientific purposes. All participants completed the survey voluntarily and received no compensation.

2) *Participants*: We intentionally targeted our survey to hardware experts, as our primary goal is to test the applicability of our framework, and they have the necessary technical knowledge to help us achieve this goal compared to laypersons. This recruitment choice reduces the likelihood of misinterpretation and ensures a sufficient baseline of understanding. Since hardware experts are difficult to reach, we started recruitment in our own professional networks, followed by snowball sampling. The majority of our 18 participants worked as hardware designers and manufacturers.

3) *Survey Questions*: After an introduction and explanation of the purpose of our survey, we asked participants for demographic information. For professional background, we asked participants to indicate their current occupation, professional branch/industry, and field of study.

TABLE III  
THE FIVE HARDWARE EXPLAINABILITY APPROACHES CONSIDERED IN OUR FRAMEWORK AS WELL AS A SHORT DESCRIPTION AND LITERATURE REFERENCES FOR EACH OF THEM.

Approach	Description	Explainability Benefit	Sources
Trusted Manufacturing	Retains microchip manufacturing capabilities at domestic facilities to prevent targeted manipulations by a malicious (foreign) fab.	Provides information and assurances about the manufacturing process of a microchip that may not be available about such processes in foreign fabs.	[4], [5], [29], [39], [40]
Standards & Certifications	Prescribe a lower bar to be met by microchip components and respective manufacturing processes and verify compliance thereof.	Provide information about the functionality and quality of a microchip, e.g., its use of standardized communication protocols.	[41]–[43]
Open Hardware	Strives for transparency throughout the supply chain, e.g., with open-source designs, tools, and manufacturing techniques.	Provides information about the architecture, implementation, and manufacturing process of a microchip for others (mostly experts) to inspect.	[20]–[22]
Testing & Verification	Evaluates the correctness of a microchip through design and manufacturing using simulation, testing, and formal verification.	Provides information about passed tests and verification procedures of a microchip and thereby also about its functionality and correctness.	[44]–[48]
Physical Analysis	Verifies correctness and checks for undesired information leakage of a microchip using invasive or even destructive physical analysis to rule out malicious modifications.	Provides information about the low-level architecture of a microchip and its regular behavior, as well as its behavior in settings outside the specified operating conditions.	[49]–[53]

The core part of our survey consists of three matrix questions in which we asked participants to indicate their perceptions of the relations between stakeholders, desiderata, and explainability approaches. We provided definitions of all stakeholders, desiderata, and explainability approaches in the survey to ensure that participants’ interpretation coincides with our understanding of the terms. All three matrix questions asked participants to rate the relationships on a 5-point Likert-like scale. The first question asked participants to rate the *importance* of each desideratum for any given stakeholder from “1 – not at all important” to “5 – extremely important”. The second question asked participants to rate *applicability* of every explainability approach to each of the stakeholders from “1 – not at all applicable” to “5 – extremely applicable”. The third question asked participants to rate the *usefulness* of each approach to satisfy the respective requirement from “1 – not at all useful” to “5 – extremely useful”.

4) *Data Analysis*: Since our primary goal of conducting the survey is to demonstrate the usefulness of our framework, we gathered exploratory data on the relationships between stakeholders, desiderata, and explainability approaches. To this end, we focused our analysis on descriptive statistics of participants’ responses to the three matrix questions. We calculated the mean values of each cell for the three questions and visualized them as heatmaps (see Figures 4, 5, and 6).

### B. Survey Results

The exploratory survey with hardware experts enables us to uncover potentially satisfied or unmet desiderata for different stakeholders in XHW. Specifically, Figure 4 shows the perceived importance of desiderata to stakeholders, Figure 5 shows the perceived applicability of explainability approaches to stakeholders, and Figure 6 shows the perceived usefulness of explainability approaches for satisfying desiderata.

Our framework can help guide future research by systematically identifying research gaps. We did this through three

steps: ❶ identify desiderata relevant for a stakeholder, ❷ for a given desideratum, determine whether existing approaches can be used to satisfy it, and ❸ consider the applicability of the respective approaches to the stakeholder. Below, we outline the most salient gaps we identified.




Fig. 4. Mean values of participants’ responses on the perceived importance of individual desiderata for each stakeholder on a scale from “1 – not at all important” to “5 – extremely important”.

1) *Low Explainability Needs Among Hardware Manufacturers*: According to our participants’ assessment, manufacturers are not overly concerned with XHW (at least not for the desiderata we considered). To contextualize the findings, manufacturers do not need to *understand* the hardware they manufacture. Instead, they apply quality control measures to verify whether the produced microchips exhibit correct functionality. Although not directly needed for hardware manufacturing, promoting explainability during manufacturing could help other stakeholders understand the hardware.

2) *Potential Misalignment of Regulatory Initiatives*: Open hardware and trusted manufacturing align with and are part




Fig. 5. Mean values of participants’ responses on the applicability of existing explainability approaches for each stakeholder category on a scale from “1 – not at all applicable” to “5 – extremely applicable”.

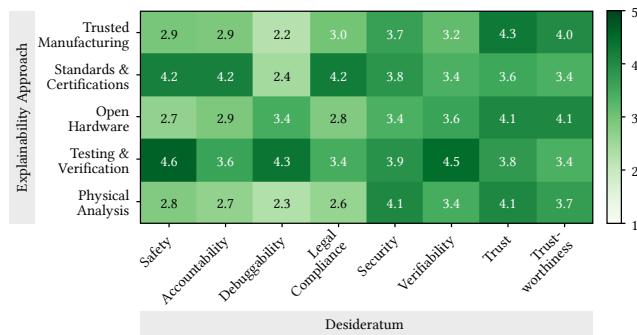


Fig. 6. Mean values of participants’ responses on the usefulness of existing explainability approaches to satisfy each desideratum on a scale from “1 – not at all useful” to “5 – extremely useful”.

of the strategy outlined in the European Chips Act [4]. Indeed, both approaches seem to be particularly well-suited for evoking trust and ensuring trustworthiness. However, our participants assessed that open hardware and trusted manufacturing have limited applicability for most stakeholders except manufacturers (see Figure 5). This indicates a potential mismatch between regulatory aspirations and the reality.

3) *Lack of Proper Explainability Approaches for End Users:* A potential gap for future research emerges as our participants think that the needs of end users are not adequately covered by existing explainability approaches. While safety, trust, and trustworthiness are rated as relevant desiderata for end users, there has not been any explainability approach that is at least moderately applicable to them. Of all the inapplicable approaches, standards and certifications perform best at a mean value of  $A_M = 2.6$  (see Figure 5). While this approach is certainly not ideal to cover all three relevant desiderata (see Figure 6), it may still be a good starting point to develop novel techniques that could satisfy end users’ needs.

### C. Limitations

As our primary objective with this exploratory survey was to demonstrate that our framework is suitable for identifying potential research gaps, it has some limitations. Most importantly, we only had hardware experts as participants, which

limits the generalizability of our findings. We are aware that statements made by hardware experts about end users should be taken with caution, as this is one stakeholder group making statements about a completely different one. In particular, the potential misalignment of regulatory initiatives (Section V-B2) and the lack of proper explainability approaches for end users (Section V-B3) may be impacted by this effect. With this exploratory study, however, we primarily aim to demonstrate that our framework is suitable for identifying potential research gaps. Our findings lay the foundation for future work to replicate the same protocol with other stakeholders.

## VI. DISCUSSION AND FUTURE WORK

To date, the understanding of hardware matters has garnered little attention in societal discourse. However, as becomes evident from media coverage (e. g., [12], [27]) and regulatory initiatives (e. g., [4], [5]), this is likely about to change. Against this backdrop, the primary goals of our work are to introduce the concept of XHW, motivate why it is needed, and offer a starting point for future research in RE and other fields [16] to thoroughly engage with all stakeholders involved. This lays the groundwork for requirements engineers of hardware system to formulate XHW requirements that suit all stakeholders’ needs.

An important direction of future research is to involve laypersons in studies and discourses about hardware. Such studies could focus on the general XHW needs of end users or their understanding of microchips. Our framework provides a useful starting point for such investigations. Moreover, studies with expert populations in the form of interviews could provide the qualitative insights currently lacking, e. g., by probing deeper into the relevance of XHW for regulatory initiatives.

Overall, we hope that our framework opens the debate on XHW and others will join us in studying the needs of different stakeholders in this new domain more broadly. Below, we discuss how our work informs concrete directions for future research, noting specific relevance to requirements engineers.

### A. Research Direction 1: Filling in Explainability Gaps

One potential explainability gap identified through our survey results concerns the end users, as the hardware experts who participated in our study think that none of the proposed XHW approaches would apply to end users. This does not mean that end users do not need XHW at all, as hardware experts speculate that certain desiderata (viz., security, safety, and trust) would still be important to end users.

Nonetheless, research on end users’ understanding of hardware is extremely limited. While it is reasonable to expect that end users have little understanding of hardware due to a lack of technical expertise, it is important to uncover their specific mental models of how hardware works. Drawing from work on end users’ mental models of other topics such as computer security [54] and the internet [55], we expect that research on end users’ mental models of hardware could shed light on the information needed in explaining hardware to end users.

Based on these deliberations, we derive our first two research questions (RQs) for future work:

**RQ1:** What are end users’ mental models of hardware?

**RQ2:** What information about hardware is required to align end users’ mental models with system models?

In order to answer RQ1, a qualitative user study could elicit end users’ mental models of hardware. For RQ2, we would need to verify, again in a user study, whether the approaches of our framework are really insufficient to produce information that can reconcile end users’ mental model with the actual system models. If the results for RQ2 show that the existing XHW approaches are indeed insufficient for end users, new approaches must be developed for their needs.

Research on usable security and privacy has been successful in developing and advocating standardized labels for IoT devices [56] and smartphone apps [57], drawing inspiration from food nutrition labels. Against this background, we imagine that hardware labels could function as an explainability approach for hardware that is particularly suitable for end users.

Accordingly, we formulate a third RQ for future research:

**RQ3:** What information shall hardware labels provide to meet the XHW requirements and needs of end users?

There are already standards and norms in the hardware community that requirements engineers can draw from to answer this RQ, such as ISO 26262 [41] and IEC 61508 [42] for hardware safety and FIPS PUB 140-3 [43] for hardware security. A comprehensive requirements elicitation and interpretation, as well as validation (e.g., through prototyping), would be next steps here.

### B. Research Direction 2: Devising New XHW Approaches

We see the potential of adapting existing research on XAI (and explainability in general) to generate new approaches for XHW. In line with the above idea of hardware labels, research on model cards [58] and datasheets for datasets [59] can already be found in the debates on XAI. This research could serve as a basis for adapting results from XAI and, in particular, XAI approaches for XHW.

Further XAI approaches we envision to be adaptable are *counterfactual explanations* [60] and *LIME* [61]. Counterfactual explanations allow system integrators to check whether a given microchip conforms to its original schematic by examining what needs to change in the microchip’s input for a given deviation in its output. Employing the idea behind LIME, hardware designers could create surrogate models of a microchip for specific input regions by varying its input and looking at the changes in the output. These models could be used to simulate and check the microchip’s behavior.

Against this background, another avenue for future research would be to analyze the extent to which (ideas behind) XAI approaches can be transferred to XHW to devise new explainability approaches for hardware, leading to the following RQ:

**RQ4:** How can XAI approaches be adopted for XHW?

A possible starting point for answering this RQ would be to distill features of XAI approaches and see whether these

are useful for hardware. Such features could be taken from XAI taxonomies or reviews (e.g., [3], [16], [62]), and their suitability for XHW could be discussed in a workshop with hardware experts.

### C. Research Direction 3: The Right to Repair

Given the considerable ecological footprint of producing modern electronics, repairability has become a concern for both legislators and end users. In fact, a *right to repair* is enshrined in law in more and more countries (e.g., the United States, the United Kingdom, and India), requiring, among others, that a device should be constructed and designed in a manner that allows repairs to be made easily (see, e.g., [63]). Providing details on microchip functionality and interfaces, which aligns with the goals of XHW, could empower end users to properly judge causes for microchip breakdown, determine the suitability of replacement parts, and learn how to replace affected components. Furthermore, information on utilized materials and compositions could aid recycling efforts of products broken beyond repair.

Thus, the final avenue for future research that we want to emphasize is to find out if and how XHW could support a right to repair. We derive our last RQ for future research:

**RQ5:** How can XHW help to provide and obtain the information necessary to exercise the right to repair?

The hardware labels mentioned above (see Section VI-A) could be a first step to answer this RQ, requiring future research based on our framework. Overall, however, a thorough requirements elicitation and interpretation, especially from system integrators, would be valuable in answering this RQ. The collected requirements would then have to be reviewed by legal scholars to determine whether they cover the law.

## VII. CONCLUSION

In this article, we propose Explainable Hardware (XHW) as a new concept and motivate its necessity through legislative initiatives and existing findings on XAI as well as software explainability. Motivated by these findings, we conceive a framework comprising different stakeholders, their desiderata, and existing approaches towards XHW. In an exploratory survey among 18 hardware experts, we demonstrated its applicability and usefulness to identify research gaps.

Our framework paves the way for further studies, e.g., to explore end users’ needs concerning XHW more concretely. In particular, it is plausible that end users will interact differently with a system, depending on the installed hardware and their level of information about it. In this line of thought, prospective XHW approaches could go hand in hand with right-to-repair legislation to facilitate repair and recycling. Inspirations for such approaches could be taken from XAI, e.g., in the form of labels, or by adopting other existing XAI techniques. Overall, XHW represents a vital new research direction to which RE researchers can contribute significantly.



## ACKNOWLEDGMENTS

We would like to thank Gabriel Lima and all participants of the EIS colloquium for feedback on early versions of this article. Furthermore, we would like to thank three anonymous reviewers and Hannah Deters for helpful comments on a more advanced version of this article. Special thanks go to Jakob Droste and Markus Langer for feedback on all versions.

Work on this paper was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—EXC 2092 CASA—390781972, through the DFG grant 389792660 as part of TRR 248, and by the Volkswagen Foundation grants AZ 9B830, AZ 98509, and AZ 98514 “Explainable Intelligent Systems” (EIS).

The Volkswagen Foundation and the DFG had no role in preparation, review, or approval of the manuscript; or the decision to submit the manuscript for publication. The authors declare no other financial interests.

## REFERENCES

- [1] T. Speith, J. Speith, S. Becker, Y. Zou, A. Biega, and C. Paar, *Supplementary Material for Research Preview “Explainability as a Requirement for Hardware: Introducing Explainable Hardware (XHW)”*, 2024. [Online]. Available: <https://doi.org/10.5281/zenodo.10991465>
- [2] A. Panesar, “Ethics of intelligence,” in *Machine Learning and AI for Healthcare: Big Data for Improved Health Outcomes*. New York, NY, USA: Apress, 2019, ch. 6, pp. 207–254.
- [3] L. Chazette, W. Brunotte, and T. Speith, “Exploring explainability: A definition, a model, and a knowledge catalogue,” in *Proceedings of the 29th IEEE International Requirements Engineering Conference (RE 2021)*, J. Cleland-Huang, A. Moreira, K. Schneider, and M. Vierhauser, Eds. Piscataway, NJ, USA: IEEE, 2021, pp. 197–208.
- [4] European Commission. (2022) European Chips Act: Staff Working Document. [Online]. Available: <https://digital-strategy.ec.europa.eu/en/library/european-chips-act-staff-working-document>
- [5] Senate of the United States. (2022, August) CHIPS and Science Act 2022 (P.L. 117-167). [Online]. Available: <https://www.congress.gov/bills/117/congress-house-bill/4346/text>
- [6] S. Adey, “The hunt for the kill switch,” *IEEE Spectrum*, vol. 45, no. 5, pp. 34–39, 2008.
- [7] J. Clements and Y. Lao, “Hardware trojan attacks on neural networks,” 2018. [Online]. Available: <https://arxiv.org/abs/1806.05768>
- [8] W. Li, J. Yu, X. Ning, P. Wang, Q. Wei, Y. Wang, and H. Yang, “Hu-Fu: Hardware and software collaborative attack framework against neural networks,” in *Proceedings of the 2018 IEEE Computer Society Annual Symposium on VLSI (ISVLSI 2018)*, W. Zhang, C. J. Xue, and Z. Shao, Eds. Piscataway, NJ, USA: IEEE, 2018, pp. 482–487.
- [9] Y. Liu, L. Wei, B. Luo, and Q. Xu, “Fault injection attack on deep neural network,” in *Proceedings of the 2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD 2017)*, S. Parameswaran, Ed. Piscataway, NJ, USA: IEEE, 2017, pp. 131–138.
- [10] G. T. Becker, F. Regazzoni, C. Paar, and W. P. Burleson, “Stealthy dopant-level hardware trojans,” in *Proceedings of the 15th International Workshop on Cryptographic Hardware and Embedded Systems (CHES 2013)*, ser. Lecture Notes in Computer Science, G. Bertoni and J. Coron, Eds., vol. 8086. Berlin/Heidelberg, Germany: Springer, 2013, pp. 197–214.
- [11] S. Shankland. (2022) M1 Ultra: Apple just unveiled its most powerful mac chip yet. CNET. [Online]. Available: <https://www.cnet.com/tech/computing/m1-ultra-apple-just-unveiled-its-most-powerful-chip-yet/>
- [12] P. Mozur, J. Liu, and R. Zhong. (2022) “The Eye of the Storm”: Taiwan Is Caught in a Great Game Over Microchips. The New York Times. [Online]. Available: <https://www.nytimes.com/2022/08/29/technology/taiwan-chips.html>
- [13] L. Chazette and K. Schneider, “Explainability as a non-functional requirement: Challenges and recommendations,” *Requirements Engineering*, vol. 25, no. 4, pp. 493–514, 2020.
- [14] W. Brunotte, L. Chazette, V. Klös, and T. Speith, “Quo vadis, explainability? – A research roadmap for explainability engineering,” in *Proceedings of the 28th International Working Conference on Requirements Engineering: Foundation for Software Quality (REFSQ 2022)*, ser. Lecture Notes in Computer Science, V. Gervasi and A. Vogelsang, Eds., vol. 13216. Cham, Switzerland: Springer International Publishing, 2022, pp. 26–32.
- [15] M. A. Köhl, K. Baum, M. Langer, D. Oster, T. Speith, and D. Bohlender, “Explainability as a non-functional requirement,” in *Proceedings of the 27th IEEE International Requirements Engineering Conference (RE 2019)*, D. E. Damian, A. Perini, and S. Lee, Eds. Piscataway, NJ, USA: IEEE, 2019, pp. 363–368.
- [16] M. Langer, D. Oster, T. Speith, H. Hermanns, L. Kästner, E. Schmidt, A. Sesing, and K. Baum, “What do we want from explainable artificial intelligence (XAI)? – A stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research,” *Artificial Intelligence*, vol. 296, 2021.
- [17] J. Burrell, “How the machine ‘thinks’: Understanding opacity in machine learning algorithms,” *Big Data & Society*, vol. 3, no. 1, pp. 1–12, 2016.
- [18] S. Mann, B. Crook, L. Kästner, A. Schomäcker, and T. Speith, “Sources of opacity in computer systems: Towards a comprehensive taxonomy,” in *Proceedings of the 31st IEEE International Requirements Engineering Conference Workshops (REW 2023)*, F. Dalpiaz, J. Horkoff, and K. Schneider, Eds. Piscataway, NJ, USA: IEEE, 2023, pp. 337–342.
- [19] J. Bonvoisin, R. Mies, J.-F. Boujut, and R. Stark, “What is the ‘source’ of open source hardware?” *Journal of Open Hardware*, vol. 1, no. 1, pp. 5–5, 2017.
- [20] OpenRISC Community. (2018) OpenRISC. [Online]. Available: <https://openrisc.io/>
- [21] lowRISC CIC. (2019) OpenTitan. [Online]. Available: <https://opentitan.org/>
- [22] The Linux Foundation. (2019) The CHIPS Alliance. [Online]. Available: <https://chipsalliance.org>
- [23] R. S. Portnoff, L. N. Lee, S. Egelman, P. Mishra, D. Leung, and D. A. Wagner, “Somebody’s watching me?: Assessing the effectiveness of webcam indicator lights,” in *Proceedings of the 33rd ACM Conference on Human Factors in Computing Systems (CHI 2015)*, B. Begole, J. Kim, K. Inkpen, and W. Woo, Eds. New York, NY, USA: Association for Computing Machinery, 2015, pp. 1649–1658.
- [24] E. Zeng, S. Mare, and F. Roesner, “End user security and privacy concerns with smart homes,” in *Proceedings of the Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. Berkeley, CA, USA: USENIX Association, 2017, pp. 65–80. [Online]. Available: <https://www.usenix.org/conference/soups2017/technical-sessions/presentation/zeng>
- [25] J. Lau, B. Zimmerman, and F. Schaub, “Alexa, are you listening? privacy perceptions, concerns and privacy-seeking behaviors with smart speakers,” *Proceedings of the ACM on Human-Computer Interaction*, vol. 2, no. CSCW, 2018.
- [26] D. Y. Huang, N. Aphorpe, F. Li, G. Acar, and N. Feamster, “IoT inspector: Crowdsourcing labeled network traffic from smart home devices at scale,” *Proceedings of the ACM on Interactive, Mobile, Wearable, and Ubiquitous Technologies*, vol. 4, no. 2, pp. 1–21, 2020.
- [27] D. Clark. (2021) The tech cold war’s ‘most complicated machine’ that’s out of China’s reach. The New York Times. [Online]. Available: <https://www.nytimes.com/2021/07/04/technology/tech-cold-war-chips.html>
- [28] United States District Court, Southern District of Texas, Houston Division. (2010, January) United States of America vs. Ehab Ashoor. No. 10-20354 (5th Cir. Apr. 29, 2011). [Online]. Available: <https://assets.bwbx.io/documents/users/iqjWHBFdfxIU/r9dKMMM0G5I/v0>
- [29] Defense Science Board Task Force, “Report of the defense science board task force on high performance microchip supply,” Office of the Under Secretary of Defense, Tech. Rep., February 2005. [Online]. Available: <https://dsb.cto.mil/reports/2000s/ADA435563.pdf>
- [30] D. Wang, Q. Yang, A. Abdul, and B. Y. Lim, “Designing theory-driven user-centric explainable AI,” in *Proceedings of the 37th ACM Conference on Human Factors in Computing Systems (CHI 2019)*, S. A. Brewster, G. Fitzpatrick, A. L. Cox, and V. Kostakos, Eds. New York, NY, USA: Association for Computing Machinery, 2019, pp. 1–15.
- [31] Q. V. Liao, D. M. Gruen, and S. Miller, “Questioning the AI: Informing design practices for explainable AI user experiences,” in *Proceedings of the 38th ACM Conference on Human Factors in Computing Systems (CHI 2020)*, R. Bernhaupt, F. F. Mueller, D. Verweij, J. Andres, J. McGrenere, A. Cockburn, I. Avellino, A. Goguy, P. Bjørn, S. Zhao,



- B. P. Samson, and R. Kocielnik, Eds. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1–15.
- [32] W. Brunotte, L. Chazette, V. Klös, E. Knauss, T. Speith, and A. Vogelsang, “Welcome to the first international workshop on requirements engineering for explainable systems (RE4ES),” in *Proceedings of the 29th IEEE International Requirements Engineering Conference Workshops (REW 2021)*, T. Yue and M. Mirakhorli, Eds. Piscataway, NJ, USA: IEEE, 2021, pp. 157–158.
- [33] L. Kästner, M. Langer, V. Lazar, A. Schomäcker, T. Speith, and S. Sterz, “On the relation of trust and explainability: Why to engineer for trustworthiness,” in *Proceedings of the 29th IEEE International Requirements Engineering Conference Workshops (REW 2021)*, T. Yue and M. Mirakhorli, Eds. Piscataway, NJ, USA: IEEE, 2021, pp. 169–175.
- [34] M. Langer, K. Baum, K. Hartmann, S. Hessel, T. Speith, and J. Wahl, “Explainability auditing for intelligent systems: A rationale for multidisciplinary perspectives,” in *Proceedings of the 29th IEEE International Requirements Engineering Conference Workshops (REW 2021)*, T. Yue and M. Mirakhorli, Eds. Piscataway, NJ, USA: IEEE, 2021, pp. 164–168.
- [35] Z. C. Lipton, “The myths of model interpretability,” *ACM Queue*, vol. 16, no. 3, p. 31–57, 2018.
- [36] R. R. Hoffman, S. T. Mueller, G. Klein, and J. Litman, “Metrics for explainable AI: Challenges and prospects,” 2018. [Online]. Available: <https://arxiv.org/abs/1812.04608>
- [37] A. Tomlinson, S. Parkin, and S. A. Shaikh, “Drivers and barriers for secure hardware adoption across ecosystem stakeholders,” *Journal of Cybersecurity*, vol. 8, no. 1, 2022.
- [38] T. Speith, “Building bridges for better machines: From machine ethics to machine explainability and back,” Ph.D. dissertation, Saarland University, 2023.
- [39] K. Vaidyanathan, B. P. Das, H. E. Sumbul, R. Liu, and L. T. Pileggi, “Building trusted ICs using split fabrication,” in *Proceedings of the 7th IEEE International Symposium on Hardware-Oriented Security and Trust (HOST 2014)*, F. Koushanfar, M. Hsiao, and M. Potkonjak, Eds. Piscataway, NJ, USA: IEEE, 2014, pp. 1–6.
- [40] F. Imeson, A. Emtenan, S. Garg, and M. V. Tripunitara, “Securing computer hardware using 3D integrated circuit (IC) technology and split manufacturing for obfuscation,” in *Proceedings of the 22nd USENIX Security Symposium (USENIX Security 2013)*, S. T. King, Ed. Berkeley, CA, USA: USENIX Association, 2013, pp. 495–510. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/presentation/imeson>
- [41] International Organization for Standardization. (2018) Road vehicles – functional safety – part 11: Guidelines on application of ISO 26262 to semiconductors. [Online]. Available: <https://www.iso.org/obp/ui/#iso:std:iso:26262:-11:ed-1:v1:en>
- [42] S. Brown, “Overview of IEC 61508 design of electrical/electronic/programmable electronic safety-related systems,” *Computing and Control Engineering Journal*, vol. 11, no. 1, pp. 6–12, 2000.
- [43] National Institute of Standards and Technology. (2019) FIPS PUB 140-3 - Security Requirements for Cryptographic Modules. [Online]. Available: <https://doi.org/10.6028/NIST.FIPS.140-3>
- [44] N. Nayani and M. Mollaghasemi, “Validation and verification of the simulation model of a photolithography process in semiconductor manufacturing,” in *Proceedings of the 30th Conference on Winter Simulation: Simulation—A Bridge to the Future (WSC 1998)*, D. J. Medeiros, E. F. Watson, J. S. Carson II, and M. S. Manivannan, Eds. Piscataway, NJ, USA: IEEE, 1998, pp. 1017–1022.
- [45] R. L. Geiger, P. E. Allen, and N. R. Strader, *VLSI Design Techniques for Analog and Digital Circuits*. New York, NY, USA: McGraw-Hill Publishing Company, 1990.
- [46] W. Maly, “Realistic fault modeling for VLSI testing,” in *Proceedings of the 24th ACM/IEEE Design Automation Conference (DAC 1987)*, A. O’Neill and D. Thomas, Eds. New York, NY, USA: Association for Computing Machinery, 1987, pp. 173–180.
- [47] A. Kuehlmann, A. Srinivasan, and D. P. LaPotin, “Verity - A formal verification program for custom CMOS circuits,” *IBM Journal of Research and Development*, vol. 39, no. 1-2, pp. 149–166, 1995.
- [48] M. Althoff, S. Yaldiz, A. Rajhans, X. Li, B. H. Krogh, and L. T. Pileggi, “Formal verification of phase-locked loops using reachability analysis and continuization,” in *Proceedings of the 2011 IEEE/ACM International Conference on Computer-Aided Design (ICCAD 2011)*, J. R. Phillips, A. J. Hu, and H. Graeb, Eds. Piscataway, NJ, USA: IEEE, 2011, pp. 659–666.
- [49] R. Torrance and D. James, “The state-of-the-art in IC reverse engineering,” in *Proceedings of the 11th International Workshop on Cryptographic Hardware and Embedded Systems (CHES 2009)*, ser. Lecture Notes in Computer Science, C. Clavier and K. Gaj, Eds., vol. 5747. Berlin/Heidelberg, Germany: Springer, 2009, pp. 363–381.
- [50] L. Azriel, J. Speith, N. Albartus, R. Ginosar, A. Mendelson, and C. Paar, “A survey of algorithmic methods in IC reverse engineering,” *Journal of Cryptographic Engineering*, vol. 11, no. 3, pp. 299–315, 2021.
- [51] P. C. Kocher, “Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems,” in *Proceedings of the 16th Annual International Cryptology Conference (CRYPTO 1996)*, ser. Lecture Notes in Computer Science, N. Kobitz, Ed., vol. 1109. Berlin/Heidelberg, Germany: Springer, 1996, pp. 104–113.
- [52] D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, “The EM side-channel(s),” in *Proceedings of the 4th International Workshop on Cryptographic Hardware and Embedded Systems (CHES 2002)*, ser. Lecture Notes in Computer Science, B. S. Kaliski, c. K. Koç, and C. Paar, Eds., vol. 2523. Berlin/Heidelberg, Germany: Springer, 2002, pp. 29–45.
- [53] M. Hsueh, T. K. Tsai, and R. K. Iyer, “Fault injection techniques and tools,” *Computer*, vol. 30, no. 4, pp. 75–82, 1997.
- [54] R. Wash, “Folk models of home computer security,” in *Proceedings of the Sixth Symposium on Usable Privacy and Security (SOUPS 2010)*, L. F. Cranor, Ed. New York, NY, USA: Association for Computing Machinery, 2010.
- [55] R. Kang, L. Dabbish, N. Fruchter, and S. B. Kiesler, ““My data just goes everywhere”: User mental models of the internet and implications for privacy and security,” in *Proceedings of the Eleventh Symposium On Usable Privacy and Security (SOUPS 2015)*, L. F. Cranor, R. Biddle, and S. Consolvo, Eds. Berkeley, CA, USA: USENIX Association, 2015, pp. 39–52. [Online]. Available: <https://www.usenix.org/conference/soups2015/proceedings/presentation/kang>
- [56] P. Emami-Naeini, Y. Agarwal, L. F. Cranor, and H. Hibshi, “Ask the experts: What should be on an IoT privacy and security label?” in *Proceedings of the 41st IEEE Symposium on Security and Privacy (SP 2020)*, G. F. Ciocarie, H. Shacham, and A. Oprea, Eds. Piscataway, NJ, USA: IEEE, 2020, pp. 447–464.
- [57] L. F. Cranor, “Mobile-app privacy nutrition labels missing key ingredients for success,” *Communications of the ACM*, vol. 65, no. 11, pp. 26–28, 2022.
- [58] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, “Model cards for model reporting,” in *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency (FAT\* 2019)*, D. Boyd and J. H. Morgenstern, Eds. New York, NY, USA: Association for Computing Machinery, 2019, pp. 220–229.
- [59] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. M. Wallach, H. Daumé III, and K. Crawford, “Datasheets for datasets,” *Communications of the ACM*, vol. 64, no. 12, pp. 86–92, 2021.
- [60] S. Wachter, B. Mittelstadt, and C. Russell, “Counterfactual explanations without opening the black box: Automated decisions and the GDPR,” *Harvard Journal of Law & Technology*, vol. 31, no. 2, pp. 841–887, 2017.
- [61] M. T. Ribeiro, S. Singh, and C. Guestrin, ““Why should I trust you?”: Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2016)*, B. Krishnapuram, M. Shah, A. J. Smola, C. C. Aggarwal, D. Shen, and R. Rastogi, Eds. New York, NY, USA: Association for Computing Machinery, 2016, pp. 1135–1144.
- [62] T. Speith, “A review of taxonomies of explainable artificial intelligence (XAI) methods,” in *Proceedings of the 5th ACM Conference on Fairness, Accountability, and Transparency (FAccT 2022)*, C. Isbell, S. Lazar, A. Oh, and A. Xiang, Eds. New York, NY, USA: Association for Computing Machinery, 2022, pp. 2239–2250.
- [63] S. Svensson, J. L. Richter, E. Maitre-Ekern, T. Pihlajarinne, A. Maigret, and C. Dalhammar, “The emerging ‘right to repair’ legislation in the EU and the US,” in *Proceedings of the 7th International Symposium and Environmental Exhibition on Going Green—CARE INNOVATION*, B. Kopacek, Ed., 2018, pp. 27–29.