



Build & Share

A Sandbox for Enabling the AI Ecosystem

This report summarises the results of the second implementation phase (2024–2026) of the Innovation Sandbox for Artificial Intelligence (AI). Building on the insights from the first phase, the Department for Economic Affairs of the Canton of Zurich further developed the sandbox in a targeted manner and established it as a key instrument for enabling the AI ecosystem. The sandbox allows AI providers to test their applications under real-world conditions in collaboration with public sector partners, to clarify regulatory issues at an early stage, and to generate concrete impulses for policy, administration, and the economy. In the second phase, the sandbox team successfully implemented five AI use cases: AI in medical documentation, digital eye clinic, AI for building permits, sensor-based bridge monitoring, and autonomous inspection robots. With two use cases, the sandbox delivered tangible value in the healthcare sector for the first time. Another project supported robotics companies in entering the EU market. While the first phase (2022–2024) focused on AI applications in the field of computer vision, the second phase primarily tested generative AI. The results of the second phase confirm the AI sandbox as an effective instrument for learning and implementation that reduces regulatory uncertainty, develops best practices based on real-world use cases, and thereby strengthens the AI ecosystem. At the policy level, the insights gained served as a key foundation for the Government Council's decision to strengthen the AI ecosystem for the 2026–2029 period and for the establishment of sandboxes as an instrument for promoting AI innovation at the national level.

Contents

01.

*Confirmation and
Further Development*

Page 4

03.

*Implementation of the
Sandbox Projects*

Page 8

05.

*Impact and Scaling
of Results*

Page 20

07.

Glossary

Page 26

02.

*Selection of
Use Cases*

Page 6

04.

*Insights from the
AI Use Cases*

Page 16


06.

Conclusion and Outlook

Page 23

01.

Confirmation and Further Development



In 2022, the Canton of Zurich, in collaboration with the Metropolitan Area of Zurich, launched the Innovation Sandbox for Artificial Intelligence (AI). Based on concrete AI use cases, the sandbox addresses regulatory questions and enables AI providers to access suitable test partners and data sources from the public sector. In contrast to many international approaches, the focus lies on **real-world testing** rather than purely regulatory assessments. The sandbox team systematically prepares all results, publishes them transparently, and makes them available for knowledge transfer to organisations facing similar challenges. At the same time, the insights gained feed into the further development of regulatory frameworks. The strategic objectives, the sandbox design, and the insights from Phase I are documented in the final report [Play & Learn – How to strengthen an AI hub with a sandbox.](#)

Resonance and Impact of the First Phase

The first phase clearly confirmed the effectiveness of the approach: the sandbox team implemented five use cases under real-world conditions and further developed AI solutions in close collaboration with public partners. The extensive use of the results (e.g. thousands of report downloads and broad uptake across the ecosystem) underscores the re-

levance of the sandbox. In addition, the insights fed directly into regulatory and policy processes, providing a concrete basis for the further development of AI-related frameworks. At the same time, the approach evolved from a cantonal initiative in Zurich into an internationally recognised reference model, attracting significant interest beyond Switzerland.

«Our sandbox approach deliberately leverages the strengths of the Swiss AI ecosystem: collaboration, openness, and a strong focus on practical application.»

*Raphael von Thiessen,
Programme Lead AI, Canton of Zurich*

* Highlighted words are explained in more detail in the glossary.

¹ The sandbox team consists of an interdisciplinary core team comprising a programme lead (generalist), a regulatory expert (legal specialist), a technical expert (data scientist), and a communications expert.

01. Confirmation and Further Development

Roll-out and Sustainable Continuation of Projects

Beyond the immediate project outcomes, the impact of the sandbox is reflected in the continued development and operational deployment of individual use cases. For example, the [Smart Parking project](#) is in productive use in the City of Frauenfeld and has been further expanded. Following successful testing with the Canton of Schwyz, the [machine translation solution](#) for commercial register extracts has been rolled out to several cantons. Continued implementation is not an explicit objective of the sandbox; it remains the responsibility of the respective project partners. However, when use cases are further developed or scaled after the testing phase, this demonstrates their tangible value and the practical relevance of the sandbox. In addition, sustainable effects arise through the provision and reuse of data: within the [drone inspection project](#), IBM Research developed a structured, annotated dataset for infrastructure inspection, which serves as a benchmark for future AI models and is publicly accessible via scientific publications and platforms such as [Hugging Face](#).

Maintaining Success Factors

In addition to the substantive outcomes, structural factors also proved critical to the success of the sandbox. A key success factor is the high degree of personnel continuity. The members of the steering committee² have supported the programme since 2021 with minimal changes in composition. The project team has likewise remained stable over the years. This continuity strengthens institutional knowledge, shortens coordination processes, and enables the consistent further development of the instrument. As in the first phase, the Department for Economic Affairs (personnel resources) and the Metropolitan Area of Zurich (budget for programme implementation) jointly ensured the financing of the sandbox.

For the second implementation phase, the sandbox team deliberately focused on continuity combined with targeted further development. In addition to personnel stability, proven success factors were retained:

- Open project calls across all sectors
- A clear focus on regulatory clarification and access to test partners and data sources
- No financial compensation for participating organisations
- Implementation of all projects within the existing legal framework, without the possibility of regulatory exemptions
- A strong emphasis on public communication to facilitate knowledge transfer

Strategic Development of the Sandbox

Building on these stable foundations, Phase II introduced targeted strategic enhancements by incorporating feedback from project partners and the steering committee:

- Strengthened networking among sandbox participants through events and workshops
- Implementation of projects in new sectors (especially the healthcare system)
- Targeted engagement with EU regulatory requirements to support market access for local AI providers


Technological Developments

While the first phase focused primarily on **computer vision**, the sandbox in Phase II also integrated **generative AI**. This expansion was deliberately timed after the initial widespread market introduction of chatbots, enabling a more reflective and practice-oriented application of the technology.

² The steering committee comprises representatives from the Offices for Economic Affairs of the Cantons of Zurich and Schwyz, the Statistical Office and the Digital Administration unit of the Canton of Zurich, the ETH AI Center, the UZH Digital Society Initiative (DSI), swissICT, and ZHAW Entrepreneurship.

02.

Selection of Use Cases



In the second implementation phase, the project call took place under changed conditions: the AI sandbox was already established as a recognised instrument for promoting innovation. Accordingly, numerous enquiries from potential project partners were received prior to the official call. The project team therefore introduced a new step, engaging proactively with interested parties before submission to clarify open questions and assess relevance for the sandbox at an early stage.

Submissions and Selection Process

The number of submissions remained moderate and comparable to the first phase. In total, 24 project proposals were received, compared to 21 in the initial call. Selection followed the same [criteria](#) and process as in Phase I:

- **Submission:** Written submission of a use case via a low-threshold online form
- **Validation:** A 30-minute interview to assess motivation and domain expertise
- **Prioritisation:** Evaluation and prioritisation by the steering committee

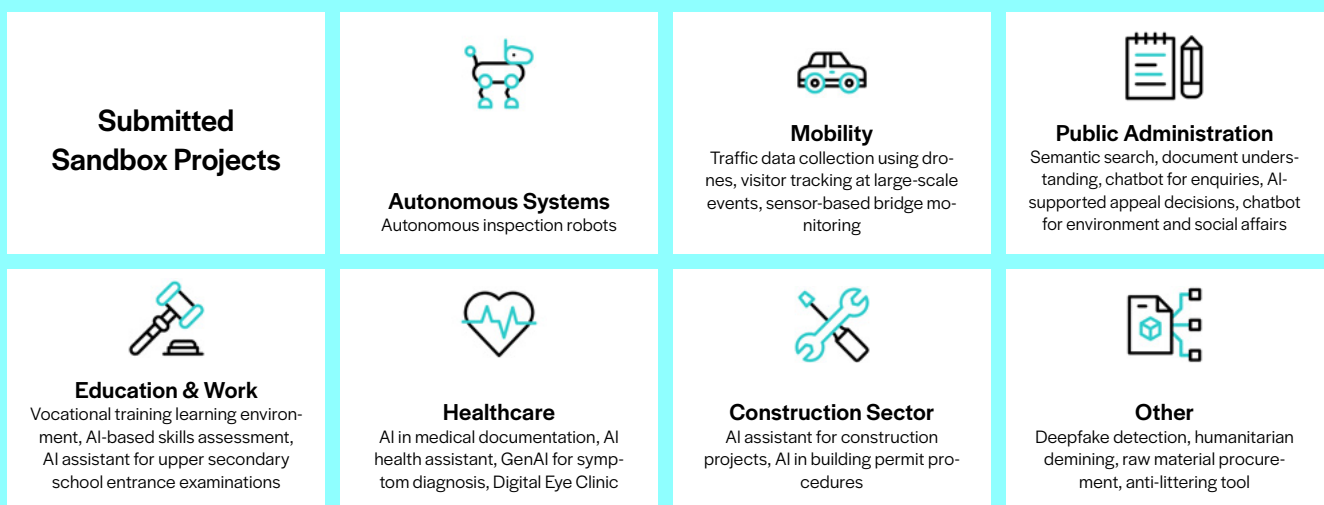


Figure 1: The 24 submitted projects illustrate the breadth of potential AI applications and underscore the added value of a cross-sectoral sandbox approach.

02. Selection of Use Cases

Strategic Focus: Healthcare Sector

Prior to the second project call, the sandbox team deliberately established a strategic focus on the healthcare sector. The objective was to systematically integrate a domain characterised by high regulatory complexity and significant innovation potential into the sandbox. To target suitable AI providers, the project team leveraged the digital health center Bülach as an existing partner. This approach proved effective: a total of four project proposals were submitted from the healthcare sector.

Findings and Recommendations

The evaluation of the submissions confirmed the importance of clearly defined use cases. For healthcare projects, precisely specified AI services and demonstrated experience in the healthcare domain proved to be key success factors. Less suitable were open-ended or problem-oriented proposals, such as generic healthcare chatbots or applications from AI service providers without a concrete use case.


Furthermore, the short online interviews once again proved to be a central element of the selection process. They enable a better understanding of the motivation behind project submissions, allow for a realistic assessment of the maturity of use cases, and provide a means to verify competencies. A purely written selection process would not be sufficient for this purpose. Going forward, the focus should remain consistently on clearly defined use cases, complemented by structured interviews as an essential component.

«Particularly in the healthcare sector, controlled testing environments are essential to effectively align regulatory requirements with innovation»

*Stefan Lienhard,
CEO, digital health center bülach (dhc)*

03.

Implementation of the Sandbox Projects

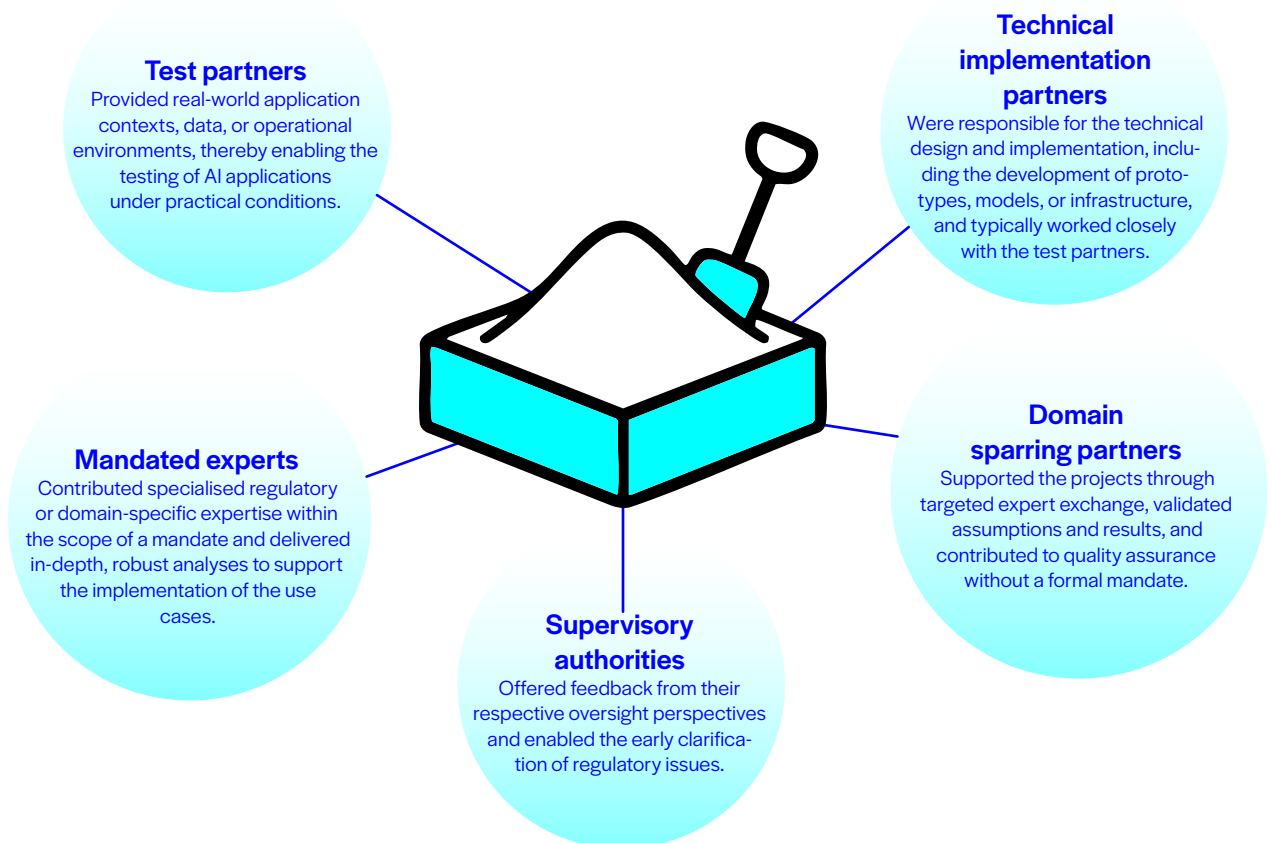


The sandbox team once again implemented the projects using a portfolio approach: in collaboration with participating organisations, it initiated eight use cases and realised five of them in parallel with slight time staggering; the reasons for the non-implementation of individual projects are outlined in the section «Additional Projects.»

team to leverage synergies in a targeted manner and to allocate resources flexibly through rolling budgeting. The portfolio approach accounted for varying levels of project maturity and allowed resources to be deployed where they generated the greatest value. At the same time, the project team remained flexible in accommodating the day-to-day operations of participating organisations.

The implementation window spanned 18 months, with most use cases completed within a duration of 9 to 12 months. This approach enabled the project

Implementation took place in close collaboration with a wide range of partners, each assuming clearly defined roles and responsibilities according to their expertise.



03. Implementation of the Sandbox Projects

Management of IP and Confidentiality

The handling of intellectual property (IP) and trade secrets is an integral part of project implementation and is defined jointly with all project partners at an early stage. The participating stakeholders identify protected content (e.g. proprietary algorithms, models, training data, or process logic), clarify responsibilities, and define the handling of sensitive data and results. A fundamental distinction is made between pre-existing IP contributed by organisations to the sandbox and results jointly developed within the project. Contributed IP remains with the respective partners and is not disclosed. By contrast, jointly developed, non-competitive insights (e.g. methodological approaches, regulatory classifications, reference architectures, or selected prototypes) are to be made accessible in a targeted manner. The published outputs of the sandbox, including best practices and selected artefacts such as datasets or prototypes, contain only content that does not infringe third-party rights and has been approved by all participants; confidential or competitively sensitive aspects are excluded.

Use Cases with Real-World Testing

In Phase II, the project team implemented a total of five use cases, which can be grouped into two categories. Three projects combined regulatory questions with concrete real-world testing of AI systems: the digital eye clinic, AI in building permit procedures, and sensor-based bridge monitoring. In these cases, the project team tested AI applications under real operating conditions in the public sector and generated practical insights for operations, integration, and governance.

Focus on Regulatory Analysis

Two additional projects deliberately focused on in-depth regulatory analysis, as the prerequisites for operational testing were not met. In the area of medical documentation, the stringent requirements for handling highly sensitive health data precluded deployment in a public hospital. In the project on autonomous inspection robots, no suitable test partner could be identified. Nevertheless, both projects delivered substantial value by establishing regulatory clarity, deriving best practices, and enabling the scalability of findings beyond the specific use case. The following sections provide a concise overview of the five use cases and their key results. Detailed information is presented in the linked project reports.

03. Implementation of the Sandbox Projects




1. AI in Medical Documentation

Medical documentation entails a high administrative burden and consumes valuable time of healthcare professionals. AI-supported speech recognition and *large language models* offer the potential to efficiently transcribe, structure, and improve the quality of reports. Based on the use case of the AI provider [MPAssist](#), this sandbox project developed a practice-oriented guideline for the safe use of AI in medical documentation, particularly with regard to data protection and the delineation from medical devices. Due to the stringent data protection requirements governing data processing, operational deployment with patient data in a public hospital was not feasible. Instead, the project focused on the systematic clarification of regulatory issues. To scale the results, the project team organised a regulatory roundtable with more than ten start-ups in the field of AI-supported medical documentation, together with the competent data protection supervisory authority. This format proved particularly effective, as AI providers rarely have direct access to supervisory authorities. Collaboration with additional stake-

holders, including Swissmedic, the Department of Health of the Canton of Zurich, and the eHealth working group of swissICT, was a key success factor in deepening the use case from multiple perspectives. The recommendations developed in the project quickly established themselves as a reference within the healthcare system and support providers, hospitals, and medical practices in the responsible adoption of AI-enabled documentation solutions. Their broad implementation in the form of proof-of-concepts, as well as pilot and roll-out phases, demonstrates that the recommendations are being translated directly into practice and provide concrete guidance for ongoing implementation processes.

[AI in medical documentation – Legal frameworks and recommendations](#)

03. Implementation of the Sandbox Projects



2. Digital Eye Clinic

Diabetic retinopathy is one of the most common ocular diseases and poses increasing challenges to healthcare systems due to rising patient numbers, shortages of skilled professionals, and escalating costs. AI-enabled diagnostics offer the potential to automate screening processes, relieve medical staff, and improve efficiency. In this sandbox project, conducted with the [Ophthalmology Clinic](#) of Stadtspital Zürich and the [Spröss Foundation for Ophthalmology](#), the medical, technological, regulatory, and ethical prerequisites for the clinical deployment of AI diagnostics were systematically assessed.

The project delivered three key outcomes with immediate practical relevance. First, it developed a practice-tested evaluation framework for assessing AI diagnostic solutions, which is also applicable to other medical specialisations. Second, the project team built a modular AI diagnostic platform based on open-source technologies, as no commercial provider met the stringent requirements of a public hospital. In doing so, the project demonstrated the tangible potential of *intrapreneurship* and in-house

development within the public healthcare sector. Third, the project clarified key regulatory, data protection, and ethical requirements and derived concrete best practices for the implementation of AI diagnostics in hospitals. The results are transferable beyond the specific use case and support the safe and responsible scaling of AI in clinical practice.

[Digital Eye Clinic – Bringing AI diagnostics into practice](#)

03. Implementation of the Sandbox Projects




3. AI for Building Permits

Building permit procedures are complex, and the regulatory requirements are highly fragmented. Incomplete or erroneous applications result in significant additional effort for municipalities, specialised authorities, and applicants. Against this backdrop, this sandbox project examined how AI can be used in a targeted manner to relieve and optimise the process. In collaboration with practice and technology partners, the project team tested concrete AI use cases and developed a prototype for the application of AI in the submission and review of applications for solar installations, heat pumps, and charging stations. The prototype is modular and designed as an open system, enabling municipalities or other cantons to adopt it, adapt it to their procedures, and further develop it as needed. The work produced clear results: rule-based systems already enable substantial efficiency gains for well-structured requirements. The AI-supported pre-checks developed within the project provide additional value, particularly in assessing plausibility and completeness of applications. At the same time, a systematic evaluation of different AI models revealed significant differences

in quality and consistency, especially in the interpretation of plan drawings. For model validation, the project used real building applications from the City of Kloten. Technical implementation was carried out in collaboration with the partners [Nokema](#), [Gossweiler Ingenieure](#), and [Byte Studio](#). The project results provide a robust basis for decision-making regarding the future use of AI in building permit procedures. They were directly incorporated as technical input into the response to the [postulate on AI in building permit procedures](#) and demonstrate in a practice-oriented manner at which stages of the process AI can already provide meaningful support, where rule-based approaches remain preferable, and where current feasibility limits exist. At the same time, the realistic efficiency potential was analysed, highlighting the conditions under which AI can, in the future, contribute to reducing the workload of specialised authorities and shortening processing times.

[AI for Building Permits – Use Cases, Prototypes and Practical Insights](#)

03. Implementation of the Sandbox Projects




4. Sensor-based Bridge Monitoring

Rising traffic volumes and an ageing infrastructure base are significantly increasing the effort required for monitoring and maintaining bridges. This sandbox project, submitted by [irmos technologies](#), examined how data-driven monitoring using AI can provide improved decision-making foundations for operations, maintenance, and investment. Following the withdrawal of the originally planned implementation partner for a municipal pedestrian bridge, the project was reoriented and carried out in collaboration with [Schweizerische Südostbahn AG \(SOB\)](#). At the Reidholz railway bridge in Wädenswil, sensors captured real loads from train traffic, providing for the first time a precise, data-based picture of the structural condition. Building on this, the project developed scenarios for the remaining service life of the bridge, supporting concrete maintenance and investment decisions. In parallel, the project team clarified key legal issues relating to liability, data protection, and the operation of sensor-based monitoring systems. The pilot project demonstrates in a practical manner how data-driven approaches can

make the monitoring of critical infrastructure safer, more efficient, and more sustainable, particularly in the heavily utilised infrastructure of the Metropolitan Area of Zurich.

[Sensor-based Bridge Monitoring – Data Collection and Forecasting for Infrastructure Maintenance](#)

03. Implementation of the Sandbox Projects



5. Autonomous Inspection Robots

Autonomous inspection robots automate hazardous and resource-intensive inspections in critical infrastructure, thereby enhancing safety and efficiency. The objective of this sandbox project was to systematically clarify the regulatory prerequisites for EU market access for such systems. Based on a specific use case by [ANYbotics](#), the project examined how autonomous inspection systems can be classified in compliance with the requirements of the **EU AI Act** and the **EU machinery regulation** and brought to the EU market. As no suitable test partner for operational deployment could be identified, the focus was placed on in-depth regulatory analysis. In addition, the project team, in collaboration with [Modulos](#), tested the use of an AI governance platform and prepared the application of the **ISO/IEC 42001** standard on an exemplary basis. For the first time, the sandbox project also engaged in cross-border collaboration with competent supervisory authorities, e.g. the Federal Network Agency in Germany and the Office for Digital Innovation of the Principality of Liechtenstein. A particular challenge was posed by the dynamically

evolving EU requirements and the interdependencies between the AI Act and the machinery regulation. The project established a structured regulatory classification of autonomous inspection systems within the interplay of the AI Act and machinery regulation and developed practice-oriented compliance strategies. In doing so, the sandbox made a concrete contribution to facilitating EU market access for Swiss robotics companies.

[Autonomous inspection robots - Approaches to the AI Act and EU machinery legislation](#)

03. Implementation of the Sandbox Projects

Additional Projects

In Phase II, the sandbox team initiated a total of eight projects. Three of these were not pursued further by mutual agreement. The discontinued use cases included an AI-supported platform for preparing for upper secondary school examinations, chatbots for highly regulated domains, and the detection of **deepfakes** in the context of identity fraud. In two cases, the project team was unable to identify suitable implementation partners for operational testing. In another case, the project's orientation as a research initiative and the required financial contribution did not align with the objectives of the sandbox. Despite their non-implementation, the sandbox team generated value in all cases: it facilitated access for participating companies to relevant public authorities, initiated exchanges with potential users, and provided guidance on regulatory and organisational requirements.

Dependencies on Test Partners


Collaboration with test partners is a critical success factor, but also represents the greatest dependency within the sandbox approach. Availability, access to data, and internal decision-making processes on the partner side are often only partially predictable and directly affect feasibility. Sandbox projects with a focus on practical implementation therefore require a strong openness to outcomes, as well as flexible and iterative prioritisation by the project team. Planning certainty can only be ensured to a limited extent; the key lies in the ability to identify dependencies at an early stage and to adjust timelines and project focus accordingly.

«Although our project was not implemented, the AI sandbox enabled us to access public authorities and significantly accelerated our impact study on AI in education»

Jan Bühlmann, Founder, GoGymi

04.

Insights from the AI Use Cases



After two phases with a total of ten implemented AI projects, clear and recurring patterns have emerged across the dimensions of technology, law, and organisation. The insights are based on concrete AI use cases from a range of sectors, from smart parking and medical documentation to building permit procedures and infrastructure monitoring.




Figure 2: Across two phases, the sandbox programme successfully completed a total of ten AI projects between 2022 and 2026.

04. Insights from the AI Use Cases

Technology

- **AI is rarely the primary challenge:** in many cases, the deployed AI technologies are already available and mature (e.g. computer vision in smart parking or speech recognition for medical transcription). The more significant challenges lie in infrastructure, data access, integration, governance, and adoption in day-to-day operations.
- **Rule-based systems outperform generative AI for well-defined decision logic:** generative models demonstrate strengths in processing unstructured information, such as summarisation, transcription, information structuring, or preliminary checks (e.g. medical documentation or pre-checks in building permit procedures). However, they reach their limits in clearly defined, traceable, and legally relevant decision-making processes. In such cases, rule-based approaches deliver more robust and transparent results.
- **AI extends beyond large language models:** many high-impact use cases in the public sector are based on computer vision or sensor data rather than language models (e.g. smart parking, drone inspections, sensor-based bridge monitoring).
- **Model selection is task-specific:** different models exhibit distinct strengths and weaknesses depending on the task (e.g. visual plan interpretation versus text extraction in building applications). Unlike generic **LLM benchmarks**, practical experience shows that a single «best» model often does not exist, particularly in diverse use cases involving multiple data inputs (e.g. images and text).
- **Model-agnostic architecture proves advantageous:** given the rapid pace of technological development, systems should be designed to allow models to be exchanged and continuously evaluated (e.g. benchmarking in the «AI for building permits» project).
- **Combine *frontier models* and open-source solutions strategically:** requirements relating to data protection, performance, cost, and scalability vary significantly across use cases. Organisations can combine commercial frontier models from the cloud with locally deployed open-source models and should select the most appropriate configuration based on specific requirements (e.g. in healthcare contexts).
- **Data access is the most powerful lever:** even seemingly simple datasets can generate substantial value for the development and optimisation of AI systems. When made accessible, they enable a wide range of providers to train, validate, and improve their models (e.g. high-resolution drone imagery of a runway).



04. Insights from the AI Use Cases

Legal

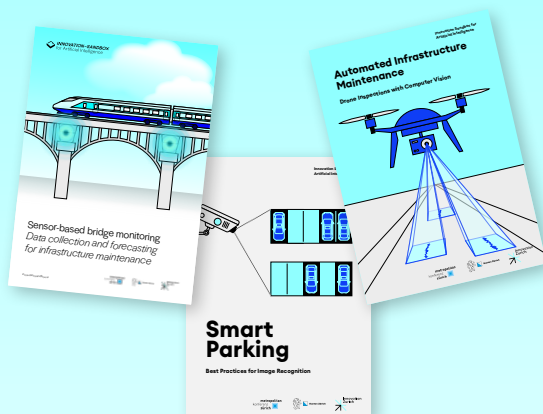
- **AI providers should integrate regulatory considerations early into product development:** if regulatory requirements are addressed too late, flawed assumptions may arise in the business model. Early clarification is essential to avoid pursuing an incorrect development path. For example, the «Autonomous Systems» project showed that retrofitting existing agricultural vehicles for automation may require full recertification of the product.
- **AI does not operate in a legal vacuum:** existing legal frameworks, e.g. in data protection and sector-specific regulation, are central in all projects and already address many of the key issues. AI-specific regulation, such as the implementation of the AI Convention into Swiss law, complements this framework in a targeted manner, especially in areas where AI systems affect fundamental rights.
- **Interdisciplinary legal issues are the norm:** AI applications typically span multiple areas of law (e.g. data protection and medical device regulation in the digital eye clinic). Isolated assessments are insufficient. AI providers require a holistic legal evaluation across the entire use case, rather than isolated compliance assessments of individual components.
- **Rapid regulatory guidance is essential:** supervisory authorities must be able to provide timely guidance and legal certainty for new use cases, particularly for technologies that spread rapidly across a sector (e.g. AI in medical documentation).
- **Continuous regulatory learning becomes critical:** traditional regulatory approaches reach their limits with dynamic, adaptive systems. New formats for exchange and feedback between authorities and providers are gaining importance beyond rigid approval mechanisms.



04. Insights from the AI Use Cases


Organisation

- **Strong stakeholder engagement is decisive:** successful projects are characterised by committed implementation partners who contribute time, resources, networks, and domain expertise (e.g. extensive testing in the smart parking project in Frauenfeld). Expectations towards implementation partners must be clearly defined from the outset.
- **Value must be clearly articulated:** the benefits of AI are not always immediately apparent to domain experts (e.g. sensor-based monitoring of bridge vibration behaviour as a complement to visual inspections). The key is a clear link to concrete decisions and operational value in day-to-day activities. AI should be deployed where it creates tangible value, not merely because it is technologically feasible.
- **The interface between domain expertise and technology is critical:** the greatest value emerges where domain knowledge and technical expertise are closely integrated. Without this translation layer, AI systems remain either technically sound but operationally irrelevant or vice versa.
- **Quality depends on systematic testing and domain validation:** reliable results require structured testing, clear metrics, and repeated evaluation (e.g. thousands of test runs and LLM-as-a-judge in the building permit project). At the same time, domain-specific assessment is essential. Technology providers are often unable to fully evaluate outputs; this must be carried out by subject-matter experts (e.g. experienced ophthalmologists in the digital eye clinic). Ad hoc testing is insufficient.
- **Scaling often fails at integration, not at AI performance:** the transition from pilot to operational deployment frequently fails due to insufficient integration into existing systems, processes, and governance structures, rather than limitations of the models themselves.
- **New ways of working require active management:** AI reshapes existing processes and roles. Acceptance is fostered through transparent communication, training, and the involvement of employees (e.g. drone inspections as a complement to manual runway inspections).




05.

Impact and Scaling of Results



The sandbox aims to systematically scale the insights generated in individual projects. It transfers project outcomes in a structured manner to additional application domains, policy processes, and public discourse in order to achieve sustainable structural impact. The impact unfolds across three interrelated levels:



1. Impact within the Sandbox Projects


This level concerns the immediate benefits for the participating AI providers and public sector partners.

- **Concrete innovation:** Support for AI providers in the development, testing, and validation of AI applications under real-world conditions, particularly within the public sector context.
- **Access to test environments and data:** Facilitation of suitable test partners as well as real application contexts, including access to public sector datasets.
- **Regulatory clarification:** Structured classification of legal requirements and reduction of uncertainty.
- **Capacity building in public administration:** Development of practical AI expertise within specialised units, public authorities, and other public institutions through direct involvement in concrete use cases.

«The AI sandbox systematically translates concrete projects into broader impact: from practical application and knowledge transfer to the shaping of AI policy.»

*Raphael von Thiessen,
Programme Lead AI, Canton of Zurich*

05. Impact and Scaling of Results



2. Knowledge Transfer within the AI Ecosystem

At this level, project-specific insights are made accessible beyond individual use cases.

- **Systematic knowledge dissemination:** Publication of project reports including best practices, videos, and methodological insights for public administration, industry, and the professional community.
- **Disclosure and reuse:** Release of selected prototypes and platforms (e.g. AI in building permit procedures, Digital Eye Clinic) to enable further development by third parties.
- **Transfer formats:** Organisation of regulatory roundtables, expert events, and webinars.
- **International networking:** Collaboration with regulatory authorities (including Germany and Liechtenstein) and exchange with partners across Europe, North and South America, and Asia.



3. Systemic Impact and Public Embedding

At this level, project insights are translated into structural developments.

- **Direct policy impact:** Technical input into policy processes, such as the [postulate](#) on the use of AI in building permit procedures.
- **Policy impulses at the national level:** Contributions to the further development of AI policy instruments, particularly in the context of the [Federal Council's decision on AI regulation](#), which established sandboxes as an instrument for promoting innovation, as well as providing a basis for a parliamentary initiative from the Council of States on [AI testing environments](#).
- **Transfer to other sectors:** Dissemination of practical experience to sandbox or innovation initiatives in education ([Educa](#)), healthcare ([Swissmedic](#)), and mobility ([MODI](#)).
- **Public visibility and legitimacy:** Recipient of the [Digital Economy Award 2025](#) (Government category) and strong media presence, including interviews in [Neue Zürcher Zeitung](#) and on [Tele Z](#). Accessible formats such as podcasts are also well suited to communicating the sandbox to a broader audience ([AI Talk Live-Podcast](#)).

05. Impact and Scaling of Results

Limitations of Individual Use Cases

The sandbox deliberately focuses on clearly defined use cases. This bottom-up approach enables practice-oriented insights but remains limited to individual applications. At the same time, this focus entails a risk: the more the selection is concentrated on a small number of use cases, the greater the impact if individual projects cannot be implemented as planned. To maximise impact, careful selection of strategically relevant and robust use cases is therefore essential. Projects such as AI in medical documentation demonstrate that use cases with high systemic relevance can generate particularly strong leverage.

Resource Intensity

Operating a sandbox is resource-intensive. Scaling from five to, for example, fifty projects would require substantial personnel and organisational investment. The sandbox team addresses this challenge through a strong partner network; nevertheless, natural limits to scalability remain. A further limitation lies in the timing of project calls, which currently take place approximately every 1.5 years. No new projects can be initiated between these calls. If the submission window is missed, participation is only possible in the subsequent call. Given the rapid pace of AI development, this can lead to delays in generating insights.

Limited Scope of Insights within Existing Legal Frameworks

The sandbox operates deliberately within existing legal frameworks and does not allow for the suspension of regulatory requirements. While this strengthens the practical applicability of the results, it also limits the scope of insights regarding potentially innovative but not (yet) legally permissible approaches. Certain questions therefore cannot be fully explored and remain constrained by the current legal framework.

Implementation and Policy Uptake

The sandbox generates impulses for AI innovation in the public sector and for regulatory development. At the same time, the public sector holds significant potential: as an early adopter of AI systems, it can actively drive demand for innovation, gain practical experience, and thereby strengthen the local AI ecosystem. However, large-scale implementation lies with the respective institutions. Public authorities, municipalities, federal agencies, educational institutions, and infrastructure operators must independently deploy, operate, and integrate AI systems into their processes. Similarly, regulatory insights feed into democratic processes; whether and how they are translated into legislation is determined through political deliberation. The Innovation Sandbox demonstrates what works and how AI can be deployed responsibly. Broad-scale adoption, however, remains a shared responsibility across the entire AI ecosystem.

06.

Conclusion and Outlook



The second implementation phase has demonstrated that sandboxes are an effective instrument for bringing together innovation, regulation, and economic development in the field of AI. The projects delivered tangible results, created regulatory clarity, and enabled learning under real-world conditions, both through operational testing and regulatory analysis.

Local AI providers and public institutions benefited through increased legal certainty, supervisory authorities gained practice-oriented foundations for proportionate regulation, and policymakers received evidence-based decision-making inputs. The portfolio approach proved effective, as it enables flexibility, openness to outcomes, and targeted prioritisation.

Several areas for further development of the sandbox are emerging:

- **Agile project intake:** More frequent project calls combined with continuous onboarding of projects to respond more rapidly to new ideas, technologies, and regulatory questions.
- **Differentiation of project types and support needs:** Early distinction between projects focused on regulatory analysis and those involving real-world testing, complemented by a staged «funnel» approach based on the level of clarification required, from self-service to targeted support via a regulatory desk, and through to in-depth guidance within the AI sandbox.
- **Thematic calls with a portfolio logic:** Targeted calls aligned with prioritised thematic areas, such as the financial sector, which is of particular relevance to the Zurich economic region, to leverage synergies between projects and reduce the risks associated with individual initiatives.
- **Systematisation and accessibility of insights:** Structured preparation, consolidation, and targeted dissemination of best practices and learnings from sandbox projects (e.g. thematic playbooks, a one-stop shop), complemented by an AI-supported interaction component (e.g. chatbot) to maximise discoverability, usability, and reuse across different target audiences.
- **Dialogue with the public and civil society:** Expansion of formats that bring sandbox activities closer to the general public and civil society actors (including NGOs), incorporating participatory elements to strengthen transparency, acceptance, and trust in AI.
- **Expanded testing scope:** Policy clarification on how testing can be enabled under clearly defined conditions beyond existing legal frameworks, under strict supervision and with robust safeguards.
- **International collaboration:** Enhanced cooperation and coordination in the field of AI with international partners, for example through the development of an international database of sandbox projects and insights to strengthen transparency and knowledge transfer, as well as the establishment of a sandbox summit for cross-border comparison of national and regional approaches.

06. Conclusion and Outlook

Overall, Phase II demonstrates that the Innovation Sandbox for AI not only enables individual projects but also functions as an environment for learning and implementation. It establishes the foundation for testing AI responsibly, contextualising it within policy frameworks, and shaping its long-term deployment in the public interest.

From the AI Sandbox to the Systematic Strengthening of the AI Ecosystem

The successful implementation and further development of the AI sandbox formed a key foundation for the [Government Council of the Canton of Zurich's decision to strengthen the AI ecosystem](#). This development was made possible in large part through the funding provided by the Metropolitan Area of Zurich, which established the basis for the creation

and institutionalisation of the instrument. The experience gained within the sandbox demonstrated that practice-oriented testing environments not only enable innovation but also create regulatory clarity, build trust, and provide concrete inputs for policy decision-making. In this way, the AI sandbox acted as a catalyst for a broader, strategic approach to AI policy.

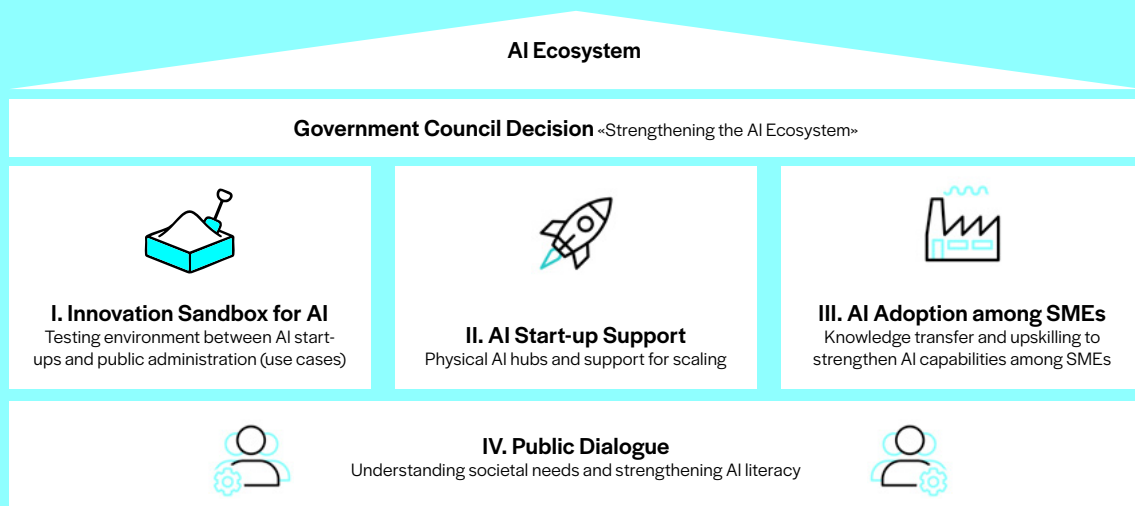


Figure 3: The Government Council's decision to strengthen the AI ecosystem builds on the AI sandbox and the AI dialogue and establishes four key areas for action.

06. Conclusion and Outlook

With the Government Council's decision, the AI ecosystem in Zurich is now being systematically advanced and scaled more broadly. Building on the insights from the sandbox, the canton is pursuing an integrated approach along four pillars: the further development of the AI sandbox as a testing and learning environment, targeted support for AI start-ups, enabling SMEs to adopt and use AI, and the expansion of a structured AI dialogue with the public. This shifts the focus from individual pilot projects to a sustainable ecosystem that addresses innovation, application, and societal integration in equal measure.

The AI sandbox remains a central component and will be further developed in a targeted manner. At the same time, new offerings are being created for SMEs to facilitate capability building and the adoption of concrete AI applications, while existing AI start-up networks and physical AI clusters are being strengthened. The AI dialogue with the public, already established in more than ten municipalities across the Canton of Zurich, will be systematised and expanded to incorporate diverse perspectives and to further strengthen acceptance and trust in the responsible use of AI.

The AI sandbox thus marks not a conclusion, but the starting point of a long-term AI ecosystem strategy. It has demonstrated how AI can be responsibly tested, appropriately contextualised within regulatory frameworks, and made practically usable, thereby laying the foundation for the next phase of development of the AI ecosystem in Zurich through 2029.

Glossary

Computer Vision

A subfield of artificial intelligence that automatically analyses and interprets visual data (e.g. images or videos), for example for object detection, classification, or damage identification.

Deepfakes

Audio, image, or video content generated or manipulated using AI to create realistic but artificially produced representations of individuals or events.

EU AI Act

EU regulation governing AI systems based on a risk-based approach. It defines requirements for high-risk AI systems, transparency obligations, and prohibitions of certain applications.

Frontier Models

Highly advanced, large-scale AI models (e.g. large language or multimodal models) representing the current state of the art. They are trained on extensive datasets using substantial computational resources and underpin many generative AI applications, while also introducing heightened requirements for safety, governance, and regulation.

Generative AI (Generative Artificial Intelligence)

AI systems that generate new content such as text, images, code, or audio, typically based on large pre-trained models.

Intrapreneurship

Entrepreneurial activity within established organisations, where employees initiate and drive innovation projects, often with an exploratory focus.

ISO/IEC 42001

An international standard for management systems in the field of artificial intelligence. It defines requirements for governance, risk management, and organisational processes to ensure the responsible development and use of AI.

Large Language Models (LLMs)

Large neural network-based language models trained on extensive text corpora to understand and generate natural language; they form the foundation of many generative AI applications.

LLM Benchmarks

Standardised tests used to evaluate and compare language models based on defined tasks (e.g. text comprehension, knowledge questions, or logical reasoning). Their explanatory power is limited, as real-world use cases are often more complex and involve diverse data formats and requirements.

Machinery Regulation (EU Machinery Regulation)

The EU legal framework governing the safety of machinery and related products. Its relevance to AI arises particularly in the context of autonomous or safety-critical systems such as robots.

Real-World Testing

The evaluation of an AI system under real operating conditions using actual data, processes, and organisational contexts, in order to validate technical, regulatory, and operational aspects in a practical setting.

Authors



Raphael von Thiessen
Programme Lead AI, Canton of Zurich



Lukas Willi
Project Lead AI, Canton of Zurich

With expert support from

Dr. Aurelia Tamò-Larrieux, Associate Professor, UNIL

Dr. Clara Guerra, Head of Office, Office for Digital Innovation, Principality of Liechtenstein

Demian Niemeyer, Senior AI Regulatory Expert, appliedAI Institute for Europe

Ivo Ruckstuhl, Director, Swiss Centre for Responsible AI

Melhem Daoud, Senior Project Manager Digital Infrastructure and AI, digitalswitzerland

Dr. Sebastian Singler, Director, PwC

Yvonne Finger, Head of Division, Federal Network Agency (Bundesnetzagentur), Federal Republic of Germany

Impressum

Publisher

Division of Business and Economic
Development, Canton of Zurich
Metropolitan Area Zurich Association
Innovation Zurich

Project conception and coordination

Raphael von Thiessen
Standortförderung Kanton Zürich
8090 Zürich
raphael.vonthiessen@vd.zh.ch

Authors

Raphael von Thiessen
Lukas Willi

Design

here we are gmbh, here-we-are.ch

Publication

This report is published exclusively in digital format
and is available in German and English.

Copyright

All content of this publication, in particular texts and
graphics, is protected by copyright. The copyright is
held by the Office for Economic Development of the
Canton of Zurich. The publication may be shared
with proper attribution and may be quoted from,
provided the source is fully cited.

Project-Steering

- Office for Economy, Canton of Zurich
- Statistical Office, Canton of Zurich
- Department for Digital Administration,
Canton of Zurich
- Office for Economy, Canton of Schwyz
- Zurich Metropolitan Area Association
- ETH AI Center
- Center for Information Technology, Society,
and Law (ITSL), University of Zurich
- ZHAW entrepreneurship