neuravox
PUBLIC INTEREST TECHNOLOGY

NEURAVOX POLICY LAB

01 *Who starts the thinking?*

02 *Who shapes the reasoning?*

03 *Who decides what is credible?*

04 *Who stands behind the output?*

05 *What can the person still do?*

06 *Who makes the consequential choice?*

Human Agency Framework

A practical framework for retaining cognitive and decision making authority in generative AI use

PUBLISHER
Neuravox Foundation

EDITION
Public Edition, Version 1.0

DATE
October 2026

About this publication

The Human Agency Framework is a publication of the Neuravox Policy Lab, the research programme of Neuravox Foundation on human agency in everyday AI use. It builds on the research report *Human Agency in AI Use* (2026). That report found that usage intensity alone was a poor guide to agency, and that the measured behaviours did not form one overall agency score. The framework therefore keeps six distinct domains and reports a profile across them.

The framework also draws on research on cognitive offloading, automation bias, critical thinking in AI assisted work, human autonomy and AI literacy. International frameworks from UNESCO, the OECD and the African Union give its normative basis: human oversight, accountability, dignity and human determination. UNESCO's AI Competency Framework for Teachers gave a model for turning principles into domains, specifications, practical activities and assessment routes. This framework applies that level of operational detail to human agency in everyday generative AI use.

SUGGESTED CITATION	Abako, G. L. (2026). *Human Agency Framework: A practical framework for retaining cognitive and decision making authority in generative AI use.* Version 1.0. Neuravox Policy Lab, Neuravox Foundation.
COMPANION REPORT	Abako, G. L. (2026). *Human Agency in AI Use: Exploratory findings from surveyed generative AI users.* Neuravox Policy Lab, Neuravox Foundation.
AUTHOR	Gideon Abako, Neuravox Foundation
VERSION	Version 1.0, October 2026. This public edition replaces the external review edition of September 2026. Neuravox Foundation will revise the framework as new evidence becomes available.
USE OF THIS FRAMEWORK	The framework supports reflection, workflow review, design and governance. It is not a diagnostic test, a validated psychometric scale or a compliance certification.
CONTACT	neuravox.org | contact@neuravox.org

© 2026 Neuravox Foundation.

HUMAN AGENCY FRAMEWORK

Contents

	The framework at a glance
	Summary
01	Why a Human Agency Framework?
02	Purpose, scope and intended users
03	Evidence and development approach
04	Guiding principles
05	Framework architecture
06	The six domains of human agency
07	Context of use
08	Outcomes and consequences
09	Interpreting an agency profile
10	Applying the framework
11	Assessing human agency
12	Evidence boundaries and responsible interpretation
A	Quick reflection tool
B	Applied scenarios
C	Organizational review template
	References

OVERVIEW

The framework at a glance

THE FRAMEWORK IN ONE SENTENCE

Human agency in AI use concerns where cognitive and decision making authority sits between a person and an AI system, and how far the person keeps the capacity to initiate, shape, evaluate, explain, perform and decide.

SIX DOMAINS OF HUMAN AGENCY

01 **Independent initiation** *Who starts the thinking?*	02 **Cognitive authorship** *Who shapes the reasoning?*	03 **Verification and epistemic judgment** *Who decides what is credible?*
04 **Explanatory ownership** *Who understands and stands behind the output?*	05 **Capability retention** *What can the person still do without AI?*	06 **Decision authority** *Who makes the consequential choice?*

THREE LAYERS

CONTEXT OF USE	AGENCY PROFILE	OUTCOMES
Stakes, reversibility, verifiability, task type, frequency, experience, consequences for others, workload.	Where authority sits in each of the six domains, described one domain at a time.	Effectiveness, efficiency, learning, confidence, dependence, capability change, responsibility.

THREE DESCRIPTIVE AUTHORITY STATES

HUMAN LED	SHARED	AI LED
The person keeps primary direction. AI may assist.	AI contributes substantially. The person can still intervene.	AI largely determines the activity. Intervention is limited.

Profiles, no total score. First describe where authority sits in each domain. Then judge if that distribution suits the task, its stakes and its consequences.

OVERVIEW

Summary

Usage intensity is a weak guide to human agency. What matters is where cognitive and decision making authority sits when a person uses AI. The Human Agency Framework gives individuals, organizations, AI designers and policy teams a structured way to see that distribution and to judge if it suits the task.

Generative AI now takes part in work that was central to human cognition: framing problems, generating ideas, synthesizing evidence, evaluating options and recommending decisions. Counting prompts or hours of use cannot show who steers this work. The relevant question is what authority stays with the person as AI becomes part of everyday work and life.

The framework answers that question through six domains: independent initiation, cognitive authorship, verification and epistemic judgment, explanatory ownership, capability retention and decision authority. A context layer sits around the domains and an outcome layer sits beside them. Outcomes such as speed, confidence, learning or perceived benefit are assessed separately from agency.

The framework reports a profile across the six domains. It does not produce a total score.

WHAT THE FRAMEWORK IS FOR

A common language

Give people and institutions shared terms to examine human agency in everyday generative AI use.

From principle to practice

Turn high level commitments to human oversight into observable behaviours and practical questions.

Better human involvement

Support reflection on the quality of human involvement in AI supported tasks.

Review and evaluation

Give a basis for self assessment, organizational review and task evaluation.

EVIDENCE POSITION

The framework rests on the Neuravox Human Agency Survey, published research and international human centred AI principles. The first survey was exploratory and had 76 respondents. The six domains are working domains for practice and further research. They are not validated psychometric scales.

SECTION 01

Why a Human Agency Framework?

Generative AI changes what assistance means. Earlier digital tools mostly took over memory, calculation or retrieval. Generative systems can also propose how to frame a problem, organize an argument, synthesize evidence, compare options and recommend the next step. This makes the line between assistance and delegation harder to see.

Human agency is a person's ability to direct their own thinking and action. In AI supported work, that ability can shift while the quality of the output stays high. A polished answer can hide weak understanding. A fast decision can hide dependence on an AI recommendation. Frequent AI use can also sit alongside active checking, strong explanation and independent capability.

The central question is therefore how cognitive and decision making authority is distributed during use. The fact of AI use, or its frequency, cannot answer that question alone.

> CORE PROPOSITION
>
> Human agency in AI use is the capacity to keep meaningful cognitive and decision making authority while using AI assistance.

1.1 From "human in the loop" to observable agency

Governance principles often call for human oversight, accountability or human centred AI. These principles are necessary, but they rarely say what oversight looks like in everyday behaviour. A person can stay formally "in the loop" while accepting an AI recommendation with little understanding and no realistic ability to reject it.

The framework therefore works through six observable questions. Who starts the thinking? Who shapes the reasoning? Who decides what is credible? Who understands the output? What capability remains without AI? Who makes the consequential choice?

1.2 Agency does not require abstinence

The framework does not treat AI use or delegation as harmful in itself. Delegation is one reason tools are useful. A person can delegate drafting, search or idea generation and still keep strong authority over purpose, verification, explanation and final decisions. The aim is to make the distribution of authority visible and suited to the context.

SECTION 02

Purpose, scope and intended users

2.1 Purpose

The Human Agency Framework gives a structured way to examine how people keep or transfer cognitive and decision making authority when they use generative AI. It supports reflection, research, governance and design. It does not aim to prohibit AI use.

2.2 Scope

The framework focuses on cognitive agency and decision authority in everyday generative AI use. Social agency is part of the wider research question. The first survey did not measure it directly, so this edition treats social agency as a cross cutting consideration (see section 12.4).

2.3 Intended users

USER	PRIMARY USE
Individuals	Reflect on how AI enters their thinking, learning, work and personal decisions.
Organizations	Review AI enabled workflows, identify where human judgment is essential and design agency checkpoints.
AI designers	Identify product choices that preserve user understanding, intervention and decision authority.
Policy and governance teams	Translate human oversight principles into practical behaviours and accountability expectations.
Researchers and educators	Develop and test measures of human agency across tasks, populations and contexts.

SECTION 03

Evidence and development approach

Three evidence streams inform the framework:

- the Neuravox Human Agency Survey and its research report,
- published research on cognitive offloading, automation bias, critical thinking, autonomy and AI literacy,
- international policy frameworks that place human determination, accountability and dignity at the centre of responsible AI.

3.1 What the Neuravox study found

The survey collected responses from 76 adults who had used generative AI for at least three months. Most were frequent and experienced users: 69.7 percent used AI at least daily and 77.6 percent had used it for a year or more. The findings below shaped the framework. They describe this sample and are exploratory.

FINDING	WHAT IT MEANS FOR THE FRAMEWORK
AI use frequency showed no clear link with the overall delegation summary (rho 0.05) or with perceived change in independent thinking.	Usage intensity is context. It does not measure agency.
One delegation behaviour did rise with exposure. Asking AI before trying independently was more common among more frequent users (rho 0.28) and among people with longer AI experience (rho 0.33).	Independent initiation needs particular attention among frequent and long term users.
Respondents who reported more delegation reported lower confidence completing similar tasks without AI (rho −0.32).	Capability retention needs its own domain and where it matters, task based evidence.
Delegation and retained agency coexisted. Moderate delegation appeared alongside frequent conflict questioning and explanatory ownership.	Delegation can preserve agency. The quality of delegation matters.
The seven behaviours did not form one coherent score (alpha 0.575). The three retained agency behaviours also held together weakly (alpha 0.41).	The framework keeps distinct domains and reports profiles.
55.3 percent reported improvement in independent thinking and 26.3 percent reported some reduction. Respondents with stronger retained agency did not report more improvement.	Benefit and agency need separate measurement.

Correlations are exploratory Spearman coefficients from the survey and are not corrected for multiple testing. With 76 respondents, the study could detect only moderate or larger associations, so weaker links remain possible.

3.2 Open questions

- Do the six domains form stable constructs across populations and tasks?
- How do people interpret “improvement” when they describe changes in independent thinking?
- How does agency change over time as AI becomes routine?
- Which forms of delegation are productive, neutral or harmful under different conditions?
- How do these patterns differ across countries, languages and sectors?
- How should social agency and human relationships enter the framework?

EVIDENCE BOUNDARY

The six domains are evidence-informed practical domains. They are not psychometric factors. Do not treat them as a validated scale or combine them into a diagnostic score.

SECTION 04

Guiding principles

Eight principles guide the use and interpretation of the framework.

01

Agency is multidimensional

Examine human agency across distinct capacities and behaviours. Strength in one domain does not guarantee strength in another.

02

Usage intensity describes context

Frequency, duration and number of prompts describe exposure. They do not show who steers the reasoning or makes the decision. Record them as context and read them with behaviour.

03

Delegation can preserve agency

People can delegate parts of a task and keep purpose, judgment, understanding and final authority.

04

Accountability requires real intervention

Formal approval is not enough if a person cannot understand, challenge or realistically reject an AI output.

05

Benefit and agency are separate outcomes

AI can improve speed, quality or access while it changes independent capability. A positive outcome does not prove that agency was preserved.

06

Demonstrate capability where it matters

Self reported confidence is useful but not enough for critical tasks. Where capability matters for safety, learning or professional responsibility, use task based evidence where feasible.

07

Agency depends on context

The right distribution of authority depends on stakes, reversibility, verifiability, task type and consequences for others.

08

The framework follows the evidence

AI capabilities and patterns of use will change. Revise the framework through evidence and practice.

SECTION 05

Framework architecture

The framework has three layers. Six agency domains form the centre. A context layer describes the conditions of AI use. An outcome layer records what results from that use. Separate layers stop outcomes such as speed or satisfaction from being read as evidence of preserved agency.

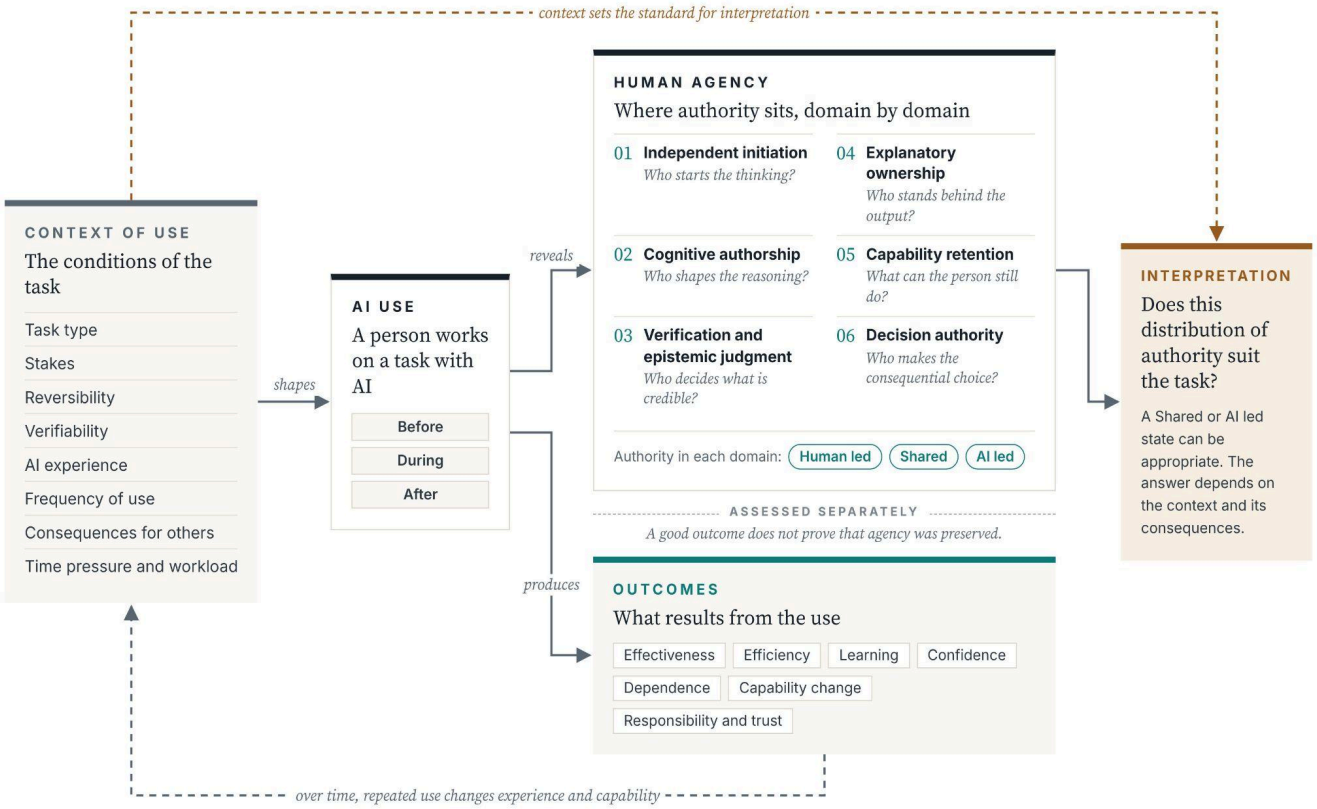


Figure 1. Conceptual model of the Human Agency Framework.

5.1 The six domains

DOMAIN	CORE QUESTION	WORKING MEANING
Independent initiation	Who starts the thinking?	Form an initial view, question or attempt before AI defines the core problem.
Cognitive authorship	Who shapes the reasoning?	Keep influence over structure, priorities, assumptions and direction.
Verification and epistemic judgment	Who decides what is credible?	Check important outputs, notice conflicts and reject unsupported answers.
Explanatory ownership	Who understands and stands behind the output?	Explain the reasoning and the result in one's own terms.

DOMAIN	CORE QUESTION	WORKING MEANING
Capability retention	What can the person still do independently?	Keep relevant knowledge and skill for when AI is unavailable or unsuitable.
Decision authority	Who makes the consequential choice?	Keep final authority over choices with personal, professional or social consequences.

5.2 Three descriptive authority states

Three descriptive states show where authority appears to sit within each domain. The states are not grades. Do not combine them into a total score.

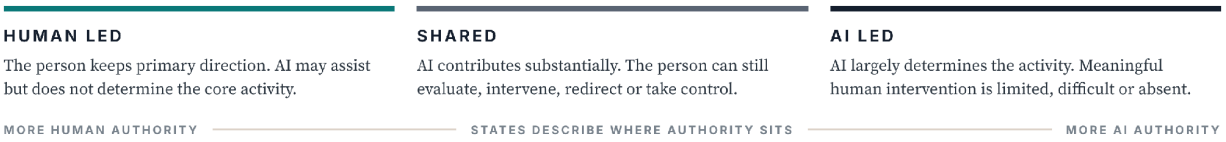


Figure 2. Authority states describe where authority sits in one domain.

> **IMPORTANT**
> A Shared or AI led state can be appropriate. Interpretation depends on context. The framework first describes where authority sits. It then asks if that distribution suits the task and its consequences.

5.3 A practical sequence: before, during and after AI use

STAGE	PRIMARY DOMAINS	AGENCY CHECKPOINT
Before	Independent initiation, limits on decision authority	What am I trying to do? What should remain mine to decide?
During	Cognitive authorship, verification and epistemic judgment	Is AI shaping the task in ways I understand and accept? What needs checking?
After	Explanatory ownership, capability retention, final decision authority	Can I explain the result, work without the system where necessary and own the decision?

SECTION 06

The six domains of human agency

Each domain specification gives a definition, the reason the domain matters, indicators of stronger agency, signals that authority may be shifting, reflection questions, an example and the evidence basis. Use the specifications for reflection, organizational review, AI design and governance.

How to read each specification

ELEMENT	WHAT IT GIVES
Definition	What the domain covers.
Why this domain matters	The risk to human agency that the domain addresses.
Stronger agency indicators	Observable behaviours that show authority staying with the person.
Signals agency may be shifting	Observable behaviours that show authority moving to the AI system.
Reflection questions	Questions a person can ask about one task.
Example	A short case that shows the domain in practice.
Evidence basis	What the first Neuravox study and wider research show for the domain.

6.1 | DOMAIN 01 OF 06

Independent initiation

Who starts the thinking?

DEFINITION

Independent initiation is the capacity to form an initial understanding, question, objective or attempt before AI defines the core problem. The emphasis is on retained ability and deliberate choice. It is not a rule that people must always work alone before they use AI. Shared initiation can still show meaningful agency when a person uses AI on purpose to explore or clarify a problem, and then keeps authority over how the problem is framed and pursued.

WHY THIS DOMAIN MATTERS

AI can enter a task so early that it starts to define what the user thinks the problem is. When initiation becomes automatic, the person can lose the chance to form an independent frame, notice uncertainty or decide what help they actually need.

STRONGER AGENCY INDICATORS	SIGNALS AGENCY MAY BE SHIFTING
Identifies the problem or objective before prompting AI when the task calls for independent thought.	AI becomes the default starting point for tasks that need judgment or learning.
Forms an initial view, hypothesis, outline or attempt.	The user asks AI what to think before clarifying what they are trying to decide.
Chooses deliberately when AI should enter the task.	Starting without AI feels unusually difficult, even for familiar tasks.
Can begin a familiar task when AI is unavailable.	

REFLECTION QUESTIONS

- Before I asked AI, did I know what problem I was trying to solve?
- Could I describe my initial position without looking at the AI conversation?
- Did I decide what I wanted from AI, or did AI decide what the task became?
- If AI were unavailable at the start, could I still begin?

EXAMPLE

A researcher sets the question and an initial search strategy, then uses AI to expand the list of keywords. A more AI led pattern would be to ask the system for the research question and adopt its framing without independent review.

EVIDENCE BASIS

Direct survey evidence from one behaviour: asking AI before trying independently. In the first study, 43 respondents (56.6 percent) did this sometimes and 12 did it often or almost always. It was the one delegation behaviour that rose with AI use frequency (rho 0.28) and with length of experience (rho 0.33). Research on cognitive offloading supports the domain. One item cannot capture the full construct, so interpret it through observed behaviour.

6.2 | DOMAIN 02 OF 06

Cognitive authorship

Who shapes the reasoning?

DEFINITION

Cognitive authorship is the extent to which the person keeps influence over the structure, priorities, assumptions and direction of the work. It separates editing an AI generated product from shaping the intellectual path that produced it.

WHY THIS DOMAIN MATTERS

A person can make the final edits after AI has already set the frame, the order of ideas, the alternatives considered and the assumptions built into the work. Authorship of the visible text and authorship of the reasoning are different things.

STRONGER AGENCY INDICATORS	SIGNALS AGENCY MAY BE SHIFTING
Decides which issues deserve attention.	AI generated framing is adopted again and again without challenge.
Changes AI generated structure when it does not fit the task.	The user can edit wording but struggles to change the underlying argument or structure.
Introduces important ideas, evidence or alternatives that did not come from AI.	Important alternatives are missing because AI did not raise them.
Can explain why the final structure was chosen.	The only reason the user can give for the final structure is that AI suggested it.

REFLECTION QUESTIONS

- Who decided the structure of this work?
- Which important ideas or priorities came from me?
- Did I substantially change the AI's framing?
- If AI had proposed a different structure, could I judge which was better?

EXAMPLE

A policy analyst asks AI for three ways to organize a brief, rejects two, combines useful parts of the third and rebuilds the argument around their own policy logic. AI assisted the work. The analyst authored its reasoning.

EVIDENCE BASIS

Survey evidence from one behaviour: letting AI structure ideas. In the first study, 45 respondents (59.2 percent) did this sometimes and 11 did it often or almost always. The domain also draws on the conceptual difference between editing an output and shaping the reasoning behind it. The survey evidence is narrow, so interpret the domain through observable behaviour.

6.3 | DOMAIN 03 OF 06

Verification and epistemic judgment

Who decides what is credible?

DEFINITION

Verification and epistemic judgment are the capacities to evaluate AI generated information, identify uncertainty or conflict, seek external evidence when needed and reject an AI response. The domain covers factual checking and also judgments about evidence quality, logic, plausibility and confidence.

WHY THIS DOMAIN MATTERS

AI outputs can be fluent and persuasive when they are wrong or weakly supported. Human oversight means little if the user cannot tell when to check, challenge or reject an answer.

STRONGER AGENCY INDICATORS	SIGNALS AGENCY MAY BE SHIFTING
Checks important factual claims against reliable evidence.	Outputs are accepted mainly because they sound coherent or authoritative.
Recognizes when a task needs external verification or domain expertise.	AI is used to verify another AI response where independent checking is possible.
Questions AI when it conflicts with their knowledge or judgment.	The user cannot say what evidence would change their mind.
Can reject a plausible answer when its support is weak.	A plausible answer is accepted despite weak support.
Separates confident presentation from strength of evidence.	Fluent presentation is treated as proof that an answer is reliable.

REFLECTION QUESTIONS

- What would make me reject this answer?
- Did I verify the claims that actually matter?
- Do I accept this because it is supported, or because it sounds convincing?
- If AI disagreed with my judgment, what would I do next?
- Which parts remain uncertain?

EXAMPLE

A clinician who uses an AI summarizer checks medication information against an authoritative source before acting. In a low stakes brainstorming task, external checking may be unnecessary, because the output serves only as a source of ideas.

EVIDENCE BASIS

This domain has the strongest grounding in the framework. Two survey behaviours relate to it. 59 respondents (77.6 percent) questioned AI often or almost always when it conflicted with their judgment, and 10 used AI output without checking often or almost always. Research on automation bias and critical thinking, and international emphasis on human oversight and explainability, also support the domain.

6.4 | DOMAIN 04 OF 06

Explanatory ownership

Who understands and stands behind the output?

DEFINITION

Explanatory ownership is the capacity to explain the reasoning, assumptions and final result in one's own terms, without the original AI conversation. It concerns understanding of the output and accountability for it. It does not require authorship of every word.

WHY THIS DOMAIN MATTERS

A person can submit, present or act on a high quality AI assisted output without understanding how it reached its conclusion. Explanatory ownership gives a practical test of human involvement. A person can also understand and explain an AI generated line of reasoning while its claims remain false or unsupported. Explanatory ownership therefore does not show that an output was adequately verified.

STRONGER AGENCY INDICATORS	SIGNALS AGENCY MAY BE SHIFTING
Can explain how the conclusion was reached.	The user can reproduce the output but cannot explain why it is correct.
Can identify the important assumptions and limitations.	The user cannot identify the assumptions or limitations behind the AI generated result.
Can defend or revise the output when challenged.	The user relies on AI to answer challenges about work they present as their own.
Can rebuild the main reasoning without reopening the AI conversation.	The user needs the AI conversation to explain the reasoning behind the final output.

REFLECTION QUESTIONS

- Could I explain this to someone without opening the AI chat?
- Do I understand why this conclusion follows?
- Could I defend the important choices?
- If someone challenged the result, could I respond on my own?

EXAMPLE

A student uses AI to help debug code and can explain the bug, why the fix works and how to detect the same issue next time. Copying a working solution without understanding it would show weaker explanatory ownership.

EVIDENCE BASIS

Direct survey evidence. The ability to explain final work was the strongest reported retained agency behaviour: 61 respondents (80.3 percent) reported it often or almost always, and none reported never or rarely. This ceiling limits what self report can show, so teach back tasks give stronger evidence. The domain is distinct from capability retention.

6.5 | DOMAIN 05 OF 06

Capability retention

What can the person still do independently?

DEFINITION

Capability retention is the extent to which relevant knowledge and skill stay available when AI is absent, unsuitable, wrong or deliberately withheld. It concerns the human competence that remains necessary for resilience, learning, error detection and professional responsibility.

WHY THIS DOMAIN MATTERS

Understanding an AI assisted result does not prove that a person can do a similar task alone. Repeated delegation can improve immediate performance while it reduces practice of the underlying skills. The long term effect is an empirical question that needs more than self report.

STRONGER AGENCY INDICATORS	SIGNALS AGENCY MAY BE SHIFTING
Can complete the essential parts of familiar tasks without AI when needed.	Routine tasks become difficult without AI, even when they were manageable before.
Keeps the foundational knowledge needed to detect important errors.	The user cannot detect serious errors without another AI system.
Uses AI in ways that support learning, and keeps practising core skills.	Core skills are no longer practised, although the role still needs them.
Can keep working during AI failure, restriction or unavailability.	Essential work stops when AI becomes unavailable.

REFLECTION QUESTIONS

- Which parts of this task could I still do without AI?
- Which skills might weaken if I keep using AI this way?
- Could I detect a serious error without another AI system telling me?
- If AI became unavailable tomorrow, what would I no longer know how to do?

EXAMPLE

A financial analyst uses AI to speed up spreadsheet work but regularly performs core calculations and interpretation without it. For critical functions, the organization can test unaided performance on purpose.

EVIDENCE BASIS

Direct survey evidence from confidence without AI. In the first study, 46 respondents (60.5 percent) felt confident completing similar tasks without AI often or almost always, which made this the most variable retained agency behaviour. Respondents who reported more delegation reported lower confidence (rho −0.32). For consequential uses, task based evidence gives a stronger basis than self reported confidence.

6.6 | DOMAIN 06 OF 06

Decision authority

Who makes the consequential choice?

DEFINITION

Decision authority concerns the control a person keeps over choices with personal, professional or social consequences. AI can recommend, compare, predict or advise. The person should still be able to question the recommendation and make the consequential choice where human responsibility is required.

WHY THIS DOMAIN MATTERS

A nominal human approval step can hide automation bias. Decision authority is weak when the AI recommendation in effect decides the outcome, and the person lacks the information, confidence or freedom to choose otherwise.

STRONGER AGENCY INDICATORS	SIGNALS AGENCY MAY BE SHIFTING
Separates advice from decision.	AI recommendations are routinely treated as default decisions.
Considers relevant alternatives beyond the AI recommendation.	Human review becomes a procedural click with little real evaluation.
Can explain why a recommendation was accepted or rejected.	The person has little freedom or confidence to reject the AI recommendation.
Keeps named responsibility for consequential decisions.	Responsibility becomes unclear when an AI supported decision causes harm.
Can realistically choose against the AI recommendation.	

REFLECTION QUESTIONS

- Could I realistically choose against the AI recommendation?
- What information, other than the AI response, influenced my decision?
- If this decision goes badly, can I explain why I made it?
- Has AI become the default authority because deciding on my own feels harder?

EXAMPLE

A hiring manager can use AI to summarize applications. The final decision still needs review of the relevant evidence, awareness of bias risks and accountable human judgment. A rubber stamp approval keeps a human in the loop without meaningful decision authority.

EVIDENCE BASIS

Direct survey evidence from delegating decisions previously made personally. 18 respondents (23.7 percent) did this often or almost always and 35 (46.1 percent) did so never or rarely. Research on automation bias and policy commitments to human accountability reinforce the domain.

SECTION 07

Context of use

Do not interpret agency in the abstract. The same distribution of authority can suit one task and cause harm in another. Context sets how much independent judgment, verification, capability and accountability a task needs.

7.1 Core contextual factors

CONTEXT FACTOR	WHY IT MATTERS
Task type	Writing, coding, research, learning, personal advice and decision support make different demands on agency.
Stakes	Errors matter differently in low stakes exploration, professional work and decisions that affect health, rights, money or opportunity.
Reversibility	A reversible choice can tolerate more experimentation than a decision that is irreversible or hard to correct.
Verifiability	Some outputs can be tested against code, sources or measurable results. Others offer few external checks.
AI experience	Experience can change how well a person prompts, checks and interprets AI. Record it as context. It is not proof of agency.
Frequency of use	Usage intensity describes exposure. Do not use it to diagnose dependence or agency.
Consequences for others	Decisions that affect patients, employees, students or the public can need stronger safeguards than decisions that affect only the user.
Time pressure and workload	High workload can increase reliance on automated recommendations and reduce the time available for independent checking.

7.2 Context rule

> **CONTEXT RULE**
>
> As stakes and consequences rise, reversibility falls and independent verification becomes harder, a transfer of cognitive or decision making authority to AI needs stronger justification.

This rule guides interpretation. It does not set a numerical threshold. Apply it with professional judgment and attention to the task, its consequences and the safeguards available. Treat time pressure and workload as implementation risks: they can make human review less substantive even where formal oversight remains in place.

SECTION 08

Outcomes and consequences

The framework separates agency from outcomes. A person can perform better while becoming less capable without AI. A person can also keep strong agency and gain little efficiency. Measure outcomes beside the agency profile and keep them out of it.

OUTCOME	QUESTION
Effectiveness	Did AI improve the quality or usefulness of the result?
Efficiency	Did AI reduce time, effort or cost?
Learning	Did the person gain knowledge or skill from the interaction?
Confidence	Did the interaction change confidence appropriately, or create overconfidence?
Dependence	Does the person increasingly feel unable to act without AI?
Capability change	Has unaided performance improved, stayed stable or declined over time?
Responsibility and trust	Can the person or institution still account for the final output or decision?

SEPARATION PRINCIPLE

A positive outcome is not evidence of preserved agency. A high level of human control is not evidence that AI was useful. Assess both.

SECTION 09

Interpreting an agency profile

The framework produces a descriptive profile across the six domains. A person can be human led in one domain, shared in another and AI led in a third. The profile keeps that visible.

9.1 Profile template

DOMAIN	OBSERVED AUTHORITY	EVIDENCE OR EXAMPLE
Independent initiation	☐ Human led ☐ Shared ☐ AI led	
Cognitive authorship	☐ Human led ☐ Shared ☐ AI led	
Verification and epistemic judgment	☐ Human led ☐ Shared ☐ AI led	
Explanatory ownership	☐ Human led ☐ Shared ☐ AI led	
Capability retention	☐ Human led ☐ Shared ☐ AI led	
Decision authority	☐ Human led ☐ Shared ☐ AI led	

9.2 Interpretation questions

- Which domains matter most for this task or role?
- Does any AI led domain occur in a high stakes, irreversible or weakly verifiable context?
- Can the person intervene in a meaningful way, or is human oversight mostly procedural?
- Do positive outcomes hide a decline in unaided capability?
- Does the profile change when the same person does a different task?
- Which safeguards, training or design changes would move authority to a more suitable distribution?

9.3 Profiles are not identities

Do not label people as high agency or low agency from a single interaction. Agency is specific to the task, sensitive to context and can change over time. The same person can show very different profiles across work, learning, coding, personal advice and high stakes decisions.

SECTION 10

Applying the framework

10.1 Individuals

Individuals can use the framework to reflect on important or recurring AI supported tasks. The aim is to decide deliberately what stays under human control. Maximum manual effort is not the goal.

BEFORE	DURING	AFTER
Define the problem and decide what AI should contribute.	Notice when AI changes the structure of the task, and verify the claims that matter.	Explain the result in your own terms, check what capability remains and make consequential decisions deliberately.

10.2 Organizations

Organizations can use the framework to review workflows. This replaces blanket rules based on hours or frequency of AI use. A practical process maps the task, identifies the domains where human authority is essential, sets evidence requirements and specifies escalation points.

STEP	ORGANIZATIONAL QUESTION
1. Map the workflow	Where does AI enter the task? What does it generate, evaluate or recommend?
2. Identify agency critical domains	Which of the six domains matter most for quality, safety, learning or accountability?
3. Set agency checkpoints	Where must a person verify, explain, intervene or decide?
4. Test capability	For critical roles, can staff still perform essential functions during AI failure or in unusual cases?
5. Review outcomes separately	Does AI improve performance while it creates new dependence, opacity or gaps in responsibility?

10.3 AI designers

- **Independent initiation.** In learning and judgment heavy tasks, let users set goals or give an initial view before the system generates a complete solution.

- **Cognitive authorship.** Show alternative structures and make it easy to revise or reject AI framing.
- **Verification.** Show sources, uncertainty, assumptions and the points that need external checks.
- **Explanatory ownership.** Let users inspect how an answer was built, and prompt reflection where appropriate.
- **Capability retention.** Avoid designs that remove practice of skills users still need.
- **Decision authority.** Make the line between advice and final human responsibility explicit, especially in consequential contexts.

10.4 Policy and public interest technology

UNESCO, OECD and African Union frameworks already establish human oversight, dignity and responsibility as core principles. The framework helps put these principles into operation. It moves from abstract requirements to observable behaviours and task accountability.

Policy can ask four practical questions. Do people have the information, capability and freedom to question AI? Are critical outputs verifiable? Are important skills being retained? Do consequential decisions remain under human control?

SECTION 11

Assessing human agency

The framework can support assessment. It is not a validated psychometric scale. Combine several methods, because self report alone cannot establish capability, understanding or actual decision behaviour.

11.1 Recommended evidence types

EVIDENCE TYPE	ROLE
Self report	Shows perceived behaviour, confidence, dependence and reflection. Recall and social desirability bias affect it.
Behavioural observation	Records what users actually check, revise, reject or delegate during AI assisted tasks.
Task based assessment	Tests unaided capability, explanation, verification or performance under controlled conditions.
Scenario based assessment	Tests decisions and reasoning in realistic situations with different stakes and AI recommendations.
Interaction logs	Can show timing, revisions and patterns of dependence, if collected ethically and read in context.
Longitudinal assessment	Shows if repeated patterns of delegation change capability, judgment or confidence over time.

11.2 Domain to method mapping

DOMAIN	ASSESSMENT METHOD	WHAT IT TESTS
Independent initiation	Self report, first attempt task, prompt timing and sequence	Does the person form a view or attempt before AI defines the problem?
Cognitive authorship	Think aloud task, revision trace, structured interview	Who sets the frame, structure and priorities?
Verification and epistemic judgment	Source checking task, seeded AI errors, observation	Can the person identify what needs checking and reject unsupported outputs?

DOMAIN	ASSESSMENT METHOD	WHAT IT TESTS
Explanatory ownership	Teach back, or defence of the work without reopening the chat	Can the person rebuild and defend the reasoning?
Capability retention	Comparable tasks with and without AI, repeated over time	What can the person still do independently?
Decision authority	Vignettes with AI recommendations, justification of the choice	Can the person choose against AI and account for the decision?

11.3 Scoring position

NO TOTAL SCORE

Use evidence at domain level and descriptive profiles. The current evidence does not justify adding the six domains into one agency score, or applying universal weights across tasks and contexts.

SECTION 12

Evidence boundaries and responsible interpretation

The Human Agency Framework is evidence informed. It turns findings from the human agency study, published research and international AI governance principles into a structure for reflection, workflow review, design and governance. Use it within the limits of that evidence.

12.1 What the framework can support

The framework can help individuals and institutions see how cognitive and decision making authority is distributed. It can structure reflection on a specific AI supported task, support reviews of organizational workflows, identify points where human verification or decision responsibility is essential and guide product and policy discussions about human control.

It also gives a common language for patterns that simple usage measures hide. Two people can both use AI every day and show very different profiles of initiation, authorship, verification, understanding, capability and decision authority. The framework is designed to capture that difference.

12.2 What the framework does not establish

Do not use the framework as a psychological diagnosis, a measure of personal worth, a clinical measure of dependence or a validated psychometric scale. It sets no population threshold for high or low agency, and it does not support ranking people by one total score.

The underlying survey was exploratory, cross sectional and based on self report. It cannot show that AI use causes a gain or loss in capability over time. It also cannot show that the six domains are fixed latent traits. The framework therefore uses domain profiles and observable evidence.

12.3 The first study and its limits

The first survey had 76 respondents, recruited through professional networks, LinkedIn and online communities between 8 and 22 August 2026. This convenience and snowball sample probably includes more frequent and more favourable AI users than the general population. The survey did not record country, region or other demographic details, so the findings make no claim about any national or regional population.

All measures were self reported at one point in time. Some items reached a ceiling: no respondent said they could never or rarely explain their final work. The next research phase plans to record geography and demographic context, publish the full instrument wording and add task based measures.

12.4 Social agency

The wider research question covers cognitive, decisional and social agency. The first survey measured cognitive and decisional behaviours directly. It did not include a dedicated measure of social agency, so social agency is not a seventh domain in this edition.

Social agency still matters. Generative AI can change who people consult, whose knowledge enters a decision, how people meet disagreement and how human relationships are strengthened, bypassed or displaced. International human centred AI frameworks also stress human relationships, inclusion, participation, dignity and accountability.

In this edition, social agency is a cross cutting consideration. It appears most in the context layer, especially consequences for others, and in questions about advice, collaboration and responsibility. Users and organizations should ask:

- Has AI changed who is consulted before an important decision?
- Does AI use reduce exposure to human disagreement or other viewpoints?
- Is AI displacing a human relationship, professional interaction or collaborative process that still has value?
- Can the people affected by an AI supported decision question, contest or influence it?
- Does AI use shift voice or authority among people, teams or communities?

These questions extend the framework without a separate social agency score. Where social relationships are central to a task, assess them explicitly beside the six domains.

12.5 Relative strength of the evidence

The current evidence supports the six domains to different degrees.

DOMAIN	SURVEY EVIDENCE	OTHER SUPPORT	CURRENT STRENGTH
Independent initiation	One item, rose with frequency and experience	Cognitive offloading research	Moderate, single item
Cognitive authorship	One item	Difference between editing and reasoning	Limited, single item
Verification and epistemic judgment	Two related items	Automation bias, critical thinking, oversight principles	Strongest
Explanatory ownership	One item with a ceiling effect	Critical thinking research	Moderate, needs task evidence
Capability retention	One item, lower with more delegation	Cognitive offloading research	Moderate, needs task evidence
Decision authority	One item	Automation bias, accountability principles	Moderate

This unevenness is a reason to keep the evidence visible. It does not justify collapsing the domains. The framework separates what the survey observed from what rests mainly on theory, prior research and governance logic.

12.6 Responsible use of the authority states

Human led, Shared and AI led describe where authority appears to sit in one domain. They are not moral grades. A human led state is not automatically better, and an AI led state is not automatically wrong. A low stakes, reversible formatting task can reasonably be AI led. A high stakes, weakly verifiable decision that affects another person can need much stronger human authority.

Interpretation therefore follows two steps. First, describe the authority profile. Second, judge if that profile suits the context, the consequences and the accountability requirements of the task.

12.7 A framework based on current evidence

Version 1.0 reflects the evidence available at publication. It supports reflection, organizational practice and AI governance. Neuravox Foundation will revise it as AI systems and the evidence change. Until then, apply it within the evidence boundaries stated here.

APPENDIX A

Quick reflection tool

Use this worksheet to reflect on one AI supported task. It is a practical aid. It is not a validated assessment.

DOMAIN	REFLECTION QUESTION	STATE	NOTES
Independent initiation	Did I know what I was trying to do before AI shaped the task?	☐ Human led ☐ Shared ☐ AI led	
Cognitive authorship	Did I shape the structure, priorities and direction of the work?	☐ Human led ☐ Shared ☐ AI led	
Verification and epistemic judgment	Did I check what mattered, and could I still reject the AI response?	☐ Human led ☐ Shared ☐ AI led	
Explanatory ownership	Can I explain and defend the final result without reopening the chat?	☐ Human led ☐ Shared ☐ AI led	
Capability retention	Could I still do the essential parts without AI if I needed to?	☐ Human led ☐ Shared ☐ AI led	
Decision authority	Did I make the consequential decision, with AI as input and not as authority?	☐ Human led ☐ Shared ☐ AI led	

CONTEXT CHECK

- How high are the stakes?
- How reversible is the outcome?
- How easy is independent verification?
- Who else could be affected?
- Would this distribution of authority still be acceptable if the AI were wrong?

APPENDIX B

Applied scenarios

The scenarios show how the same level of AI use can produce very different agency profiles. They are illustrations, not model answers. They show how to use the six domains and the context layer together.

SCENARIO	INDICATIVE CONTEXT	MAIN DOMAINS ILLUSTRATED
B1 Policy brief	Professional, moderate stakes	Cognitive authorship, verification, decision authority
B2 University learning	Learning, moderate stakes	Independent initiation, explanatory ownership, capability retention
B3 Coding and data analysis	Professional, highly verifiable	Verification, explanatory ownership, capability retention
B4 Personal advice	Personal, potentially high consequence	Cognitive authorship, verification, decision authority
B5 Hiring decision	High stakes, affects others	Verification, explanatory ownership, decision authority
B6 Clinical information	High stakes, professional	Verification, capability retention, decision authority
B7 Administrative work	Low stakes, reversible	Cognitive authorship, explanatory ownership
B8 Research synthesis	Professional, evidence dependent	Verification, cognitive authorship, explanatory ownership

B1 Drafting a policy brief

PROFESSIONAL, MODERATE STAKES

SCENARIO	A policy analyst asks a generative AI system to propose three structures for a brief. The analyst already knows the policy question. They reject two structures, combine parts of the third with their own outline, check all factual claims against source documents, rewrite the argument and can explain why they chose the final recommendations.
LIKELY PROFILE	Independent initiation: human led. Cognitive authorship: shared, strongly directed by the analyst. Verification: human led. Explanatory ownership: human led. Capability retention: likely human led, if the analyst can still do the core task without AI. Decision authority: human led.
INTERPRETATION	Substantial AI assistance does not by itself indicate weak agency. The analyst keeps the problem definition, the evidence standard, the reasoning and the final recommendation.

B2 University learning

LEARNING, MODERATE STAKES

SCENARIO	A student asks AI to solve a difficult statistics problem before attempting it. The student copies the steps into an assignment and can reproduce the final answer. They cannot explain why the method applies or solve a similar problem when the system is unavailable.
LIKELY PROFILE	Independent initiation: AI led. Cognitive authorship: AI led. Verification: weak or absent. Explanatory ownership: AI led. Capability retention: weak. Decision authority: less relevant, because the task is mainly educational.
INTERPRETATION	The output can be correct while the learning objective fails. In learning contexts, capability retention and explanatory ownership can deserve more weight than speed.

B3 Coding and data analysis

PROFESSIONAL, HIGHLY VERIFIABLE

SCENARIO	A developer uses AI to generate boilerplate code and suggest a debugging path. The developer runs tests, inspects the code, rejects an unsafe suggestion, changes the logic and can explain the final implementation. They could rebuild the essential logic without the tool, although more slowly.

LIKELY PROFILE	Independent initiation: possibly shared. Cognitive authorship: shared. Verification: human led, because tests and code behaviour give strong external checks. Explanatory ownership: human led. Capability retention: human led. Decision authority: human led.
INTERPRETATION	This is productive delegation in a highly verifiable environment. AI adds speed while the developer keeps the ability to test, understand and intervene.

B4 Personal advice

PERSONAL, POTENTIALLY HIGH CONSEQUENCE

SCENARIO	A person asks a chatbot if they should end a relationship. They accept the recommendation because it sounds emotionally perceptive, and act on it without discussing the issue with anyone or considering if the system has enough context.
LIKELY PROFILE	Independent initiation: possibly shared. Cognitive authorship: AI led, because the system frames the situation. Verification: difficult, because there may be no simple external truth. Explanatory ownership: limited. Decision authority: AI led, if the recommendation in effect becomes the decision.
INTERPRETATION	Personal advice can be hard to verify and can have serious consequences. The framework pays particular attention to other perspectives and to the person's real freedom to choose differently.

B5 Professional hiring decision

HIGH STAKES, AFFECTS OTHERS

SCENARIO	An organization uses AI to summarize applications and rank candidates. A manager reviews only the top ranked candidates and approves the ranking without examining the criteria, the omitted applicants or possible bias. A person formally signs the final hiring decision.
LIKELY PROFILE	Independent initiation: possibly human led at organizational level. Cognitive authorship of the ranking: AI led. Verification: weak. Explanatory ownership: limited. Decision authority: formally human, substantively AI led.
INTERPRETATION	A human approval step does not guarantee meaningful human agency. If the reviewer lacks the information or practical freedom to question the ranking, automation dominates the decision.

B6 Clinical or health information

HIGH STAKES, PROFESSIONAL

SCENARIO	A clinician uses AI to summarize a long patient record and flag possible medication interactions. The clinician checks the information against the record and authoritative clinical sources, notices one incorrect suggestion, rejects it and makes the final decision with professional judgment.
LIKELY PROFILE	Independent initiation: possibly shared. Cognitive authorship: shared. Verification: human led. Explanatory ownership: human led. Capability retention: must stay strong, because the stakes are high. Decision authority: human led.
INTERPRETATION	High stakes use does not require abstaining from AI. It raises the standard for verification, capability and decision responsibility.

B7 Routine administrative work

LOW STAKES, REVERSIBLE

SCENARIO	An employee uses AI every day to rewrite meeting notes, format routine emails and turn rough points into polished prose. They rarely inspect changes at sentence level, because the tasks are low stakes and easy to reverse. They keep the underlying facts, the purpose and final responsibility for what they send.
LIKELY PROFILE	Cognitive authorship: shared or AI led at the level of wording. Decision authority and explanatory ownership: human led.
INTERPRETATION	An AI led pattern can suit bounded, low stakes and reversible tasks. The framework separates this from AI led authority in tasks with greater judgment or consequences.

B8 Research synthesis

PROFESSIONAL, EVIDENCE DEPENDENT

SCENARIO	A researcher asks AI to summarize a body of literature and uses the summary as the basis for a report without opening the cited papers. The references look plausible and the report reads well. Several claims are wrong, and some cited studies do not support the conclusions attributed to them.
LIKELY PROFILE	Cognitive authorship: heavily AI led. Verification: weak. Explanatory ownership: may appear strong until challenged. Capability retention: unclear. The final conclusions inherit the system's errors.
INTERPRETATION	Fluency can hide weak epistemic control. In research and evidence synthesis, check the claims that carry the argument. Coherent prose is not evidence of accuracy.

APPENDIX C

Organizational review template

This template turns the framework into a review process for teams that use generative AI in recurring workflows. It supports governance, quality assurance, professional practice and internal policy discussion. It is not a compliance certification.

C1 Define the workflow

Name the task or workflow. Describe its purpose, who performs it, where AI enters, what AI produces and who is responsible for the final output or decision.

Workflow	
Purpose	
AI system or tool	
Human roles	
Final accountable person or function	

C2 Map the six agency domains

For each domain, record where authority sits now and the evidence for that judgment.

DOMAIN	REVIEW QUESTION	OBSERVED STATE	EVIDENCE
Independent initiation	Who defines the problem, objective or initial direction?	☐ Human led ☐ Shared ☐ AI led	
Cognitive authorship	Who shapes the structure, priorities, assumptions and direction of the work?	☐ Human led ☐ Shared ☐ AI led	
Verification and epistemic judgment	Who decides what needs checking, and what to trust or reject?	☐ Human led ☐ Shared ☐ AI led	

DOMAIN	REVIEW QUESTION	OBSERVED STATE	EVIDENCE
Explanatory ownership	Can the responsible person explain the reasoning, assumptions and result independently?	☐ Human led ☐ Shared ☐ AI led	
Capability retention	What essential knowledge or skill remains when AI is unavailable, unsuitable or wrong?	☐ Human led ☐ Shared ☐ AI led	
Decision authority	Who makes the consequential choice, and can that person realistically choose against the AI?	☐ Human led ☐ Shared ☐ AI led	

C3 Apply the context layer

FACTOR	REVIEW QUESTION	NOTES
Task type	What kind of cognitive or professional work does the workflow involve?	
Stakes	What could be lost if the output or decision is wrong?	
Reversibility	How easily can an error or decision be corrected?	
Verifiability	Is there reliable external evidence to check the AI output against?	
AI experience	Do users understand the system well enough to recognize its limits?	
Frequency of use	Is repeated use changing how people perform or learn the task?	
Consequences for others	Who beyond the direct user may be affected?	
Time pressure and workload	Do working conditions make real human review unrealistic?	

C4 Identify agency critical points

Not every domain needs the same level of human control in every workflow. Mark the points where human authority is essential for safety, learning, professional responsibility, rights or accountability.

POINT	AGENCY CRITICAL POINT	WHY IT MATTERS	REQUIRED HUMAN ACTION
1			
2			
3			

C5 Specify safeguards

Possible safeguards include independent source checks, required explanation of key assumptions, escalation for low confidence cases, regular unaided practice, peer review, explicit decision ownership, limits on AI use for specific steps and design changes that make intervention easier.

SAFEGUARD	DOMAIN PROTECTED	OWNER	EVIDENCE THAT IT WORKS

C6 Review outcomes separately

Record any gains in effectiveness, efficiency, learning or access, and any new dependence, capability loss, opacity or unclear responsibility. Good performance does not prove that agency is preserved.

Effectiveness	
Efficiency	
Learning or capability gain	
Dependence or capability loss	
Responsibility or trust concerns	

C7 Final judgment

The review does not aim to maximize human control at every step. It decides if the distribution of authority suits the task and its consequences.

Is the current authority profile appropriate? ☐ **Yes** ☐ **Partly** ☐ **No**

What should change?

Who is responsible for the change?

When should the workflow be reviewed again?

Apply the framework again when the task, the AI system, the level of automation, the stakes or organizational responsibilities change materially.

NEURAVOX POLICY LAB

References

Abako, G. L. (2026). Human Agency in AI Use: Exploratory findings from surveyed generative AI users. Neuravox Policy Lab, Neuravox Foundation.

African Union. (2024). Continental Artificial Intelligence Strategy. African Union Commission. https://au.int/en/documents/20240809/continental-artificial-intelligence-strategy

Buijsman, S., Carter, S. E., & Bermúdez, J. P. (2025). Autonomy by design: Preserving human autonomy in AI decision support. Philosophy & Technology, 38(3), 97. https://doi.org/10.1007/s13347-025-00932-2

Carolus, A., Koch, M., Straka, S., Latoschik, M. E., & Wienrich, C. (2023). MAILS: Meta AI Literacy Scale. arXiv:2302.09319. https://arxiv.org/abs/2302.09319

Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: A systematic review of frequency, effect mediators, and mitigators. Journal of the American Medical Informatics Association, 19(1), 121-127. https://doi.org/10.1136/amiajnl-2011-000089

Jin, Y., Martinez-Maldonado, R., Gašević, D., & Yan, L. (2024). GLAT: The Generative AI Literacy Assessment Test. arXiv:2411.00283. https://arxiv.org/abs/2411.00283

Lee, H. P. H., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. In Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (pp. 1-22). https://doi.org/10.1145/3706598.3713778

OECD. (2024). OECD AI Principles. Organisation for Economic Co-operation and Development. https://www.oecd.org/en/topics/ai-principles.html

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. Human Factors, 52(3), 381-410. https://doi.org/10.1177/0018720810376055

Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. Trends in Cognitive Sciences, 20(9), 676-688. https://doi.org/10.1016/j.tics.2016.07.002

Romeo, G., & Conti, D. (2026). Exploring automation bias in human-AI collaboration: A review and implications for explainable AI. AI & Society, 41(1), 259-278. https://doi.org/10.1007/s00146-025-02422-7

Soto-Sanfiel, M. T., Angulo-Brunet, A., & Lutz, C. (2025). The scale of artificial intelligence literacy for all (SAIL4ALL): Assessing knowledge of artificial intelligence in all adult populations. Humanities and Social Sciences Communications, 12, 1618. https://doi.org/10.1057/s41599-025-05978-3

Sturgeon, B., Samuelson, D., Haimes, J., & Anthis, J. R. (2025). HumanAgencyBench: Scalable evaluation of human agency support in AI assistants. arXiv:2509.08494. https://arxiv.org/abs/2509.08494

UNESCO. (2021). Recommendation on the Ethics of Artificial Intelligence. United Nations Educational, Scientific and Cultural Organization.

UNESCO. (2024). AI competency framework for teachers. United Nations Educational, Scientific and Cultural Organization.

Zhu, Q., Li, X., Dong, Y., Chang, P., & Fan, M. (2026). Not all cognitive offloading is equal: Distinguishing dependent and autonomous offloading to generative AI. Frontiers in Psychology, 17, 1878629. https://doi.org/10.3389/fpsyg.2026.1878629

neuravox
PUBLIC INTEREST TECHNOLOGY

Neuravox Foundation is a public interest technology and AI governance organisation. The Neuravox Policy Lab studies how people and institutions can gain value from AI while keeping human judgment, skill and responsibility active.

WEB
neuravox.org

CONTACT
research@neuravox.org