CLAIR

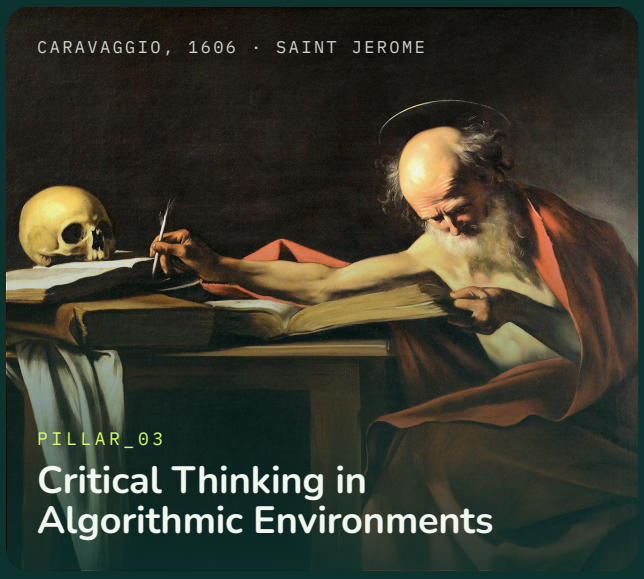
International Conference on
Leadership & AI in Research.

Research Agenda

Leadership and AI in Research









First Edition, 21st October 2026 COPENHAGEN 2026 TO BERLIN 2027

IN PARTNERSHIP WITH

Stiftelsen IMIT
Institute for Management of Innovation and Technology

westwerk. events

00 Contents

MAIN AGENDA

– Executive summary

01 1. Preface and purpose

02 2. Framework and method

03 3. Focus area: Scientific and engineering method

04 4. Focus area: Leadership governance and accountability

05 5. Focus area: Critical thinking in algorithmic environments

06 6. Focus area: Organisational and cultural transformation

07 7. Integrated synthesis

08 8. Recommendations for collaborative research

09 9. Development towards Berlin 2027

10 10. Sources and background materials

APPENDICES

A1 Appendix 1. Scientific and engineering method

A2 Appendix 2. Leadership governance and accountability

A3 Appendix 3. Critical thinking in algorithmic environments

A4 Appendix 4. Organisational and cultural transformation

HOW TO USE THIS DOCUMENT

The main agenda presents the purpose, method, research priorities and proposed next steps. Appendices 1–4 provide the underlying workshop syntheses for scrutiny and participant review. Each uses the same five requested outputs and includes an evidence trail, cross-area connections and unresolved issues.

Question references: **M** = Method; **G** = Governance; **C** = Critical Thinking; **O** = Organisational Transformation. Appendix section labels such as A1.2 and A2.2 identify sections within an appendix; they are distinct from research question codes.

Executive summary

The CLAIR Research Agenda builds on CLAIR's distinctive focus on leadership and AI in research. Following the inaugural conference in Copenhagen on 16–18 September 2026, it establishes a shared basis for investigating how AI changes research and the leadership it requires. Its next ambition is to translate these priorities into a practical CLAIR Leadership Framework, with CLAIR 2027 in Berlin providing the setting to develop and critically review a first version.

The initial empirical foundation comprises four facilitated group discussions involving 25 participants across sectors. Transcripts and the available group synopsis were synthesised into 16 research questions addressing four interconnected focus areas: scientific and engineering methods; leadership, governance and accountability; critical thinking; and organisational and cultural transformation. These are collaboratively generated priorities, formulated for participant review, rather than statements of formal consensus.

The central message is that realising AI's value in research requires scientific methods, human judgement, governance and organisational capabilities to develop together. Leaders need practical ways to decide where AI adds value, how its contributions should be evaluated, who remains accountable, and how teams retain the expertise and confidence to question its outputs.

Four proposed research directions translate these priorities into collaborative inquiry:

- **Research workflow evaluation:** comparing AI-supported workflows with existing practice to assess research quality, reproducibility, resource use and human oversight.
- **Governance in practice:** examining how responsibility, decision authority and oversight operate in everyday research.
- **Expertise and critical thinking:** tracking how AI use affects learning, independent judgement and the ability to recognise and challenge unreliable outputs over time.
- **Organisational implementation:** investigating how organisations move from successful pilots to sustained use, including changes in capabilities, working practices and realised value.

These directions will underpin the Leadership Framework, combining workflow evaluation checklists, responsibility and escalation guidance, critical scrutiny practices, and implementation approaches with measures of progress and value. Development will draw on research, organisational cases and exchange between researchers, leaders and practitioners.

CLAIR 2027 in Berlin will advance the transition from research agenda to leadership practice, focusing initially on life sciences and healthcare. Talks and interactive sessions will examine leadership decisions, supporting evidence, implementation challenges and transferable practices. The intended outcome is a first framework for application, evaluation and refinement in participating organisations, distinguishing evidence and experience from propositions requiring testing.

The agenda remains relevant across R&D sectors, with wider contributions exploring transferability and adaptation. It will evolve through participant review, organisational experience and research. This first edition, dated October 2026, will be publicly disseminated in November 2026 through CLAIR and partner channels and published on the CLAIR website. Projects, partnerships and resources remain to be agreed.

Invitation to contribute: We invite researchers, leaders and practitioners to refine the questions, share cases and decision tools, propose evaluation settings, and help shape the CLAIR Leadership Framework towards Berlin 2027.

01 Preface and purpose

This agenda is shaped by the contributions of the CLAIR community. We thank workshop participants and facilitators, speakers, dialogue and panel contributors, programme organisers and those supporting synthesis and review. Acknowledgement recognises contribution without implying endorsement of every interpretation or recommendation.

CLAIR occupies a distinctive position as a focused conference on leadership and AI in research. It brings together researchers, leaders and practitioners across public and private R&D to examine how leadership must evolve alongside AI. Its four connected perspectives - scientific and engineering methods, critical thinking, governance and organisational transformation - provide a shared framework for that inquiry. CLAIR translates the dialogue into a living research agenda and opportunities for collaboration.

Why a shared research agenda

Organisations face shared questions about AI delegation, evaluation, accountability and expertise. The agenda connects practical experience with academic inquiry to identify research needs and support collaboration. It serves researchers, public and private R&D organisations, educators and leaders, and aims to inform decisions by research funders and policymakers. Findings require contextual interpretation before transfer to other settings.

The workshop priorities should be considered alongside the European Commission's Group of Chief Scientific Advisors' opinion on AI uptake in science (2024) and the OECD's *Artificial Intelligence in Science* (2023). CLAIR contributes questions grounded in participants' experience, examining how leadership, governance and organisational conditions support responsible AI use in research. [11, 12]

A living document with several purposes

The agenda is a living framework for research priorities, public–private collaboration, organisational learning and policy dialogue. It will directly inform the programme, themes and collaborative activities of CLAIR 2027 in Berlin. This first edition, dated October 2026, will be publicly disseminated in November 2026 through CLAIR and partner channels and published on the CLAIR website.

Each edition should record what has changed, what evidence informed the change and which questions remain open. Contributions may challenge an existing interpretation, add a missing perspective or propose a study. Participation in the process does not imply endorsement of every statement. Differences should be documented where they affect the research questions or their application.

Scope and practical value

CLAIR's remit is leadership and AI in research across R&D sectors, including health and life sciences, energy and power, automotive, telecommunications and defence. The initial workshop material draws substantially on health and life sciences, reflecting the experience represented at CLAIR 2026.

Building on this foundation, CLAIR 2027 in Berlin, September 22-23, will focus primarily on leaders in life sciences and healthcare. Its ambition is to translate the Research Agenda into a practical CLAIR Leadership Framework, helping leaders evaluate AI-supported research, assign responsibility, sustain critical thinking and guide organisational change.

Research findings, concrete organisational cases and interactive workshops will inform a first version for subsequent application, evaluation and refinement. Proposed outputs include decision tools, evaluation checklists and implementation guidance, with a clear distinction between evidence-informed practices and propositions requiring further testing.

Contributions from other R&D sectors will help examine which leadership practices transfer across settings and which require adaptation. The Research Agenda will remain the foundation for continued inquiry, while the Leadership Framework develops its practical application.

Contributors and acknowledgement

We gratefully acknowledge those who contributed to the Copenhagen workshops, the wider conference dialogue, and the development and review of this agenda. Inclusion recognises contribution and does not imply endorsement of every interpretation or proposal.

Natalie Aleksić, Stockholm School of Economics; **Alexia Athanasopoloulou**, Erasmus University Rotterdam; **Hans Björkman**, Karlstad University; **Niklas Blomberg**, Innovative Health Initiative; **Iain Buchan**, University of Liverpool; **Peter Casteleyn**, European Institute of Innovation through Health Data (i-HD); **Robert Demir**, IMIT; **Ola Engkvist**, AstraZeneca; **Alexander Flaig**, Linköping University; **Tobias Fredberg**, Chalmers University of Technology; **Evgeny Galimov**, Phastar Inc.; **Fiona Gerster**, Roche; **Martin Ingvar**, Karolinska Institute; **Dipak Kalra**, i-HD; **Tom Metcalf**, i-HD; **Hyewon Nina Koo**, AstraZeneca; **Zoi Kolitsi**, i-HD; **Charlotta Kronblad**, Gothenburg University; **Alexander Krupp**, Bayer AG; **Nathan Lea**, i-HD; **Johan Luthman**, Lundbeck; **Magnus Lundbäck**, Getinge; **Luciano Lykkebo**, Westwerk Events; **Martin Löwendahl**, Gartner; **Claes Magnusson**, Malmö Higher Vocational College; **Aaron Mann**, Clinical Research Data Sharing Alliance; **Jörg Pareigis**, Karlstad University; **Johanna Pregmark**, Chalmers University of Technology; **Villim Prpić**, KTH Royal Institute of Technology; **Claudia Richter**, Bayer AG; **Anca Scortariu**, European Commission; **Thomas Senderovitz**, Danish Association of the Pharmaceutical Industry; **Thomas Steger-Hartmann**, Bayer AG; **Nakshathra Suresh**, EIRIS; **Aneta Tyszkiewicz**, EFPIA; **Petra Wilson**, Health Connect Partners.

OCTOBER 21, 2026

Mats Sundgren, Chair of the CLAIR Steering Committee

02 Framework and method

How the four focus areas connect

The four focus areas address interdependent conditions for credible and valuable AI-supported research. *Scientific and engineering methods* establish how knowledge is generated and validated. *Critical thinking* enables researchers to question evidence and assumptions. *Leadership and governance* establish direction, responsibility and authority to intervene. *Organisational culture and structures* shape whether people have the capabilities, incentives and confidence to exercise that judgement.

These relationships become tangible in everyday research: an AI-generated hypothesis requires a suitable validation method, researchers capable of challenging it, clear responsibility for accepting its conclusions, and a working environment that supports scrutiny. Changes in any one area therefore have implications for the others.

The proposed collaborative research directions investigate these connections through comparative evaluation of research workflows, studies of governance in practice, longitudinal studies of expertise and critical thinking, and research on organisational implementation. Each brings several focus areas together to examine how methods, judgement, responsibility and organisational conditions jointly influence research quality and value. Figure 1 illustrates these interrelationships; Sections 7 and 8 develop the synthesis and proposed research directions.

Generation of the workshop material

On 18 September 2026, 25 participants from across sectors took part in four groups, distributed as evenly as possible according to interest. Each group addressed one focus area and was supported by a facilitator. These actual arrangements, supplied by the conference coordinator, supersede the larger group numbers anticipated in earlier planning documents.

The common workshop design progressed from individual reflection through an exchange of experiences to joint synthesis. Groups were asked to identify three to five research questions or unresolved challenges, leadership and governance implications, methodological and organisational consequences, systemic barriers and enabling conditions, and directions for continued inquiry.

Discussions were recorded for transcription under the Chatham House rule arrangements reported by the coordinator. The supplied records use speaker codes in several transcripts but retain some identifying details. This First edition reports themes without attributing statements to individuals or affiliations. Source transcripts require a further check for residual identifiers before wider circulation.

Approach to the first synthesis

The first synthesis reviewed the four transcripts and the leadership synopsis against the organisers' five requested outputs. Related contributions were grouped into themes; tensions and differing views were retained. Themes were formulated as research questions and compared across focus areas to identify shared challenges. AI assisted this editorial process. The resulting questions are collaboratively generated research priorities, synthesised from the workshop dialogue and formulated as research questions for participant review. M, G, C and O identify areas; numbering does not imply rank.

Appendices 1–4 trace discussion themes to questions, implications, barriers and proposed studies through source locators. They distinguish workshop contributions from subsequent editorial proposals. Initial contributor feedback has begun to refine the agenda; wider participant review will check fidelity, omissions and differences of interpretation.

Evidence boundaries

Self-selection, uneven coverage and transcription errors limit interpretation. Workshop claims identify matters for inquiry rather than verified findings; wider-sector relevance requires further investigation. Slides establish framing, not plenary conclusions; no closing-panel transcript was used. Figure 1 maps the emerging priorities and shared concerns. The following sections develop the questions; Appendices 1–4 provide the evidence trail.

Four focus areas and emerging research priorities

PILLAR_01 · M

AI and the Scientific & Engineering Method

M1	Human–AI responsibilities
M2	Validation and reproducibility
M3	Expertise and scientific judgement
M4	Incentives and research quality

PILLAR_02 · G

Leadership, Governance & Accountability

G1	Governance embedded in practice
G2	Accountability in AI-supported work
G3	Assurance and transparency
G4	Leadership and shared purpose

Shared concerns

Expertise • Trust • Accountability • Learning

PILLAR_03 · C

Critical Thinking in Algorithmic Environments

C1	AI and critical thinking
C2	Expertise and decision context
C3	Practices for critical evaluation
C4	Learning and independent reasoning

PILLAR_04 · O

Organisational & Cultural Transformation in R&D

01	Trust, values and participation
02	From pilots to sustained use
03	Purpose and value assessment
04	Integration of expertise and roles

Collaboratively generated research priorities, synthesised from the workshop dialogue and formulated as research questions for participant review.
Short theme labels; full research questions appear in the agenda. Codes identify questions, not rankings.

FIGURE 1. Four focus areas and emerging research priorities. Collaboratively generated research priorities, synthesised from the workshop dialogue and formulated as research questions for participant review.



03 Focus area: Scientific and engineering method

AI can participate in data curation and preparation, question formulation, hypothesis generation, experimentation, analysis and interpretation. This focus area examines how scientific and engineering methods should develop as researchers delegate activities to AI while retaining responsibility for validation, reproducibility and scientific judgement. Appendix 1 documents the workshop synthesis and source trail; subsequent contributor input extends the inquiry to AI-assisted data preparation.

Workshop synthesis

Participants considered different roles for AI, ranging from a tool within a defined task to a more autonomous participant in a research workflow. These roles raise different questions about control, justification and evaluation. The discussion repeatedly returned to whether assurance should depend on understanding the process, testing the output, or a combination appropriate to the research context.

A related concern was the continued capacity of humans to oversee research. Delegating a task may alter opportunities to acquire expertise through practice. The group also questioned whether existing publication and career incentives support rigorous inquiry and meaningful outcomes, and how these incentives interact with AI-enabled increases in output.

Priority research questions

M1 **Human–AI responsibilities:** How should responsibilities be allocated between researchers and AI systems across the research process, taking account of context, capability and consequences?

M2 **Validation and reproducibility:** What validation and reproducibility practices are needed for AI-assisted and agentic research workflows, particularly when behaviour varies with context and user input?

M3 **Expertise and scientific judgement:** What expertise must researchers retain and develop to assess AI contributions, recognise limitations and intervene effectively?

M4 **Incentives and research quality:** How do research incentives influence the balance between AI-enabled speed, output volume and methodological quality, and what changes would support rigorous practice?

Leadership and governance implications

Leaders need to make delegation an explicit research decision, with clarity about purpose, validation and responsibility. Oversight should address the whole workflow as well as individual outputs. A proposed research priority is to identify when additional review, escalation or interruption improves quality, and how leaders can recognise those situations without unnecessarily restricting useful experimentation.

Question retained for review: How should strong empirical results be weighed against limited visibility into the process that produced them? The discussion did not resolve a single answer applicable to all research settings.

Methodological and organisational consequences

The discussion suggests that AI use should be examined within a complete research workflow. Contributor review highlighted AI-assisted data curation and preparation as a specific application of M2. Research should examine whether automated transformations improve data quality, semantic consistency and interoperability while preserving provenance and fitness for the intended research purpose. Evaluation should consider errors or biases introduced during preparation and their effects on subsequent findings.

Relevant features include the task delegated, the data and tools used, validation steps, changes made during the work and the basis for accepting a conclusion. The appropriate documentation will differ between a bounded analytical task and an adaptive sequence of research decisions.

Reproducibility also requires a clear object of assessment. Researchers may seek to reproduce an output, an analytical procedure or a broader scientific finding. A proposed study should specify which of these matters and how variation will be interpreted. Apparent consistency alone cannot establish validity; evaluation must relate to the scientific purpose.

Organisations also need to understand how changes in work affect the development of expertise. If early-career researchers perform fewer foundational tasks, alternative ways of building the judgement needed for later oversight may be necessary. This is a research question about learning and capability, not an assumption that all delegation causes skill loss.

Systemic barriers and enabling conditions

The group raised publication and performance incentives, concentration of access to data and computing, and the capacity of institutions to change. Measures centred on output volume may not capture scientific contribution. Equally, replacing them with new measures can create unintended incentives. Research should examine behaviours and outcomes rather than assume that additional metrics improve quality.

Within M4, a specific line of inquiry is how funding, research assessment and publication policies reward methodological rigour, validation, data stewardship and reproducibility as AI increases the volume and speed of research outputs.

Potential enabling conditions include access to suitable evaluation settings, interdisciplinary expertise and opportunities to compare workflows over time. These conditions need to be investigated across organisations with different resources, risk profiles and research objectives.

Proposed directions for collaboration

Comparative workflow studies: Public and private research teams could compare selected AI-supported workflows with relevant existing approaches. Evaluation could examine scientific validity, reproducibility, error detection, time, resource use and the quality of human intervention. These measures are proposals for study design, not workshop-endorsed standards.

Studies of researcher development: Universities and R&D organisations could investigate how different patterns of AI use affect the acquisition of domain knowledge, methodological judgement and the ability to detect failure. Longitudinal designs would help distinguish initial performance gains from changes in capability.

Connections to the other focus areas

Validation depends on critical thinking and appropriate expertise. Delegation raises governance questions about decision authority and accountability. Incentives, training and the opportunity to challenge a result are organisational conditions. Research on scientific method should therefore include these relationships rather than hold them outside the study boundary.



04 Focus area: Leadership governance and accountability

The discussion in Workshop Group 2 examined leadership when no individual can understand all the information or technical processes involved. Its emphasis was on governance as a practical, evolving part of organisational work, connected to purpose, trust and the ability to evaluate AI in context. The accompanying synopsis reinforces these themes. Appendix 2 documents the supporting synthesis and source trail.

Priority research questions

G1 **Governance embedded in practice:** How can AI governance be co-created, embedded in everyday organisational work and adapted through experience while supporting responsible innovation?

G2 **Accountability in AI-supported work:** How should responsibilities and accountability be allocated and maintained across those who select, introduce, use and oversee AI applications?

G3 **Proportionate assurance and transparency:** What combinations of explanation, documentation, validation and ongoing review provide sufficient assurance for different AI applications and organisational contexts?

G4 **Leadership capabilities and shared purpose:** What capabilities and practices enable leaders to establish shared purpose, recognise the limits of their knowledge and support responsible experimentation?

Leadership and governance implications

Participants questioned the expectation that leaders must personally know every answer. Building teams, drawing on distributed expertise and setting direction were presented as important complements to technical understanding. The practical challenge is to make responsibility clear while allowing decisions to draw on several forms of expertise.

Methodological and organisational consequences

Governance arrangements need to be studied where people use them. Facilitator feedback emphasised the need for dynamic governance models that evolve as organisational conditions, capabilities and risks change. This need extends beyond AI, although AI increases its urgency. Within G1, research could examine how organisations revise decision rights, oversight arrangements and working rules through experience while maintaining clear accountability.

The group discussed organisational sandboxes and evaluation of behaviour in context, including after deployment. Research could examine how concerns are raised, how exceptions are handled and how learning changes rules or working practices.

Barriers and enabling conditions

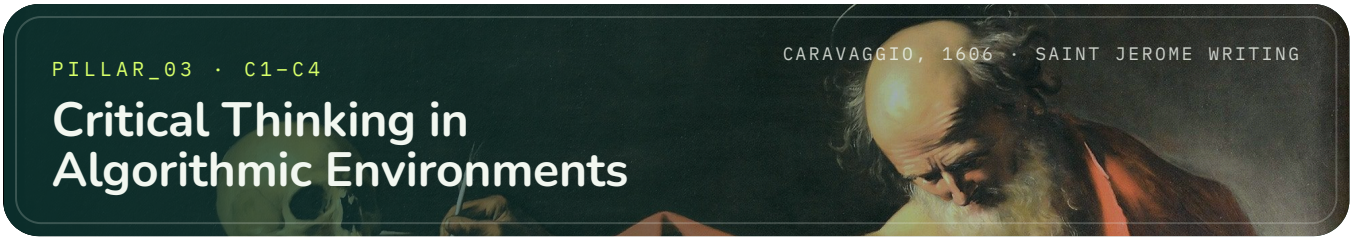
Rigid structures, fragmented responsibilities and limited interdisciplinary engagement may constrain adaptation. The synopsis also identifies funding and the need to understand the underlying organisational problem before adding AI. Useful enabling conditions may include shared purpose, participation in governance design and access to expertise beyond technical disciplines.

Proposed inquiry and unresolved tension

Comparative case studies could examine governance arrangements across organisations; bounded experiments could evaluate alternative review and escalation processes. Outcomes should include practical use, accountability and the effect on innovation. These proposed studies would investigate the tension between assurance and flexibility, including when outcome testing is adequate and when further explanation is necessary.

Contributor review identified a further line of inquiry within G1 and G3: how organisational governance interacts with external requirements, including funder conditions, professional standards and applicable regulation. Comparative studies across R&D sectors could examine where common guidance, standards or assurance approaches are useful and where differences in context require tailored arrangements.

A further question, sharpened through facilitator review, concerns the balance between AI-enabled production and human decision-making capacity. As AI makes analysis and other research outputs faster and less costly to produce, leaders may face more material and choices than they can meaningfully assess. Research within G2–G4 could examine decision-making models for this environment: what can be delegated, which decisions require human judgement, how attention is prioritised and when escalation is necessary. Studies should assess decision quality and leaders' capacity to exercise oversight, alongside speed.



05 Focus area: Critical thinking in algorithmic environments

The discussion in Workshop Group 3 considered how people judge increasingly accessible and convincing information. Participants explored both the risk of accepting AI outputs too readily and the possibility of using AI to encounter alternative perspectives. They also identified difficulties in agreeing what critical thinking means across disciplines. Appendix 3 documents the supporting synthesis and source trail.

Priority research questions

C1 **AI and the quality of critical thinking:** Under what conditions does AI strengthen or weaken critical thinking, scientific judgement and the quality of research decisions?

C2 **Expertise context and decision consequences:** How do domain knowledge, experience, time constraints and the consequences of a decision influence people's ability and willingness to question AI-generated outputs?

C3 **Practices for critical evaluation:** Which combinations of source verification, reflection, alternative perspectives and AI-supported challenge improve the evaluation of evidence and interpretation?

C4 **Learning and independent reasoning:** How can education and organisational practice develop and sustain independent reasoning as researchers increasingly work with AI?

Leadership and governance implications

Organisations need to make room for questioning and explain what level of scrutiny a task requires. Treating review as an individual responsibility alone overlooks workload, incentives and authority. A further research question is whether people can challenge an AI-supported conclusion when it is presented with organisational or professional authority.

Methodological and organisational consequences

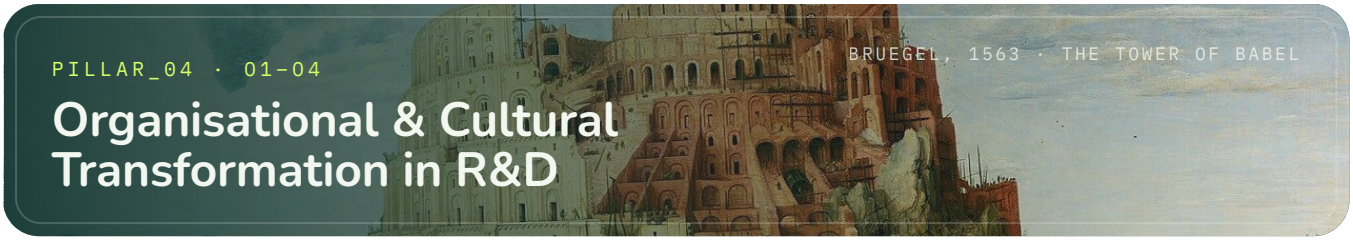
Research should define the aspect of critical thinking being assessed and distinguish it from confidence, fluency or task completion. The workshop's discussion of sources and explanations points to the need to test whether users can evaluate evidence, including when an explanation sounds plausible. Self-reported confidence would need to be considered alongside demonstrated judgement.

Barriers and enabling conditions

The discussion raised limited checking capacity, assumptions about AI reliability, confirmation of existing beliefs and disciplinary differences in terminology. Potential enablers include shared concepts, opportunities for reflection, access to underlying evidence and deliberate exposure to competing interpretations.

Proposed inquiry and unresolved tension

Studies could compare reasoning with and without AI or evaluate structured reflection and alternative-perspective prompts. Longer-term work could assess whether improvements persist when assistance is removed. These are proposed designs. The central tension remains that AI can expand the perspectives available while also making it easier to reinforce an existing view. Connections to expertise development, governance and learning culture should be examined explicitly.



06 Focus area: Organisational and cultural transformation

The discussion in Workshop Group 4 connected AI implementation with organisational values, trust, psychological safety and how people are treated through change. It also considered how organisations move beyond pilots, define value and integrate expertise that is often separated across disciplines or professional groups. [4] Appendix 4 documents the supporting synthesis and source trail.

Priority research questions

O1 **Trust values and participation:** How do organisational values, trust, psychological safety and employee participation shape the responsible adoption and use of AI in R&D?

O2 **From pilots to sustained implementation:** What combinations of capabilities, organisational structures and workflow changes enable AI initiatives to progress from successful pilots to sustained use?

O3 **Purpose and value assessment:** How should organisations define and assess the value of AI-enabled transformation, taking account of intended outcomes, context and realistic comparisons with existing practice?

O4 **Integration of expertise and organisational roles:** How should scientific, technical and organisational expertise be brought together, and responsibilities arranged, to support AI-enabled transformation?

Leadership and governance implications

Participants linked organisational values with concrete decisions about inclusion and development. A key issue is whether people are supported to contribute as work changes. Leaders also need to clarify the purpose of implementation and how competing short-term and longer-term objectives will be assessed.

Methodological and organisational consequences

Implementation research should examine roles, working relationships and learning processes alongside the technology. The discussion highlighted the difficulty of judging outputs when the basis for a recommendation is unclear. Trust therefore needs to be studied as a relationship among people, systems and institutional practices, rather than reduced to willingness to adopt a tool.

Barriers and enabling conditions

Professional silos, uncertain measures of success and the gap between pilots and implementation were recurring concerns. Participants valued settings in which different disciplines could work on a shared problem. They also questioned how organisations with limited resources can achieve useful change and avoid activities that produce little practical value.

Proposed inquiry and unresolved tension

Comparative implementation studies could track adoption, employee experience, working practices and realised value over time. Research could test whether structured exchanges across professional boundaries improve learning and implementation. The appropriate scope remains open: some questions may transfer across sectors, while others require close attention to health, life sciences or another specific R&D setting. CLAIR should examine this boundary through concrete cases.

07 Integrated synthesis

Taken together, the discussions suggest that scientific quality, human judgement, governance and organisational conditions are mutually dependent. The following synthesis identifies connections across the material. It is an editorial interpretation for participant review, rather than a record of agreement among all four groups. The question codes link this synthesis to the four appendices.

SHARED THEME	INTEGRATED RESEARCH CHALLENGE	QUESTIONS
Expertise and agency	Maintaining the capability and authority to question, redirect and take responsibility for AI-supported work.	M3 G4 C2 C1 C4
Trust and assurance	Combining validation, explanation, judgement and organisational trust in ways suited to the task.	M2 G3 C3 O1
Accountability in practice	Making decision authority and responsibility workable as tasks become distributed.	M1 G1 G2
Incentives and value	Rewarding credible knowledge and useful outcomes while recognising unintended effects of metrics.	M4 O3
Learning and adaptation	Updating methods, governance and capabilities as experience and technologies change.	G1 C4 O2
Interdisciplinary collaboration	Connecting different forms of expertise around concrete research and implementation problems.	G4 O4

TABLE 1. Shared themes and integrated research challenges

Connections that deserve explicit study

Retaining expertise depends on opportunities to practise, incentives to question and authority to intervene. Assurance depends both on technical evaluation and on people recognising the limits of that evaluation. Organisational value depends on what is measured and on whether learning can alter the original implementation plan. Studies that isolate only one of these elements may miss the conditions that explain their results.

Differences and unresolved tensions

The material contains competing views on the extent to which understanding a process is necessary when outputs can be tested. It also raises tensions between acceleration and reflection, experimentation and oversight, and common principles and context-specific practice. These differences should inform study design rather than be removed during consolidation.

A further question concerns whose outcomes count. Immediate productivity, long-term expertise, scientific contribution and wider public value may point towards different choices. Collaborative projects should make these objectives explicit and examine trade-offs. The agenda therefore needs both focused studies and comparisons that test how findings change across settings. Figure 2 connects the questions and shared challenges with four proposed directions for collaborative research.

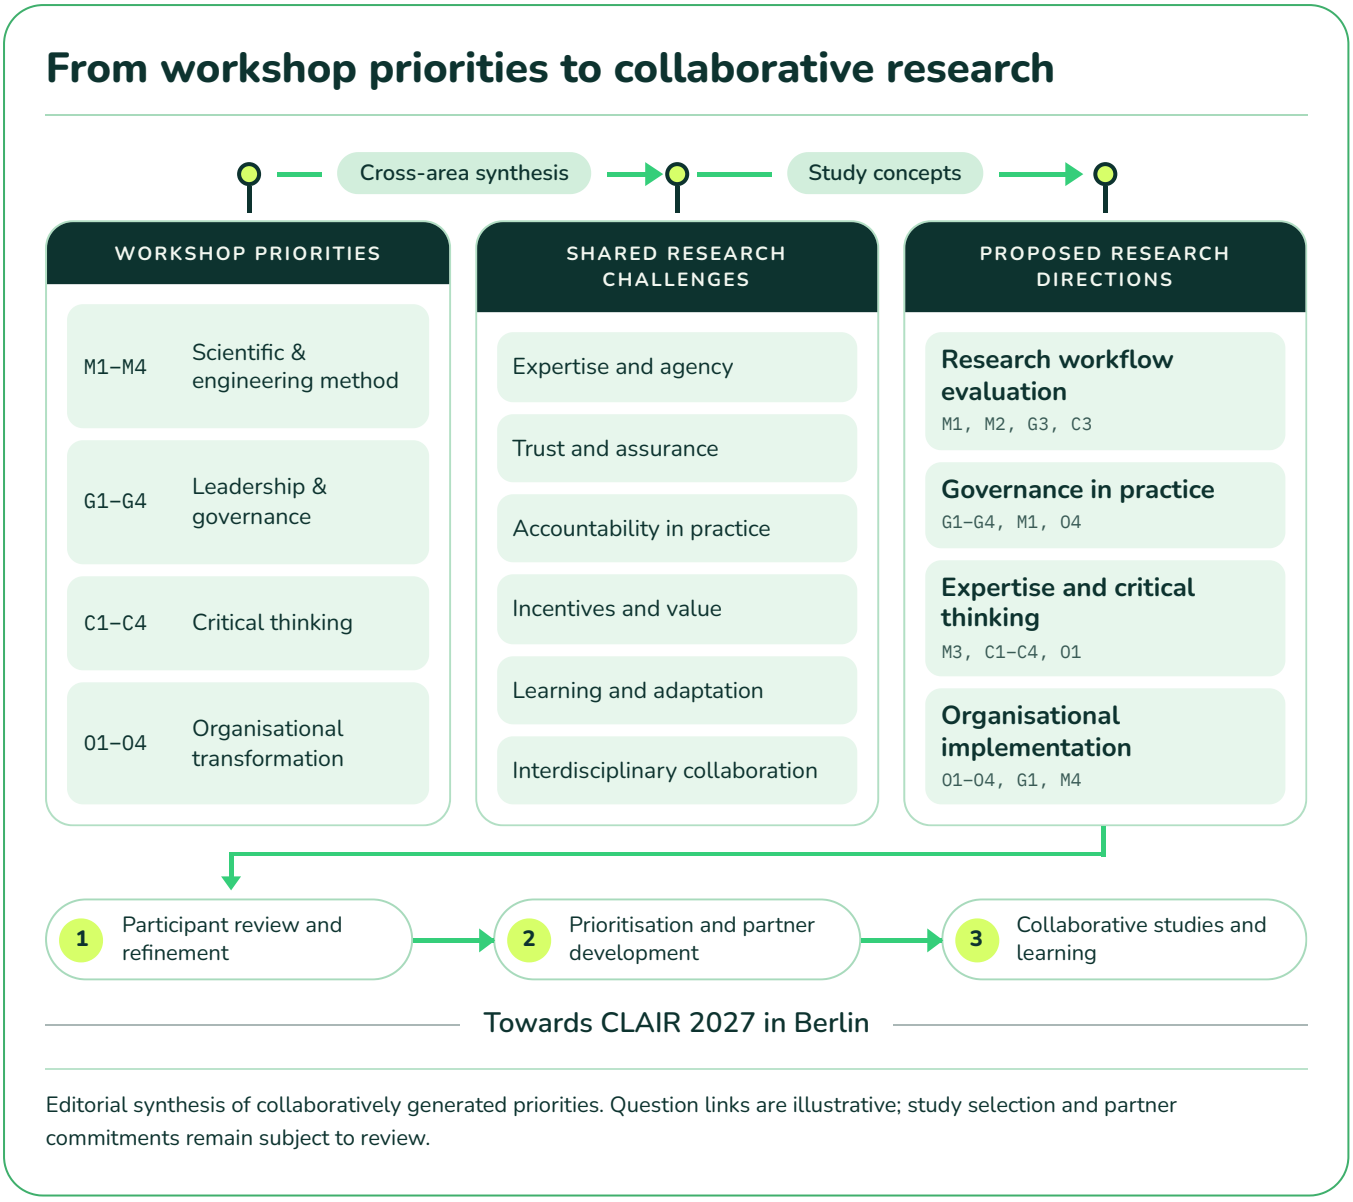


FIGURE 2. From workshop priorities to collaborative research. Links illustrate the editorial synthesis; proposed studies and programme directions remain subject to review.

08 Recommendations for collaborative research

The workshop synthesis provides a basis for developing a focused programme of collaborative research. The four directions below bring together questions across CLAIR's focus areas and translate them into possible studies. They are proposals developed through the synthesis, subject to participant review and prioritisation; they do not constitute agreed projects or institutional commitments. Figure 2 illustrates these connections.

Validate and prioritise the questions

Each workshop group should review its synthesis for accuracy, omissions and differences of interpretation. Wider consultation with CLAIR participants and representatives from additional R&D sectors should then test the relevance of the questions and identify opportunities for collaboration.

An initial research portfolio should be selected according to significance, researchability, feasibility, access to suitable settings and opportunities for public–private collaboration. The rationale for selection should be documented. Important questions that cannot yet be investigated should remain visible, together with the evidence, resources or partnerships needed to pursue them. Figure 3 summarises these research directions and their intended practical outputs.

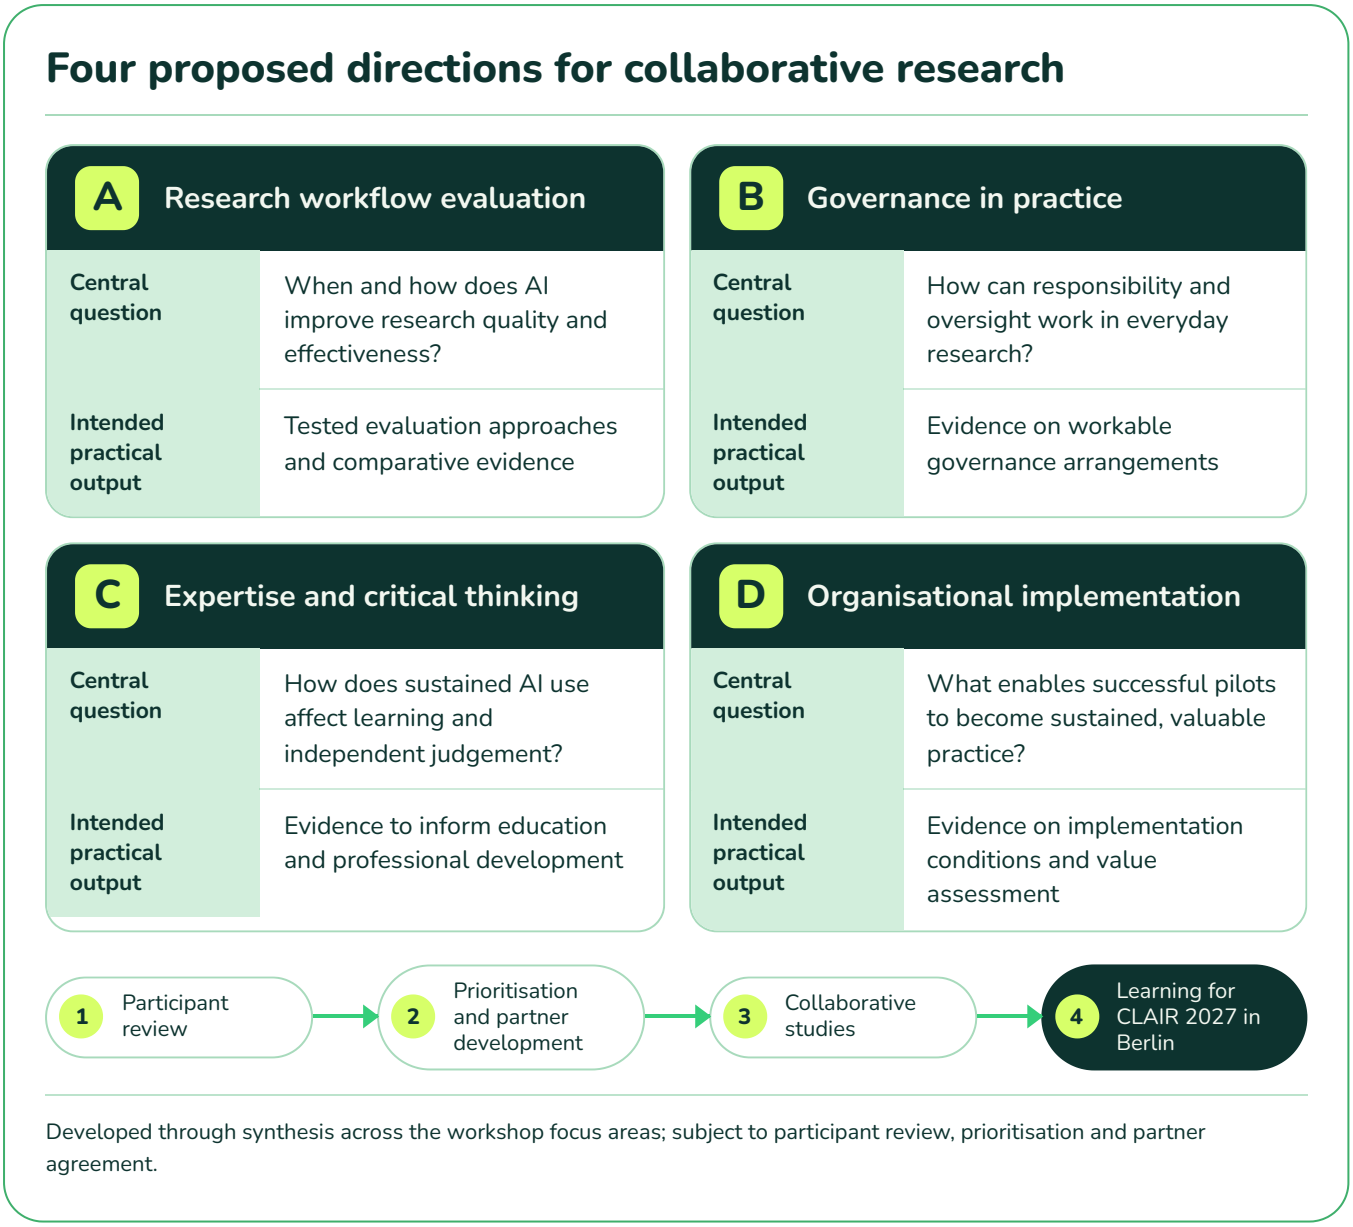


FIGURE 3. Four proposed directions for collaborative research, their intended practical outputs and the pathway towards CLAIR 2027 in Berlin. Research directions and subsequent activities remain subject to participant review, prioritisation and partner agreement.

A. Comparative research workflow evaluation: This direction examines how AI-supported research workflows compare with existing practice and under what conditions they improve research quality and effectiveness. Relevant activities include data curation and preparation, question formulation, hypothesis generation, experimentation, analysis and interpretation.

Academic and public or private R&D teams could compare defined workflows, specifying the research task, AI application, allocation of responsibilities and validation procedures. Evaluation should consider data quality and provenance, validity, reproducibility, error detection, time and resource use, alongside the quality of human oversight and intervention.

Intended practical output: Tested evaluation approaches and comparative evidence to help researchers and leaders decide where AI adds value, what validation is needed and when human intervention remains essential.

Illustrative links: M1, M2, G3 and C3.

B. Governance in practice: This direction examines how governance and accountability operate in everyday research. It asks how responsibilities are allocated, who has authority to question or interrupt AI-supported work, and how organisations respond to uncertainty, errors and unexpected outcomes.

Comparative case studies and bounded organisational experiments could examine decision authority, documentation, escalation, monitoring and the adaptation of governance through experience. Studies should assess whether these arrangements support both credible assurance and useful experimentation.

Studies should also examine how internal governance responds to funder conditions, professional standards and applicable regulation, identifying opportunities for shared guidance and assurance across sectors. Studies could also compare decision-making models under different levels of AI-generated workload, examining how delegation, prioritisation and escalation affect decision quality, leaders' workload and the timely exercise of human judgement.

Intended practical output: Evidence on workable governance arrangements, including approaches to assigning responsibility, escalating concerns and adapting oversight to different research contexts.

Illustrative links: G1–G4, M1 and O4.

C. Expertise and critical thinking: This direction investigates how sustained AI use affects the development and exercise of expertise, independent reasoning and scientific judgement. It addresses both opportunities to strengthen learning and the possibility that delegation reduces opportunities to acquire essential capabilities.

Longitudinal studies involving researchers, educators and employers could follow changes in knowledge, reasoning and decision-making over time. Evaluation should include the ability to assess unfamiliar outputs, recognise uncertainty, challenge recommendations and act when AI assistance is unreliable. Comparisons could examine different learning practices, levels of experience and patterns of AI use.

Intended practical output: Evidence to inform researcher education, professional development and working practices that sustain independent judgement and effective scrutiny of AI-supported research.

Illustrative links: M3, C1–C4 and O1.

D. Organisational implementation: This direction investigates what enables organisations to move from successful AI pilots to sustained and valuable use. It considers how capabilities, workflows, organisational roles, incentives and culture interact during implementation.

Management researchers and participating organisations could follow initiatives over time, comparing settings with different resources and institutional conditions. Studies should examine employee participation, trust, learning, workflow redesign and organisational ownership, alongside the value realised against the original research purpose and existing practice.

Intended practical output: Evidence on the conditions that support sustained implementation, together with approaches for assessing value and identifying why promising initiatives fail to become established practice.

Illustrative links: O1–O4, G1 and M4.

Develop feasible studies and partnerships

Each selected study concept should specify its research question, setting, proposed design, intended output, partner contributions and resource requirements. Partners should agree access to data and research settings, publication expectations and the handling of sensitive information before work begins.

Where relevant, research funders and policymakers should be invited as partners to help identify decision needs, implementation constraints and opportunities for comparative study. Policy-relevant outputs could include evidence to inform funding conditions, investment in shared data and research infrastructure, and proportionate assurance approaches. Partner roles and arrangements for independent analysis and reporting should be explicit.

Collaboration should draw on the expertise required by the question and, where useful, compare experiences across R&D sectors. Such comparisons should examine how findings depend on context, including research objectives, organisational resources and the consequences of failure. Specific projects, partners and funding remain to be agreed.

Connect research with shared learning and Berlin 2027

A light coordination process should enable participating teams to exchange methods, findings and implementation difficulties. Common terminology and selected comparable outcomes can support learning while preserving meaningful differences between settings.

The November 2026 publication and invitation to contribute will help identify additional research settings, expertise and potential partners. Reviewed priorities and emerging study concepts will inform possible tracks, panels and collaborative workshops at CLAIR 2027 in Berlin. Evidence and lessons from subsequent work will support further revisions of the living Research Agenda.



09 Development towards Berlin 2027

The CLAIR Research Agenda provides a foundation for continued inquiry and practical leadership development. CLAIR 2027 in Berlin, on 22–23 September 2027, will advance this work by bringing researchers, leaders and practitioners together to develop a first CLAIR Leadership Framework, with an initial focus on life sciences and healthcare.

The framework will translate the agenda's four interconnected focus areas into practical support for leadership decisions: evaluating AI-supported research, assigning responsibility, sustaining critical thinking and guiding organisational transformation. Research findings, concrete organisational cases and structured workshops will inform its development.

The ambition for Berlin is to develop and critically review an initial set of decision tools, evaluation checklists and implementation guidance for subsequent application and evaluation in participating organisations. These outputs will distinguish evidence-informed practices from propositions requiring further testing, making their intended use and limitations explicit.

The Research Agenda and Leadership Framework will develop together. Research will inform practical approaches, while organisational experience will reveal new questions and priorities. Contributions from other R&D sectors will help examine which practices transfer across settings and which require adaptation.

PROPOSAL · A LEARNED SOCIETY FOR LEADERSHIP AND AI IN RESEARCH

To provide continuity beyond individual conferences, CLAIR proposes the establishment of a learned society for leadership and AI in research. The society would provide direction for future meetings, steward the living Research Agenda and Leadership Framework, and support collaborative learning and research across R&D sectors. Its proposed membership, governance structure and stewardship arrangements will be developed in consultation with the CLAIR community and presented for discussion at CLAIR 2027 in Berlin.

Proposal contributed by Martin Ingvar, Karolinska Institute.

Researchers, leaders and practitioners are invited to contribute evidence, cases, practical tools and opportunities for evaluation, helping turn the Copenhagen dialogue into a shared resource for leadership practice.

10 Sources and background materials

The following sources and background materials informed the Research Agenda. Numbers in the text refer to this list. Titles are lightly normalised for readability; the original supplied files remain the underlying records.

[1] CLAIR Day 3 Scientific Method Table. Table 1 discussion transcript, 18 September 2026.

[2] CLAIR Day 3 Leadership and Governance. Table 2 transcript, with Synopsis Day 3 Leadership Table, 18 September 2026.

[3] CLAIR Day 3 Critical Thinking Table. Table 3 discussion transcript, 18 September 2026.

[4] CLAIR Day 3 Organisational Culture Transcript. Table 4, recording duration 1 hour 20 minutes 38 seconds, 18 September 2026.

[5] Day 3 Workshops. Presentation setting out the workshop process and five expected group outputs.

[6] CLAIR Day 3 Facilitator Pre Read. September 2026 materials describing facilitation, recording and synthesis arrangements.

[7] CLAIR Research Agenda Executive Summary. Pre-conference statement of purpose, scope and intended development.

[8] CLAIR 2026 Agenda Overview. Conference programme, Copenhagen, 16–18 September 2026.

[9] AI in the Backseat and supplied Day 1 and Day 2 dialogue materials. Conference framing and discussion prompts.

[10] Day 3 Opening Dialogue and Closing Panel decks. Framing questions for agenda development, feasibility and continued collaboration.

[11] European Commission, Group of Chief Scientific Advisors (2024). Successful and timely uptake of artificial intelligence in science in the EU. Scientific opinion.

[12] OECD (2023). Artificial Intelligence in Science: Challenges, Opportunities and the Future of Research. OECD Publishing, Paris.

APPENDIX 1 · M1–M4

Appendix 1: Theme - AI and the Scientific and Engineering Method

Workshop Group 1• Preliminary synthesis for participant review



A1.1 Purpose and analytical basis

This appendix examines how AI reshapes scientific inquiry, experimentation, knowledge creation and the relationship between human and machine-generated evidence. The source is the Table 1 discussion transcript. The analysis brings together related contributions within the five outputs requested by the organisers. It is a preliminary interpretive synthesis; the proposed questions and activities have not been formally ranked or endorsed by the group.

Its purpose is to show how discussion contributions inform the agenda and to support corrections, additions and refinement by participants. It presents themes without attributing statements to individuals or organisations. The common output framework is documented in source [5].

A1.2 Priority research questions and unresolved challenges

Collaboratively generated research priorities, synthesised from the workshop dialogue and formulated as research questions for participant review.

M1–M4 are question references, introduced during synthesis, not a ranking. The questions were neither prescribed before the workshop nor formally agreed in this wording by the group.

The central challenge is how scientific practice should develop when humans and AI systems jointly formulate questions, conduct research and evaluate evidence. Participants considered adapting research arrangements while retaining scientific rigour and avoiding assumptions that machines reason or behave like humans. [1, opening discussion]

M1 **Human–AI responsibilities.** How should responsibilities be allocated between researchers and AI systems across the research process, taking account of context, capability and consequences?

M2 **Validation and reproducibility.** What validation and reproducibility practices are needed for AI-assisted and agentic research workflows, particularly when behaviour varies with context and user input?

M3 **Expertise and scientific judgement.** What expertise must researchers retain and develop to assess AI contributions, recognise limitations and intervene effectively?

M4 **Incentives and research quality.** How do research incentives influence the balance between AI-enabled speed, output volume and methodological quality, and what changes would support rigorous practice?

A1.3 Leadership and governance implications

Resource validation as part of research. Participants questioned whether AI applications are sufficiently tested before deployment and observed that designing appropriate tests can require substantial effort. The synthesis proposes explicit resources, ownership and time for validation.

Clarify responsibility as roles change. Human–AI collaboration raises questions about who evaluates evidence, authorises decisions and responds to unexpected consequences.

Protect competence and quality. Concerns about learning scientific principles and incentives favouring output suggest examining how leaders sustain opportunities for critical assessment and rigorous practice.

These implications derive from the workshop material and require further investigation. [1]

A1.4 Methodological and organisational consequences

Three proposed directions follow: evaluate complete research workflows, including human review; specify what should be tested, documented and reproducible; and assess effects on expertise and unintended outcomes over time. These extend the discussion of contextual testing, reproducibility and researcher education. Their effectiveness remains to be investigated. [1]

A1.5 Systemic barriers and enabling conditions

The left column summarises discussion concerns; the right proposes corresponding conditions to investigate.

BARRIER OR TENSION		POTENTIAL ENABLING CONDITION
Application-specific testing is demanding.	→	Allocate validation resources and develop shared evaluation approaches.
Speed and output pressures can displace rigour.	→	Recognise testing, reproducibility and substantive scientific contribution.
Delegation may reduce opportunities to learn.	→	Provide supervised practice that demonstrates independent understanding.
Initial tests may miss wider consequences.	→	Continue evaluation and review unexpected outcomes.

Detailed financing, policy and regulatory mechanisms remain to be developed where the workshop material does not establish a specific proposal.

TABLE A1.1. Systemic barriers and potential enabling conditions

A1.6 Directions for continued inquiry and collaboration

PROPOSED ACTIVITY	PURPOSE	QUESTIONS
Comparative workflow studies	Compare allocation of responsibilities, research quality, reproducibility and human intervention.	M1 M2
Shared validation cases	Develop and compare contextual tests, documenting limitations and unexpected results.	M1 M2
Studies of expertise and incentives	Follow how AI use, training and organisational expectations influence methodological competence.	M3 M4

TABLE A1.2. Proposed directions for inquiry and collaboration

These activities are proposals developed through synthesis, not agreed projects. Public- and private-sector research organisations and relevant academic disciplines could develop them jointly, with emerging findings informing CLAIR 2027 in Berlin.

A1.7 From workshop dialogue to research questions

DISCUSSION EVIDENCE AND SOURCE LOCATOR	AGGREGATED THEME	REFERENCE
Humans and machines jointly formulate and execute scientific questions. `Source: Opening discussion on the scientific method as a human construct`	**Human–AI research arrangements**	M1
Agentic applications may be deployed before adequate contextual testing. `Source: Discussion of independent test data and customised systems`	**Validation and reproducibility**	M2
Young researchers need to learn core scientific principles. `Source: Passage on educating young researchers and PhD students`	**Development of expertise**	M3
Output incentives can compete with quality. `Source: Discussion of time constraints and incentives for output`	**Incentives for rigour**	M4

TABLE A1.3. From workshop dialogue to research questions

Evidence is paraphrased from source [1] in Section 10. Locators identify relevant transcript passages for participant review.

A1.8 Connections unresolved issues and participant review

Validation depends on critical thinking; delegation raises governance questions; and incentives and training are organisational conditions. Open issues include how far scientific methods require adaptation, what constitutes sufficient validation, and how assurance can be resourced. These questions connect M1–M4 with the other three focus areas.

Participant review will check fidelity, differences of view and omissions, and distinguish corrections from new proposals. M1–M4 remain provisional formulations of the collaborative dialogue pending that review.

Contributor feedback extended M2 to the validation of AI-assisted data curation and preparation, including data quality, semantic consistency, interoperability and provenance. It also sharpened M4 to examine funding, research assessment and publication policies that support rigour, data stewardship and reproducibility. These are consultation refinements to the agenda, rather than separately agreed workshop outputs.

APPENDIX 2 · G1–G4

Appendix 2: Theme - Leadership Governance and Accountability

Workshop Group 2 • Preliminary synthesis for participant review



A2.1 Purpose and analytical basis

This appendix examines leadership, governance and accountability as AI becomes embedded in research-intensive organisations. The sources are the Table 2 transcript and accompanying group synopsis. The synopsis provides a contemporaneous consolidation of the discussion; the research questions below are subsequent editorial formulations. Both sources are organised against the five requested outputs. [2]

Its purpose is to show how discussion contributions inform the agenda and to support corrections, additions and refinement by participants. It presents themes without attributing statements to individuals or organisations. The common output framework is documented in source [5].

A2.2 Priority research questions and unresolved challenges

Collaboratively generated research priorities, synthesised from the workshop dialogue and formulated as research questions for participant review.

G1–G4 are question references, introduced during synthesis, not a ranking. The questions were neither prescribed before the workshop nor formally agreed in this wording by the group.

The central challenge is how leaders establish purposeful, accountable AI use when they cannot personally command all the relevant knowledge. Participants emphasised shared inquiry, organisational context and recognition of knowledge limits. The synopsis proposed governance embedded in everyday work and adapted through experience. [2]

G1 **Governance embedded in practice.** How can AI governance be co-created, embedded in everyday organisational work and adapted through experience while supporting responsible innovation?

G2 **Accountability in AI-supported work.** How should responsibilities and accountability be allocated and maintained across those who select, introduce, use and oversee AI applications?

G3 **Proportionate assurance and transparency.** What combinations of explanation, documentation, validation and ongoing review provide sufficient assurance for different AI applications and organisational contexts?

G4 **Leadership capabilities and shared purpose.** What capabilities and practices enable leaders to establish shared purpose, recognise the limits of their knowledge and support responsible experimentation?

A2.3 Leadership and governance implications

Establish the purpose of AI use. Participants questioned broad encouragement to adopt AI without understanding the problem, application or consequences. Leadership needs a clear connection between AI use and organisational purpose.

Make governance useful in everyday work. The synopsis proposed lived, co-created governance. The research task is to understand how guidance can support decisions and help resolve uncertainty.

Combine direction with humility and visibility. Leaders need to acknowledge knowledge limits and understand actual AI use, including the effects of deadlines and local working practices.

These implications derive from the workshop material and require further investigation. [2]

A2.4 Methodological and organisational consequences

Proposed research should study governance where decisions are made, investigate bounded experimentation through organisational sandboxes, and bring social-science and humanities perspectives into governance design. The synopsis explicitly identifies sandboxes and interdisciplinary participation. Different applications require attention to their own context, risks and responsibilities. [2]

A2.5 Systemic barriers and enabling conditions

The left column summarises discussion concerns; the right proposes corresponding conditions to investigate.

BARRIER OR TENSION		POTENTIAL ENABLING CONDITION
Governance can feel like an additional task.	→	Co-create guidance integrated into workflows and revised through experience.
Leadership may lack visibility into actual use.	→	Maintain dialogue about applications, working practices and emerging problems.
Rigid structures and professional boundaries restrict adaptation.	→	Create interdisciplinary settings for bounded experimentation.
Limited explainability creates uncertainty.	→	Investigate contextual combinations of output testing, documentation and validation.
Funding and institutional incentives constrain change.	→	Resource evaluation and examine publication and reward practices.

TABLE A2.1. Systemic barriers and enabling conditions

Detailed financing, policy and regulatory mechanisms remain to be developed where the workshop material does not establish a specific proposal.

A2.6 Directions for continued inquiry and collaboration

PROPOSED ACTIVITY	PURPOSE	QUESTIONS
Governance in everyday work	Compare how organisations embed governance, allocate responsibility and respond to problems.	G1 G2
Collaborative governance sandboxes	Test arrangements for documentation, explanation, validation and oversight.	G1–G3
Leadership learning studies	Study contextual understanding, shared purpose, knowledge limits and responsible experimentation.	G4

TABLE A2.2. Directions for continued inquiry and collaboration

These activities are proposals developed through synthesis, not agreed projects. Public- and private-sector research organisations and relevant academic disciplines could develop them jointly, with emerging findings informing CLAIR 2027 in Berlin.

A2.7 From workshop dialogue to research questions

DISCUSSION EVIDENCE AND SOURCE LOCATOR	AGGREGATED THEME	REFERENCE
Governance should be dynamic, lived and co-created. `Source: Synopsis item 1`	**Governance as evolving practice**	G1
Leadership encouragement can outpace understanding of actual uses. `Source: Transcript discussion of AI use under organisational deadlines; synopsis item 3`	**Visibility and accountability**	G2
Acceptable explainability depends on use, assumptions and validation. `Source: Synopsis item 5`	**Contextual assurance**	G3
Leaders cannot know everything and need a common purpose. `Source: Opening transcript discussion; synopsis item 2`	**Leadership for shared inquiry**	G4

TABLE A2.3. From workshop dialogue to research questions

Evidence is paraphrased from source [2] in Section 10. Locators identify relevant transcript passages and synopsis items for participant review.

A2.8 Connections unresolved issues and participant review

Assurance links to scientific validation, recognition of limitations to critical thinking, and governance in everyday behaviour to organisational culture. Open tensions concern innovation and oversight, sufficient explanation, and humility alongside decisive action. Participants also questioned whether formal training alone could address leadership gaps. [2]

Participant review will check fidelity, differences of view and omissions, and distinguish corrections from new proposals. G1–G4 remain provisional formulations of the collaborative dialogue pending that review.

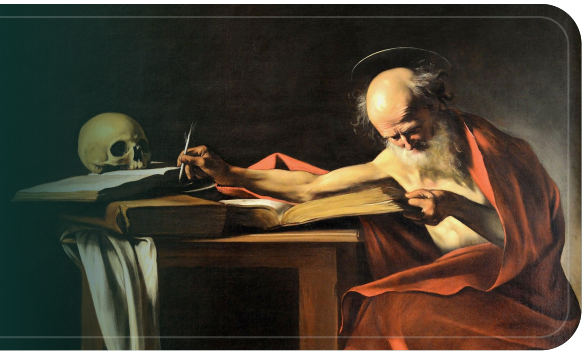
Contributor feedback extended the inquiry within G1 and G3 to the relationship between organisational governance and external requirements, including funder conditions, professional standards and applicable regulation. This includes examining opportunities for common guidance and assurance across sectors. This extension arose during consultation and remains subject to wider review.

Facilitator review also emphasised dynamic governance as a broader organisational need intensified by AI. It proposed further inquiry into decision-making models where AI-enabled production may outpace human capacity for assessment and judgement. This refinement draws on the facilitator's reflection on both the group discussion and wider conference dialogue and remains subject to participant review.

APPENDIX 3 · C1–C4

Appendix 3: Theme - Critical Thinking in Algorithmic Environments

Workshop Group 3 • Preliminary synthesis for participant review



A3.1 Purpose and analytical basis

This appendix examines how scientific judgement, intellectual independence and critical thinking can be maintained and strengthened as AI becomes embedded in research. The source is the Table 3 discussion transcript. Related contributions are grouped within the five requested outputs. Contrasting views are retained, particularly on generational differences and the potential for AI itself to support critical challenge. [3]

Its purpose is to show how discussion contributions inform the agenda and to support corrections, additions and refinement by participants. It presents themes without attributing statements to individuals or organisations. The common output framework is documented in source [5].

A3.2 Priority research questions and unresolved challenges

Collaboratively generated research priorities, synthesised from the workshop dialogue and formulated as research questions for participant review.

C1–C4 are question references, introduced during synthesis, not a ranking. The questions were neither prescribed before the workshop nor formally agreed in this wording by the group.

The central challenge is how people evaluate both AI-generated information and its significance for decisions. Participants distinguished checking correctness from understanding meaning, questioned the capacity available for verification and explored AI-supported alternatives and challenge. These possibilities remain matters for investigation. [3]

C1 **AI and the quality of critical thinking.** Under what conditions does AI strengthen or weaken critical thinking, scientific judgement and the quality of research decisions?

C2 **Expertise context and decision consequences.** How do domain knowledge, experience, time constraints and the consequences of a decision influence people's ability and willingness to question AI-generated outputs?

C3 **Practices for critical evaluation.** Which combinations of source verification, reflection, alternative perspectives and AI-supported challenge improve the evaluation of evidence and interpretation?

C4 **Learning and independent reasoning.** How can education and organisational practice develop and sustain independent reasoning as researchers increasingly work with AI?

A3.3 Leadership and governance implications

Make questioning possible. Participants suggested that scepticism outside work does not always transfer into professional settings. Research should examine whether people feel able to challenge information and established positions.

Relate scrutiny to decision consequences. The significance of a decision matters for the level of attention and review it warrants. Appropriate allocation of scrutiny remains a research question.

Recognise human limitations and avoid age assumptions. Defensiveness and overconfidence can restrict challenge. Concerns about younger researchers were countered by observations about scepticism developed through digital experience.

These implications derive from the workshop material and require further investigation. [3]

A3.4 Methodological and organisational consequences

Proposed studies should distinguish verification from interpretation, compare individual review with peer discussion and AI-supported critique, and assess reasoning alongside final answers. Relevant outcomes include recognition of uncertainty and revision of judgement when evidence changes. Participants discussed requesting sources and explanations, but their availability alone does not establish effective evaluation. [3]

A3.5 Systemic barriers and enabling conditions

The left column summarises discussion concerns; the right proposes corresponding conditions to investigate.

BARRIER OR TENSION		POTENTIAL ENABLING CONDITION
Ready-made answers may be accepted without questioning.	→	Build reflection on evidence and assumptions into work.
Checking requires time and relevant knowledge.	→	Allocate review capacity according to the decision and context.
Defensiveness and existing beliefs restrict challenge.	→	Use structured consideration of alternative explanations.
Critical skills may not transfer into professional practice.	→	Connect existing scepticism with scientific evaluation through learning.
Generational assumptions obscure actual capability.	→	Assess demonstrated reasoning, experience and task context.

TABLE A3.1. Systemic barriers and enabling conditions

Detailed financing, policy and regulatory mechanisms remain to be developed where the workshop material does not establish a specific proposal.

A3.6 Directions for continued inquiry and collaboration

PROPOSED ACTIVITY	PURPOSE	QUESTIONS
Comparative evaluation studies	Compare human-only, AI-assisted and peer-review approaches to evidence assessment and interpretation.	C1 C3
Studies across decision contexts	Examine expertise, time, expectations and consequences as conditions for scrutiny.	C2 C3
Longitudinal learning studies	Follow how AI use and learning practices affect independent reasoning over time.	C1 C4

TABLE A3.2. Directions for continued inquiry and collaboration

These activities are proposals developed through synthesis, not agreed projects. Public- and private-sector research organisations and relevant academic disciplines could develop them jointly, with emerging findings informing CLAIR 2027 in Berlin.

A3.7 From workshop dialogue to research questions

DISCUSSION EVIDENCE AND SOURCE LOCATOR	AGGREGATED THEME	REFERENCE
AI may encourage acceptance while also introducing alternative perspectives. `Source: Discussion of ready-made answers and AI agents assigned a critical role`	**Conditions for critical thinking**	C1
Knowledge, checking capacity and decision consequences matter. `Source: Source-checking discussion and later comparison of decision stakes`	**Context for judgement**	C2
Correctness and meaning are distinct questions. `Source: Exchange asking whether information is correct and what it means`	**Evidence and interpretation**	C3
Generational views differ; scepticism may need to transfer into work. `Source: Exchange contrasting concerns about younger researchers with digital-life experience`	**Learning independent reasoning**	C4

TABLE A3.3. From workshop dialogue to research questions

Evidence is paraphrased from source [3] in Section 10. Locators identify relevant transcript passages for participant review.

A3.8 Connections unresolved issues and participant review

Critical thinking supports scientific validation, the exercise of accountable judgement and organisational learning. Open questions concern when AI-supported challenge improves judgement, when it produces false confidence and how much scrutiny different decisions require. Contrasting views about generations and AI's contribution should be investigated rather than resolved by assumption. [3]

Participant review will check fidelity, differences of view and omissions, and distinguish corrections from new proposals. C1–C4 remain provisional formulations of the collaborative dialogue pending that review.

APPENDIX 4 · 01–04

Appendix 4: Theme - Organisational and Cultural Transformation in R&D

Workshop Group 4 • Preliminary synthesis for participant review



A4.1 Purpose and analytical basis

This appendix examines the capabilities, structures, behaviours and cultural conditions needed to integrate AI responsibly and productively into research and development. The source is the Table 4 transcript, with a recorded duration of 1 hour 20 minutes 38 seconds. Contributions are grouped within the five requested outputs. Timestamps in the evidence table provide locators without identifying speakers. [4]

Its purpose is to show how discussion contributions inform the agenda and to support corrections, additions and refinement by participants. It presents themes without attributing statements to individuals or organisations. The common output framework is documented in source [5].

A4.2 Priority research questions and unresolved challenges

Collaboratively generated research priorities synthesised from the workshop dialogue and formulated as research questions for participant review.

O1–O4 are question references, introduced during synthesis, not a ranking. The questions were neither prescribed before the workshop nor formally agreed in this wording by the group.

The central challenge is translating AI capabilities into sustained improvements while maintaining trust, involving people and adapting working practices. Participants connected organisational values with how people are treated, questioned why successful pilots fail to become lasting practice and argued for starting with the intended transformation. [4, 00:01–07:13]

O1 **Trust values and participation.** How do organisational values, trust, psychological safety and employee participation shape the responsible adoption and use of AI in R&D?

O2 **From pilots to sustained implementation.** What combinations of capabilities, organisational structures and workflow changes enable AI initiatives to progress from successful pilots to sustained use?

O3 **Purpose and value assessment.** How should organisations define and assess the value of AI-enabled transformation, taking account of intended outcomes, context and realistic comparisons with existing practice?

O4 **Integration of expertise and organisational roles.** How should scientific, technical and organisational expertise be brought together, and responsibilities arranged, to support AI-enabled transformation?

A4.3 Leadership and governance implications

Begin with intended change. Participants proposed asking what the organisation wants to achieve and which AI capabilities would support it. This connects technology choices with purpose.

Express values through decisions about people. Inclusion, learning and support during change were discussed as practical expressions of values. Their effects on trust and participation require examination.

Connect trust and responsibility with actual work. Participants explored the ability to question analysis and clarify accountability. They also questioned arrangements that separate AI strategy from research needs, people and processes.

These implications derive from the workshop material and require further investigation. [4]

A4.4 Methodological and organisational consequences

Proposed research should follow implementation beyond the pilot, examine skills, processes and tools together, and evaluate value against actual practice and user needs. Combining implementation measures with evidence about people's experience could explain how organisational outcomes arise. The discussion also identified workflow redesign and reuse of data and models as relevant. [4, 05:08–19:14]

A4.5 Systemic barriers and enabling conditions

The left column summarises discussion concerns; the right proposes corresponding conditions to investigate.

BARRIER OR TENSION		POTENTIAL ENABLING CONDITION
Successful pilots do not become lasting practice.	→	Examine workflow changes, organisational ownership and continued support.
Unfamiliar analysis can undermine trust.	→	Enable questioning, examination of evidence and clear responsibility.
People can experience change as imposed.	→	Support participation and learning consistent with organisational values.
AI responsibility can become detached from research.	→	Connect scientific teams, technical specialists and implementation roles.
Unchanged workflows and repeated work reduce value.	→	Redesign processes and support appropriate data and model reuse.
Unrealistic comparisons can distort evaluation.	→	Assess against actual practice, intended purpose and user needs.

TABLE A4.1. Systemic barriers and enabling conditions

Detailed financing, policy and regulatory mechanisms remain to be developed where the workshop material does not establish a specific proposal.

A4.6 Directions for continued inquiry and collaboration

PROPOSED ACTIVITY	PURPOSE	QUESTIONS
Longitudinal implementation studies	Follow pilots into routine use, examining capabilities, workflows, ownership and outcomes.	O2 O4
Studies of trust and participation	Compare how involvement, learning and communication influence trust and implementation.	O1 O2
Collaborative value evaluation	Develop and test assessments against intended purpose, working practices and users' needs.	O3

TABLE A4.2. Directions for continued inquiry and collaboration

These activities are proposals developed through synthesis, not agreed projects. Public- and private-sector research organisations and relevant academic disciplines could develop them jointly, with emerging findings informing CLAIR 2027 in Berlin.

A4.7 From workshop dialogue to research questions

DISCUSSION EVIDENCE AND SOURCE LOCATOR	AGGREGATED THEME	REFERENCE
Treatment of people expresses values; trust influences participation. `Source: 00:01–05:08`	**Human and cultural conditions**	O1
Pilots may not become sustained practice; skills, processes and tools interact. `Source: 05:08; 11:17–16:51`	**Implementation capability**	O2
Clarify what AI should achieve and evaluate quality in its use case. `Source: 06:03–09:12`	**Purpose and value**	O3
Implementation involves people, processes, technology and organisational roles. `Source: 17:12–19:14`	**Integration of expertise**	O4

TABLE A4.3. From workshop dialogue to research questions

Evidence is paraphrased from source [4] in Section 10. Locators identify transcript passages or recording timestamps for participant review.

A4.8 Connections unresolved issues and participant review

Workflow changes connect with scientific standards; implementation roles with governance; and learning and questioning with critical thinking. Open issues concern organisational arrangements, trust with appropriate scrutiny, the pace of analysis versus human review capacity, and sustained change. Proposals about where AI responsibilities should sit are alternatives for investigation, not an agreed organisational prescription.

Participant review will check fidelity, differences of view and omissions, and distinguish corrections from new proposals. O1–O4 remain provisional formulations of the collaborative dialogue pending that review.