

Sandboxing Artificial Intelligence

Balancing Innovation, Regulation, and Stakeholder Needs

About FARI

FARI is an independent, not-for profit Artificial Intelligence initiative led by the Vrije Universiteit Brussel (VUB) and the Université libre de Bruxelles (ULB). The aim of the research institute is to enable, promote and perform excellent cross-disciplinary research on Artificial Intelligence in Brussels, inspired by the humanistic values of freedom, equality and solidarity that lay at the foundations of both the Vrije Universiteit Brussel (VUB) and the Université libre de Bruxelles (ULB), internationally acclaimed and with a local impact.

For more information visit <https://www.fari.brussels/>

Acknowledgements

This paper was authored by Stig A. Due, Hira Shah, Thiago Moraes, Dr. Nathan Genicot, and Dr. Martin Canter. We would especially like to thank all the participants in the workshop for their valuable insights and discussions, which greatly informed the development of this paper. Particular thanks go to those who contributed their expertise and feedback throughout the process.

Table Of Content

Introduction	5
The Role of Testing in AI Sandboxes	6
Legal and Regulatory Considerations	8
Stakeholder Engagement	12
AI Innovators.....	12
Regulators	14
Civil society	16
Societal and Scientific Experts	17
Knowledge Generation and Harmonization of Sandboxes.....	19
Conclusion	21
Bibliography	23

Executive summary

As artificial intelligence (AI) systems become increasingly integrated into society, the need for effective regulation that balances innovation with public interest is more urgent than ever. The European Union's AI Act introduces a novel approach to this challenge through the establishment of AI regulatory sandboxes, controlled environments where AI systems can be developed and tested under regulatory supervision.

This whitepaper draws on insights from workshops hosted by CAIRNE and FARI in 2024, bringing together diverse stakeholders across the AI ecosystem. It explores the practical design and implementation challenges of AI regulatory sandboxes, with a particular focus on stakeholder engagement, legal frameworks, technical testing, and cross-border harmonization.

Key themes include the importance of taking a use-case-driven approach to sandbox operation, the need for clarity around legal mandates and data protection, and the essential role of trust-building to secure active participation from innovators, regulators, civil society, and experts. The report also emphasizes the sandbox's role in generating regulatory knowledge and how that knowledge can be leveraged to support EU-wide harmonization without sacrificing local adaptability.

Rather than providing final answers, this whitepaper aims to stimulate a productive dialogue among stakeholders and contribute to the evolving practice of responsible AI governance in Europe.

Introduction

2024 marks an important year for the field of AI governance. The EU passed the AI Act, which is the first comprehensive effort to regulate AI worldwide. The AI Act not only seeks to mitigate risks related to the rapid development of this technology, but it also has provisions aimed at boosting AI innovation in the EU. The main tool to be employed to achieve this balance is AI regulatory sandboxes, as laid out in Article 57 of [the AI Act](#). The provision requires each member state to establish at least one national AI regulatory sandbox within August 2026.

Regulatory sandboxes are not a new regulatory tool and have already been applied to domains like data privacy like the [Norwegian privacy sandbox](#). However, this tool is new within AI regulations in the EU. This has led to several questions arising about how to best establish the sandbox and comply with the provisions of the AI Act. As a result, there has been a high demand for cooperation to plan how regulators should proceed. This involves conferences, workshops, reports, etc.

One of the organizations that is working on promoting collaboration in answering the questions is the [Confederation of Laboratories for Artificial Intelligence Research in Europe](#) (CAIRNE). In January 2024, they hosted a workshop in Zurich, where various stakeholders from across Europe. This proved productive and was followed up by a workshop in collaboration with [AI For the Common Good Institute Brussels](#) (FARI) on the 19th of November, 2024. This workshop gathered 30 participants amongst various stakeholder groups from across the EU, with the goal of providing a forum for organizations engaged in establishing AI-related sandboxes across Europe. FARI facilitated the discussions and generated actionable recommendations to enhance innovation environments that ensure responsible and sustainable development of AI.

The workshop took place over a full day. In the morning, participants shared pain points they had identified in their sandboxing experience. A presentation and Q&A with representatives of the European Commission AI Office followed. In the afternoon, four breakout groups were made, and each group worked on finding a strategy to answer a series of questions related to different topics on AI Regulatory Sandboxes. In this white paper, we will highlight the key issues that were identified by the participants and some possible solutions. The goal of this report is not to give definitive answers, but rather to stimulate a discussion that will benefit all stakeholders regarding the implementation of AI regulatory sandboxes.

The views and opinions expressed in this paper are those of the participants and do not necessarily reflect the official position of FARI or CAIRNE. The content is based on discussions held during a workshop and is intended to capture a range of perspectives shared in that context.

The Role of Testing in AI Sandboxes

Two key questions that underpinned the workshop were how sandboxes should be structured, and what the sandboxes should do. It was pointed out that there is a distinction between a sandbox that focuses more on the regulatory aspect (giving legal guidance to the participants in the sandbox) and a sandbox that provides technical testing (provides a testing environment, provides datasets, etc.) It is also possible that a sandbox can facilitate both legal guidance and technical testing. We can find several examples of sandboxes that have been operational since before the AI Act was enacted that focus only on regulatory guidance. This means that we have an idea of what this entails and looks like in practice. However, there are not as many existing examples of a sandbox providing technical testing. This means that the discussion on testing in a sandbox focuses largely on how a sandbox can facilitate technical testing.

The sandbox can facilitate technical testing by providing resources, data, and testing infrastructure. Technical testing can uncover potential risks that legal assessment alone cannot identify. What this looks like in practice, however, is not clear. It was pointed out that testing AI systems is a complex endeavor because systems can be targeted toward a variety of sectors, as diverse as health, law, or transportation, to name a few. All these systems would have different needs, for example, access to data, computational resources, or something else. The stakeholders present mentioned the difficulty of designing a sandbox that can cover all participants' needs, as well as provide the necessary resources.

Another issue that arises when we consider technical testing in the sandbox is who should facilitate the testing. As stated in the AI Act, the National Competent Authority (NCA) is responsible for operating the sandbox, but should they facilitate this technical testing in addition to testing compliance with the regulations? Providing technical testing can be resource-intensive, and the regulatory sandbox could act as a connector hub to avoid this problem. The sandbox could assess the needs of the participants and put them in contact with relevant partners, like Test and Experimentation Facilities (TEFs), data factories, or other testing environments.

It became clear through this discussion that a one-size-fits-all solution will not suffice. Different sandbox participants may have different needs and expectations from the sandbox depending on their specific case. The solution that could benefit the participants the most is a use case-driven approach. This would mean that instead of having a standardized plan in the sandbox, each case would be assessed, and their needs would be identified. The sandbox process would then be tailored to their needs, to maximize the benefit of the sandbox for all stakeholders.

If the sandbox provides technical testing by acting as a hub, that will free up their resources to better provide regulatory guidance to the participants. This can take the form of the legal team of the sandbox and the participant discussing how to interpret how the law applies to their case, where it isn't clear if their system is compliant. Previous experience from [fintech sandboxes](#) have shown that the legal discussion in a sandbox provides SMEs and start-ups legal certainty, which in turn has resulted in the product and service being able to increase their go-to-market speed.

As we've seen, testing—particularly technical testing—is a core component of the regulatory sandbox. Yet for testing efforts to have a real impact, they must be supported by a solid legal and institutional foundation. A sandbox must be embedded within the national regulatory framework to ensure its legitimacy, scalability, and long-term relevance. The following section explores key legal and regulatory considerations that remain open and require further clarification.

Key takeaways on the role of testing in AI regulatory sandboxes	
(Potential) Challenges	Actions
<ul style="list-style-type: none"> • Uncertain what the regulatory sandbox is supposed to test. • Several sandboxes operational today only provide regulatory guidance. • A one-size-fits-all approach fails to meet the needs of all the sandbox participants. • The operators of the regulatory sandbox may lack the resources to provide comprehensive technical testing. 	<ul style="list-style-type: none"> • Policymakers must clarify whether sandboxes focus on AI compliance assessment, technical aspects, or both. • Regulators need to acknowledge the necessity of technical testing as it can uncover weaknesses and risks in AI systems. • Use a case-driven approach: Testing requirements vary across cases, some projects may require rigorous testing, while others may only need regulatory guidance. • The regulatory sandbox can function as a hub, connecting companies to relevant collaborators like TEFs.

Legal and Regulatory Considerations

The AI Act leaves considerable freedom to Member States regarding the institutional setup to implement the regulation. Article 57 of the AI Act states that each Member State must designate at least one National Competent Authority (NCA), which can be a market surveillance authority, a notifying authority, or both. Member States also have the option to create a new authority to serve as the NCA. The role of the NCA is primarily to operate the AI regulatory sandbox and act as the main contact point between the Member State and the AI Office. This flexibility offers advantages but also raises practical concerns. One key issue is how AI sandboxes should be integrated into existing regulatory policies and administrative procedures. Another is whether it is more effective to establish a new authority or to designate an existing one as the NCA.

The AI Act also mandates that national sandboxes be horizontal in nature, meaning they are not sector specific. However, there is a strong case for also developing sector-specific sandboxes, especially for high-risk AI systems. These could be tailored to areas like healthcare or manufacturing and operate under existing sectoral regulators such as health agencies or financial supervisors. Doing so would require coordination between sector-specific and national authorities. Sector-specific sandboxes may also facilitate cross-border collaboration, offering more targeted testing environments for high-risk use cases. Beyond national efforts, sandboxes should aim to integrate with broader EU AI infrastructure such as TEFs, AI factories, living labs, and EDIHs to ensure coherence and scalability.

On the question of which authority should operate the sandbox, we can look at examples across Member States. Denmark has designated its Data Protection Authority (DPA) as the NCA. This makes sense, as DPAs must be involved under the AI Act due to the frequent processing of personal data. However, DPAs are often perceived as less innovation-friendly, with a strong focus on privacy that may not align with the more flexible, experimental aims of a sandbox. If innovation is a priority, another type of authority might be more suitable. Italy, for example, has assigned this role to its cybersecurity authority. This body may have stronger subject matter expertise in areas relevant to AI risk and may be more aligned with innovation goals. That said, a potential challenge here is funding — innovation-friendly authorities may lack the resources to effectively run a sandbox. Spain offers an alternative model: it created a new authority by royal decree. This allowed for a rapid setup, but such a mechanism may not be available in other Member States, and creating a new authority from scratch is often costly and time-consuming.

Regardless of which authority is chosen, it's important to clarify what powers the NCA will hold. These include access to key components of the AI system under development — such as data, source code, and models — as well as administrative exemptions that protect participants from certain penalties while still holding them accountable under civil and criminal law.

Another layer of complexity involves data protection. The AI Act requires that DPAs be “associated and involved with” data-related aspects of sandboxes, implying a role beyond consultation. However, the scope of this involvement is still unclear. Sandboxes also require a controlled testing environment, though it remains to be defined who is responsible for this — the sandbox provider or the participating developers. Additionally, it is not yet clear whether exit reports, which are confidential, will be shared with DPAs.

The issue of further processing of personal data was also discussed. Article 59 of the AI Act allows further processing only in limited cases where the system serves a substantial public interest, such as in health, safety, or the environment. Even then, this processing must take place in a functionally separate, controlled environment, and only when anonymous or synthetic data cannot be used. These are expected to be rare cases, and there was broad agreement on the need for clear risk mitigation strategies. Data protection, therefore, must be integrated into the design and operation of the sandbox from the outset.

Finally, while the AI Act provides the foundation for national sandboxes, it does not itself create a sufficient legal basis for implementation at regional or local levels. Member States will need to pass specific national legislation to operationalize sandboxes ahead of the Act's application in 2026. Given the frequent need to process personal data, these laws must align with existing frameworks like the GDPR. To avoid fragmentation and ensure coherence, there is a strong argument for the European Commission to issue an implementing act or guidance that defines common parameters and provides operational clarity.

While the legal and regulatory framework sets the foundation for how AI regulatory sandboxes are established and operated, their success ultimately depends on how effectively they engage stakeholders. The next section examines how stakeholder participation can be structured to ensure the success of the sandbox.

Key takeaways from legal and regulatory considerations	
(Potential) Challenges	Actions
<ul style="list-style-type: none"> • Member states have wide discretion under the AI Act on how to set up sandboxes, which can lead to inconsistent implementation. • Ambiguity over whether to create new institutions or assign responsibilities to existing ones. • Lack of clarity on the powers and responsibilities of NCAs. • Ambiguity around the role of DPAs in sandbox operation. • Legal uncertainty on further processing of personal data within sandboxes. • National horizontal sandboxes might not adequately address the needs of high-risk or sector-specific AI systems. 	<ul style="list-style-type: none"> • Encourage the AI office to issue an implementing act or common guidelines to promote harmonized sandbox practices across the union. • Weigh trade-offs between establishing new AI authorities or expanding mandates of existing institutions based on national administrative capacity. • Define minimum competencies for NCAs, including access to data, systems, and scope of exemptions. • Clarify the meaning of “associated and involved” under the AI Act through EU-level guidance. • Develop clear criteria and safeguards for further processing aligned with GDPR Art. 6(4). • Establish sector-specific sandboxes (e.g., in healthcare or manufacturing) alongside the national sandbox, ideally operated by domain-relevant authorities.

<ul style="list-style-type: none"> • Legal uncertainty at the regional/local level as the AI Act only provides a legal basis for the national sandbox. 	<ul style="list-style-type: none"> • Member states should adopt national laws that enable sandbox operation early and clarify legal bases.
---	---

Stakeholder Engagement

When we think about the stakeholders involved in a regulatory sandbox, we might have many different ideas about who can be involved. The most basic conception of the stakeholders is the sandbox operator and the AI innovators. However, if we start mapping out all the potential stakeholders, the list becomes quite extensive. Broadly, we are talking about four types of stakeholders in this context, namely: AI innovators, regulators, civil society, and experts. All these stakeholders bring their own perspectives and strengths, however, that also includes their unique challenges and concerns.

There are four questions at the core of understanding stakeholders' engagement with the sandbox. The first question is about the stakeholders' incentives to invest time and resources into the sandbox. On the flip side of this, the second question we can ask is what the stakeholder contributes to the sandbox by participating. The third question that we need to address is the concerns of the different stakeholders for the sandbox. The final question we need to look at is about how we engage with the different stakeholders to maximize the output of the sandbox.

The first stakeholder we will explore are AI innovators. As the driving force behind the development of AI, they occupy a unique and influential position. These actors possess the technical expertise underpinning AI systems and are at the forefront of innovation. The AI regulatory sandbox is designed with these innovators in mind, offering a space to test their technologies for weaknesses, risks, and legal compliance. In the following section, we examine their incentives for participating in the sandbox, their concerns, the value they bring to the process, and strategies for effective engagement.

AI Innovators

The stakeholder we will explore is the AI innovators. In this category, we consider AI providers and AI deployers, both from the public and private sector, including start-ups, SMEs, public administration, and knowledge institutions. It must be emphasized that the first target for AI

regulatory sandboxes under the AI Act are SMEs, which play a major role in the EU on innovation. To ensure their participation in the sandbox process, AI innovators, especially SMEs, require some form of incentive to join.

The first incentive is legal certainty. Sandbox participation supports AI innovators in understanding legal requirements to ensure compliance with existing regulations. Specifically, as the AI Act is quite new, not a lot of practical examples exist yet on how to apply the regulation. A second incentive is building trust in AI innovators. While the AI Act only requires conformity self-assessment, a successful participation in a sandbox can still lead to an increase in credibility for the AI innovators, potentially attracting clients and investors for their innovative solutions. This can also lead to reduced time to market, for the benefit of the AI innovators.

The AI innovators contribute directly to the value of the sandbox process by bringing expertise on AI systems, and state-of-the-art innovation in their domain of expertise. This enables evidence-based regulatory learning for the regulators directly, compared to traditional methods of regulating innovation which are slower in their adaptability. The regulatory sandbox framework fosters discussions between regulators and innovators to develop better regulation faster and mitigate risks while promoting innovation.

However, AI innovators also have concerns that need to be addressed to ensure their participation in the sandboxing process. Most concerns relate to trust. Can companies trust that strict confidentiality will be upheld? Can they trust that their intellectual property will be protected? Can they trust the process itself? For example, how do the participants know that the information they share during the sandbox process will not be used against them? In some instances, the sandbox process might even be viewed as a hidden audit process, meaning that innovators would be investigated by the sandbox to check whether they are compliant, also on regulation which is not covered by the sandbox agreement. Addressing those concerns is also crucial to ensure AI innovators full participation to the process.

The key focus for mitigating the concerns of the AI innovators is trust building. All stakeholders need to understand that the sandbox process is a regulatory tool that is there to serve all affected parties. To build trust with the innovators, it is important to start early, ideally before the sandbox process starts. This can be done through workshops, open Q&As, open office hours, etc. The goal is to help the innovators understand what the process involves, and what to expect as they participate. Unclear issues like what legal exemptions are relevant, what type of

guidance they can expect, and what their liability in the sandbox process is can be answered in this pre-sandboxing phase.

To aid the companies in this initial trust-building process, it is crucial that the sandbox strives to be open and transparent. One measure that was discussed was reporting out learnings from the sandbox. This should consist not only of the exit reports as required by the AI Act, but also continuous reports about what happens in the sandbox. There are many ways of building this trust, however, the point here is to highlight the importance of prioritizing the trust-building process.

Key takeaways from engagement with stakeholders: AI Innovators	
(Potential) Challenges	Actions
AI providers, AI deployers, SMEs and Start-ups <ul style="list-style-type: none">• Risk of IP leaks and confidentiality breaches.• Lack of trust and open communication channels.• Limited transparency in sandbox activities.	AI providers, AI deployers, SMEs and Start-ups <ul style="list-style-type: none">• Implement strict confidentiality and IP protection procedures.• Foster trust through open communication (e.g. workshops, Q&As, office hours).• Provide ongoing reports on sandbox activities.

Regulators

Having explored the perspective of AI innovators within AI regulatory sandboxes, we now turn to another core stakeholder: the regulators. These actors are responsible for operating the sandbox and ensuring that AI systems comply with the AI Act. Unsurprisingly, their incentives, contributions, concerns, and modes of engagement differ significantly from those of AI innovators. In this section, we will examine the regulators' perspective on the sandbox.

The regulators are as important as the AI innovators and are responsible for the sandbox functioning properly, facilitating both regulatory guidance and technical testing. In addition to making sure the AI innovators get the most out of the sandbox, they are the ones who carry out the regulatory learning. The NCA is the member state's contact point for the AI office, ensuring that the regulatory learning from the sandbox, including best practices and lessons learned, is communicated to the AI office through the exit report. The objective of "regulatory learning" is to be able to create, amend and adapt regulations when necessary.

The regulators have two primary incentives to participate in the sandbox process. First, they want to improve compliance with existing regulations. In this context, that means giving regulatory guidance and supporting the AI innovators with compliance. Secondly, they want to understand what is happening regarding the evolution of both technology and the market. The sandbox enables them to follow both these developments.

The regulators have several potential concerns or limitations that should be addressed. First, they can lack resources in the form of financial, technical, and human resources. If they are not allocated the necessary resources, they might have to outsource some parts of the sandbox process. This problem can lead to diverging quality of the sandbox across the EU member states, potentially undermining goals of harmonization on the EU level and potentially leading to forum shopping.

Another challenge regulators may face is the risk of providing excessive guidance or unintentionally distorting the market by favoring one company over another. If regulators appear biased or disproportionately support a particular company's efforts, it could result in an unfair competitive advantage. The regulators must be very careful in their role as neutral actors in the process. This problem could be mitigated by having transparent and neutral selection criteria, ensuring that the selection process is fair and equal.

It is clear that the regulators play a key role in a regulatory sandbox. They are responsible for coordinating and operating the sandbox. This gives them the responsibility of balancing the interest of the stakeholders involved, to make sure that the sandbox is as productive as possible. At the same time the regulators face challenges relating to their access to resources, which needs to be addressed.

Key takeaways from engagement with stakeholders: Regulators	
(Potential) Challenges	Actions
<ul style="list-style-type: none"> • Inadequate resources to operate sandboxes. • Limited access to expertise. • Risk of market distortion because of biases. • Lack of fair selection process. 	<ul style="list-style-type: none"> • Allocate necessary resources for effective sandbox operations. • Ensure access to relevant expertise (internal or external). • Focus on maintaining market neutrality by avoiding favoritism or biased guidance. • Develop transparent and unbiased selection criteria.

Civil society

Civil society can have an important role in the sandboxes if it is actively facilitated. It is important to include civil society in the sandbox process because they can help promote inclusive and human-centered AI services. This increases the positive social impact of the various projects (or decreases the negative social impact). Often, civil society groups, like NGOs, represent the interests of groups that can sometimes be neglected. This is valuable to minorities, or groups that do not have representation through other channels. The incentives of these stakeholders to participate in the sandbox process align with the reason for which they often exist in the first place. They want to represent their members and protect their rights. An extension of that is to gain influence over the AI Agenda, so they can maximize the positive social impact. The reason for wanting influence over the AI agenda is to get better and more trustworthy AI services.

The concerns or limitations of the civil society stakeholders are similar to those of the regulators. The concerns that arise are a lack of resources, AI literacy, and access. It is important to facilitate the involvement of these actors in the sandboxes to make sure sandboxes maximize the benefit for the whole of society. The potential benefit these stakeholders can have may be illustrated through the past example of the Dutch children benefit scandal where approximately 26,000 parents were wrongfully accused of making fraudulent benefits claims, resulting in families being driven into severe financial hardship. When an investigation was launched, it

was found that the cause of this scandal was algorithmic bias in the system that was employed. The system racially profiled minorities and discriminated against them. It was through the report *Xenophobic machines* that civil society actors were able to uncover this scandal, leading to positive change. This case illustrates how important civil society actors are in general, and it is important to have them represented in the sandbox process to protect the rights of different groups in society.

What can we do to best facilitate their involvement in a sandbox, or what do we have to keep in mind when designing the sandbox? First, it is important to give these actors real influence, so that they are not just consulted superficially, checking off the box. Secondly, measure that can prove itself useful is providing training on AI literacy so the civil society actors can meaningfully engage with a sandbox. Thirdly, subsidizing their participation costs or supporting them financially to be involved could be considered. The goal here is to create mechanisms that can give these actors an economic incentive to participate, while also highlighting them as a key actor in the sandbox process.

Key takeaways from engagement with stakeholders: Civil Society	
(Potential) Challenges	Actions
<ul style="list-style-type: none">• Superficial stakeholder engagement.• Low AI literacy and awareness.• Financial barriers to participation.	<ul style="list-style-type: none">• Ensure meaningful involvement by giving real influence.• Promote AI literacy and accessibility.• Offer financial support to resource-limited organizations.

Societal and Scientific Experts

Involving experts in the sandbox is important to ensure the quality of the operations. The distinction we make between scientific experts and societal experts in this context is that societal experts are individuals with personal experience from a minority community. This brings important perspectives that the scientific experts or other stakeholders are not able to capture. Societal experts may work with the civil society actors involved in the sandbox to ensure that vulnerable minority groups are represented in the sandbox. They may have economic incentives as well as interests in increasing their visibility. However, the ultimate goal of these actors in participating in the sandbox is to promote safer and fairer AI systems.

The scientific experts, on the other hand, must serve a unique role in the sandbox. The reason for this is that often these experts, for example university professors, do not have any personal stake in the case. That allows them to consult a developer, regulator, or civic society, allowing them to be a resource to everyone. They already tend to have similar involvements in innovation funding agencies to provide expertise whenever required. Another benefit of scientific experts is that they can bring a perspective that considers the big picture, and the possible consequences that may arise in the future. Experts can bring context, highlight risks, and apply critical thinking to the sandbox process, making them key stakeholders.

Experts face some hurdles in participating in the sandbox, and some of them are similar to the ones encountered by civil society actors. The main hurdle is the limited available resources of the sandbox. Without proper resource planning and access facilitation, experts might lack visibility and access to the sandbox. Their role as advisors, and the knowledge they can provide, is a crucial element to the sandboxing process. One solution is the creation of a panel of experts that can be consulted whenever required. This allows for a wider range of experts to be available, while limiting the consumption of financial resources that having full-time experts would require.

There are several things to keep in mind when trying to engage experts in the sandbox. First off, it is important to acknowledge their importance in the sandbox. As an extension of that, they often want visibility and a real impact. This is quite like civil society engagement. The point identified that can differ is how they want long-term engagement and future opportunities.

Key takeaways from engagement with stakeholders: Societal and scientific experts	
(Potential) challenges	Action
<ul style="list-style-type: none"> • Minority perspectives are often overlooked by traditional stakeholders. • Incentivizing societal experts with lived experience to participate. • Scientific experts often have no direct stake, which may limit their engagement or prioritization. 	<ul style="list-style-type: none"> • Involve societal experts with lived experience to represent vulnerable or minority groups. • Facilitate participation through funding, visibility initiatives, and collaboration with civil society actors. • Highlight their value as neutral, big-picture thinkers and ensure they are engaged as cross-cutting resources.

<ul style="list-style-type: none"> • Experts lack structured ways to contribute consistently. • Lack of recognition and long-term opportunities may discourage expert participation. 	<ul style="list-style-type: none"> • Establish a dedicated expert panel for consultation during key sandbox stages. • Acknowledge their contributions, ensure visibility, and provide avenues for sustained engagement beyond the immediate sandbox.
--	--

This overview of the different stakeholders in an AI regulatory sandbox highlights the complex network of incentives, contributions, concerns, and engagement methods required for effective sandboxing. There are many moving parts when it comes to effectively involving all stakeholders to get the most out of the sandbox process. It requires careful coordination by the sandbox operators to utilize the full potential of the sandbox. In the next section, we will look at the knowledge generated in the sandbox, its important role, and how it can contribute to the harmonization of sandboxes. Then we will look at to what extent we want to harmonize the different AI regulatory sandboxes.

Knowledge Generation and Harmonization of Sandboxes

A key aspect of AI regulatory sandboxes is regulatory learning. This concept needs to be detailed and specified to ensure that stakeholders understand what this entails. In the sandbox process, several types of resources are created that could be transferred to other sandboxes. In our discussion, three types of knowledge outputs were identified. The first one the mandatory exit report covering specific use cases and projects, a crucial element in the sharing of best practices and regulatory learning. The second type of resource output is methodology, which can be viewed as a more structural form of knowledge compared to case-specific knowledge.

This methodology relates to procedures in the sandbox, such as legal guidance and testing procedures, and can be analyzed across different phases of the sandbox.

Another important element of knowledge generation in the regulatory sandbox is competence building. This includes the skills and expertise gained by individuals involved in the process, who can then share their knowledge with other stakeholders. This transfer of expertise strengthens the overall regulatory learning ecosystem, ensuring that insights gained from one sandbox experience contribute to broader improvements in AI governance.

The knowledge generated in the sandbox process can be the basis for the harmonization of sandboxes. The experiences and best practices are documented in exit reports that can be used by other sandbox operators to improve their operations. This can be in the form of improving internal procedures or avoiding repeating the tests on use cases solved in other sandboxes. The generation of knowledge and sharing of exit reports is important to strengthen collaboration between different sandboxes.

An important question that arises in relation to collaboration between the regulatory sandboxes is to which extent AI regulatory sandboxes should be harmonized. While harmonization can be beneficial, leaving flexibility to member states can often be useful. For example, the eligibility criteria because different countries may have different needs. These needs can also encompass cultural elements that make it difficult to harmonize eligibility criteria. For example, general trust in regulators may vary depending on whether you are running a sandbox in a Nordic country compared to an Eastern European country. The difference in trust can impact how strict the eligibility criteria can be, where rigid criteria can work in a Nordic country with high trust, it can discourage companies in countries with lower trust. A solution that works in one region might not work in another. This illustrates that despite harmonization being important, there must be room for divergence.

Within a member state, the issue of alignment between sandboxes at different administrative levels also requires attention. A key consideration is the extent to which a national sandbox should allow flexibility for regional or local sandboxes to adapt the framework to their specific needs. As with alignment between member states, it is important to define broad, overarching principles for the sandbox while preserving space for local initiatives to tailor their approach. One possible model could involve the national sandbox focusing primarily on regulatory aspects, while regional or local sandboxes emphasize technical testing. This setup would give companies the opportunity to engage with the sandbox that best fits their needs.

Key takeaways on knowledge generation and harmonization of sandboxes	
(Potential) Challenges	Actions
<ul style="list-style-type: none"> Limited sharing of insights between sandbox initiatives. National and regional differences hinder uniform sandbox criteria. Lack of trust in cross-border sandbox outcomes. 	<ul style="list-style-type: none"> Regulators should prioritize knowledge-sharing mechanisms to ensure regulatory learning is disseminated across jurisdictions. Allow some flexibility to accommodate local contexts regarding eligibility criteria. Develop a high-quality, mutually recognized exit report to validate results across jurisdictions.

Conclusion

AI regulatory sandboxes represent a promising, though complex, tool to support both innovation and accountability in the AI ecosystem. As illustrated throughout this whitepaper, sandboxes have the potential to address a wide range of stakeholder needs. These include AI innovators, who are seeking clarity and go-to-market speed. They also include regulators, who aim to build robust, risk-based oversight mechanisms. Finally, civil society groups benefit as well, as they advocate for fairness and inclusion.

The key to unlocking this potential lies in thoughtful design and careful implementation. Sandboxes must strike a balance between flexibility and structure, harmonization, and contextual sensitivity. A one-size-fits-all model is unlikely to work. Instead, a use-case driven,

modular approach, supported by strong collaboration between stakeholders, can help ensure that sandboxes are fit for purpose.

At the same time, the sandbox is not an isolated initiative — it is part of a broader AI governance ecosystem that must be continuously refined. The challenges of resource allocation, stakeholder trust, regulatory consistency, and meaningful participation will not solve themselves. Addressing them will require ongoing commitment, transparent processes, and deliberate knowledge sharing at both national and EU levels.

Ultimately, regulatory sandboxes are a chance to learn by doing. If we get it right, they can serve as important tools, not only for smarter regulation, but also for more trustworthy, inclusive, and impactful AI across Europe.

Bibliography

Sources on sandboxes:

Buocz, T., Pfotenhauer, S., & Eisenberger, I. (2023). Regulatory sandboxes in the AI Act: Reconciling innovation and safety? *Law, Innovation and Technology*, 15(2), 357-389. <https://doi.org/10.1080/17579961.2023.2245678>

Bagni, F. & Seferi, F. (2025). *Regulatory sandboxes for AI and cybersecurity: Questions and answers for stakeholders*. Cybersecurity National Lab. ISBN 9788894137378.

Datasphere Initiative. (2025). *Sandboxes for AI: Tools for a new frontier*. Datasphere Initiative.

EUSAiR (2025, March). *EU regulatory sandboxes for AI (EUSAiR): Roadmap*. https://eusair-project.eu/app/uploads/2025/04/EUSAIR_RoadMap_v1_final.pdf

Federal Ministry for Economic Affairs and Energy. (2019, July). Making space for innovation: The handbook for regulatory sandboxes. BMWi.

Markussen, T. E. (2023). *Evaluation of the Norwegian Data Protection Authority's Regulatory Sandbox for Artificial Intelligence*. Norwegian Data Protection Authority.

Attrey, A., M. Leshner and C. Lomax (2020), "The role of sandboxes in promoting flexibility and innovation in the digital age", OECD Going Digital Toolkit Notes, No. 2, OECD Publishing, Paris, <https://doi.org/10.1787/cdf5ed45-en>.

OECD. (2023). *Regulatory sandboxes in artificial intelligence*. OECD Digital Economy Papers, No. 3562. https://www.oecd.org/en/publications/regulatory-sandboxes-in-artificial-intelligence_8f80a0e6-en.html

Von Thiesen, R. (2024). *Play & Learn: How to strengthen an AI hub with a sandbox*. Division of Business and Economic Development, Canton of Zurich Metropolitan Area Zurich Association Innovation Zurich. https://www.zh.ch/content/dam/zhweb/bilder-dokumente/themen/wirtschaft-arbeit/wirtschaftsstandort/dokumente/sandbox_en.pdf

Sources on the EU AI Act:

Future of Life Institute. (2025). *The AI Act Explorer: EU Artificial Intelligence Act*. <https://artificialintelligenceact.eu/ai-act-explorer/>

European Commission, [Artificial Intelligence – Questions and Answers](#), 2023

EP Legislative Observatory, [Procedure file on Artificial Intelligence Act](#), 2021/0106(COD)

Council, [Artificial intelligence \(AI\) act: Council gives final green light to the first worldwide rules on AI](#), Press release, 21 May 2024

European Commission, [Decision establishing the European Artificial Intelligence Office, C\(2024\) 390 final](#)

European Parliament, EPRS, [Artificial intelligence act](#), Legislative briefing, September 2024

European Parliament, EPRS, [Generative AI and watermarking](#), December 2023

European Parliament, EPRS, [Parliament's negotiating position on the artificial intelligence act](#), At a glance, June 2023

European Parliament, EPRS, [General-purpose artificial intelligence](#), At a glance, March 2023

European Parliament, EPRS, [Artificial intelligence act and regulatory sandboxes](#), Briefing, June 2022

European Parliament, EPRS, [Artificial intelligence in healthcare: Applications, risks, and ethical and societal impacts](#), Study, June 2022

European Parliament, EPRS, [Regulating facial recognition in the EU](#), In-depth analysis, September 2021



UNIVERSITÉ
LIBRE
DE BRUXELLES



VRIJE
UNIVERSITEIT
BRUSSEL



THE BRUSSELS-CAPITAL REGION AND EUROPE INVEST IN YOUR FUTURE!