



D2.2

Workshop Report

“Experimental Design & Policy Support”

Status: DRAFT

Dissemination Level: Public

Note: When viewed inside Microsoft Teams, formatting and layout may not display correctly. The document formatting is optimised for PDF export, which is the official submission format.

Abstract

Key words Experimental Design; Policy Support; Strategic Research and Innovation Agenda (SRIA); Resource for AI Science in Europe (RAISE); Interoperability; Causal Inference; Trustworthy AI; Open Science

This report presents the findings and co-creation outcomes of the European SCIANCE workshop on **"Experimental Design & Policy Support,"** a key initiative under the Resource for AI Science in Europe (RAISE) project framework. Organized by DFKI, the workshop convened cross-disciplinary domain experts to identify, refine, and prioritize strategic AI research and innovation needs. Through an iterative approach combining initial quantitative screening, qualitative Miro-board mapping, and a formal Multi-Criteria Decision Analysis (MCDA), the panel evaluated 34 distinct research areas. The results reveal a clear shift away from routine, legacy statistical components like simple Bayesian optimization toward complex, system-level challenges including uncertainty representation, causal inference, and ethical, legal, and social aspects (ELSA) compliance. These insights establish a unified, open, and trustworthy infrastructure framework designed to directly inform the upcoming Strategic Research and Innovation Agenda (SRIA) and guide European Commission science funding tracks.

Revision History

Version	Date	Description	Author/Reviewer
V 0.1	29/10/2026	First draft	Dr. Andrea Lösch, Jessica Otto
V0.2		First draft review	
V1.0		Submitted Version	

Document Description

Work Package 6	AI in Science Working Groups Terms of Reference		
Document Type	Deliverable		
Document Status	Draft	Version	0.1
Dissemination Level	Public		
Copyright Status	 This material by Parties of the SCIANCE Consortium is licensed under a Creative Commons Attribution 4.0 International License.		
Lead partner	DFKI		
Document Link			
DOI			
Author(s)	Dr. Andrea Lösch, Jessica Otto		
Reviewer(s)	To be added based on feedback provided		
Approved by	SCIANCE Management Committee		

Table of contents

Executive Summary	4
1 Introduction	5
1.1 Focus of the workshop	5
1.2 Methodology	5
1.3 Workshop attendance	6
1.4 Data analysis	6
2 Workshop contents	6
2.1 Workshop agenda	6
2.2 Welcome & introduction	7
2.3 Review of relevant AI research & innovations	9
2.4 Initial survey	13
2.5 Interactive discussion	14
2.6 Consensus session	15
2.7 Summary & closing statements	16
3 Workshop results	17
3.1 Initial survey	17
3.2 Interactive discussion	21
3.3 Consensus session (MCDA)	21
4 Conclusions and recommendations	24
5 Glossary of abbreviations	26

SCIANCE Project

The [SCIANCE](#) project is a European coordination initiative establishing the foundations for the [Resource for AI Science in Europe](#) (RAISE). The project promotes interdisciplinary collaboration to integrate AI into scientific research by enabling streamlined access to trustworthy data repositories, computing resources, and research infrastructures across Europe. Bringing together leading research infrastructures, scientific organisations, and AI centres of excellence, SCIANCE develops a Strategic Research and Innovation Agenda (SRIA) through an iterative co-creation process.

Building on a comprehensive AI-in-science landscape analysis, the SRIA identifies long-term research challenges where AI can accelerate discovery and defines cross-domain AI research and innovation priorities, focusing on open, trustworthy, and frugal AI methods for science. A roadmap for infrastructure upgrades and cooperation mechanisms supports its implementation, translating scientific research and innovation priorities into coordinated action. To consolidate the European AI-in-science ecosystem, SCIANCE pilots the RAISE Secretariat, supported by the RAISE Digital Hub, community-building events at annual AIS Summits, and the RAISE Academy, providing a continuous framework for engagement, capacity building, knowledge sharing, and cross-domain collaborations.

By combining scientific domain expertise with AI research and infrastructure knowledge, SCIANCE delivers robust, evidence-based guidance that is scientifically credible and strategically aligned with European Commission objectives, laying the foundations for future collaborative, AI-enabled science in Europe.

Executive Summary

Context and Objective

The workshop “**Experimental Design & Policy Support**” was organized as a core co-creation track under the European Union’s SCIANCE initiative, establishing the strategic blueprint for the Resource for AI Science in Europe (RAISE). Operating on the principle that individual laboratories rarely possess the computational or data scale required for self-contained, reliable AI discovery, the event gathered leading domain experts to map cross-cutting AI-in-science priorities. The ultimate objective of this exercise is to feed an evidence-based, expert-validated portfolio directly into the upcoming Strategic Research and Innovation Agenda (SRIA) to accelerate European technological sovereignty and scientific competitiveness.

Methodological Workflow

The expert session followed a bottom-up, expert-driven co-creation workflow consisting of four interconnected methodological phases:

1. **Evidence-Based Preparation:** Baseline desk research pre-mapped initial research and innovation priority areas to provide a structured starting platform.
2. **Expert Validation and Extension:** On-site and remote experts evaluated the preliminary collection on a 5-star Likert scale via an automated online survey to identify baseline anomalies.
3. **Interactive Co-Creation:** Participants utilized a digital Miro board to dynamically expand the scope, debate unconventional ideas, group semantic proximities, and map conceptual dependencies.
4. **Multi-Criteria Decision Analysis (MCDA):** The final prioritized areas were formally scored via a Formbricks survey based on Scientific Impact (double weight), Economic Impact (double weight), and European Feasibility/Accountability (single weight).

Key Findings and Definitive Trends

The mathematical consensus and qualitative clustering tracks demonstrated a profound shift away from abstract, high-level ethics manifestos and standalone optimization formulas toward testable verification methods and open software infrastructure:

- **Rigorous Validation Trumps Routine Optimization:** Driven by an acute need for validation, “Uncertainty representation” secured the absolute number 1 ranking (Weighted Score: 4.00), closely supported by “Causal inference in explainable AI” at rank 2 (3.88; highest specific Scientific Impact score: 4.40).
- **The De-prioritisation of Established Formula Baselines:** Mature, legacy mathematical baselines like standalone *Bayesian optimization* (rank 30) and *ML-DOE* (rank 23) dropped to the bottom tiers. As such, they appear to be commercially mature commodities that do not require strategic public research funding.
- **Actionable Verification Over Abstract Principles:** While abstract entries like general *Ethics and governance technologies* dropped to the bottom (rank 34), the expert-driven priority area “**ELSA-compliant experimental design**” tied for rank 5 overall), establishing a clear mandate to build ethical check lines directly into the software code block.
- **Human Competency as a Core Strategic Pillar:** Dynamic expert note additions inserted entirely new socio-technical dimensions into the architecture, with “**Training scientists in the era of Agentic science**” securing a high 7th rank.

Top 15 Strategic Priorities Overview

The final ranked portfolio synthesizes the 34 raw entries into 15 re-grouped, action-oriented research and innovation vectors distributed across four core tiers:

- **Tier I: Trustworthy Optimization, Validation & Methodology Fitness (Items 1–4):** Encompasses statistical uncertainty representation, causal inference explainability, programmatic algorithmic fitness checks, and open European evaluation benchmark platforms.
- **Tier II: Interoperability, Cyber-Physical Interfaces & Lab Design (Items 5–8):** Focuses on *in silico*-real world abstraction protocols, next-generation autonomous closed-loop lab architectures, semantic search space annotation, and multi-agent swarm system simulation frameworks.
- **Tier III: Translational AI, Decision Support & Compliance (Items 9–12):** Focuses on actionable ELSA-compliant code boundaries, translational public policy decision support systems, multi-modal procedural audit sampling pipelines, and modular policy engines requiring automated justification trails.
- **Tier IV: Infrastructure Strategy, Human Capital & Security (Items 13–15):** Prioritizes multi-site zero-trust lines for confidential data processing, agentic science human capacity training, and energy-efficient, low-cost custom computing hardware (chips).

Immediate Next Steps

The report will be published via the RAISE Forum for public review. The feedback will be used to inform and develop the final expert survey for identifying the strategic research and innovation priorities across all SCIANCE areas.

1 Introduction

1.1 Focus of the workshop

The workshop “**Experimental Design & Policy Support**” focused on identifying and ranking cross-cutting AI research and innovation priorities to accelerate the scientific discovery lifecycle and strengthen public policy decision-making networks across Europe. Organized under the SCIANCE co-creation process for the development of the Strategic Research and Innovation Agenda (SRIA), the event addressed critical systemic challenges, including statistical uncertainty calibration, black-box data models, and the need for testable risk governance.

Participants explored the strategic intersection of adaptive closed-loop laboratories, causal inference, and translational AI platforms. Key technological and methodological areas highlighted included machine learning-driven Design of Experiments (ML-DOE), conversational hypothesis-generation agents, and Ethical, Legal, and Social Aspects (ELSA) compliance guardrails embedded directly into the software code layer. Special emphasis was placed on cross-domain interoperability protocols and shared open evaluation platforms to connect virtual simulation layers with physical, remote laboratory execution grids. Through interactive discussions and a formal multi-criteria decision analysis managed by DFKI moderators, the session evaluated and ranked 34 research priorities based on their scientific impact, economic benefit, and European feasibility. Ultimately, the workshop aimed to consolidate these expert-validated priorities into a structured portfolio to help shape Europe’s open, shared, and trustworthy AI-in-Science infrastructure roadmap.

1.2 Methodology

The SCIANCE WP2 workshops follow a **bottom-up, expert-driven co-creation approach** to identify and prioritise strategic AI research and innovation (R&I) needs across selected scientific domains and cross-cutting areas. The methodology combines **evidence-based preparation, structured expert elicitation, interactive discussion, and multi-criteria prioritisation**. Its purpose is not simply to identify existing AI solutions, but to determine where future research and innovation efforts could have the greatest potential to advance AI-enabled science at European level. Overall, the methodology consists of the following blocks:

1. **Evidence-based preparation of the workshops and initial mapping:** Prior to the workshops, SCIANCE consolidated relevant evidence from state-of-the-art reviews, landscape analyses and good practices. Based on this evidence, a preliminary collection of relevant AI research and innovations — including tools, applications, models — was prepared for each workshop topic. This provided a common evidence base while keeping the process open to contributions from the participating experts. In the corresponding preparatory session preceding the actual workshop, experts were asked to prepare individually by identifying the **five AI research and innovations they consider most important** for the respective topic and to assess their perceived relevance. They were explicitly encouraged to identify emerging developments and “outliers” that may not yet be mainstream.
2. **Expert validation and extension:** The workshop started with a plenary review of the evidence and the preliminary collection of AI research and innovations. Participants then were asked to assess the relevance of the identified items using a five-point relevance scale. The assessment took place in form of an initial online survey that was analyzed on the spot. Most importantly, the initial list provided to participants was deliberately kept open: experts could add missing research and innovations and indicate their perceived relevance. This allowed the methodology to combine a structured starting point with bottom-up expert input used to extend the original findings while at the same time reducing the risk of the initial evidence base determining the outcome in advance. The result was a **complete, expert-validated collection of relevant AI research and innovations**, accompanied by an initial assessment of their perceived relevance.
3. **Interactive co-creation and consolidation of priorities:** The initial collection then formed the basis for an interactive plenary discussion. Experts jointly examined the identified items, addressed open questions and differences in interpretation or relevance, and discussed particularly novel or unconventional ideas in greater depth. The explicit inclusion of such “outliers” or “crazy ideas” was intended to ensure that the process does not merely reproduce established technologies and research directions, but also captures potentially important future developments. Through this discussion, the broad collection was consolidated into a **final set of thematic AI research & innovation priorities** for the respective scientific or cross-cutting area.
4. **Multi-criteria prioritisation and consensus:** The final set of priorities were subsequently ranked by experts using a **lightweight multi-criteria decision analysis (MCDA)** in the form of an online survey based on formbricks. Specifically, three criteria were applied:
 - **Scientific impact** – the extent to which addressing the R&I need could unlock or significantly accelerate AI-enabled research.

- **Economic impact** – the potential effect on the economic performance of European companies and organisations, including productivity, cost reduction, value creation, competitiveness and revenues.
- **Feasibility/accountability** – the extent to which there is both a realistic pathway and a need to address the R&I need at European scale, for example through research, policy, infrastructure or coordination.

In the analysis, scientific impact and economic impact received **double weighting**, while feasibility/accountability receives a **single weighting**. The weighting is therefore designed to balance scientific significance and potential economic relevance with the practical question of whether there is a credible rationale and pathway for European-level action. The outcome of this consensus session was a **jointly validated and ranked portfolio of 34 AI research and innovation priorities**.

1.3 Workshop attendance

Overall, seven experts attended the workshop on-site:

- Prof. Masahiro Ryo - Leibniz Centre for Agricultural Landscape Research (ZALF)
- Prof. Jean-Pierre Vilotte – Institut de physique du globe de Paris (IPGP), CNRS [representative for Jalal Fadili - AISSAI Center, CNRS]
- Corina Pascu - European Union Agency for Cybersecurity (ENISA) [representative for Yonah Welker - MIT CSAIL, EMEA Campus + EU/EIT projects]
- Dr. Dragos Ovidiu Tofan - Alexandru Ioan Cuza University [representative for Prof. Mihaela Tofan - Alexandru Ioan Cuza University]
- Dr. phil. nat. Sebastian Lobentanzer – Computational Health Center, Helmholtz Munich
- Prof. Ricardo da Silva Torres – Wageningen University & Research
- Mark Dietrich – Bloodstone Consulting OU

Furthermore, feedback to the surveys was provided remotely from 3 additional experts:

- Yonah Welker - MIT CSAIL, EMEA Campus + EU/EIT projects]
- Jilles Vreeken - CISPA
- Tommaso Dorigo - Italian institute for nuclear physics (INFN)

1.4 Data analysis

In the context of the **initial survey**, we used a structured 5-star Likert relevance scale, where 5 denoted "Extremely Relevant" and 1 denoted "Extremely Irrelevant". The objective of this initial quantitative tracking was to identify potentially irrelevant items (relevance rating 2 or below) and give experts a chance to add additional relevant items before initiating open co-creation tracks in the interactive discussion round.

For the analysis of the final **Miro board**, we first extracted every individual text element (sticky notes) and analyzed the text cards using semantic proximity mapping (affinity grouping) looking at underlying systemic requirements, shared operational goals, dependencies etc. For instance, cards addressing “data governance”, “access control”, “secure sharing”, and “automated audit trails” were mapped together because they all represent systemic requirements for building trust in multi-site digital collaborations. Once items were sorted into loose groups, we synthesized these groups into corresponding thematic clusters, capturing the strategic innovation and research priorities established during the live session. Finally, we verified the clusters by cross-referencing them with the formal results of the multi-criteria decision analysis survey.

In the final **multi-criteria decision analysis survey**, experts ranked all research and innovation priorities identified previously in terms of three criteria (scientific impact – rated double, economic impact – rated double, feasibility/accountability – rated single) based on a 5-point Likert scale.

2 Workshop contents

2.1 Workshop agenda

Figure 1 below provides the details of the workshop agenda and sequence of the different sessions:

09:00 - 09:15	Welcome & introduction by SCIANCE
09:15 – 10:00	Review of relevant AI research and innovations (i.e. tools, applications, models, infrastructures) (plenary session)
10:00 – 10:15	Initial survey for checking relevance and enabling extension (plenary session)
10:15 – 10:45	Coffee Break
10:45 – 11:30	Interactive discussion: Identification of AI research & innovation priorities for enabling and/or improving experimental design & policy support
11:30 – 12:00	Consensus session: Ranked list of thematic AI for Science research priority areas for experimental design & policy support
12:00 – 12:30	Summary and closing statements by SCIANCE
12:30 – 13:30	Lunch

Figure 1 – Agenda for the Workshop “Experimental Design & Policy Support”

2.2 Welcome & introduction

The workshop opened with a welcome and introduction by Dr. Andrea Lösch representing the SCIANCE team. This initial session set the stage by introducing DFKI's workshop team, outlining the day's agenda, and explaining the SCIANCE co-creation process used to develop the Strategic Research and Innovation Agenda (SRIA). The corresponding workshop slides are provided below.



DFKI's Workshop Team



Dr. Tilman Becker
Senior Researcher
Director RICAIP
Centre of Excellence



Dr. Andrea Lösch
Senior Consultant
SCIANCE Team Lead



Dr. André Meyer-Vitali
Senior Researcher
Scientific Director
CERTAIN




Jessica Otto
Assistant Researcher
Applied Computer Science



RAISE – Resource for AI Science in Europe





2.3 Review of relevant AI research & innovations

The preparatory review for this workshop established the strategic framework for integrating AI into the scientific discovery cycle and the legislative framework at a European scale. The session explicitly framed **RAISE** as an essential mechanism to pool data, computing resources, and talent required for adaptive and closed-loop experimentation, noting that individual laboratories rarely possess the scale needed for self-contained, reliable AI-driven research.

The thematic overview covered three main areas:

- **The Shift to Adaptive Experimentation:** The presentation identified **Bayesian Optimisation (BO)** as the most mature AI component currently deployed for experimental design, particularly in drug discovery, materials science, and chemistry. BO models quantify experimental uncertainty probabilistically to select optimal future actions rather than relying on exhaustive trial-and-error. This approach forms the foundation for **Machine Learning-driven Design of Experiments (ML-DOE)**, which continually adapts experimental plans as new results arrive, dramatically reducing material costs, execution time, and environmental waste.
- **AI-Supported Hypothesis Generation:** The framework detailed emerging capabilities where Large Language Models (LLMs) and domain-specific AI agents act as cognitive assistants. These systems identify latent knowledge gaps in literature, draft templates, generate experimental designs, and tailor analysis plans. However, the session emphasized caution: LLM-based hypothesis generation remains unstandardized and must serve as a tool to support, rather than replace, narrow domain expertise.
- **Translational AI and Policy Support:** The final component explored how AI connects basic science with public decision-making networks. By combining multi-modal LLMs, knowledge graphs, and digital twins, translational research creates interactive advisory systems, risk assessments, and impact forecasts. The session concluded that embedding **Ethical, Legal, and Social Aspects (ELSA)** compliance, testable risk governance, and clear decision logs directly into the scientific workflow can turn European regulatory mandates into competitive assets for technological sovereignty.

All slides of this session are provided below.



European strategic relevance: Experimental design & policy support

- **RAISE** is positioned as a mechanism to pool data, compute and talent needed for adaptive and closed-loop experimentation.
- Federated research data and computing infrastructure are important because individual laboratories rarely have the scale needed for reliable AI-driven experimentation.
- Through the expert workshops, **SCIENCE** will provide the **research and innovation priorities for SRIA** for the different scientific disciplines and cross-cutting topics.
- Today's expert workshop will link **AI-guided experimentation and policy support** to European competitiveness and technological sovereignty.
- Overall, Europe's emphasis on trustworthy AI, risk governance and human oversight can become a competitive asset if embedded directly into SRIA and the upcoming research infrastructure.
- **Priority-setting question** today: Which research capacities and innovations should Europe develop as shared, open, auditable infrastructure in the context of **experimental design and policy support**?



1. Experimental design: The shift to adaptive experimentation

- **Bayesian optimisation (BO)** is presented as the most mature/deployed AI component for experimental design, especially in drug discovery, materials science and chemistry.
- BO models experimental outcomes probabilistically, quantifies uncertainty and selects promising conditions instead of relying on exhaustive trial-and-error.
- Active learning / adaptive design repeatedly trains a surrogate model and selects the next high-value experiment using acquisition functions such as uncertainty or expected improvement.
- **ML-driven Design of Experiments (ML-DOE)** extends classical DOE by adapting the experimental plan as new results arrive.
- Core scientific value: fewer experiments, lower time/material cost and potentially less hazardous or environmental waste





2. From optimization to AI-supported hypothesis generation

Emerging capacities

- LLMs identify knowledge gaps and suggest research directions.
- Domain-specific AI agents can automate parts of experiment setup.
- Systems can generate experiment designs from literature, templates and hypotheses.
- AI-generated analysis plans can be tailored to user data and objectives.

Maturity & caution

- LLM-based hypothesis generation is newer and less standardised than BO / active learning.
- Powerful in narrow, well-specified domains; not a general replacement for domain expertise.
- Generated hypotheses, protocols and plans require scientific validation.
- Potential is greatest when connected to evidence, data and experimental feedback.



3. Closed-loop experimentation: towards self-driving laboratories

- **Closed-loop workflow:** experimental data → AI model → next experiment → new data → model update.
- Autonomous agents can analyse outcomes and recommend changes to reaction conditions or protocols.
- Self-driving laboratories **combine AI planning with robotic execution**, potentially minimising human intervention between cycles.
- This creates a **bridge between experimental design and physical execution** — and makes the quality of laboratory automation a strategic enabler.
- The key direction identified by the review is **trusted autonomy**: independent action within defined boundaries, with safeguards and human accountability





4. Experimental design: key research challenges

Scientific / methodological

- Surrogate models depend on sufficient and representative data.
- Uncertainty estimates can be poorly calibrated in early-stage research.
- Adaptive designs complicate conventional pre-registration and fixed-protocol approaches.
- Best practices for documenting and auditing adaptive experimental trajectories are not yet converged


Operational / human

- AI recommendations can be enacted directly by robots, raising safety and cost risks.
- Human sanity-checking remains important, especially for hazardous experiments.
- Researchers need combined domain, statistics and ML literacy.
- Skills scarcity can become a bottleneck even where technology is available.



5. Interdisciplinary collaboration & policy support: from finding to impact

- AI for policy support sits at the interface between scientific results and action: treatment, product, regulation, policy or public decision.
- **Translational AI** can synthesise complex evidence and strengthen links between basic and applied disciplines — but depends on data quality, governance and sustained investment.
- **Knowledge translation and decision support** combine LLMs, multimodal LLMs, retrieval systems, knowledge graphs and digital twins.
- Typical **applications** include recommendations, interactive advisory systems, policy information, impact forecasts, risk assessments and scenario-based decisions.
- AI therefore has a potential role not only in producing knowledge, but in **connecting scientific knowledge across disciplines and into decision processes**





6. Policy support: governance becomes part of the scientific workflow

Emerging innovation needs

- Science-adapted AI risk assessment and verification tools.
- Structured mapping from legal risk categories to concrete technical checks.
- Staged ethics oversight aligned with AI-system maturity and potential impact.
- Decision logs, human approval and escalation paths for consequential uses.

Challenges for Europe

- Research increasingly spans academia, public bodies and industry; boundaries between research and application can be blurred.
- Abstract principles such as fairness or robustness must become testable practices.
- AI-assisted analysis at scale can create new statistical risks, including more false positives.
- Regulatory requirements remain dynamic, creating uncertainty for institutional governance investment.



Our initial AI research and innovations at a glance

AI research and innovations relating to experimental design

- Optimal Experimental Design (OED) / Bayesian Optimisation
- Sequential & batch active learning for surrogate models
- Adaptive Random Experiment Design (ARED) / Error-driven approaches
- ML-driven Design of Experiments (ML-DOE)
- LLM-based hypothesis generation and planning
- Autonomous/self-driving labs (closed-loop workflow)
- Data encoding & preprocessing methods

AI research and innovations relating to policy support

- Translational research
- Risk assessment technologies
- Ethics and governance technologies



2.4 Initial survey

Following the review, Jessica Otto introduced an online survey designed to extend and validate the initial research findings. Participants were asked to provide direct feedback by rating the relevance of identified technologies across

categories like Bayesian optimization, active learning and ML-driven Design of Experiments. This interactive step ensured that community insights and potential "outlier" innovations were captured to build a complete baseline list. The corresponding workshop slides are provided below.



Initial survey for extending and validating our findings




Review of relevant AI research and innovations

- Major findings so far and relevant AI research and innovations we have identified so far

Initial survey – your feedback matters now!

- Goal:** Extending and validating AI research and innovations we have identified so far
- Tool:** Online survey that names all relevant AI research and innovations identified for experimental design & policy support
- Survey conduct:**
 - Please rate the relevance (5 stars) of each AI research & innovation identified (5* highest relevance)
 - Please add any new AI research and innovations (including "outliers") they feel fit and also indicate perceived relevance
- Result:** Complete list of relevant AI research and innovations for experimental design & policy support with an indication of their perceived relevance









Initial survey for extending and validating our findings



2.5 Interactive discussion

Moderated by Dr. André Meyer-Vitali, the interactive discussion transitioned the workshop from a broad overview to a focused evaluation of priorities. Utilizing a [Miro board](#) managed by Jessica Otto and moderated by Dr. Andrea Lösch, participants jointly reviewed the compiled collection of innovations, analyzed and refined them, and discussed "crazy

ideas” in detail. This collaborative session successfully extended the initial findings into a complete set of relevant research and innovation priorities. The corresponding workshop slides are provided below.



Interactive discussion and consensus session

•Interactive discussion

- Starting point:** Complete collection of relevant AI research and innovations categorized by their perceived relevance
- What we will do now:** Joint review of this collection (i) checking for open issues/questions (e.g. relevance) and (ii) discussing outliers/“crazy ideas” in detail
- Result:** Final list of 15 research and innovation priorities based on relevance

•Consensus session

- Starting point:** Final list of 15 research and innovation priorities based on relevance
- To dos:** Lightweight multi-criteria decision analysis → Rank each AI research and innovation based on
 - Scientific impact (double weight)
 - Economic impact (double weight)
 - Feasibility/accountability (single weight)
- Result:** Final ranked list of top 15 research and innovation priorities relating to experimental design & policy support, taking into account their scientific impact, economic impact and feasibility/accountability




Bayesian optimization
Please, rate on each of the following criteria:

- Scientific impact: extent to which addressing this research and innovation will unlock or significantly accelerate AI-enabled research
- Economic impact: Net change in European companies' and organisations' economic performance
- Feasibility/accountability: Extent to which there is a realistic path through policy, infrastructure, or coordination

	1 (Very low)	2 (Low)	3 (Moderate)	4 (High)	5 (Very high)
Scientific impact	☐	☐	☐	☐	☐
Economic impact	☐	☐	☐	☐	☐
Feasibility/accountability	☐	☐	☐	☐	☐

Back Next

Powered by Formbricks





2.6 Consensus session

In the final consensus session, the top priorities were formally ranked using a lightweight multi-criteria decision analysis managed by Jessica Otto. Participants evaluated each research area through a structured Formbricks survey based on its scientific impact, economic impact, and feasibility at a European scale. The process resulted in a definitive, ranked list of strategic AI for Science priorities to feed directly into the upcoming SRIA and implementation roadmap. The relevant workshop slides are provided below.



Criteria definition & final survey


- Scientific impact:**
 - Definition:** Extent to which addressing this research and innovation will unlock or significantly accelerate AI-enabled research
 - Scores:** 1=very low, 2=low, 3=moderate, 4=high, 5=very high
- Economic impact:**
 - Definition:** The net change in European companies' and organisations economic performance – measured through productivity gains, cost reductions, value creation, and increases in competitiveness and revenues – resulting from this particular research and innovation
 - Scores:** 1=very low, 2=low, 3=moderate, 4=high, 5=very high
- Feasibility/accountability**
 - Definition:** Extent to which there is a realistic path and a need – through policy, infrastructure, or coordination – to address this particular research and innovation at European scale (e.g. there is no funding need if corresponding technologies/solutions do already exist)
 - Scores:** 1=very low, 2=low, 3=moderate, 4=high, 5=very high








Final survey for ranking research and innovations






2.7 Summary & closing statements


The workshop concluded with a final session led by Dr. Andrea Lösch to consolidate the day's findings and outline the immediate next steps. The team detailed plans to document the top 15 identified research and innovation priorities in an official workshop report and share them with the wider community via the RAISE Forum. Ultimately, these validated community insights will be integrated into the broader SCIANCE co-creation process to shape the upcoming AI in Science Strategic Research and Innovation Agenda (SRIA) and implementation roadmap.



Travel arrangements & reimbursement

-  Your travel costs to Ghent will be reimbursed based on the SCIANCE Travel policy for DFKI workshops
-  Reimbursement requires actual attendance - please do not forget to sign the attendance list today
-  Simply use the SCIANCE Reimbursement form for DFKI workshops to claim back your travel costs








SCIANCE **Summary and next steps**

- **Workshop results:** Joint definition of the top research and innovation priorities **for experimental design and policy support**
- **Next steps:**
 - **Consolidate workshop results** and document the identified Top Research & Innovation Priorities in a corresponding **workshop report**
 - Share the findings with the wider AI-in-Science community through the **RAISE Forum**
 - **Validate and further refine** the priorities based on **feedback from the broader community**
 - Feed the validated results into the SCIANCE process towards the **AI in Science SRIA** and implementation roadmap

**RAISE
Forum**




forum.raise4science.eu

3 Workshop results

3.1 Initial survey

Figure 1 below provides a screenshot of the initial survey.

Figure 1 – Initial survey extract



SCIANCE

Workshop "Experimental Design & Policy Support"

Thank you for your participation. The purpose of this questionnaire is to collect your feedback to assess the relevance of AI research and innovations identified for literature review.

Wenn Sie dieses Formular übermitteln, werden Ihre Details wie Name und E-Mail-Adresse nicht automatisch erfasst, es sei denn, Sie geben es selbst an.

* Erforderlich

1. Rate the relevance of Bayesian optimization * 

Extremely irrelevant ☆ ☆ ☆ ☆ ☆ Extremely relevant
2. Rate the relevance of Active learning/adaptive design * 

Extremely irrelevant ☆ ☆ ☆ ☆ ☆ Extremely relevant
3. Rate the relevance of ML-driven Design of Experiments (ML-DOE) * 


Extremely irrelevant ☆ ☆ ☆ ☆ ☆ Extremely relevant

As indicated earlier (see 1.2 Workshop methodology) the first part of the expert workshop focused on establishing an analytical baseline by assessing and extending a preliminary collection of artificial intelligence and laboratory automation technologies. Experts were asked to evaluate eight pre-mapped research and innovation priorities that had derived from desk research. For this initial evaluation, experts used a structured 5-star Likert relevance scale, where 5 denoted "Extremely Relevant" and 1 denoted "Extremely Irrelevant". The objective of this initial quantitative tracking was to identify potentially irrelevant items (relevance rating 2 or below) and give experts a chance to add additional relevant items before initiating open co-creation tracks in the interactive discussion round. Table 1 below illustrates the results of the initial survey.

Table 2 - Ranked Results from the initial survey

Rank	Topic & Score
1	Data encoding & preprocessing methods 4.14 Durchschnittliche Bewertung  Ebene 5  4 Ebene 4  1 Ebene 3  1 Ebene 2  1 Ebene 1
1	Translational research 4.14 Durchschnittliche Bewertung  Ebene 5  4 Ebene 4  1 Ebene 3  1 Ebene 2  1 Ebene 1
1	Ethics and governance technologies 4.14 Durchschnittliche Bewertung  Ebene 5  4 Ebene 4  2 Ebene 3 Ebene 2 Ebene 1  1


4 LLM-based hypothesis generation and planning



4 Autonomous/self-driving labs (closed-loop workflow)




4 Risk assessment technologies




7 Adaptive Random Experiment Design (ARED) / Error-driven approaches




8 ML-driven Design of Experiments (ML-DOE)



9 Optimal Experimental Design (OED) / Bayesian Optimisation



10 Sequential & batch active learning for surrogate models



As shown in Table 1 above, while there were no items that had to be removed because of low relevance (2 or below), four core tiers seemed to emerge with regard to the initial research and innovations items proposed, namely:

- **The Consensus Top Tier (Rank 1 Tiers):** Three distinct topics tied for the highest baseline relevance, each achieving an identical average score of 4.14 out of 5.00. Both **"Data encoding & preprocessing methods"** and **"Translational research"** captured identical vote splits, with 4 experts selecting max-relevance and the remainder distributed across the remaining levels. Crucially, **"Ethics and governance technologies"** matched this tier with 4 max-relevance votes and a tighter clustering at Level 4 (2 votes) and Level 1 (1 vote). This high rating shows that participants initially prioritized the data cleaning and administrative compliance layers as essential requirements for modern scientific research.
- **The Mid-Tier Clusters (Ranks 4–7):** A secondary cluster focused on automated generation and software loops. **"LLM-based hypothesis generation and planning," "Autonomous/self-driving labs (closed-loop workflow),"** and **"Risk assessment technologies"** all tied with a solid relevance average of 3.86. The data for self-driving labs showed a particularly strong consensus, with 6 out of 10 experts choosing high-relevance (Level 4). Mid-tier algorithmic variations followed closely, with **"Adaptive Random Experiment Design (ARED)"** scoring 3.57.
- **The Low-Tier Algorithmic Baselines (Ranks 8–10):** The lowest-rated items in the baseline survey were the core mathematical methodologies. **"ML-driven Design of Experiments (ML-DOE)"** averaged 3.43, while **"Optimal Experimental Design (OED) / Bayesian Optimisation"** dropped to 3.14. Tying at the bottom of the spectrum was **"Sequential & batch active learning for surrogate models"** with a score of 2.86. The expert comments indicated that these core statistical components are viewed as established algorithmic infrastructure. Consequently, they


scored lower because they represent standard, available tools rather than novel research priorities requiring new European funding frameworks.

The findings show that the priorities identified in each tier belong to both groups – experimental design as well as policy support – so there is no dominance of one or the other aspect of the research process in terms of priority.

3.2 Interactive discussion

As explained earlier (see 1.2 Workshop methodology), the interactive discussion was designed to expand the scope of the preliminary findings by allowing the seven experts to propose new concepts, identify structural challenges, and debate unconventional ideas that were all collected, tracked and jointly evaluated using a digital Miro board. Figure 2 below provides the outcomes of the interactive discussion with the yellow sticky notes representing the new items brought forward by the experts and the red sticky notes representing the initial research and innovations presented above (see 3.1 Initial survey).

Figure 2 – Miro board results of the interactive discussion



The actual grouping on the Miro board showed a clear subordination of specific mathematical methods beneath broader validation and governance layers. Core red baseline items from the literature review like *Optimal Experimental Design (OED)*, *active surrogate learning*, and *ML-DOE* were moved to the right part of the board. They were grouped directly under a comprehensive new yellow cluster focused on **"Uncertainty representation," "Design harness software infrastructure as a new common for AI,"** and **"Evaluation and benchmark platform as a common."** This indicates that individual machine learning loops are ineffective unless the EC funds shared software harnesses to track and evaluate statistical uncertainty across different laboratories.

The yellow sticky notes also introduced entirely new domains focused on cross-domain translation and socio-technical safety, which were missing from the initial desk baseline. Experts introduced critical yellow cards such as **"Causal inference in explainable AI," "ELSA-compliant experimental design,"** and **"Develop trustworthy science-driven agents."**

Rather than treating policy support as a simple text-prompting task, the panel added yellow notes focusing on structural methodology, including *"Adequacy of the chosen AI methods to the problem"* and *"Establish characterizations of the scientific 'search space'."*

3.3 Consensus session (MCDA)

Figure 3 below provides a part of the survey used for a multi-criteria decision analysis (MCDA) to conclude the final ranking of the research and innovation priorities.

Figure 3 – Extract from final survey used for multi-criteria decision analysis

Bayesian optimization

Please, rate on each of the following criteria:

- Scientific impact:** extent to which addressing this research and innovation will unlock or significantly accelerate AI-enabled research
- Economic impact:** Net change in European companies' and organisations' economic performance
- Feasibility/accountability:** Extent to which there is a realistic path through policy, infrastructure, or coordination

1 (Very low) 2 (Low) 3 (Moderate) 4 (High) 5 (Very high)

Scientific impact	☐	☐	☐	☐	☐
Economic impact	☐	☐	☐	☐	☐
Feasibility/accountability	☐	☐	☐	☐	☐

Back
Next

Powered by Formbricks

Table 2 below summarizes the outcomes of the final survey that was conducted as part of the consensus session, focussing on the final ranking of the research and innovation priorities identified in the context of the preceding interactive discussion.

Table 3 – Ranked results from the multi-criteria analysis (MCDA) survey

Rank	Topic	Responses (n)	Avg Scientific	Avg Economic	Avg Feasibility	Weighted score
1	Uncertainty representation	10	4.20	3.90	3.80	4.00
2	Causal inference in explainable AI	10	4.40	3.60	3.40	3.88
3	Adequacy of the chosen AI methods to the problem	10	3.90	3.90	3.70	3.86
4	Protocols for in silico-real world interoperability	10	4.10	3.90	3.20	3.84
5	Autonomous/self-driving labs	10	4.20	3.80	2.90	3.78
5	ELSA-compliant experimental design	10	3.90	3.50	4.10	3.78
7	Protocols for Human-AI interaction	10	3.80	3.60	3.90	3.74
7	Training scientists in the era of Agentic science	10	3.80	3.60	3.90	3.74
9	Low-cost computing incl. hardware (chip)	10	3.90	3.90	2.90	3.70
9	Develop trustworthy science-driven agents	10	3.90	3.70	3.30	3.70

11	Data encoding & preprocessing methods	10	3.90	3.30	3.90	3.66
12	Evaluation and benchmark platform as a common	10	4.00	3.30	3.60	3.64
13	Experiment (Design) semantic annotation (documenting design choices etc.)	10	3.70	3.40	3.90	3.62
13	Work in a sophisticated way with confidential data	10	3.70	3.80	3.10	3.62
15	Establish characterizations scientific "search space"	10	4.00	3.30	3.40	3.60
16	Physical AI incl. Robotics & AR for human assistance in Human-AI collaboration	10	3.80	3.60	2.90	3.54
17	Reduce complexity overall	10	3.30	3.80	3.20	3.48
17	Building scientific knowledge infrastructure	10	3.70	3.20	3.60	3.48
19	Establish characterizations scientific "search space"	10	4.00	3.00	3.30	3.46
20	AI solutions that fit with fixed procedural audit sampling or HR compliance work	10	3.00	3.80	3.60	3.44
21	Design harness software infrastructure as a new common for AI	10	3.40	3.30	3.70	3.42
22	AI tool that can handle financial audit & HR altogether (so that one does not need to apply several tools), such as statistical sampling, risk assessed frameworks, legal and regulatory procedures	10	2.80	3.90	3.50	3.38
23	ML-driven Design of Experiments (ML-DOE)	10	3.60	3.40	2.80	3.36
24	Risk assessment technologies	10	2.90	3.80	3.20	3.32
25	The value of pooling should be tested	10	3.80	2.80	3.30	3.30
26	LLM-based hypothesis generation and planning	10	3.50	3.20	3.00	3.28
27	Active learning/adaptive design	10	3.50	3.10	3.10	3.26
28	Policy support-modularization to allow use in many context	10	2.90	3.40	3.60	3.24
29	Audit sampling and risk assessment use of fixed statistical framework	10	2.80	3.40	3.70	3.22

30	Bayesian optimization	10	3.30	2.90	3.40	3.16
31	Multi-agent & swarm intelligence for representing complex systems	10	3.50	3.00	2.60	3.12
32	Translational research	10	3.10	3.20	2.70	3.06
33	Policy Support: methods/prompts to force justification of policy recommendation	10	2.60	2.90	3.30	2.86
34	Ethics and governance technologies	10	2.70	2.60	3.30	2.78

As becomes apparent from Table 2 above, the mathematical rankings confirmed the shift seen in the interactive discussion: **newly introduced expert priorities (originally yellow notes) swept the top spots, while the desk-review baseline items (originally red notes) dropped significantly in priority.** In addition to this general pattern, three other important patterns emerged from the survey results:

- A Shift from Pure Optimisation to Rigorous Validation (Ranks 1–3):** The expert panel’s emphasis on reliability was validated by the number 1 overall ranking of **"Uncertainty representation"** (Weighted Score: 4.00), which scored a high 4.20 for Scientific Impact and 3.80 for Feasibility. **"Causal inference in explainable AI"** ranked 2 with a weighted score of 3.88 and a maximum Scientific Impact rating of 4.40. **"Adequacy of the chosen AI methods to the problem"** secured place 3 (Weighted Score: 3.86). This top tier indicates that experts rejected black-box optimization loops and instead prioritized research into causal explanation models and uncertainty tracking.
- Decoupling of Hardware from Algorithmic Intelligence (Ranks 4–5):** The baseline red item **"Autonomous/self-driving labs"** tied for rank 5 overall with a weighted score of 3.78, matching the newly introduced yellow priority **"ELSA-compliant experimental design"**. However, the core mathematical models that drive these labs dropped drastically in the rankings: **"ML-driven Design of Experiments (ML-DOE)"** fell to rank 23 (Weighted Score: 3.36), and **"Bayesian optimization"** fell to rank 30 (Weighted Score: 3.16). This divergence indicates that standalone optimization formulas are viewed as mature tools. Strategic funding must instead target the high-level interoperability frameworks (**"Protocols for in silico-real world interoperability"**, ranked 4 at a weighted score of 3.84) and systemic ethical guardrails that govern them.
- Drop of Broad Compliance and General Baseline Tiers:** The most striking change occurred within the administrative compliance tracks. Although **"Data encoding & preprocessing methods"**, **"Translational research"**, and **"Ethics and governance technologies"** initially tied for first place in the baseline survey, they fell significantly during the comprehensive MCDA evaluation, landing at rank 11 (Weighted Score: 3.66), rank 32 (Weighted Score: 3.06), and rank 34 (Weighted Score: 2.78) respectively. Overall it hence seems that in the context of this expert workshop, abstract ethics statements (red notes) provide little practical value. Funding should instead be directed toward actionable, testable technical checks embedded directly into the code layer, such as active causal inference and verifiable uncertainty boundaries.

4 Conclusions and recommendations

Based on the above results (3 Workshop results), in particular the weighted MCDA scores, semantic proximity mapping of the Miro board, and the structural interaction between the baseline methodologies and newly contributed expert requirements, the following top 15 research and innovation priorities can be defined and grouped into four tiers:

Tier I: Trustworthiness	1: Rigorous Statistical Uncertainty Representation and Propagation across Discovery Loops
	<u>Strategic Focus:</u> Researching advanced mathematical frameworks to model, quantify, and track epistemic and aleatoric uncertainty throughout the entire data generation cycle, preventing downstream failure propagation in automated workflows. <u>MCDA Grounding:</u> Ranked as top priority (Rank 1) overall (Weighted Score: 4.00); driven by its foundational importance for scientific validation.

	2: Causal Inference and Structural Explainability in Scientific Artificial Intelligence <u>Strategic Focus:</u> Moving beyond predictive correlation models by funding research into causal machine learning architectures (such as structural causal models) to provide testable, explainable hypotheses. <u>MCDA Grounding:</u> Ranked 2nd overall (Weighted Score: 3.88); achieved the highest individual Scientific Impact score (4.40) in the entire portfolio.
	3: Methodological Adequacy and Algorithmic Fitness Assessment Framework(s) <u>Strategic Focus:</u> Developing automated diagnostic tools, benchmarking suites, and mathematical selection frameworks to evaluate whether a chosen AI model's structural assumptions natively match the specific physical parameters of a scientific problem. <u>MCDA Grounding:</u> Ranked 3rd overall (Weighted Score: 3.86); establishes a systematic safety check against the arbitrary application of black-box models.
	4: Shared Open Evaluation Platforms and Standardized Benchmark Commons <u>Strategic Focus:</u> Architecting continent-spanning, open-access platform infrastructures containing reference datasets, baseline simulators, and hold-out validation pipelines to benchmark new AI-driven discovery loops under strict, reproducible conditions. <u>MCDA Grounding:</u> Positioned at raw rank 12 (Weighted Score: 3.64); provides the operational platform required to scale the top three priorities.
	5: Cross-Domain Abstraction Protocols for In Silico-Real World Interoperability <u>Strategic Focus:</u> Formatting standardized programmatic APIs, semantic models, and hardware translation layers that allow purely virtual experiment designs to execute flawlessly on heterogeneous, remote laboratory robotics. <u>MCDA Grounding:</u> Ranked 4 th overall (Weighted Score: 3.84); bridges the operational gap between abstract data generation and physical runtime.
Tier II: Interoperability, Cyber-Physical Interfaces & Lab Design (Priorities 5–8)	6: Next-Generation Autonomous Closed-Loop Orchestration Systems (Self-Driving Laboratories) <u>Strategic Focus:</u> Integrating advanced Bayesian optimization, adaptive batch active learning, and surrogate model tuning into unified, self-governing software architectures that run end-to-end experimental design loops without manual human intervention. <u>MCDA Grounding:</u> Ranked 5 th overall (Weighted Score: 3.78); groups and upgrades the mature mathematical baseline elements (ML-driven Design of Experiments ML-DOE, Bayesian Optimisation BO, and active learning) into an active system orchestration layer.
	7: Semantic Annotation, Search Space Characterization, and Metadata Documentation <u>Strategic Focus:</u> Formulating semantic web structures, ontologies, and automated metadata recorders to explicitly document experiment design choices, structural boundaries, and tracking metrics within the discovery trail. <u>MCDA Grounding:</u> Combines raw ranks 13, 15, and 19; moves laboratory documentation from passive text logs into machine-actionable, discoverable search environments.
	8: Multi-Agent Swarm Intelligence for Representing and Simulating Complex Systems <u>Strategic Focus:</u> Researching decentralized, collaborative multi-agent simulation frameworks to model and map highly complex, emergent phenomena across multi-variable systems (such as climate models, cellular biology networks, or material degradation). <u>MCDA Grounding:</u> Ranked at raw rank 31 (Weighted Score: 3.12); repositioned from an isolated modeling application into a core simulation component for virtual experimentation.
	9: ELSA-Compliant Experimental Design and Algorithmic Guardrails <u>Strategic Focus:</u> Translating abstract Ethical, Legal, and Social Aspects (ELSA) guidelines into verifiable code modules and technical checks embedded directly into autonomous laboratory loops to guarantee compliance with European legal norms. <u>MCDA Grounding:</u> Ranked 5 th overall (Weighted Score: 3.78); fundamentally prioritizes testable compliance checking over passive, high-level ethics documentation.
Tier III: Translational AI, Decision Support & Compliance (Ranks 9–12)	10: Translational Decision Support Architectures and Complex Evidence Synthesis <u>Strategic Focus:</u> Combining large-scale retrieval systems, domain knowledge graphs, and digital twins into unified translational platforms capable of gathering distributed basic research and converting it into impact forecasts for public policymakers. <u>MCDA Grounding:</u> Regroups raw rank 32 with the initial slide baseline; shifts translational AI from retrospective literature reviews into active, forward-looking policy scenario evaluation.
	11: Multi-Modal Procedural Audit Sampling and Compliance Validation Loops <u>Strategic Focus:</u> Designing specialized AI verification tools optimized to perform compliance checks, legal assessments, and technical auditing against fixed procedural or operational guidelines. <u>MCDA Grounding:</u> Combines raw ranks 20, 22, and 29; groups fragmented compliance, financial, and HR audit requirements under a single, unified algorithmic checking tier.
	12: Modular Policy Support Engines with Automated Justification Prompting

	Strategic Focus: Building modular, highly context-aware policy recommendation modules that require the underlying AI algorithms to provide explicit mathematical, scientific, and empirical justifications for every proposed policy directive. MCDA Grounding: Consolidates raw ranks 28 and 33; forces algorithmic accountability inside public decision-making frameworks.
Tier IV: Infrastructure Strategy, Human Capital & Security (Priorities 13–15)	13: Advanced Protocols for Multi-Site Collaboration with Confidential Data <u>Strategic Focus:</u> Engineering secure, zero-trust network lines, federated data processing networks, and secure multi-party computation layers to enable collaborative, cross-border research without exposing sensitive intellectual property or restricted data sets. <u>MCDA Grounding:</u> Ranked at raw rank 13 (Weighted Score: 3.62); addresses the essential security requirements for distributed European research.
	14: Human-Capacity Capacity Building and Training for the Era of Agentic Science <u>Strategic Focus:</u> Formulating specialized, cross-disciplinary university curriculums and structural training modules to equip domain scientists with combined literacy in statistics, data tokenization, and multi-agent interaction. <u>MCDA Grounding:</u> Tied for rank 7 overall (Weighted Score: 3.74); backed by a high Feasibility rating of 3.90, showing immediate readiness for implementation across European nodes.
	15: Low-Cost Hardware Customization and Dedicated Science-Driven Computing Architectures <u>Strategic Focus:</u> Customizing energy-efficient edge processors, low-cost domain-specific chips, and open hardware architectures specifically tailored to execute complex surrogate modeling directly at the laboratory bench. <u>MCDA Grounding:</u> Combines raw ranks 9, 17, and 21; balances scientific capability with the strategic imperative to reduce total compute complexity, computing costs, and environmental carbon footprints.

As the research and innovation priorities emerging from the Miro board and the MCDA were partially merged and/or re-designed, the validation of the emerging research and innovation priorities presented above still needs to be carried out by the larger community.

5 Glossary of abbreviations

AI: Artificial Intelligence

API: Application Programming Interface

ARED: Adaptive Random Experiment Design

BO: Bayesian Optimisation

DFKI: German Research Centre for Artificial Intelligence (*Deutsches Forschungszentrum für Künstliche Intelligenz*)

ELN: Electronic Lab Notebook

ELSA: Ethical, Legal, and Social Aspects

EOSC: European Open Science Cloud

HPC: High-Performance Computing

LIMS: Laboratory Information Management System

LLM: Large Language Model

MCDA: Multi-Criteria Decision Analysis

ML-DOE: Machine Learning-driven Design of Experiments

OED: Optimal Experimental Design

RAISE: Resource for AI Science in Europe

RaaS: Robotics-as-a-Service

SRIA: Strategic Research and Innovation Agenda

XAI: Explainable Artificial Intelligence