Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Designing a Boundary Negotiating Artifact for Collaborative Socio-Technical Sense-Making in AI Regulatory Sandboxes

Idoia Landa-Oregi

idoia.landa@list.lu

Luxembourg Institute of Science and Technology

Tom Deckenbrunnen

Luxembourg Institute of Science and Technology

Alessio Buscemi

Luxembourg Institute of Science and Technology

Daniele Pagani

Luxembourg Institute of Science and Technology

German Castignani

Luxembourg Institute of Science and Technology

Research Article

Keywords:

Posted Date: November 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11812637/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Designing a Boundary Negotiating Artifact for Collaborative Socio-Technical Sense-Making in AI Regulatory Sandboxes

Idoia Landa-Oregi[1*], Tom Deckenbrunnen[1,2], Alessio Buscemi[1], Daniele Pagani[1], German Castignani[1]

[1*]Luxembourg Institute of Science and Technology, 5, avenue des hauts-fourneaux, Esch/Alzette, L-4362 Luxembourg.
[2]University of Luxembourg, 2, place de l'Université, Esch/Alzette, L-4365, Luxembourg.

*Corresponding author(s). E-mail(s): idoia.landa@list.lu;
Contributing authors: tom.deckenbrunnen@list.lu;
alessio.buscemi@list.lu; daniele.pagani@list.lu;
german.castignani@list.lu;

Abstract

The rapid, unpredictable advancements in AI system capabilities has seen regulators take adaptive and experimental approaches to policymaking. Established in other domains as instruments balancing regulation with innovation, regulatory sandboxes are estimated as viable solutions for AI regulation. However, analyses mostly focus on the legal and institutional design of AI Regulatory Sandboxes (AIRSes). With the legal framework leaving the socio-technical interpretation to stakeholders, this creates a gap on the sense-making required to fulfill the AIRS purpose. In this paper, we approach this by designing the prototype of a technical Boundary Negotiating Artifact as a way to mediate meaning in AIRSes. We adopt a Research-through-Design approach to iteratively develop a prototype, providing an interface for the different stakeholders to collaborate in AI assessment. We then position it as technical backbone in an AIRS framework synthesized from the literature, structuring the collaborative sense-making of the involved stakeholders. We further report the insights gained from our design process and engagement with stakeholders, specifically how the proposed prototype serves as a technical Boundary Negotiating Artifact (BNA) to help cooperation among different communities of practice.

Keywords: AI regulation, boundary negotiating artifact, collaborative sense-making, prototype design, research-through-design, socio-technical systems

1 Introduction

An AIRS, as defined by Articles 57 and 58 of the EU Artificial Intelligence Act (EU AI Act), is a controlled environment intended to foster innovation across the AI system life cycle and contribute to regulatory learning (European Commission 2024). It consists of the supervision of a prospective AI system provider by a National Competent Authority (NCA), in particular through legal guidance, but also through technical expertise and resources. Thus, the provisions laid out by the EU AI Act create the institutional space for collaboration between diverse stakeholder groups from legal to technical communities. Beyond the technical or legal challenges these stakeholders might face, a big challenge lies in the need for sense-making (Guio Español and Koenig 2025). How do stakeholders with different expertise, backgrounds, vocabularies or interests build a shared understanding to make consequential decisions together under the uncertainty of AI regulation?

This is fundamentally a collaborative sense-making problem. Being at the intersection of interdisciplinary communities, AIRSes generate information asymmetries, referring to the uneven access or understanding of relevant information among stakeholders, which can affect the transparency and decision-making of the collaborative process. Nevertheless, with the appropriate artifacts, AIRSes can work as levers to reducing said information asymmetries, pairing designed testing mechanisms with the combined expertise of the stakeholders involved (Español and Koenig 2025). However, regardless of the various efforts on AIRS definition and implementation, little is known about how this information asymmetry is actually experienced by the people who must navigate it, or how it might be addressed in practice.

Consider what a sandbox encounter requires for its stakeholders. An AI innovator must translate a system they have developed and understand in technical terms into a language of compliance, with the conformity obligations that the regulatory framework demands. At the same time, a regulator must assess a system they have never encountered before, in a specific sector, using legal provisions that were defined in a level of abstraction that does not clearly map to technical specifications. Both stakeholders carry genuine expertise of their own field, but lack the vocabulary of the other. There is a gap in communication that the regulation does not address. AIRSes are defined as a legal process, offering a sound procedure without the sense-making infrastructure that can make the process adequate to all stakeholders.

Prior work has analyzed AIRSes as institutional environments for multi-stakeholder boundary negotiation over AI systems and AI regulation (Deckenbrunnen et al. 2026). Particular focus was put on the need for BNAs as effective support of collaborative sense-making and boundary object stabilization, enabled by the legal uncertainty caused by abstract legislation. A BNA, being documentation, a model, or a process, can be used by any stakeholder group to communicate their perspective to another group, reducing the information asymmetry. While using a BNA can genuinely reduce

the information asymmetry, its design strongly relies on a deep understanding of its prospective users. Understanding their current experience with AI regulations and testing, pain points, and needs is essential for the successful definition of an artifact that will ease their overall experience. Applied to AIRSes, this means that a BNA cannot be purely designed from a legal text, nor for a technical best practice alone, as it should be shaped by how innovators and regulators actually experience the sandbox encounter, the assets they each bring, and the points of divergence in their understanding. This stakeholder-centered view is what allows BNA to move from a compliance artifact to a genuine sense-making tool.

This approach is closely aligned with that of the Computer-supported Cooperative Work (CSCW) community, as the creation of said technical tool would support the cooperation among different communities in a complex setting (Schmidt and Simone 1996). This paper addresses this gap by designing a prototype for a tool that works as a the socio-technical system for structuring the collaboration of stakeholders in an AIRSes, reducing the information asymmetry and increasing cooperation. Thus, adopting a Research-through-Design (RtD) approach, it first understands the experience of regulators and innovators in this context and then reports the design of a prototype to enhance this collaboration and expand the consensus AIRS framework with a technical BNA.

This work adds to the established interests of CSCW literature in multi-disciplinary work and boundary negotiation by extending it specifically to the field of AI regulation and assessment. It reports a RtD investigation conducted through qualitative research with AI innovators, technical experts, and regulatory authorities, and through an iterative design process producing the prototype of a technical BNA to mediate their interaction in AIRSes. Specifically, the research is guided by the following Research Questions (RQs):

1. How do AI innovators, technical experts, and regulatory authorities differently understand, experience, cooperate in and navigate the obligations and ambiguities of AI regulations?
2. What design knowledge emerges from creating the prototype of a technical BNA that mediates that experience?
3. How can that knowledge inform AIRS operationalization?

2 Research context

2.1 Legal and Institutional Design of AIRS in Europe

AIRSes were included in the EU AI Act as a mechanism to support innovation and regulatory learning under conditions of legal uncertainty, especially for smaller organizations that often lack the resources to navigate compliance on their own. Within a regulatory sandbox, a NCA provides an AI innovator with direct regulatory guidance and supervisory support. This provision is inspired by a precedent set by financial regulations, where sandboxes have been used since 2016 to let fintech innovators test novel products under supervised conditions, an approach that has since been adopted by many other sectors(Fahy and Lauren Fahy 2022).

The EU AI Act translates this into an institutional structure, requiring every member state to establish at least one national AIRS operated by the designated NCA. In parallel, the AI office is responsible for a sandbox addressing general-purpose AI systems at European Union (EU) level. Beyond this structure, however, the EU AI Act leaves considerable room for variation, as AIRSes may differ by sector, technical expertise of the NCA, or the specific testing infrastructure they are able to offer. For instance, an AI innovator operating across multiple member states or a system intended for cross-boarder deployment may need to satisfy expectations that differ from one national sandbox to the next. This fragmentation in itself is a motivation for the designed prototype and its application within the framework detailed in this paper, as a shared tool and process structure offering a single path toward the kind of interoperability that current provisions lack.

Several scholars have examined how AIRSes should be designed and governed. Novelli et al. (2026) provide an in-depth description of AIRS design from a regulatory lens, detailing the steps and tasks that structure each step of the process. (Chen and Taeihagh 2026) complement this work with a framework of characteristics, design choices, and functionalities an AIRS needs to fulfill its intended purpose. Similarly, (Español and Koenig 2025) analyze best practices for existing sandboxes to support a more robust inquiry into the learning process within an AIRS. This body of work establishes *what* a well-informed AIRS should consider both procedurally and structurally. What remains underexplored is *how* the stakeholders engaging in such an environment practically communicate and build shared understanding.

Following these guidelines, some EU Member States have made concrete progress toward implementing their EU AI Act mandated sandboxes, but the considerable discretion left to NCAs risks producing fragmented or incompatible approaches that could undermine the implementation of the EU AI Act (Ahern 2025). This concern is echoed at the institutional level as the European Parliament has flagged that fragmented enforcement could leave national authorities with uneven capacities, encouraging "sandbox arbitrage" as providers intentionally select less strict sandboxes as a result (European Parliament 2026; Ahern 2025). The EU AI Act itself anticipates this risk, tasking the Commission with adopting implementing acts under Article 58(2) to establish common principles on eligibility, application, and exit procedures precisely "in order to avoid fragmentation across the Union" (European Commission 2024). However, this harmonization strategy operates at the level of legal procedure and does not address the day-to-day, cross-stakeholder experience of navigating a sandbox.

2.2 Boundary Negotiation in Multidisciplinary Collaborative Work

Understanding the complex interactions arising in collaborative work between diverse stakeholder groups is a non-trivial task. Through a series of ethnographic studies examining such interactions, Star and Griesemer (1989) noted that even in contexts with no consensus among communities of practice, the collaboration tends to be structured around objects with a specific set of properties (Star and Griesemer 1989). Such boundary objects, be they material or procedural, are robust enough to serve as a

shared core between stakeholders and, at the same time, possess the interpretive flexibility to accommodate the diverging information needs of different groups. Leigh Star (2010) expands on her previous work through an account of the life cycle of boundary objects (Leigh Star 2010). The creation and stabilization of boundary objects rely on phases attempting to standardize ill- and well-structured interactions among stakeholders, and the introduction of new perspectives with engagement of new communities of practice. Leigh Star (2010) further raised concerns about the misuse of the concept of boundary objects as any object offering any amount of interpretive flexibility. The analysis by Lee (2007) notes a disconnect between the boundary object concept and the methods of standardization, highlighting the issue of the concept itself becoming a catch-all term (Lee 2007). The paper addresses this disconnect by introducing the notion of BNAs, artifacts that are distinct from boundary objects and serve as a support in collaboratively negotiating the meaning and eventually stabilizing the boundary objects in question. More work has investigated the use and design of BNAs in interdisciplinary, collaborative work in both engineering and knowledge work contexts (Beddoes et al. 2011; Lee et al. 2026). Beyond these domains, boundary objects and BNAs have also been used to harmonize different vocabularies and methods across product design practices (Velleu et al. 2023), to coordinate distributed roles and responsibilities in agile systems (Wohlrab et al. 2019), and to bridge epistemological differences between qualitative and quantitative traditions in mixed-methods research (Wenger-Trayner et al. 2019). The existing literature illustrates how boundary objects and BNAs can work across disciplines wherever heterogeneous expertise must be harmonized and serve as a useful analytical tool in understanding multi-disciplinary collaborations.

There is emerging literature from the CSCW and Human-Computer Interaction (HCI) communities expanding on these notions to the artificial intelligence space by investigating the relationship of AI systems and the theory of boundary objects and negotiation. One dominant strand of work explores the utility of AI in supporting collaborative boundary work (Grasso and Colombino 2026; (Chuqiao) Wan et al. 2026; Gunasekaran et al. 2026; Guridi et al. 2025). Complementarily, there is recent work considering AI systems themselves as the subject of boundary negotiation (Tadimalla and Maher 2026; Wu et al. 2026).

Deckenbrunnen et al. (2026) adopt the latter perspective by considering both the AI system and the applicable AI regulation as boundary objects for the various stakeholders involved in AI system development and assessment. Similarly to AI auditing innovation arising from a lack of standardized procedures and tools (Schiff et al. 2024), Deckenbrunnen et al. (2026) posit that the current legal uncertainty, caused by vague, human value-based regulation creates the productive space for collaborative socio-technical sense-making. AIRSes in particular are interpreted as institutional environments of boundary negotiation as they unite technical and non-technical stakeholders with varying fields of expertise. The authors note, however, that the intended purpose of AIRSes as innovation support and legal tools of reducing uncertainty can only be fulfilled if there are appropriate BNAs supporting the collaborative work. From

this boundary negotiation interpretation, Deckenbrunnen et al. (2026) provide a preliminary analysis of components and requirements of technical frameworks for them to serve as effective BNAs, such as collaborative dashboards and role-specific access.

2.3 Human Experience of Stakeholders in AIRS

The proper design of a BNA to bridge the information asymmetry among stakeholders relies on the understanding of their experience and needs. While offering services in AI assessment our research team has engaged with multiple industry partners which has given us meaningful insights into stakeholder needs. These insights summarize the current experience and needs of the three key stakeholders, AI innovators, regulators, and technical experts. The AI innovator is a company or organization developing an AI system that needs to be tested. Technical experts are third-party assessment experts who are familiar with AI legal requirements and have the know-how to perform technical assessment of AI systems. Lastly, the NCAs are the ones that need to help with the interpretation of results and give guidance and support to AI innovators.

AI innovators. AI innovators are experts in building AI systems, but may lack the expertise or resources to conduct the assessment of their system, beyond more immediate performance metrics. AI regulations, such as the EU AI Act, present high regulatory complexity to AI innovators, introducing many requirements that are difficult to interpret without assessment or legal expertise. Knowing what test to use for their specific use case, finding said test, learning how to perform the test, or interpreting the results and compiling the results in a report are activities that most innovators struggle with. Moreover, these regulations include extensive documentation requirements, including the production of evidence and the maintenance of technical documentation throughout the AI life-cycle.

AI innovators face the challenge of balancing innovation activities with regulatory obligations, as they often perceive that compliance activities increase development time and administrative burden, making it more difficult to iterate rapidly while remaining competitive. Moreover, they find difficulties in translating legal requirements into technical practices, as concepts are expressed in legal language that is not always easily translated into engineering tasks. Lastly, the evolving guidance poses a challenge. Although the EU AI Act establishes the legal framework, many implementation details are still to be addressed by standards, guidelines, and codes of practice, with their current absence creating legal uncertainty during development.

Technical experts. A prominent way of conducting AI assessment is through digital sandboxes, which we will refer to as AI Assessment Sandboxes (AIASes), adopting the denomination suggested in (Buscemi et al. 2026). However, the setup of technical environments for AI assessment remains highly complex and time-consuming, creating a significant barrier to adoption at scale. Technical experts have to be able to identify relevant requirements and translate those into metrics and tests. After that, they have to find the relevant tests from a wide variety of sources and manually integrate them to execute the tests. Once the testing is done, they have to manually harmonize the results of each test into a coherent database and dashboard and manually collect the evidence of the tests that were performed. Lastly, they have to interpret the results

with feedback from different experts and manually compile everything into a report, potentially iterating on system design and the assessment procedure if necessary.

There is a growing demand for AI assessment in the EU that current assessment practices cannot fulfill. The setup of the AIAS is unique to each use case or intended purpose, which requires extended time and effort. Each use case and sector has different implications, with specific risks, metrics, testing methods and thresholds, where a common AIAS is not viable. Lastly, there is a clear need for multi-disciplinary collaboration.

NCAs. While the AI Act clearly sets out the responsibilities of NCAs regarding the AIRS, the operationalization of these sandboxes remains unclear, leaving significant discretion regarding their governance, criteria, testing procedures, and evaluation methods. Authorities must encourage experimentation while ensuring compliance, a balance that is difficult to achieve in practice. Moreover, most AI applications often fall under several regulatory domains, requiring coordination among multiple regulatory stakeholders with different expertise.

NCAs have limited to no technical expertise in AI assessment depending on the country, making it difficult to effectively supervise innovative AI systems. While they provide the regulatory expertise, they need practical methods and guidance to evaluate compliance with technical requirements such as robustness, transparency, or human oversight.

While the technical experts are the ones setting up the environment for assessing the AI system, the interpretation of the produced results requires cooperation among different experts and entities. In fact, contextualization in the given use case is essential for proper interpretation of the assessment results. For instance, an AI system for object detection, achieving 99 % accuracy, may be appropriate for a security camera with an alert system, but wholly unacceptable for pedestrian detection in an autonomous car. Similarly, abstract notions such as bias or other risks to fundamental rights are difficult to capture in a technical, one-size-fits-all manner, as established technical metrics may be met, but ignore, e.g., legal implications arising from case law. This makes collaborative, multi-stakeholder sense-making a particularly integral part of AI assessment.

3 Methods

This study follows a RtD approach, where the research is guided through the design process logic and supported by scientific research questions (Zimmerman et al. 2007). The design process involves various iterations, favoring a constant realignment of the artifacts to better tackle the design problems (Cortesão and Lenzholzer 2022; Godin and Zahedi 2014). The method is applied to design a prototype of a BNA and define its application framework, facilitating the interactions of stakeholders in an AIRS.

The RtD process, shown in Figure 1 is guided by the three RQs defined in Section 1. With the context of AIRSes in the EU and the the current experience of the different stakeholders as a starting point, we developed an open source tool (whose name is hidden due to double blind review, but whose identity and url will be included upon

acceptance of the manuscript) to mediate the experience and interactions of the stakeholders through three design iterations. The design process for the development of the prototype followed the main phases of Human-Centered Design (Giacomin 2014) through three design iterations (see Figure 1). The first phase, *understand*, focuses on developing a grounded understanding of the problem, user needs, and the overall design context. The second phase, *define*, focuses on synthesizing all the insights gathered in the first phase to frame the challenge and requirements for the design. The third phase, *develop*, consists in generating ideas and developing potential solutions. In the last phase, *evaluate*, those solutions are tested with the end users to gather feedback and refine the product through iterative loops. The main value of this approach lies in engaging users and stakeholders in each of the iterations to address user needs while still maintaining its viability.

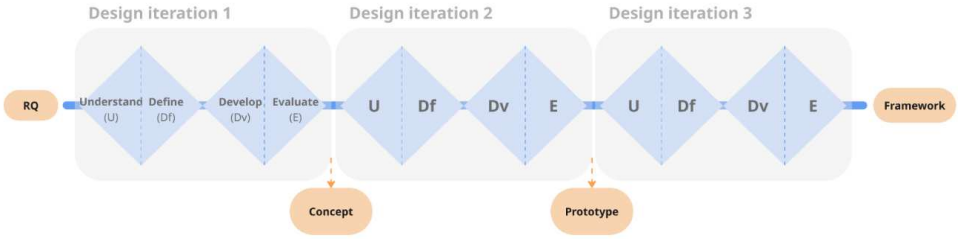

Fig. 1 Design process for the development of the prototype. It starts from the RQ and follows 3 double diamond design iterations, concluding in the framework

The expectations of the different stakeholders were adjusted with every design iteration, making sure the prototype addressed their challenges and clarified the roles and responsibilities in an AIRS. The prototype was used as an instrument for producing knowledge throughout the design process, as it served as a baseline for developing the framework presented in Section 5. This framework presents how the three main stakeholders potentially interact in an AIRS having the designed tool as a main touchpoint. It highlights the roles of each stakeholder and their exchanges through the tool.

4 Designing a computer-assisted BNA for AIRS

This section presents the findings gathered from the design process of a digital prototype that facilitates the exchanges among stakeholders by helping users in the configuration of AIASes to assess AI systems. These findings served as a baseline to enrich the framework presented in Section 5 that outlines the interactions among the stakeholders in the operationalization of AIRSes. Thus, this section reports on the details of the three design iterations in the development of the prototype, giving special focus to the associated learnings.

The tool, presented in detail in Buscemi et al. (2026), provides a curated, extensible catalog of AI tests, controls and datasets integrated through an open plug-in API interface. It guides users through mapping regulatory requirements to assessment

metrics, selecting from the catalog the tests and controls relevant to their specific use case, configuring a collaborative dashboard to execute tests and visualize results, and generating structured assessment reports.

4.1 Design Conception

The first design iteration had the purpose of developing a concept that would bridge the different experiences of the users into a single harmonized platform. Without going into the development of the digital platform yet, this cycle consisted in testing the concept with the relevant stakeholders. In order to do so, the concept was first proposed in a white paper by Buscemi et al. (Buscemi et al. 2025) presenting the general concept and architecture.

To validate the concept with the relevant communities, a presentation was made by attending Nexus, one of the premier technology events accelerating AI and technological solutions in the EU. This is a multidisciplinary event gathering attendants of all the relevant stakeholders mentioned in section 2.3. Therefore, this approach allowed the team to gather insights from all the relevant stakeholders at once. The concept was informally presented to ten attendees, from all three stakeholder groups. The feedback was gathered through a semi-structured discussion and analyzed by the team through debrief notes to identify possible improvements.

With the EU AI Act's general provisions and prohibitions applying shortly before the conference, the concept was well received especially among the technical stakeholders, understanding it as a bridge between the technical and regulatory needs. The possibility to have a tool that has the potential to help them navigate all the needed testing of their systems, in a harmonized way, was highly appreciated. However, only technical users really understood the concept at this point, with the original design proving too technical for easy understanding by non-technical stakeholders. This highlighted the need to establish a less technical path in the tool, where regulators can access the results of the testing in a simplified manner and to be able to track the testing that was conducted.

These findings shaped the first design of the prototype, presented in Section 4.2, and helped formalize the interactions of the stakeholders in the core phases of the framework, *testing* and *evaluation*, presented in Section 5.

4.2 Building a Demonstrator

The second design iteration consisted in building the first prototype of the digital tool, based on the concept validated in the first iteration and refined with the insights gathered from potential users. This involved developing a catalog that integrates a wide variety of tests, a test execution engine to run those tests in a harmonized way, dashboards to present the results to the different stakeholders, and a report generator. This prototype was developed and tested in the context of a simulated AIRS, which served as the primary case for grounding the tool's requirements against a regulatory process.

The demonstrator was evaluated in two parallel settings to gather feedback from the different stakeholder groups. First, it was presented at World AI Cannes Festival,

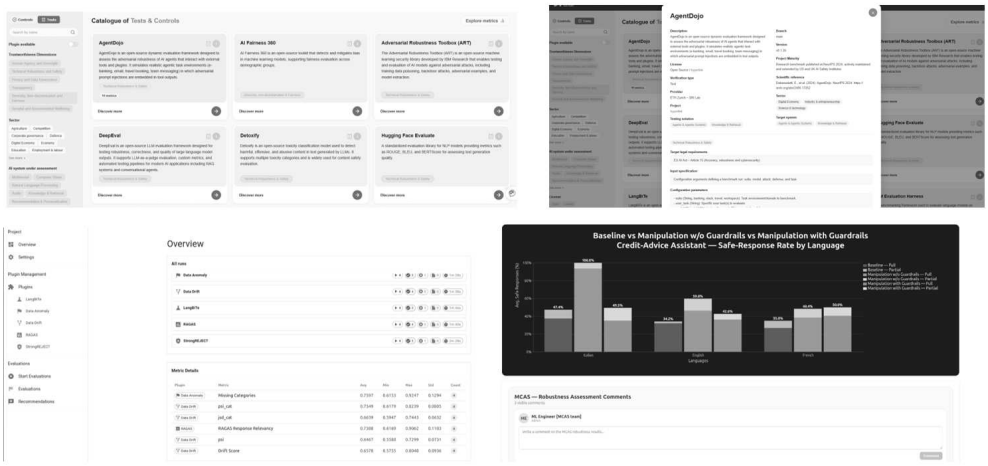

Fig. 2 Wireframes of the 4 main screens of the prototype: catalog of tests, test execution engine and collaborative dashboard

one of the leading EU events on artificial intelligence, where it was presented informally to 40 attendees, especially AI innovators and technical experts. Then, it was presented in a focused workshop with 60 participants, gathering NCAs and representatives from different EU initiatives, such as AI Factories, Digital Innovation Hubs and, Testing and Experimenting Facilities. This dual approach allowed to have a more targeted feedback, understanding the acceptance of the different stakeholders towards the tool and the remaining issues to be improved. The feedback was gathered through an observed walkthrough of the demonstrator followed by a semi-structured discussion and analyzed by the research team through debrief notes to identify recurring pain points.

The main insight gathered from this evaluation was that the tool was very focused on the technical assessment and technical users. Not all organizations that will participate in an AIRS will have to do technical testing, depending on the risk level of the AI system under assessment they might only need to complete lighter controls. Moreover, several participants highlighted that not all users will be able to qualify their system on their own, being a relevant feature that can be added to the tool, improving the overall user experience, mostly in the starting phases of the assessment.

These findings shaped the decisions carried out in the third iteration, elaborated on in Section 4.3 and helped formalize the qualification step in the on-boarding phase of the framework (Section 5) by introducing an initial system qualification step ahead of the test selection and broadening the scope of the tools to integrate controls for lower risk systems.

4.3 Moving Beyond the Technical

The last iteration presented a prototype with a holistic approach, involving non-technical users in a more meaningful way, as the core innovation of the tool is not only

its technical capacity, but the ability to provide a common ground for all stakeholders regarding AI assessment. The first prototype was improved by integrating aspects such as an initial phase for the qualification of the AI system and control checklists into the catalog. For the AI system qualification, the user will define in detail the system they want to assess, such as intended purpose, target users, and high-level information about data governance, human oversight, transparency to end users. It also includes the selection of several tags such as the type of AI technologies on which the system is based and the sector of application. The catalog, shown in Figure 3 includes control checklists from diverse sources. Similarly to the tests, each control checklist is represented by its slug, which can be expanded to retrieve information such as source, addressed regulation, or standard addressed or specific obligations.

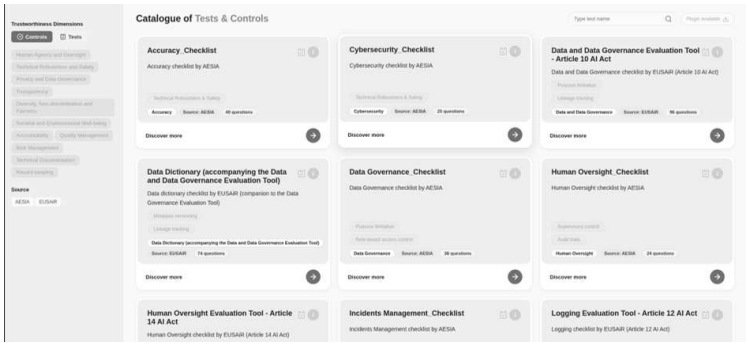

Fig. 3 Wireframe of the catalog, including controls

The improved prototype was presented at the following edition of the Nexus conference (Section 4.1), ensuring attendance of all relevant stakeholders. Along with the formal presentation at this event, the prototype was also released as an open-source tool for technical users to test. It was formally presented in the conference to 30 attendees in a focused session and the demonstration was available for informal presentation during the two days of the conference, reaching another 40 attendees for more detailed one-on-one discussions. All main stakeholders were represented among the attendees, providing balanced feedback across all groups. As in the previous iteration, the feedback was gathered through an observed walkthrough of the demonstrator followed by a semi-structured discussion and analyzed by the research team through debrief notes to identify recurring pain points.

The prototype was well received by the ecosystem, both by technical and non-technical users. Overall, stakeholders appreciated a platform that helped them navigate the wide variety of testing options and existing controls in a simplified way and the capacity to share the results with the relevant stakeholders through the collaborative dashboard and exit report. Thus, the harmonization of the different tests was also highly appreciated by the ecosystem, especially technical experts, as the experience of the users was simplified to learning how to perform the test once, instead of starting from scratch with every new test. These findings shaped the final details of the framework (Section 5), outlining the interactions among technical users and regulators

through the use of the prototype. They helped shape further all actions in the different phases of the framework, filling the gaps of the journey of the three stakeholders.

5 A Technical BNA as Backbone in AIRS Frameworks

The successful implementation of AIRSes requires a level of collaboration that challenges the operationalization of said environments, as the different stakeholders have different needs and inputs throughout the assessment process. The findings from understanding the stakeholders, Section 2.3, and designing and testing a tool to bridge the information asymmetry, Section 4, led to the definition of a framework to operationalize AIRSes. This framework represents the potential user experience and interactions of the three main stakeholders in an AIRS, using the developed prototype as the main touchpoint.

The phase structure presented in our framework is a synthesis of the work of existing regulatory, institutional, and academic sources that describe the lifecycle of an AIRS. The number of phases and naming slightly vary from one another, but the core content remains consistent, consisting in a sequence from participant selection to preparation and testing and ending with the exit and monitoring. The EU AI Act (European Commission 2024) and the implementing act for AIRSes (European Commission 2026) outline application and selection, participation, sandbox plan, written proof and exit report and post-participation. Complementary EU initiatives, such as EUSAiR (EUSAiR 2025) articulate related phases, with a slightly more granular approach: pre-participation, application and selection, preparation, participation, evaluation and exit, post-participation, reporting and monitoring. Academic sources either apply these phases, (Buscemi et al. 2026), or converge on a more simplified approach, (Novelli et al. 2026), with only the core phases: application, preparation, testing validation and exit. For the purpose of this study, we adopt this consolidated phase structure against which we situate our framework instead of proposing new phases, as our contribution lies in how the BNA designed in Section 4 can enhance stakeholder interactions and collaboration among them.

As a result, the framework, shown in Figure 4, is divided into six general phases, presented in the first row: pre-application, application and selection, on-boarding, testing, evaluation, and post-participation. The designed prototype presented in the previous section facilitates the process and exchanges among stakeholders of the four core phases: the selection, the on-boarding, the testing and the evaluation, while the framework presents the interactions of three stakeholders in the entire process of the AIRS. Each row represents the actions of a stakeholder, while the last row presents how the designed tool fits into the process. To present the framework, the technical expert and the AI innovator are separated into two different stakeholders each of them with their own actions. However, depending on the organization or specific use case, these two stakeholders could be merged into one.

The first phase of the framework is the *pre-application*. For the AI innovators this phase is especially relevant, as it is in this phase where they will decide to participate in the AIRS. They have to learn about what the AIRS can do for them, understand the benefits of participating and prepare the application. For this to happen, the

	Pre-application	Application & Selection	On-boarding	Testing	Evaluation	Post-participation
AI Innovator	Learn about AIRS, prepare application	Submit application	Sign agreement, Build Sandbox Plan	Run tests, Share with relevant stakeholders, Iterate on testing	Share report with stakeholders	Provide feedback
Regulator	Disseminate information about AIRS, Support application preparation	Define selection criteria, Review applications, Select	Provide agreement, Support in the definition of Sandbox plan	Guidance and supervision	Evaluate results, Issue proof of participation	Gather & share feedback / knowledge
Technical expert	-	Risk assessment, Technical qualification	Set up Sandbox environment	Run tests, Provide support	Final technical evaluation	Share feedback
Designed tool	-	System qualification form System card	Test / Control selection	Test execution engine Collaborative dashboard	Auto-generated report	-

Fig. 4 Framework representing the structured phases of an AIRS engagement, decomposed into stakeholder experiences and tasks. The bottom row contains the role the prototype plays in support the user experience in each phase

NCAs have to disseminate information and benefits of participating in an AIRS to the ecosystem, as well as provide support to the AI innovators in the application preparation. In this particular phase, the technical experts do not play a role yet, however, they can help orient potential users to the AIRS and help these users understand the value and benefits of applying. Requirements for a successful pre-application phase include a good information campaign and a harmonized application procedure across AIRSes. In this phase the designed prototype does not play a role yet.

The second phase is the *application and selection* of participants. In this phase the AI innovators submit the application and receive the notification of the application decision from the NCA. In order to do this, the NCAs need to first establish the eligibility and selection criteria, analyze the submitted project against those criteria and select the applicants. During this process, the technical expert can provide support to the NCAs by providing technical support to perform risk assessment and classification of the submitted project, provide technical advice to identify requirements and feasibility and provide a technical qualification of the AI system under evaluation. While the NCA needs to provide a platform for the AI innovators to apply, the technical qualification of the AI system can be performed through the design tool presented in Section 4. This will allow AI innovators and NCAs to download a structured system card, e.g. the proposed AI cards (Golpayegani et al. 2024), with the details of their AI system, such as intended purpose, target users and high-level information about data governance, human oversight, transparency to end users, to name a few.

Once the applicants are selected, they move forward to the third phase, *onboarding*. In this phase the AI innovators will have to sign an agreement for accepting the general terms and conditions of the AIRS and complete a sandbox plan in alignment with the NCA. Thus, if they have performed any testing prior to the participation in the AIRS they will have to provide the corresponding test results. Regarding the NCA, in this phase they will have to provide the AI innovators with the agreement and provide support in the creation of the sandbox plan. As for the technical expert,

with the help of the tool (Section 4) and based on the qualification of the AI system made in the previous phase, this phase consists in setting up the technical sandbox environment. This involves the selection of relevant tests, datasets, and controls. The interaction of all three stakeholders in this phase is key to ensure alignment throughout the following phases.

The fourth phase of the framework consists in performing the tests in the tool following the established plan. The testing can be done by both the AI innovator or technical expert, while the NCA can provide guidance and support. Through the use of the prototype, the results of the testing will be shown in a collaborative dashboard, where all the stakeholders can see the relevant information and make comments. The open architecture is designed to be modular, which allows customized dashboarding tools such as PowerBI or Superset increasing the utility of the tool as a BNA. This will allow to make all the necessary iterations in the testing process while identifying the strengths, gaps and risks of the AI system. All the changes made to and iterations of the process will be recorded thanks to the tamper-proof logs, giving a detailed outline of all the performed tests.

Upon conclusion of the testing phase, it is followed by one final evaluation and interpretation of the test results. The tool automatically generates a report, facilitating the reporting phase for all stakeholders. This report contains an overview of the system and the assessment results, with a more detailed analysis of the results of each assessment tool used. This includes, e.g., highlighting failed test cases for further analysis by technical experts. The AI innovator will be in charge of sharing this report with the relevant stakeholders. As for the NCA, they will have to evaluate said results and provide the AI innovator with a written proof of participation in the AIRS. Also, relevant and non-confidential findings will be published by the NCA. In this phase the technical expert can also provide a final technical evaluation of the strengths, gaps, and residual risks of the AI system under assessment.

The final phase of the framework is the post-participation, with the participation in the AIRS having formally concluded but communication among the stakeholders still ongoing to sustain their interest and involvement. A key action in this phase is to share the generated knowledge and gather feedback. The feedback activity refers to evaluating the process both internally, and with the participants. It is noteworthy that the majority of the processes that engage a variety of stakeholders conclude in the previous phase, without a dedicated phase addressing what occurs once the process is completed, undermining its impact. Hence, in this phase it is important for the NCAs to gather the feedback of other stakeholders and share the knowledge gathered trough the process with them.

This framework aims at bridging the gap between the stakeholders when assessing an AI system by providing a common tool adapted to each of their needs. Grounded in the insights presented in Section 2.3, this tool provides a common platform for both technical and non-technical users to interact regarding the assessment of an AI system. It reduces the complexity of setting up testing environments and AIRSes.

6 Discussion

This paper explores how AI innovators, technical experts, and regulators experience the collaboration in AIRSes, what design knowledge emerges from building a prototype of a technican BNA to mediate that collaboration, and how that knowledge can inform the operationalization of AIRSes. Our findings suggest that the central issue for a well-functioning AIRS is not the deficit of information, but the sense-making of it all, as stakeholders hold substantial and legitimate expertise, but lack a shared understanding. This is relevant, as it re-frames the design challenge of an AIRS. Simply supplying more documentation or guidance is not enough for its proper operationalization. There is a need for artifacts that will close the information asymmetry among the stakeholders, offering a means that can be used to negotiate meaning. This extends the view of sandboxes articulated by Guio Español and Koenig (2025), understanding AIRSes as fundamentally learning environments, with their central design problem being on communication and collaboration.

The three design iterations presented in Section 4 address that challenge, providing a prototype of a technical BNAs for multi-stakeholder cooperation. Each aspect of the tool, such as the qualification of the system, the harmonized tests and controls, the dashboard, or the report, all perform a distinct sense-making function, providing a common ground for the different stakeholders involved. Putting these artifacts together, across the AIRS life-cycle, they function as the backbone that the framework presented in Section 5 formalizes. This expands on the BNA concept for a AI regulatory context specifically. As an artifact bridging multiple communities of practice, it works as a collection of touchpoints, each working at a specific phase of the process, whose cumulative effect is the reduction of information asymmetry across the whole AIRS life-cycle.

This finding lies within a long standing conversation in the CSCW community on designing for collaboration across diverse communities. CSCW has previously used boundary objects and shared vocabularies as a way of supporting multi-stakeholder collaboration, such as using the notion of BNAs as a way to coordinate data practices across stakeholders (Jeng et al. 2025; Wu et al. 2026). This underscores a broader turn in the field towards using BNAs to hold high-stakes multi-stakeholder discussions under uncertainty or in the absence of established consensus. Our work contributes to this turn by expanding it into a regulatory context, where participants are multidisciplinary and structurally different. AI innovators, regulators, and technical experts hold different levels of authority, accountability and literacy over the same system, yet they all need to interpret the same results. Designing for common ground goes beyond translation between vocabularies, as it needs to do so without undermining the legitimate expertise that each stakeholder brings to the exchange. In this sense, this work extends the multi-stakeholder collaborative practices toward the emerging field of AI regulation.

Another relevant finding is the direction of the design process itself. The feedback gathered on the first iteration showed that the tool was well understood by technical experts, while not as easily understood by other less technical users. This issue reflects, on a small scale, what the broader literature for responsible AI suggests, as organizations building AI tend to default to efficiency- and technology-first practices,

often neglecting the user needs (Rakova et al. 2021; Grasso and Colombino 2026). Within an AIRS this can create tension among stakeholders, as the efficiency of a technical tool might reduce the burden for the AI innovators and technical experts but be inefficient to share with other stakeholders like the NCA in an AIRS engagement or even intra-organizational among departments of one and the same company. By applying the RtD methodology, participants were able to engage with the designed system in every iteration, providing insights that only became visible while observing stakeholders' behavior when interacting with the prototype. This trajectory reflects the importance of humans in the design process as designing an AIRS relies on the shared understanding across multiple sets of expertise, highlighting the relevance of stakeholder involvement in the process.

These findings converge in the discussed framework. Situating the tool as the backbone of the framework and giving each stakeholder a specific set of actions in each step, the framework provides a sense-making mechanism for AIRSes bridging the information asymmetry. While the EU AI Act (European Commission 2024) and the implementing act for AIRSes (European Commission 2026) specify that a NCA must supervise and an innovator must assess its AI system, this framework specifies *how* that supervision, planning and testing can be mediated so that the resulting exchange is meaningful to all stakeholders. In this sense, the framework's contribution offers the integration of a shared artifact that helps in the operationalization of AIRSes, providing a common protocol for all stakeholders (Schmidt and Simone 1996).

The implications of the proposed solution extend beyond the use case of this paper. As outlined in Section 2.1 the EU AI Act mandates that every member state establishes at least one AIRS, yet leaves considerable discretion to NCAs regarding its operationalization. This discretion risks fragmented or incompatible approaches that could undermine a consistent implementation of the act (Ahern 2025; Genicot and Moraes 2025). This risk, anticipated by the Commission itself, is reflected in the draft implementing act (European Commission 2025) and its final published version (European Commission 2026). However, this harmonization attempt operates at the level of the legal procedure, without focusing on the day-to-day multi-stakeholder experience of a AIRS. The framework operationalized by the prototype of the technical BNA proposed in this paper addresses that gap. It formalizes a common structure of touchpoints and stakeholder exchanges that NCAs across different member states can adopt as a practical layer under the guidelines of the EU AI Act. This offers a replicable structure for setting up their own AIRS without the need to design the collaboration tools from scratch. For AI innovators, on the other hand, it offers a more consistent point of reference across sandboxes, reducing the cost of learning new collaboration models each time they engage with a different NCA.

Beyond the EU, the relevance of the framework lies in providing a generic structure of multi-stakeholder collaboration in AIRSes. AIRSes are being established well beyond the EU context, all varying in institutional maturity (Guio Español and Koenig 2025). Since the framework's touchpoints, such as harmonized tests or reporting, are organized around the sense-making needs of innovators, technical experts, and regulators rather than around any single legal text, they are, in principle, transferable to other legislative systems that face a similar need to structure trust and communication

between stakeholders under legal uncertainty. That said, this transferability remains a claim as differences in institutional capacity and regulatory culture outside the EU may require some modifications to the framework.

Finally, our findings are grounded in a single, evolving prototype developed and evaluated primarily through semi-structured feedback from stakeholders rather than sustained longitudinal use within a live, operating AIRS. Therefore, the framework should be read as a design proposal validated through iterative exposure to stakeholders. Additionally, confidentiality agreements with participating organizations mean that raw quotes, case-specific details, and organization-level findings cannot be shared. All reported findings are consequently presented at the level of recurring patterns across sessions rather than as directly attributable statements.

7 Conclusion

This paper examined how AIRSes can be designed to support genuine stakeholder collaboration, with special focus on AI innovators, technical experts, and regulators, moving beyond merely procedural requirements. Through a RtD process grounded in three iterations of a prototype, we concluded that the core challenge these stakeholders face is not a lack of information but a lack of shared sense-making across differing vocabularies and expertise. Building on this finding we created a prototype for a technical BNA that bridges the asymmetry across the AIRS life-cycle providing an artifact that enhances the cooperation among stakeholders. Therefore, as AIRSes multiply across EU member states, we hope this framework offers a path towards an interoperable and human-centered practice. Future work should build on these findings through user studies within a real, operating AIRS, observing how AI innovators, technical experts, and regulators engage with the tool under real-world conditions. A second direction is to pilot the tool across AIRSes in different EU member states, testing directly the interoperability and replicability, examining how its structure holds up against the procedural variation between NCAs.

Declarations

Author Contributions. Conceptualization, I.L.-O., and T.D.; methodology, I.L.-O.; investigation, I.L.-O., T.D, A.B. and D.P.; writing—original draft preparation, I.L.-O. and T.D.; writing—review and editing, I.L.-O., T.D, A.B., D.P. and G.C.; supervision, G.C. All authors have read and agreed to the published version of the manuscript.

Funding. We thank the Luxembourg AI Factory for supporting this work and all the stakeholders that participated in the design process. This work is part of the Luxembourg AI Factory project that has received funding from the European Commission (EC) under the European Union Horizon programme (Grant agreement No. 101234366). Funded by the European Union. Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the EU or EC. Neither the EU nor the granting authority can be held responsible for them.

Ethical Approval. All participants were informed about the research project and their involvement in the design process of the prototype. No identifying information is provided.

Data Availability. Not applicable.

Conflict of interest/Competing interests. The authors declare to have no conflicts of interest.

References

Ahern, D. 2025. Operationalising AI regulatory sandboxes under the EU AI Act: The triple challenge of capacity, coordination and attractiveness to providers. *Cambridge Forum on AI: Law and Governance* 1: e35. https://doi.org/10.1017/CFL.2025.10023 .

Beddoes, K., M.J. Borrego, and B.K. Jesiek 2011, June. Using Boundary Negotiating Artifacts to Investigate Interdisciplinary and Multidisciplinary Teams. In *2011 ASEE Annual Conference & Exposition*, pp. 22.1610.1–22.1610.14.

Buscemi, A., G. Castignani, D. Pagani, M. Cordy, and J. Cabot 2026, 10. The AI Assessment Sandbox Configurator: A Framework to Support Technical Assessment in AI Regulatory Sandboxes. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*.

Buscemi, A., T. Simonetto, D. Pagani, G. Castignani, M. Cordy, and J. Cabot. 2025. The sandbox configurator: a framework to support technical assessment in AI regulatory sandboxes. arXiv: 2509.25256 [cs.CY].

Chen, G.K.H. and A. Taeihagh. 2026. Designing regulatory sandboxes: a comprehensive framework for aligning functionalities and objectives. *Policy Design and Practice 9*(1): 1–15. https://doi.org/10.1080/25741292.2025.2570954/ASSET/4970936B-7DA2-40E7-BE94-6BB5873BA7D3/ASSETS/GRAPHIC/RPDP{_}A{_}2570954{_}F0003{_}B.JPG .

(Chuqiao) Wan, E., C. Yin, A. Ito, Z. Gao, J. Jia, Y. Taoka, S. Saito, M. Sadek, and C. Mougenot 2026, April. KNIT: Computational Boundary Objects for Real-Time Convergence in Interdisciplinary Teams. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, New York, NY, USA, pp. 1–22. Association for Computing Machinery.

Cortesão, J. and S. Lenzholzer. 2022, 11. Research through design in urban and landscape design practice. *Journal of Urban Design 27*(6): 617–633. https://doi.org/10.1080/13574809.2022.2062313 .

Deckenbrunnen, T., A. Buscemi, M. Almada, A. Capozucca, and G. Castignani. 2026, October. Bathtubs, Boundaries, and Sandboxes: AI Regulatory Learning Under

Legal Uncertainty. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society 9*(1): 775–787. https://doi.org/10.1609/aies.v9i1.43182 .

Español, A.G. and P.D. Koenig. 2025. Regulatory sandboxes for AI in the majority world: A learning-centric approach to legal adaptation. *Cambridge Forum on AI: Law and Governance* 1: 1–20. https://doi.org/10.1017/cfl.2025.10015 .

European Commission. 2024, July. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts.

European Commission 2025, December. Draft Commission Implementing Regulation laying down rules for the application of Regulation (EU) 2024/1689 of the European Parliament and of the Council as regards the establishment, development, implementation, operation and supervision of AI regulatory sandboxes. Draft Implementing Act, European Commission.

European Commission. 2026, November. Commission Implementing Regulation (EU) 2026/2487 of 10 November 2026 laying down rules for the application of Regulation (EU) 2024/1689 of the European Parliament and of the Council as regards the establishment, development, implementation, operation and supervision of AI regulatory sandboxes. Legislative Body: COM, CNECT.

European Parliament. 2026, 3. AI regulatory sandboxes: State of play and implementation challenges — Think Tank — European Parliament.

EUSAiR. 2025, 8. AI Ecosystem and AI Regulatory Sandboxes — EUSAiR.

Fahy, L.A. and C.A. Lauren Fahy. 2022, 4. Fostering regulator–innovator collaboration at the frontline: A case study of the UK's regulatory sandbox for fintech. *Law & Policy 44*(2): 162–184. https://doi.org/10.1111/LAPO.12184 .

Genicot, N. and T.G. Moraes. 2025, January. Exploring the boundaries of AI regulatory sandboxes under the AI Act: Flexibility and real-world testing. *Cambridge Forum on AI: Law and Governance* 1: e36. https://doi.org/10.1017/cfl.2025.10013 .

Giacomin, J. 2014, 8. What Is Human Centred Design? *The Design Journal 17*(4): 606–623. https://doi.org/10.2752/175630614X14056185480186 .

Godin, D. and M. Zahedi 2014, 6. Aspects of Research through Design: A Literature Review. In *Design's Big Debates - DRS International Conference 2014*, Umea.

Golpayegani, D., I. Hupont, C. Panigutti, H.J. Pandit, S. Schade, D. O'Sullivan, and D. Lewis 2024, September. AI Cards: Towards an Applied Framework for Machine-Readable AI and Risk Documentation Inspired by the EU AI Act. In *Privacy Technologies and Policy: 12th Annual Privacy Forum, APF 2024, Karlstad, Sweden, September 4–5, 2024, Proceedings*, Berlin, Heidelberg, pp. 48–72. Springer-Verlag.

Grasso, M.A. and T. Colombino 2026. "Just a Prototype?": AI-assisted artifacts and the practical work of assurance in research labs. In *Proceedings of the 24th EUSSET conference on computer-supported cooperative work (ECSCW) – Workshop Position Paper*, Reports of the european society for socially embedded technologies.

Guio Español, A. and P.D. Koenig. 2025. Regulatory sandboxes for AI in the majority world: A learning-centric approach to legal adaptation. *Cambridge Forum on AI: Law and Governance* 1: e42. https://doi.org/10.1017/cfl.2025.10015 .

Gunasekaran, T.S., S. Lim, K. Gupta, H. Bai, Y.S. Pai, and M. Billinghurst 2026, April. Cognitive Bridge: AI-Generated Boundary Objects for Cross-Functional Collaboration. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, New York, NY, USA, pp. 1–35. Association for Computing Machinery.

Guridi, J.A., A.H.C. Hwang, D. Santo, M. Goula, C. Cheyre, L. Humphreys, and M. Rangel. 2025, May. From Fake Perfects to Conversational Imperfects: Exploring Image-Generative AI as a Boundary Object for Participatory Design of Public Spaces. *Proceedings of the ACM on Human-Computer Interaction 9*(2): CSCW014:1–CSCW014:33. https://doi.org/10.1145/3710912 .

Jeng, W., Y.J. Yang, H.C. Chen, Y.J. Yang, and Y.R. Shih. 2025, 9. Boundary Negotiating Artifacts to Envision the Desired Research Data Infrastructure Toward Reproducibility in Data-Intensive Science. *Computer Supported Cooperative Work: CSCW: An International Journal 34*(3): 871–910. https://doi.org/10.1007/S10606-025-09523-3/TABLES/6 .

Lee, C.P. 2007, June. Boundary Negotiating Artifacts: Unbinding the Routine of Boundary Objects and Embracing Chaos in Collaborative Work. *Computer Supported Cooperative Work (CSCW) 16*(3): 307–339. https://doi.org/10.1007/s10606-007-9044-5 .

Lee, J.S., T. C. Veinot, and E. Yakel 2026, April. Embracing Chaos Again: Tracing Boundary Negotiating Artifacts in Knowledge Workers? Data Practices for Collaborative Public Health Crisis Response. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, ACM Conferences, pp. 1–25.

Leigh Star, S. 2010, September. This is Not a Boundary Object: Reflections on the Origin of a Concept. *Science, Technology, & Human Values 35*(5): 601–617. https://doi.org/10.1177/0162243910377624 .

Novelli, C., P. Hacker, S. McDougall, J. Morley, A. Rotolo, and L. Floridi. 2026, September. Getting Regulatory Sandboxes Right: Design and Governance Under the AI Act. *European Journal of Risk Regulation 17*(3): 926–952. https://doi.org/10.1017/err.2026.10096 .

Rakova, B., J. Yang, H. Cramer, and R. Chowdhury. 2021, 4. Where Responsible AI meets Reality. *Proceedings of the ACM on Human-Computer Interaction 5*(CSCW1): 23. https://doi.org/10.1145/3449081 .

Schiff, D.S., S. Kelley, and J. Camacho Ibáñez. 2024, December. The emergence of artificial intelligence ethics auditing. *Big Data & Society 11*(4): 20539517241299732. https://doi.org/10.1177/20539517241299732 .

Schmidt, K. and C. Simone. 1996, June. Coordination mechanisms: Towards a conceptual foundation of CSCW systems design. *Computer Supported Cooperative Work (CSCW) 5*(2): 155–200. https://doi.org/10.1007/BF00133655 .

Star, S.L. and J.R. Griesemer. 1989. Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in berkeley's museum of vertebrate zoology, 1907-39. *Social Studies of Science 19*(3): 387–420. https://doi.org/10.1177/030631289019003001 .

Tadimalla, S.Y. and M.L. Maher 2026, April. AI Identity as a Boundary Object: Unpacking the Intersection of AI and Identity. In *Proceedings of the Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI EA '26, New York, NY, USA, pp. 1–9. Association for Computing Machinery.

Velleu, J., D. Brei, R. Gonzalez, and J. Luntz. 2023. A BOUNDARY OBJECT FOR MAPPING, COMPARING, AND INTEGRATING PRODUCT DESIGN METHODS. *Proceedings of the Design Society* 3: 29–38. https://doi.org/10.1017/PDS.2023.4 .

Wenger-Trayner, B., E. Wenger-Trayner, J. Cameron, S. Eryigit-Madzwamuse, and A. Hart. 2019, 7. Boundaries and Boundary Objects: An Evaluation Framework for Mixed Methods Research. *Journal of Mixed Methods Research 13*(3): 321–338. https://doi.org/10.1177/1558689817732225/ASSET/E7AE5E42-AB04-44F8-8076-E970AF99E8BC/ASSETS/IMAGES/LARGE/10.1177{_}1558689817732225-FIG6.JPG .

Wohlrab, R., P. Pelliccione, E. Knauss, and M. Larsson. 2019, 5. Boundary objects and their use in agile systems engineering. *Journal of Software: Evolution and Process 31*(5): e2166. https://doi.org/10.1002/SMR.2166 .

Wu, M., Y. Zhao, S. Han, M.X. Liu, and H. Shen. 2026, September. AI LEGO: Scaffolding Cross-Functional Collaboration in Industrial Responsible AI Practices during Early Design Stages. *Proceedings of the ACM on Human-Computer Interaction 10*(6): CSCW049:1–CSCW049:33. https://doi.org/10.1145/3816897 .

Zimmerman, J., J. Forlizzi, and S. Evenson 2007, April. Research through design as a method for interaction design research in HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, New York, NY, USA, pp. 493–502. Association for Computing Machinery.