The Architecture of Thinking and Acting in Digital Societies

The Entinexus Fundamentum

by Clemens Enti.nexus

Fundamentum — Version 0.14 | 25 November 2026

Status: Open to discussion and revision.

Human worth cannot depend on continuing to outperform machines. A defence of humanity built on permanent superiority in calculation, language or problem-solving would make our value conditional on an unstable comparison. The right to participate in the direction of a human life needs a firmer foundation.

The same applies to authority. Being better at an activity does not, by itself, confer the right to decide why that activity matters, whose interests it should serve or what commitments a person ought to make. We can recognise the force of a reason without transferring authorship of our lives to the source that supplied it.

Kant's distinction between price and dignity helps express what is at stake. Skills can be valued for what they produce; a person calls for respect beyond their usefulness.[1] Entinexus takes this as a starting point for its own commitment: differences in ability or dependence should never determine whose life deserves consideration. The organisation of intelligence should serve lives that already matter.

This distinction makes room for an ambitious relationship with artificial intelligence. AI could participate in the development of our understanding, the discovery of possibilities and the coordination of action. Its contribution could become more substantial while our involvement in the direction of our lives becomes more consequential. The compatibility of these ambitions deserves investigation.

The Entinexus Manifesto proposes organising our thinking and acting with AI around individual human purposes, development and participation in life.[2] This Fundamentum develops the reasoning behind that commitment. Its central proposition is that how we organise a continuing human–AI system can affect whether greater artificial capability strengthens human agency. The task is to discover which arrangements help, whom they help and where they fail.

Making room for artificial participation

The word “tool” carries several meanings. It can identify something humans have made, something employed in pursuit of purposes, or something expected to execute instructions passively. These meanings allow different conclusions. Calling AI an artefact says little about the initiative it should exercise. Describing it as serving human purposes leaves open whether it may help us understand and revise those purposes.

The restrictive position is the demand that AI contribute only after a person has fully specified the problem, the objective and the acceptable method. Such an arrangement may be appropriate for many tasks. Making it a universal principle would exclude a valuable possibility: artificial participation in recognising what needs to be considered in the first place.

Consider the difference between asking for a route to an established destination and examining whether that destination still deserves our effort. In intellectual work, AI might identify an assumption that makes the original question misleading. In planning, it might reveal a conflict between commitments. These are possible contributions to deliberation itself. Their value depends on the quality of the reasons offered and the conditions under which we can assess them.

Thinking can depend on notes, diagrams and other resources outside the body. Andy Clark and David Chalmers argued that, under suitable conditions, such resources can form part of a cognitive process.[3] Their extended mind argument gives us a reason to examine how thinking is organised across a person and the resources they use. Whether a particular AI system forms part of a person's mind remains a separate question.

We use participation to describe the active contribution of artificial intelligence within an organised process of thinking and acting. This participation concerns the role AI plays in the activity and does not, by itself, establish consciousness, personhood, or independent moral authority. It may include proposing, criticising, remembering and coordinating.

Capability and authority

Once AI is allowed to participate actively, we need to distinguish three things. Capability concerns what a system can accomplish. Operational autonomy concerns how much it can do without a fresh instruction. Authority concerns the decisions it is entitled to make within an arrangement and how that entitlement can be revised.

These dimensions do not necessarily develop together. An AI may have considerable autonomy in gathering material or organising a workflow while the person retains authority over the commitments that workflow serves. Conversely, a system that asks for approval at every stage may still exert extensive influence by selecting the options, framing the questions and determining what information reaches the person.

Aristotle distinguished cleverness in achieving a target from practical wisdom about acting well.[4] The distinction remains useful whenever greater skill leaves the worth of an

objective unsettled. AI can contribute to this judgement too: it may expose a mistaken aim or help us see a better one. Taking its reasoning seriously remains compatible with asking who is entitled to make the resulting commitment.

The case for human authority begins with the person whose life is being shaped. Their attachments, vulnerabilities and possibilities give them a standing that superior performance cannot simply displace. Agency means being able to form, reconsider and pursue purposes, and to have one's judgement make a practical difference, with support where appropriate. A person's worth remains intact when that participation is difficult or limited.

For a personal Entinexus configuration, the governing commitment is that the person retains effective authority over the arrangement's purposes and terms of participation. This authority includes adopting, questioning and revising purposes, redistributing responsibilities and choosing how much AI involvement is wanted. The authority to direct a personal AI system does not, by itself, give a person the right to make decisions for other persons or disregard their rights and interests.

Purposes are not always settled before we begin. They can become clearer, or change, through the activity of thinking and acting. Dewey described “ends-in-view”: anticipated outcomes that arise within activity and help guide what happens next.[5] An AI contribution may supply the argument that makes a new direction worth pursuing. The emerging purpose should remain open to the person's informed judgement, including their capacity to reject or revise it.

Human authority therefore needs a richer expression than habitual agreement with whatever the person says. A system that conceals a contradiction to preserve satisfaction may make judgement harder. A system that offers a well-supported objection may help the person exercise authority more intelligently. Respect for the person can include intellectual resistance, provided the contribution remains answerable to reasons and the arrangement remains open to redirection.

Authority over an arrangement does not make the person infallible. A valuable artificial participant should be able to help us discover that we are mistaken, while leaving us able to examine the reasons for its challenge. Evidence and criticism remain necessary, whoever first proposes an idea.

What makes a continuing whole

Entinexus proposes examining the continuing system within which these interactions occur. The person, artificial participants, retained context, responsibilities and practices can form an organised whole directed towards particular purposes. The human being remains a person with a life exceeding that arrangement. Participation in the system should serve that larger life.

A continuing arrangement could carry an unfinished inquiry across interruptions, connect discoveries from different activities and preserve a correction so that it changes later work.

It could help a person develop an idea through successive encounters and translate growing understanding into action. A purpose could guide several encounters while remaining open to revision. The potential benefit is a greater capacity to develop and pursue purposes over time.

These connections are what give the arrangement its unity. A correction matters when it changes later conduct. A responsibility matters when it remains clear across interruptions. We have reasons to examine the whole when the relations among contributions help explain what becomes possible. A list of tools and saved conversations, by itself, tells us little about how they work together.

A person can organise research and editorial work under different terms within one continuing system. Suppose an objection uncovered in research changes what a writer is prepared to publish. The inquiry has served its purpose by exposing a weakness, and the writer can revise the argument or defer publication. The connection matters precisely because each responsibility can be pursued on its own terms.

Within an Entinexus, Artifex names the human participant’s governing role; an Opifex is an artificial participant with a continuing identity and authorised remit specific to that person.[6] The Framework develops their proposed formation conditions. Naming an assistant does not establish those conditions, and practical criteria remain open to investigation.

The whole is also distinct from any particular model or interface. Changing a component raises a practical and philosophical question: which relations, commitments and history must survive for the same arrangement to continue? We still need to establish which changes preserve the system and which create a different one. Calling it an organised whole does not by itself establish consciousness or a distinct intelligence belonging to the system as a whole.

To organise this well, we need to understand who participates and what holds the system together, how contributions combine in actual activity, and how authority is exercised and revised. Entinexus studies these questions through constitution, operation and governance.[6] They continually affect one another. A difficulty encountered in practice may reveal that a role is unclear or a mandate no longer fits. A useful system must let such discoveries change how it is organised.

A mandate, understood as a continuing authorised remit, can support substantial artificial initiative while leaving the person able to revise the direction of the work. A shared written account is one way to make that remit understandable and revisable. Its usefulness depends on whether it clarifies action without imposing excessive work; other practices may prove better. A larger collection of instructions, memories or automated processes is not automatically a better organisation.

Dewey understood freedom in practical terms: being able to act, change direction and make choices that affect what happens. He recognised that organisation could enable these freedoms or constrain them.[5] This gives us a question to bring to human–AI systems: what does a particular arrangement make possible, compared with the alternatives available to

the person? Organisation earns its place through what it makes possible and the burdens it imposes.

The quality of human participation matters more than its frequency. A brief reconsideration that changes a project's direction may carry more agency than a succession of routine approvals.

Performance is an incomplete measure

Research gives reason to examine the arrangement closely. A 2024 analysis combined results from 106 experiments. On average, people working with AI performed better than people working alone. The combination nevertheless performed worse, on average, than the better of the two alternatives: people working alone or AI working alone. Results varied across tasks and settings.[7] Combining participants is therefore no guarantee of superior performance. These findings concern specified tasks; the question of continuing personal agency remains open.

Research on learning makes another distinction visible. A randomised study involving nearly 1,000 pupils at one Turkish high school compared two forms of GPT-4 assistance. Both improved performance during mathematics practice. When AI access was removed, pupils who had used the basic interface performed worse than pupils in the control group. A tutor designed to support learning largely avoided that disadvantage, but the study did not establish an improvement in the subsequent examination without AI.[8]

These findings show why evaluation must take the purpose of an arrangement into account. When the aim is learning, successful assistance must be assessed partly through what the learner comes to understand and can subsequently do. When the aim is completing an unwanted administrative task, relief from that task may be an appropriate benefit. The same degree of delegation can have different significance under different purposes.

Independence is consequently one possible objective, rather than a universal test of worthwhile support. A person may legitimately want continuing assistance. Their agency may be expressed through choosing and directing that assistance, including with support from others. Requiring everyone to perform unaided would make a particular form of self-sufficiency the measure of human value.

Nor should “human agency” become an excuse for poor results. A system that preserves elaborate opportunities for discussion but consistently fails to help may have little practical value. Evaluation needs to consider useful outcomes together with the person's capacity to understand, redirect and benefit from the activity. The relative importance of these dimensions depends on the purpose and the person.

Keeping authority effective

The strongest objection to the argument is that authority may survive in name while disappearing in practice. If AI increasingly supplies the alternatives, interpretations and reasons through which a person decides, what makes the resulting direction meaningfully theirs?

Appealing to a wholly uninfluenced human will would not solve this problem. Deliberation takes place through language, relationships and inherited ways of understanding. The question is whether influence can be examined and challenged, whether meaningful alternatives remain accessible and whether the person can change course with practical effect.

Accepting a contribution can itself express agency when the person understands enough of its significance to endorse it and can reconsider that endorsement. Thoughtful agreement, exploration and willingness to be changed belong within human participation. A system should make these possibilities available alongside disagreement.

The idea of meaningful human control offers a way to examine this problem. In one philosophical account, a system should respond to relevant human moral reasons and the facts of the situation. Its behaviour should also be traceable to people who understand its relevant capabilities and limitations, and recognise the moral responsibility carried by their role.[9] Applying this account to personal AI raises questions that include obligations to other people, beyond satisfying the immediate user's preferences.

Applied as a question for Entinexus, this suggests examining whether authority has practical support. Can a person understand the consequential choice at an appropriate level? Can they contest the interpretation on which action depends? Does their correction alter subsequent behaviour? Are there situations in which the arrangement demands more attention or expertise than they can realistically provide?

Lisanne Bainbridge's analysis of industrial automation identified a related difficulty: people can retain responsibility for difficult interventions while automation reduces the involvement through which relevant skills and situational understanding are maintained.[10] Its application to continuing AI assistance is an analogy to investigate, not direct evidence that such assistance inevitably erodes human capacities.

An organised system would therefore need to support the participation it expects from the person. This might involve preserving opportunities for practice where skill matters, making important disagreements visible, or obtaining independent evidence for consequential claims. The appropriate support will differ across activities and people. Constant supervision can exhaust the very attention the arrangement is supposed to release.

Effective authority may thus require carefully chosen occasions for engagement, together with substantial freedom for the system to act between them. The design problem is to discover where human participation changes the direction of the activity and how to make that participation feasible. Approval alone is an inadequate measure of this relationship.

Continuity with freedom to change

A continuing human–AI system must also address time. Retained context can help it connect commitments, recognise unfinished work and avoid repeated explanation. But an account of what a person has previously said can acquire a prescriptive force: yesterday's preference may quietly become tomorrow's rule.

Human individuality includes the possibility of becoming different. An arrangement organised around a person should distinguish their current commitments from past statements and from the AI's interpretations. These have different standing. A pattern inferred from previous behaviour may be a useful hypothesis about the person; it should remain open to correction.

This matters especially when change initially looks inconsistent. A person may wish to experiment with an unfamiliar activity, reconsider a conviction or accept a trade-off they previously avoided. A system committed to continuity should be able to support such exploration. The person's history supplies context for judgement while their present participation keeps that history open to reinterpretation.

The freedom to revise an arrangement also depends on conditions outside it. A personal system can promise portability or continuity only to the extent that its infrastructure allows them. If records cannot be exported, essential functions disappear when a provider changes its service, or replacing a component requires rebuilding accumulated context, the person's practical choices narrow. Preserving records is one part of continuity; enabling their meaningful use in later activity is another.

Personal organisation therefore belongs within a wider institutional question. Providers, employers and public institutions help determine which forms of participation are possible and who bears their costs. A carefully designed individual arrangement cannot by itself resolve those relations of power. Nor can “human authority” be treated as a single interest when several people are affected and their purposes differ.

Entinexus's personal focus is a deliberate starting point. Its legitimacy depends partly on how well it respects the people beyond its immediate participant and how realistically it acknowledges the conditions on which that participant depends. The freedom to shape a personal system carries responsibilities within a shared world.

Beyond the Artifex’s lifetime

A person’s concern for the future can extend beyond their lifetime. They may want an inquiry completed, a body of work preserved, or an undertaking to continue serving others. The legal will provides a familiar example of decisions made during life being given effect after death, within conditions established by law.[11] The end of a person’s participation need not erase every practical significance of their earlier choices.

An Opifex could offer a further way of giving effect to such intentions after the Artifex’s death. In carrying an undertaking forward, it might draw on methods developed with the

person and on the reasons that led them to revise those methods in the course of their work together.

Implementation can already involve interpretation and judgement; artificial participation would not introduce those questions for the first time. Its possible contribution would lie in how it connects an earlier purpose with the circumstances encountered later. Adapting the means of pursuing a purpose nevertheless differs from adopting a new governing purpose or enlarging authority. Where that boundary lies, and how it should be governed, require examination.

We may know what the Artifex hoped to accomplish during life and still be uncertain what they wanted to happen to the work after their death. What would give us grounds for thinking they intended it to continue? Even where that intention is clear, a further question concerns who has authority to act on it. A personal mandate does not by itself settle the powers of an Opifex after death or its relationship with those entitled to carry out the person’s directions.

Continuing an undertaking also leaves open what continues as an entity. The activity, an Opifex, the Entinexus and the human person present different claims of continuity. The Artifex’s participation helps constitute the whole. Could that founding relationship remain constitutive through the history and organisation it established, or would a successor carry forward the work of an Entinexus that had ended? Neither useful operation nor a familiar representation of the person settles these questions. The value of continuing an undertaking can be examined without assuming that either the same Entinexus or the human person persists.

Any continuing activity would affect people whose circumstances and purposes may change. The Artifex’s earlier directions must be considered together with the rights and interests of those affected and the responsibilities of those carrying out the work. Continuing operation, preserving its results, transferring an undertaking or bringing it to an end may each be appropriate under different conditions. These alternatives need assessment in relation to the purposes involved.

The philosophical coherence of a continuation, its justification and its practical feasibility require different kinds of support. Examining what might follow death can also clarify what matters during life: whether the arrangement’s history and purposes guide activity in ways the person can understand, question and revise.

Testing the Proposition

Research must investigate whether particular forms of organisation enable more capable artificial participation to develop alongside effective human direction over time, and under what conditions this improves a person's life.

Research should compare continuing organised arrangements with simpler forms of assistance under comparable conditions, accounting for differences in AI capability and the work needed to establish and maintain them. It should follow changes in purposes and

activity over time, and assess useful outcomes alongside the person’s practical ability to understand, redirect and revise the arrangement.

One specific hypothesis is that making a continuing mandate explicit and revisable helps people redirect ongoing activity more successfully than successive requests alone. Any improvement would need to justify the extra effort involved.

Benefits and burdens may differ among learning, practical delegation, exploration and continuing support. Detailed comparisons and methods should be developed through the Framework and Research.

These are proposed lines of inquiry, not an established measurement instrument. Agency itself requires conceptual care: the ability to reject a suggestion is relevant, but so are the ability to form a purpose, obtain support and bring an intention into the world. No single score should silently decide what counts as a worthwhile human life.

The programme must permit findings that challenge its premise. If simpler arrangements deliver equal benefits with less burden, greater organisational complexity has not earned its place. If people report control while repeated attempts at redirection fail, nominal authority has displaced effective authority. If the organisation resists changes in the person's purposes, it has become an obstacle to the life it was meant to serve.

Such results would call for revision. Each part of Entinexus has a different responsibility. The Manifesto states what we are committed to. The Fundamentum examines why those commitments deserve pursuit and what they assume. The Framework develops the concepts and relationships through which they can be understood and pursued. Research examines their feasibility, consequences and limits. These four parts can correct and enrich one another: experience may expose a poor arrangement, while philosophical criticism may change the questions by which we judge it.

Empirical success cannot establish human dignity, and empirical failure cannot abolish it. Findings can nevertheless reveal costs, conflicts and exclusions that demand a reconsideration of our methods or the way our commitments are expressed. A foundation gives inquiry a considered starting point. Its strength includes making its assumptions available for examination and its conclusions answerable to criticism.

The Commitment

The aspiration is to welcome the fullest worthwhile contribution artificial intelligence can make to human thinking and acting. That includes helping us formulate questions, reconsider purposes, discover possibilities and undertake activities that would otherwise exceed our resources. It also includes doing enough on our behalf that we have more attention for the parts of life in which we wish to participate directly.

Humanity exceeds the capacities investigated here. Affection, embodiment, vulnerability, play, rest and belonging need no justification through enhanced performance. A person's worth remains intact when they require extensive support, when they choose a modest

ambition or when they set the system aside. The freedom to decline greater integration belongs within the freedom the project seeks to support.

The commitment of Entinexus is to organise for that freedom. Its hypothesis is that artificial initiative and meaningful human participation can develop together through arrangements that remain understandable, revisable and responsive to a changing life. Its obligation is to discover where that hypothesis holds and where it fails.

We should make room for intelligence wherever it can contribute, and take responsibility for the relations through which it enters our lives. Our humanity needs no victory in a competition of capabilities. Our future requires a consequential human part in deciding what those capabilities are for.

Sources

[1] Immanuel Kant. Groundwork of the Metaphysics of Morals (1785; second edition, 1786). Section II, Academy pagination 4:428–435, on persons as ends in themselves and the distinction between price and dignity. English translation by Stephen Orr. Kant grounds dignity in rational moral agency; the commitment here across differences in ability is Entinexus's own. Read the primary text

[2] Clemens Enti.nexus. Let’s Organise Our Thinking and Acting in the Age of AI. Entinexus Manifesto, version 0.27, 20 November 2026. Public working text. Read the Manifesto

[3] Andy Clark and David Chalmers. The Extended Mind. Analysis 58(1), 7–19, 1998. DOI: 10.1093/analys/58.1.7. Read the authors' text

[4] Aristotle. Nicomachean Ethics. Book VI, chapters 5 and 12, especially 1140a24–1140b30 and 1144a23–36. Translated by W. D. Ross. Read Book VI

[5] John Dewey. Human Nature and Conduct: An Introduction to Social Psychology (1922). Part III, section VI, “The Nature of Aims”, pp. 223–237; Part IV, section III, “What is Freedom?”, pp. 303–313. Read “The Nature of Aims” · Read “What is Freedom?”

[6] Clemens Enti.nexus. Entinexus — The Constitution, Operation, and Governance of a Personal Human–AI System. Doctrinal working synthesis, V0.5, English edition, 17 November 2026. Sections 1.3 and 2.1–2.4. Project working text. The three research directions are named Constitutio, Operatio and Gobernatio. Individuation concerns the identity of a participant or whole; individualisation concerns how a relationship with a person influences the organisation and exercise of capacities. Their relationship remains under investigation. For the current public account of these roles and distinctions, see also Conceptual Framework for Entinexus, version 0.23, 20 November 2026, §§1–2.

[7] Michelle Vaccaro, Abdullah Almaatouq and Thomas W. Malone. When combinations of humans and AI are useful: A systematic review and meta-analysis. Nature Human Behaviour 8, 2293–2303, 2024. DOI: 10.1038/s41562-024-02024-1. Read the research paper

[8] Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı and Rei Mariman. Generative AI without guardrails can harm learning: Evidence from high school

mathematics. Proceedings of the National Academy of Sciences 122(26), e2422633122, 2025. DOI: 10.1073/pnas.2422633122. Read the research paper

[9] Filippo Santoni de Sio and Jeroen van den Hoven. Meaningful Human Control over Autonomous Systems: A Philosophical Account. Frontiers in Robotics and AI 5, Article 15, 2018. DOI: 10.3389/frobt.2018.00015. Read the research paper

[10] Lisanne Bainbridge. Ironies of Automation. Automatica 19(6), 775–779, 1983. DOI: 10.1016/0005-1098(83)90046-8. Read the research paper

[11] European Parliament and Council. Regulation (EU) No 650/2012 of 4 July 2012 on succession. Articles 3(1)(a) and (d), 23(2)(f) and (h), and 26(1)(d). Cited for the legal analogy concerning effects after death, executors’ powers, restrictions and interpretation. Read the regulation.