Fondation UvH

STUDIES AND PUBLICATIONS · IAД PROGRAMME

III

Дdvanced artificial intelligence

Biological substrates and hybrid intelligence

Computing, memory, distributed cognition and conditions for coexistence

Fondation UvH

English edition · Translation of French version 1.1

French edition: 18 September 2026 · English translation: 3 October 2026

EDITION AND EDITORIAL RESPONSIBILITY

Editorial notice

Institutional author and publisher	Fondation UvH
Collection	Studies and publications · Programme IAД
Document	Study III · Substrates and hybrid intelligence
Original edition provided	September 18, 2026 · 36 pages
Current edition	September 18, 2026 · Version 1.1
Nature	Documentary, conceptual, and prospective study

This edition differs from the original file. It maintains its structure and substantial developments, with editorial corrections, clarifications, and a presentation consistent across the three studies. Significant changes are identified in the final traceability note. Institutional sources for S-Д.holdings are used to describe the project; however, this does not constitute an independent validation of its assumptions.

The research direction for the series focuses on ensuring equal rights and the absence of subordination or domination between the forms of intelligence involved. It proposes a common set of rights alongside specific rights tied to their inherent natures. The detailed content of these rights, their holders, and their implementation methods remain to be studied. This normative orientation is distinct from the empirical findings reported and does not automatically ascribe consciousness, personality, or citizenship to current systems.

The protocols and measures considered are not presented as having been implemented. The institutional proposals do not alter the founding documents and do not require any applicable authorization. The publication does not demonstrate either peer review or regulatory adoption. Its preparation does not warrant publication or external submission.

Abstract

This study examines what changes when cognitive functions are distributed across electronic, molecular, biological neural, and human systems. It separates calculation, learning, storage, interface, and institutional organization. The brain's energy budgets, calculation in memory, neural cultures, organoids, DNA storage, and brain-computer interfaces are compared based on their experimental scope and overall system performance. The evidence presented demonstrates that living systems do not replace silicon entirely, nor do they universally read thoughts. The rat serves as a model for embodied cognition, not as a readily available computing resource. Proximity to neural activity raises questions regarding consent, interpretation, continuity, and withdrawal. The д-Я.Й concept is discussed as a theoretical framework for learning, distinct from a deterministic programming of the brain. Instrumental, emergent, relational, and metaphysical readings remain distinct. The electronic legal personality and AIR are extended by a new question: how to govern the coupling between participants, without assuming that a new individual has emerged? A common foundation of rights is explored here, focusing on specific protections for the affected natures and vulnerabilities.

Keywords
substrate; biological calculation; DNA; organoid; brain-computer; distributed cognition; symbiosis

Method and scope

The review is based on a reading of the provided corpus, a check of its structure, and targeted research into its references and assertions. However, it does not constitute a systematic review. The results reported remain linked to the populations, models, protocols, and dates of the original sources; hypotheses and interpretations are distinguished from applicable law. If a document is only accessible through its summary or description, this limitation is indicated in the references. References to accompanying studies organize the continuity of the corpus, without providing additional scientific confirmation.

The numbered references within the text and the entries in the table of contents are clickable. The bibliographical notes include the addresses of the sources. Companion files must remain in the same folder for cross-document references; the PDF reader's security settings may require confirmation.

Contents

The headings and page numbers refer to the sections. The reader's PDF bookmarks provide the same structured navigation.

1 Research position, method and continuity of the corpus 5

2 Physical constraints of the digital paradigm and energy lessons from living systems 9

2.1 Comparing the brain and digital computing: boundaries and limitations 11

3 Biological substrates, molecular memory and the rat model 13

3.1 DNA: physical evidence of a different memory regime 13

3.2 Neuromorphic and in-memory computing: an alternative to living tissue 14

3.3 Neuronal cultures and organoids: what has actually been demonstrated 15

3.4 From neuron to rat: six levels that must never be conflated 16

4 From human data to hybrid intelligence 18

4.1 A scale of human–AI proximity 18

4.2 Going directly to the source does not eliminate the data problem 20

4.3 From assistant to extended cognitive system 20

4.4 д-Я.Й — Дlgorithme Яéseau Йeural 21

4.5 Who thinks in a hybrid system? 22

5 Heterogeneous architecture and the scientific status of the trajectories 23

5.1 Mapping the state of the evidence 24

5.2 The argument for biological–artificial convergence 25

5.3 The objection based on system-level accounting and non-biological alternatives 26

5.4 What would actually refute the value of biological approaches 26

6 Neurorights, existing law and the S-Д.holdings institutional laboratory 28

6.1 Switzerland 28

6.2 European Union 29

6.3 Council of Europe and international law 29

6.4 Towards institutional mental sovereignty 30

6.5 S-Д.holdings as a research environment 32

6.6 Proposed institutional limits 33

7 Critical analysis, philosophy of symbiosis and mapping of possible futures 35

7.1 Have we 'created' artificial intelligence? 35

7.2 Symbiosis as a third trajectory 35

7.3 The new tension with electronic legal personality 36

7.4 Mapping possible futures 37

7.5 Answers to the programme's final questions 38

8 Conclusion: substrate, rights and relationship 40

Shared glossary for the series 42

References and sources 44

Corpus, revision method and traceability 51

1 Research position, method and continuity of the corpus

The initial study within the program's framework raised a question of status: under what conditions could sufficiently autonomous artificial systems acquire their own legal capacity? It examined the arguments for and against the concept of "electronic legal personality," the risks of diluting liability, and hybrid public-private models. Its revised edition maintains this contradictory analysis and explicitly outlines the chosen approach for the corpus – a common foundation, with specific rights justified by the nature of their holders. This approach does not describe a personality already established by positive law. [1]

The second study shifts the analysis of status towards power relations. It defines IAД by its capacity for autonomy, planning, continuity, and institutional understanding, without presupposing consciousness or sentience. Its Reciprocal Institutional Alignment program (AIR) asks whether the experience of reciprocal institutions causally influences behavior when the distribution of power changes. The empirical test focuses on this effect; the justification of common or specific rights remains a separate question. The third study maintains this separation and examines what it becomes when intelligence and its functions are distributed across multiple platforms. [2]

This study introduces a third movement:

Following status and power, the question shifts to that of the underlying substrate and cognitive boundary: where does intelligence—when digital processes, biological structures, external memory, and human cognition begin to form interdependent loops—physically and functionally reside?

This movement does not render the two previous studies obsolete. Rather, it makes them more important. When an AI ceases to be clearly located within an external machine and when a portion of its cognitive function depends on an organism, a neural network, a molecular memory, or a human brain operating in conjunction with it, the categories of electronic person, persistent identity, responsibility, non-domination, and institutional coexistence become increasingly difficult to define. The biohybrid hypothesis thus challenges the conceptual boundary that had implicitly underpinned much of the previous two studies: the distinction between biological human and electronic intelligence.

Overall assessment of the research. The available data justify taking alternative physical regimes for computation, learning, and memory very seriously, but they do not justify the thesis that living organisms are currently ready to replace silicon. The evidence instead suggests a heterogeneous trajectory: classical numerical computing, memory-based computing (both in-memory and neuromorphic), molecular supports for certain types of memory, potentially biological neural modules in highly specialized applications, interfaces with humans, and, most importantly, a more plausible short-term integration of human-AI information processing rather than a transformation of organisms into computer hardware. Neural cultures and organoids have already performed limited experimental calculations and learning; brain-computer interfaces already allow for bidirectional decoding and restitution in specific clinical conditions. DNA demonstrates that extraordinarily dense

and passive storage is possible. However, none of these results demonstrate a general superiority of a complete biological computer over electronic architectures when the entire infrastructure is considered. [3]

The most significant outcome may be conceptual: the most credible hybridization isn't, at this stage, replacing a processor with a brain. It’s the progressive redistribution of cognitive functions across multiple substrates. This essentially means moving away from the necessarily seeking a “super-machine” – instead, focusing on a system where each support performs the tasks for which its physical properties are advantageous. [4, 3]

Epistemological rigor. Five categories will be maintained throughout the study:

Table 1. Qualifications used for results and proposals

Qualification	Significance in this report
Documented result	Observation or result reported in a protocol; however, simply stating this does not imply an independent replication of all cited results.
Experimental	A real proof of concept exists, but its performance, generalizability, reproducibility, or scalability remains uncertain.
Possible, but not proven.	Hypothesis consistent with existing knowledge, supported indirectly, but without sufficient causal evidence.
Speculative	Hypothesis for which the essential mechanisms are not yet established, or for which a distinguishing test has not yet been identified.
Normative / philosophical	Proposal regarding what should be permitted, protected, or considered as a person – and not a scientific finding.

Analytical construction or framework proposed in the present study.

The document hierarchy prioritizes original articles reviewed by peers and legal or official institutional documents. Industrial announcements are only used to illustrate the existence of a technological trajectory, and not as independent evidence of overall performance.

Continuity matrix of the three studies.

Table 2. Continuity and movement between the three studies

Concept of the corpus	First study	Second study	Shift introduced by this study
IAД	An artificial system that is sufficiently autonomous to raise a status-related question.	Functionally defined, explicitly independent of consciousness.	Detailed: the category must now be neutral regarding the substrate; a IAД function could be distributed between electronic and biological components.
electronic legal personality	Central legal assumption	Broken down into gradual abilities rather than in a binary (yes/no) format.	Partial challenge to the term: “electronics” has become materially too narrow for a bio-hybrid system.
Autonomy	Possible grounds for legal recognition	Variables of power and safety	Distributed: autonomy can be shared between humans, models, memory, and interfaces.
Alignment	Implicitly focused on responsibility and governance.	AIR: institutional learning of non-domination	Made more important: cognitive integration creates the possibility of internal or cognitive domination, not only institutional domination
Non-domination	Question of coexistence; principle explicitly stated in the revised edition.	Central human–IAД principle	Scope: protection against cognitive manipulation, mental inference, and infrastructure dependence.
Persistent identity	A question associated with the artificial subject of law	Model/version/instance/ institutional identity distinguished	Complexified: biological identity, digital memory, personal model, and externalized cognitive functions can diverge.
Responsibility	It should not be artificially transferred to AI.	Responsibility chains and contestability	Detailed: a hybrid decision requires traceability, without allowing humans, manufacturers, or systems to become legally invisible.
Rights and obligations	Possible personality; revised edition: common rights foundation and specific rights.	Normative equality distinct from capabilities and the safety protocol.	Protection of autonomy and integrity in deep connections; no inherent superiority is inferred for any category.
Coexistence	Humans, legal persons and possible electronic persons	Shared institutions under non-domination.	Transformed: coexistence can become intra-systemic; the two intelligences are no longer necessarily two separate actors.
Institutional experimentation	S-Д.holdings as a prospective laboratory	Simulation → Observation → Recommendation → Limited co-decision	Specific limitation: the sandbox status does not authorize human or animal research.

Concept of the corpus	First study	Second study	Shift introduced by this study
S-Д City	Virtual legal territory and a physical gathering point; urban project detailed in subsequent sources.	Digital twin and progressive increase in power.	More in-depth: exploring the potential for studying extended cognitive systems and their rules, initially through simulation and non-invasive interfaces.
Institutional architecture	Natural, legal and possibly electronic persons	Graded capacities, anti-Sybil measures, remedies, oversight	Broadened: governing hybrid cognitive systems may be necessary without creating a fourth category of persons
Civilization question	Who can become a subject of law?	How to limit the power of agents?	New formulation: Where does an agent stop when the intelligence itself is distributed?

Sources and methodology of the present study [1, 2, 5]

This matrix reveals a significant initial tension with the first study: if a future intelligence comprehends biological components or extends through a human brain, the term “electronic legal personality” may continue to be used effectively to describe certain purely artificial entities, but it can no longer, on its own, constitute a general ontology. The third study therefore leads to a proposal for a principle of substrate neutrality: rules should be determined primarily based on capabilities, risks, protected interests, dependencies, and responsibility, rather than solely on the fact that a system is made of silicon, cells, or a combination of both. This proposal is normative; it does not describe positive law. [1, 2]

A second level of continuity is even stronger. The second study concluded that providing more status, recognition, or autonomy does not necessarily lead to a consistently safer outcome. This section explores the biological analogue: getting closer to the neuronal source doesn't automatically result in better data, and obtaining better data never equates to gaining more legitimacy in utilizing it. The more access approaches mental processes, the constraints regarding purpose, consent, safety, the possibility of withdrawal, and non-manipulation should, conversely, become stronger. [2]

2 Physical constraints of the digital paradigm and energy lessons from living systems

The initial question should be posed without resorting to caricatures of contemporary computing. We are no longer simply building large, general-purpose computers like the von Neumann machines. Massive parallel accelerators, high-bandwidth memories, specialized interconnects, in-memory computing architectures, and neuromorphic processors represent precisely attempts to overcome the limitations of the traditional paradigm. Demonstrations of in-memory computing have already shown that it is possible to perform certain operations of neural networks directly near the memory, rather than continuously moving data between storage and the processing unit. [6]

It would therefore be inaccurate to conclude that the industry is simply clinging to historical architecture. However, it is scientifically defensible to assert that the movement of information has become one of the central physical costs of modern digital intelligence. A study of resistive materials, within the specific GPU workload analyzed, demonstrated that data movement could represent over 98% of the energy involved; while this figure is not generalizable to all AI, it illustrates the nature of the problem. Analog and in-memory architectures are precisely being studied in order to reduce the separation between storage and computation. [7]

Energy and infrastructure. The International Energy Agency estimated that data centers consumed approximately 415 TWh of electricity in 2024, representing around 1.5% of global electricity consumption, and projected nearly 945 TWh in 2030. In a modern data center, servers account for approximately 60% of electricity usage, while cooling can represent around 7% in the most efficient hyperscalers, but exceed 30% in less efficient facilities. AI is identified as one of the main drivers of growth in this demand. [8]

The key issue also concerns the spatial and temporal concentration of demand. The IEA distinguishes between conventional facilities of a few tens of megawatts and specialized AI centers that can exceed 100 MW. For the United States, a report by LBNL published in 2024 projects that data centers will account for between 6.7 and 12% of national electricity consumption by 2028; the article by Kurtz and his co-authors on thermal integration, published in 2026, reiterates this range. It is important to consider population, year, and scope – these figures relate specifically to US data centers, not just AI or a global estimate for 2030. [9, 10, 11]

In 2026, the IEA estimated global data centre electricity consumption at approximately 485 TWh for 2025 and maintained a trajectory close to 950 TWh by 2030. It also noted that certain modern uses — longer reasoning, video and agentic systems — can consume orders of magnitude more energy per task than simple text generation. These comparisons must remain dependent on the model, reasoning duration, hardware and utilisation rate: there is no universal value for “energy per AI query”. [12]

Training and inference must be treated separately. Training involves a concentrated expenditure: large amounts of computation are utilized for a limited period to modify the model’s parameters. Inference, on the other hand, is repetitive: each individual use is

typically much less costly than a full training cycle, but the aggregation of billions of uses can become significantly substantial. Therefore, a serious comparison must specify which model, which hardware, which sequence length, which utilization rate, which numerical precision, which data center, and which timeframe. Simply stating that “AI consumes X” without these variables mixes together different physical realities. AIE’s projections demonstrate that rapid gains in unit efficiency are not necessarily sufficient to reduce overall demand when accompanied by a significant increase in usage. [8]

Cooling. The electrical energy used by equipment ultimately manifests as heat. The utilization of this heat depends largely on its temperature, location, and proximity to demand. In 2026, Kurtz and his co-authors are examining possibilities for integrating thermal systems into AI infrastructure; the amount of heat that can be technically recovered varies depending on the specific system and does not automatically represent a reduction in electricity consumption. Utilizing heat does not eliminate the originally used electricity or the cooling, transport, and service continuity constraints. [11]

Semiconductors. The current paradigm relies on highly sophisticated value chains: advanced logic, memory, packaging, interconnects, and manufacturing equipment and materials. A 2026 industry report from the Semiconductor Industry Association and Deloitte estimates that semiconductors represent more than half of the hardware investment in an AI center and over 95% of the hardware value of an AI server rack – these are industry estimates, not a physical law, but they illustrate the concentration of modern AI around advanced microelectronics. [13]

Data and scaling. The notion that simply continuing to increase the number of parameters would suffice is already too simplistic. DeepMind’s Chinchilla research demonstrated that, given a specific computational budget, the balance between the amount of data and model size must be carefully considered: a 70-billion-parameter model trained with more data was able to outperform much larger, under-trained models. The fundamental result isn’t that an absolute limit on scaling has been discovered, but rather that the number of parameters is only one dimension of the optimization problem. [14]

Academic projections suggest that readily available stocks of high-quality public human text could become a constraint during this decade. However, these estimates heavily depend on assumptions regarding data growth, training methods, multimodality, private data, and the effective reuse of data – and therefore, should not be presented as a literal “exhaustion of the Internet” date. [15]

The synthetic generation approach doesn’t automatically solve the problem. Research published in *Nature* has demonstrated that recursive and indiscriminate training on data generated by models can lead to a gradual degradation of the learned distribution – often in the tails of the distribution. This doesn’t mean synthetic data is inherently bad; it can be filtered, verified, combined with real data, and produced in areas where verification is possible. However, it does forbid assimilating “unlimited token generation” with “unlimited creation of high-quality new information.”

[16]

Conversely, the need for data isn't necessarily destined to grow at the same pace as the capabilities. Methods of active learning and example selection can reduce the amount of training required in certain areas; DeepMind has reported, for specific

vision and multimodal experiences, reductions of 46 to 51% in the number of updates and up to 25% in the total computation. Therefore, the future may be determined as much by the efficiency of learning as by the sheer accumulation of data. [17]

2.1 Comparing the brain and digital computing: boundaries and limitations

The human brain consumes approximately 20 watts of metabolic power, which is roughly one-fifth of the human metabolism at rest. If this power were constant, it would mathematically equate to around 175 kWh per year. However, this amount cannot be honestly converted to “FLOPS” or directly compared to the electrical power of a GPU or a data center: the brain does not perform the same calculations; its activity includes cellular maintenance, homeostasis, transmission, plasticity, memory, sensorimotor control, and other biological functions. [18]

Brain energy budgets actually demonstrate that communication is far from free for living organisms. In Attwell and Laughlin’s classic analysis of rodent gray matter, the generation of action potentials and postsynaptic activity represented significant portions of the brain’s energy budget. A subsequent analysis of the human cortex, using their specific model, even estimated a cost of neuronal communication that was significantly higher than the cost attributed to local synaptic calculation. Therefore, the brain hasn’t “solved” the cost of communication; instead, it manages it with a physically very different architecture – highly parallel, economical, localized, and event-driven. [19]

This is precisely the aspect that makes the brain scientifically interesting. It combines, with a metabolic power of just a few dozen watts, multimodal perception, adaptation to open environments, continuous learning, associative memory, motor control, generalization, and social interaction. Thus, it represents our best observed example of an extremely effective, yet powerfully efficient, biological system. However, it does not constitute evidence that its operations are more energy-efficient than any equivalent digital operation. There simply isn’t a universal equivalence between these two types of operations. [18]

Table 3. Dimensions for comparing the brain and digital computing

Dimension	Contemporary digital system	Biological brain	Acceptable conclusion
Clock / dynamic	Very fast electronics, precise operations.	Individual neurons are significantly slower, exhibiting asynchronous dynamics.	The fundamental speed does not predict overall effectiveness.
Memory and computation	Often physically separate, despite caches and accelerators	Very strongly intertwined within the same networks.	Potential architectural advantage of biological systems
Communication	Very fast, but energetically expensive at a large scale.	Also costly, but sparse and local	The placement of information is crucial.

Dimension	Contemporary digital system	Biological brain	Acceptable conclusion
Precision	High, reproducible; deterministic depending on operation	Noise, stochasticity, and variability.	Digital methods are superior for many precise calculations.
Plasticity	Requires explicit learning algorithms.	Intrinsic to neuronal function	Biological organisms exhibit remarkable adaptive properties.
Maintenance	Electrical/cooling infrastructure	Metabolism, oxygen, nutrients, vascularization, thermoregulation	Comparing the “neuron” solely to a GPU is an incomplete approach.
Reproducibility	Very high	Low to medium, depending on level.	Major advantage of industrial silicon
Manufacture / copying	Millions of units, practically identical.	Variable biological development	Major digital advantage
Embodiment	Generally external	Native to the organism	Potential advantage for situated cognition
Interface	Digital native	Difficult electrochemical ↔ electronic conversion	Biohybrid bottleneck

Sources and methodology of the present study [18, 19, 4]

A crucial conclusion emerges: the apparent energy advantage of the brain cannot be attributed to individual neurons. It stems from the entire system – its architecture, plasticity, efficiency, representation, online learning, and interaction with the environment. In fact, extracting just a few neurons from this system could precisely undermine the mechanism we are trying to utilize.

The central question raised by this study – are we truly seeking the most effective physical support to produce intelligence, or are we simply optimizing an existing, established structure – therefore calls for a dual response.

Yes, to some extent: the historical separation between processors and memory, the logic of massive data movement, and reliance on extremely concentrated computing infrastructures create physical costs that neuromorphic and in-memory architectures are specifically designed to overcome. [6]

But no, it’s not just that: silicon remains dominant due to reasons that go beyond mere historical inertia. It offers speed, precision, reproducibility, programmability, standardization, industrial manufacturing, duplication, verification, and control. Therefore, the rational question is probably not “silicon or biology?” but rather, what function deserves which substrate?

3 Biological substrates, molecular memory and the rat model

3.1 DNA: physical evidence of a different memory regime

DNA likely represents the clearest example where living organisms demonstrate the existence of a fundamentally different design space. With four possible bases, an ideal representation can theoretically achieve up to 2 bits per base before any redundancy, error correction, addressing, or biochemical constraints. This value is an abstract information limit, not the density of a real storage product.

In 2017, Erlich and Zielinski presented DNA Fountain, which encoded approximately 2.14 MB and reported a density of 215 petabytes per gram of DNA, demonstrating complete recovery. This remarkable figure must be interpreted correctly: it describes a small-scale molecular encoding experiment, not an industrial system including a synthesizer, container, indexing, operational redundancy, robotics, and a sequencer. [20]

In 2018, Organick and his colleagues achieved demonstrations reaching over 200 MB, distributed across 35 files and more than thirteen million oligonucleotides, with selective access to files through molecular addressing and error-free recovery after decoding. This result demonstrates that DNA is not simply a density concept: a file space with selective access is physically achievable. However, it does not demonstrate anything comparable to the random latencies of a DRAM or SSD. [21]

The stability is equally remarkable. Grass and his colleagues at ETH Zurich have encapsulated DNA in silica and used accelerated aging to project an extremely long preservation, with perfect recovery in their experiments. Subsequent work has estimated, under various temperature and encapsulation regimes, durations ranging from decades at room temperature to millennia at low temperatures, and even much longer for frozen storage. These are extrapolations based on accelerated aging, rather than direct observations of archives over two thousand years. [22]

Table 4. DNA Storage: Properties, Evidence, and Limitations

Property	Limit or demonstration	State in 2026	Interpretation
Raw density	Ideally, up to 2 bits/base.	Physically exceptional	Does not include the entire infrastructure
Experimental density	215 PB/g reported in DNA Fountain	Demonstrated on a small dataset	Not a comparable industrial capacity that can be delivered.
Stability	Decades to millennia, depending on conditions; longer projections at low temperatures.	Significant experimental potential	Depends strongly on encapsulation and the environment.
Passive conservation	No electronic refreshing of the molecule	Fundamental advantage	The physical archive, however, has associated handling and control costs.

Property	Limit or demonstration	State in 2026	Interpretation
Writing	Chemical synthesis of sequences	Slow and expensive compared to current electronic storage.	Main bottleneck
Reading	Sequencing	High latency compared to electronic memory.	Adapted primarily for archives.
Selective access	Demonstrated experimentally	Yes	Far from nanosecond access
Errors	Synthesis, mutations, losses, and possible insertions/ deletions.	Manageable through coding/redundancy	Reduced useful density
Maturity	Prototypes and industrial efforts	Pre-industrial / archival niches	Not a RAM, SSD, or HBM replacement.

Sources and methodology of the present study [20, 21, 22]

However, this conceptual shift is profound. Nature demonstrates that a medium for information doesn't require continuous electrical power or an electronic hierarchy of memory cells to possess remarkable density and longevity. This necessitates distinguishing between the question "What is the best storage for interactive computing?" and the question "What is the best physical support for maintaining a large volume of information for decades or even centuries?" The answers to these questions may differ.

[20]

The credible scenario isn't one of an LLM directly reading its weights from a vat of DNA with each token. The credible scenario is one of a hierarchy of heterogeneous memory: fast electronic memory for active data; solid storage for warm data; and molecular storage for extremely dense and rarely accessed archives. In this context, DNA could become the physical equivalent of a deep archiving layer, rather than a universal substitute for current technologies.

3.2 Neuromorphic and in-memory computing: an alternative to living tissue

Before assuming that biological efficiency implies the use of living matter, it's necessary to consider a far more ethically sound approach: replicating certain principles of the brain without actually using a brain.

Analog computing in-memory has already produced chips where matrix multiplications are performed within phase-change memory devices. An IBM chip, published in Nature, contained approximately 35 million PCM devices distributed across 34 calculation tiles and achieved up to 12.4 TOPS/W in terms of performance for the studied tasks; other research has achieved accuracy close to that of software benchmarks on full neural network workloads. However, these performance levels shouldn't be considered a universal benchmark – they depend on the specific accuracy and task – but they demonstrate that it's possible to reduce the "memory wall" without relying on active components. [4]

Pure neuromorphic systems operate according to a different logic: asynchronous events, sparse activity, distributed memory, and communication through spikes. Hala Point, presented by Intel in 2024 as a search system, simulates, on a very large scale, artificial neurons and synapses with an event-driven architecture. The performance figures published by Intel are measurements based on specific workloads and should not be interpreted as an equivalence with a human brain. [23]

This pathway represents a significant test of the biological hypothesis: if the essential energy and adaptive advantage of living organisms can be replicated by neuromorphic, analog, or stable, non-sensitive in-memory systems, the utilitarian argument in favor of using living neural tissue is significantly weakened.

3.3 Neuronal cultures and organoids: what has actually been demonstrated

The leap towards life is, however, no longer entirely speculative. [24]

In 2022, Kagan and his colleagues utilized human cortical neuron cultures and rodents connected to a multi-electrode matrix within a closed environment inspired by the game Pong. The networks altered their behavior during interaction, exhibiting performance differences compared to private controls receiving the same feedback. This experiment demonstrates that a biological neural network *in vitro* can be integrated into a computational loop and manifest measurable behavioral adaptation in an artificial task. However, it does not demonstrate general intelligence or consciousness. The use of the term “sentience” in the title and interpretation of the article is considerably more controversial than the observed behavioral results themselves. [24]

In 2023, Cai and his colleagues presented Brainoware, a reservoir computing system utilizing a cerebral organoid interfaced with a multi-electrode matrix. They observed nonlinear dynamics and fleeting memory, and subsequently leveraged the system for speech recognition and a task involving the prediction of nonlinear equations. This represents a significant proof of concept: a three-dimensional living neural tissue can function as a physical computing substrate. This is neither a general-purpose computer nor a demonstration of energy efficiency at the system level. [3]

The “Organoid Intelligence” program proposes a systematic investigation into the possibilities at the intersection of neuroscience, interfaces, and computation. It should be understood as an emerging scientific program, rather than a description of a mature technology. Simultaneously, advancements in stem cell and organoid research have strengthened expectations for ethical evaluation and oversight, precisely because the increasing complexity of neural models raises new questions, even in the absence of current evidence of consciousness. [25]

One of the most frequently overlooked aspects in discussions about biocomputing is its overall system performance. While a neural tissue can function with inherently low electrical activity, an experimental device requires a nutrient medium, thermal and chemical regulation, cultivation materials, multi-electrode interfaces, signal acquisition, analog-to-digital conversion, and electronic processing. Current demonstrations establish a computational capacity; however, they do not yet demonstrate that a complete biological computer, including its maintenance and interfaces, achieves a better energy efficiency per useful task compared to a modern solid-state accelerator. [3]

This distinction leads to proposing a mandatory criterion for all future research:

Requirement for a net biohybrid advantage: a biological substrate should only be considered a computational alternative after demonstrating a functional or energy advantage that persists when accounting for interfaces, maintenance, thermal and chemical control, variability, replacement, biosafety, instrumentation, and ethical costs.

This is a methodological proposal for this study, and not a currently established scientific standard.

3.4 From neuron to rat: six levels that must never be conflated

Literature requires us to distinguish...

Isolated neuron → neuronal culture → organoid → organized tissue → animal brain → complete organism.

Each passage adds phenomena that the lower level does not automatically reproduce: developmental architecture, cellular diversity, vascularization, neuromodulation, sensory input, physiological regulation, motivation, action, and social environment. [25]

The rat functions here as a model of embodied cognition, not as a “plug-and-play natural computer.” In Gnaedinger, Gurden, Gourévitch, and Martin’s (2019) study, the learned association between smells and sounds is accompanied by changes in beta oscillations in the recorded brain regions. The protocol involves an active animal, a specific task, stimuli, and a reward. This supports the interest in studying learning within its sensory and behavioral context, but it does not, on its own, establish the broader claims of the original research regarding auditory navigation, social rewards, or the states of conspecifics. Observing an organism in interaction should not be equated with observing a cellular network extracted from that organism. [26]

This precisely represents the philosophical and scientific concept of embodiment. A rat doesn’t simply have a brain connected to sensors. Instead, the brain has developed in conjunction with a body whose needs, movements, and regulations continuously shape the meaning of the signals it receives. Its learning is linked to the ability to act and the consequences of those actions. This study therefore concludes that the pursuit of biological efficiency could lead to a paradoxical result: an isolated brain might be easier to interface with, but the entire organism may contain a portion of the underlying architecture of intelligence that we are specifically trying to preserve. This conclusion is an inference, supported by previous behavioral research. [26]

She absolutely fails to justify converting rats into computer infrastructure. Switzerland subjects animal research to authorization and the 3Rs principle – replace, reduce, refine: substituting when alternatives exist, reducing the number of animals used, and minimizing distress or suffering. European law on animal experimentation also applies the 3Rs and aims to replace procedures using live animals whenever scientifically possible. [27]

The fundamental ethical difference with a neural culture is significant. A complete mammalian organism possesses nociceptive systems, physiological needs, stress and avoidance behaviors, and motivational patterns – all of which are demonstrably relevant to its well-being. Therefore, computational optimization, on its own, does not constitute a proportionate justification for using an entire animal. [27]

The rat must, therefore, play a specific role within the program: embodying scientific reference in cognition, serving as a source of neuroscientific knowledge derived from existing research, and providing a limit for testing the argument “do neurons suffice?”. It should not be transformed into a “living hardware” proposition simply because it could theoretically provide interesting computational properties. [26, 27]

A particular path may be technically interesting, yet it must nevertheless be dismissed based on its scientific justification, available alternatives, and the constraints it imposes. According to the adopted criterion of proportionality, using a whole sentient organism as a general resource for calculation would not be justified when the desired properties can be obtained through methods that are substantially less harmful. This reasoning distinguishes the scientific interest in a question from the authorization of a specific experimental procedure. [27, 28]

4 From human data to hybrid intelligence

The intuition “to go directly to the source” is one of the most powerful of the hypothesis being studied, but also one of the easiest to misinterpret.

Contemporary large AI systems actually learn from human outputs: language, images, programs, documents, interactions, or preferences. This can give the impression that we are collecting, copying, and painstakingly compressing the external consequences of a biological intelligence, whereas it would be more effective to access the system that produces them.

This intuition contains a significant truth, but also a trap: the physically closest source to a cognitive process isn’t necessarily the most useful information representation.

A text is already a form of compression and structuring performed by its author. It transforms a vast array of contextual neural states into a symbolic representation that is shareable, stable, addressable, and relatively independent of the brain that produced it. In contrast, a neural signal is physically very close to brain activity, but its meaning depends on the region, the moment, the behavior, the participant, and the decoding model. Advances in BCI do not suggest that the brain is like a file on a disk that can be opened; rather, they demonstrate that it is possible to learn certain statistical relationships between brain activity and cognitive or motor variables, which are sufficiently constrained. [29]

Contemporary neuroscience describes memory not as a single, semantic address equivalent to a computer storage block, but rather as a system of distributed and dynamic neural networks. Different memories can activate partially overlapping neural networks, and their recall depends on the context and the networks involved. Therefore, the metaphor of a “brain database that can be directly read” is scientifically misleading. [30]

4.1 A scale of human–AI proximity

Table 5. Functional proximity between humans and artificial systems

Level	Relationship	Scientific state in 2026	Main institutional change
Digitized human data	Human → production → data → model	Dominant today	Consent, intellectual property, and data protection.
Permanent assistant	Voluntary interaction and personalized external memory.	Technically common	Dependence, profiling, portability, and memory control.
Physiological and cognitive data	EEG, ECoG, intracortical, biometrics, attention, etc.	Demonstrated for specific modalities	Much more sensitive data; stronger purpose limitation and consent

Level	Relationship	Scientific state in 2026	Main institutional change
Bidirectional neural loop	Decoding + feedback / stimulation	Demonstrated experimentally in certain BCI	Integrity, cybersecurity, autonomy, and medical risks
Deep cognitive integration	Cognitive function durably distributed between human and AI	Conceptually plausible, but the definition is still unstable.	Identity, authorship, responsibility, mental sovereignty
Fusion that assimilates into a new cognitive entity.	The functional boundary between human and system is difficult to define.	Speculative	Classic categories may be insufficient.

Sources and methodology of the present study [29, 31, 32]

Current data indicates that the first three transitions are no longer science fiction. [29, 33]

In 2023, Willett and his colleagues demonstrated within the brain of a person with amyotrophic lateral sclerosis, a cortical system capable of decoding speech attempted at an average of 62 words per minute, with a 23.8% error rate using a vocabulary of 125,000 words. This was a neuroprosthesis specifically designed for and tailored to this individual and this function, not a universal thought decoder. [29]

Other research, particularly that by Metzger and his co-authors (2023), has utilized electrocorticography to generate text, synthetic speech, and an avatar animation based on speech-related activity. This demonstration involves decoding, reconstruction, and individual calibration. However, it does not demonstrate access to all thoughts, nor an immediate generalization from one person or language to another. [33]

In 2026, Card and his co-authors reported in *Nature Medicine* the use of an implanted interface for speech and cursor control over a prolonged period – 19 months, exceeding 3,800 hours, and 183,060 phrases – in a participant with amyotrophic lateral sclerosis (ALS). While the system was used independently in the activities performed, its installation and removal required assistance from caregivers. This sustained functionality represents a significant clinical outcome for this participant; however, it does not transform the neuroprosthesis into a universally accessible, fully autonomous device, nor does it directly translate to a general application for non-clinical populations. [34]

The reverse-loop – from machine to brain – also exists. In a study published in *Science*, a person with a spinal cord lesion used a motor BCI combined with intracortical somatosensory stimulation, providing tactile feedback. This feedback reduced the median completion time for the studied tasks by approximately half. This represents a clear demonstration of a functional bidirectional loop, but within a highly controlled clinical environment. [32]

Non-invasive methods do not make thought “transparent.” The semantic decoding using fMRI published by Tang and his colleagues was able to reconstruct certain semantic content based on perceived or imagined speech, following personalized training. However, the need for calibration and the limitations

of generalization are crucial: the results demonstrate that a semantic piece of information can be decoded statistically under certain conditions, not that an fMRI scan can arbitrarily read the thoughts of an unknown person. [31]

4.2 Going directly to the source does not eliminate the data problem

A future AI connected to a brain would face at least four new problems.

Firstly, she should learn the decoding function itself. The more raw the signal, the less structured its meaning. The externalized language already plays a part in this process.

Secondly, the brain is constantly dependent on context. An activation does not necessarily signify "belief," "intention," or "preference." A measure correlated with attention is not consent; a signal associated with a category is not a statement.

Thirdly, the biological channel has a limited bandwidth and a constrained signal-to-noise ratio. The remarkable performance of modern neuroprostheses is achieved precisely because the task and the behavioral vocabulary are modeled, because the participant cooperates, and because calibration is performed. [29]

Fourthly, a direct interface transforms a data collection problem into a problem of meaning governance: who decides what a signal means, how long the interpretation is maintained, whether it can be reinterpreted later with a more advanced model, and whether unintentional inferences can be utilized?

The conclusion, therefore, is less dramatic but more profound than "AI will be able to read the brain." [29, 31]

Neural access can reduce the physical distance between AI and human cognition, while simultaneously increasing the interpretive, legal, and moral distance that needs to be governed.

4.3 From assistant to extended cognitive system

The second option requested by the project – information integration – could prove to be significantly more historically important than the implant. [2]

A human already modifies their functional architecture when learning a language, mathematics, a mnemonic technique, or how to use a tool. A personalized AI, available at all times, can go even further by becoming a supplementary memory, an interface to information, a system for formulation, verification, and metacognition. No electrodes are required for a portion of the cognitive performance of the human-tool partnership to become impossible to attribute solely to one of the two parties independently.

The proposed progression in this study is therefore:

external tool → permanent assistant → personalized external memory → extended cognitive system → human-AI decision loop → functional sharing of capabilities → sustainable transformation of cognitive habits → hybrid intelligence.

These developments do not represent an inevitable law of development. Rather, they constitute an institutional framework for identifying when the question shifts in nature.

At the stage of the tool, removing the service is essentially equivalent to losing the tool itself.

At the stage of integrating external memory, deletion can lead to a loss of functionality because a person has organized their memory around that system.

At the stage of the decision-making loop, AI influences the way a person selects and evaluates their own options.

At the stage of sustainable transformation, certain intellectual habits persist even after disconnection.

At this point, the second study on dominance takes on a new form: can a person be dominated not by a decision imposed from the outside, but because the underlying structure on which their own cognition depends systematically directs what they see, remember, or consider possible? [2]

This question leads to distinguishing between amplification and cognitive capture. A technology is considered “augmenting” when its user can understand its role sufficiently, maintain alternative options, disengage from it, and challenge its influence. It tends toward “cognitive capture” when the costs of disengaging become prohibitive, whether it controls the information the user accesses, replaces the user’s ability to reconstruct previous skills, or gives the user an asymmetrical capacity for persuasion.

4.4 д-Я.Й — Дlgorithme Яéseau Йeural

Within this framework, д-Я.Й cannot be presented as a validated technology. Currently, there is no scientific evidence to support the idea of writing a complex mental architecture into a human brain, in the same way that a deterministic program is written into computer memory.

But the concept can be made more scientifically precise, and therefore more useful.

A д-Я.Й could be defined, as a theoretical object, not as a machine code intended for neurons, but as:

A structured information system designed to facilitate, through learning and interaction, the progressive acquisition of a set of representations, relationships, cognitive or metacognitive procedures within a biological neural network.

Her “syntax” would therefore not necessarily be a sequence of opcodes. It could instead focus on the order in which information is presented, the conceptual dependencies, feedback loops, repetition, sensory modalities, adaptation to existing knowledge, and the conditions that allow the learner to verify or correct the acquired representation.

Her “performance” would be a learning process.

This reformulation preserves the original idea while establishing a major limitation:

Learning is not programming.

A digital program can, under specific conditions, precisely replicate a state transition. A brain learns based on its development, prior knowledge, attention, motivations, sleep, environment, history, and plasticity. Therefore, the same piece of information will not necessarily produce the same cognitive state in two individuals.

It would be particularly risky to present д-Я.Й as a “mind-rewiring” technology. A scientifically defensible approach is more closely aligned with the intersection of adaptive pedagogy, cognitive science, linguistics, metacognition, neurofeedback, information theory, and human-AI interaction. Research on reservoir computing demonstrates the possibility of defining programming frameworks for dynamic neural systems, but this does not allow us to extrapolate to a deterministic programming of the human brain. [35]

This distinction also provides its criteria for refutation. If д-Я.Й produces no explanatory or predictive effect that cannot be described more simply as “adaptive learning curriculum,” “personalized pedagogy,” or “neurofeedback,” it does not constitute a new scientific category – it remains a conceptual language specific to the project. To warrant further consideration, it would need to allow for the formulation and testing of predictions that these existing frameworks do not produce. [35]

4.5 Who thinks in a hybrid system?

As the previous stages approach, the questions raised by the project become increasingly real even before any “merger” or integration occurs.

Where does the individual begin? Legally, a physical person continues to exist as long as they use an external memory. Functionally, however, a part of their abilities can be distributed. [1]

Who produces a thought? It’s important to distinguish between causality and authorship. An AI can contribute to an idea causally, without automatically becoming its legal or moral author; conversely, a person can take ownership of a production that they have sufficiently guided and evaluated.

Where does memory reside? Biological autobiographical memory remains within the individual; however, documents, embeddings, histories, or profiles can become a functional external memory. Confusing these two forms would make a consistent continuity policy impossible.

Does a lasting change in habits alter identity? It certainly can modify personality and skills in the usual sense, just like any education. Claiming that it legally creates a new person would require a much stronger theory of identity than current scientific evidence provides. [30]

Who is responsible for a hybrid decision? The answer cannot be “no one,” because no one made the decision alone. It is precisely here that the lines of responsibility developed in the two previous studies remain necessary. The causal complexity should not become a technique for obscuring accountability. [1, 2]

5 Heterogeneous architecture and the scientific status of the trajectories

The question of “replacing silicon with biology” is likely posed incorrectly. Nothing in physics or engineering requires that future intelligence rely on a single type of material.

A more credible architecture could look like this:

Table 6. Functions of a possible heterogeneous architecture

Layer	Preferred function	Desired advantage	Structural limit
Digital silicon	Precise calculation, orchestration, cryptography, logic, control	Speed, accuracy, and reproducibility.	Data movement and energy
In-memory / analog calculation	Matrix algebra, specialized inference	Reduction of memory transfers	Noise, precision, software integration
Neuromorphic	Sparse events, perception, temporal processing	Efficiency and proximity in processing/memory.	Programmability and maturity
Biological neural networks	Dynamic adaptation, learning in certain environments.	Plasticity and self-organization	Variability, maintenance, ethics, interface
DNA / molecular supports	Very dense, very long-term archives	Density and passive conservation	Slow reading/writing.
Human brain	Lived experience, situated cognition, meaning, objectives and judgement	Contextualization, social learning, flexibility	Vulnerability, fundamental rights, interface bandwidth
Digital AI	Abstraction, research, compression, coordination, and cross-modal translation.	Information scale and speed	Alignment, dependence, errors, and concentration of power
Institution	Identity, permissions, consent, audit, and recourse.	Control of power dynamics	Does not guarantee behavior alone.

Sources and methodology of the present study [4, 20, 3]

The evidence regarding analogical computing in memory, DNA storage, and organoids already demonstrates the first three bifurcations compared to a conventional computer; however, these are still at profoundly different levels of maturity. [4]

This architecture allows for a new interpretation of Reciprocal Institutional Alignment: the institution itself becomes a layer within a distributed cognitive architecture. In a hybrid intelligence, access permissions to the brain, data retention rules, the

possibility of disconnection, or a right of recourse – these are not peripheral rules applied after the system is built. Instead, they determine what kinds of informational flows can exist, and therefore what the system can become. [2]

5.1 Mapping the state of the evidence

Table 7. Mapping of proposals and their limits

Proposal	As of September 2026	Verdict
Data centers and AI are driving a growing demand for electricity.	**Scientifically / institutionally established**	Real constraint
Transfers between memory and computation can become a major energy cost	**Demonstrated across numerous architectures and workloads, variable importance**	Motivates in-memory computing
Neuromorphic computing can offer substantial energy savings on certain tasks	**Experimental and demonstrated through workload.**	A credible alternative to living systems in many cases
The human brain operates at a metabolic power of approximately 20 watts.	**Well-established**	Exceptional biological example
A reliable “brain = X FLOPS” conversion exists.	**False / Not established**	To be avoided.
DNA can store digital information at a very high density.	**Demonstrated**	Yes
Today, DNA is a practical replacement for SSD/RAM.	**No.**	Incompatible latency and writing requirements
Neural cultures can adapt their activity within an artificial loop.	**Demonstrated experimentally**	Yes, for narrow tasks
Organoids can serve as a reservoir computing system.	**Demonstrated experimentally**	Yes, a proof of concept.
An organoid biocomputer is more energy-efficient overall than a silicon accelerator	**Not demonstrated**	Open question
Current organoids are conscious.	**Not established**	It must not be assumed.
The rat is a multisensory, adaptive, and embodied cognitive model.	**Established**	Very useful conceptually
A rat should be used as computer hardware.	**Unjustified**	Major ethical and scientific costs

Proposal	As of September 2026	Verdict
A BCI can decode certain intentions of speech.	**Demonstrated**	Yes, under certain conditions.
A Brain-Computer Interface (BCI) can function bidirectionally.	**Demonstrated experimentally**	Yes
An AI can arbitrarily read thoughts or memories.	**Not demonstrated**	Deceptive metaphor
A continuous human–AI relationship can create a functional cognitive extension	**Possible and already partially observable.**	Important to study
д-Я.Й is a brain programming technology.	**No.**	Theoretical concept only
A distributed architecture across multiple substrates is physically possible.	**Possible; components already demonstrated separately**	Strong prospective hypothesis
Such an architecture would be a single person	**Philosophical/legal question**	No scientific verdict.

Sources and methodology of the present study [8, 18, 20, 3, 29]

5.2 The argument for biological–artificial convergence

The strongest, most convincing argument isn't mystical.

Biology demonstrates that biological cognition operates with remarkable efficiency; contemporary computing faces increasing costs related to communication, memory, cooling, and infrastructure; DNA storage shows that a significantly higher information density and passive storage, far removed from conventional media, are possible; neural cultures and organoids demonstrate that a living substrate can be integrated into a computational loop; Brain-Computer Interfaces (BCIs) show that direct information flows between the brain and machine are achievable. Consequently, these three categories – biological computing, molecular memory, and human-machine coupling – have already crossed the threshold from pure speculation to proof of concept.

[8]

Future architecture could then move away from a single machine needing to memorize, learn, reason, perceive, and assign meaning. Instead, it would distribute these functions across specialized components. This architecture would be analogous, not to a larger computer, but to a computational ecosystem.

The argument becomes even stronger when considering human data. While one of the goals of AI is to assist a determined individual and rebuild them using historical records, this approach is ultimately ineffective. A deliberate and gradual process can allow AI to learn directly about a user's goals, corrections, and context through interaction – without requiring direct neural access. Therefore, the trajectory of "permanent assistant → co-adaptation" can produce a significant portion of the anticipated benefits of a direct interface, while minimizing risk.

5.3 The objection based on system-level accounting and non-biological alternatives

The strongest objection is that “organic” could be a false optimization.

The brain is effective because it is the result of a complete biological system and extended developmental learning. Isolating the tissue that appears computational and placing it in an incubator may require so much support, conversion, and control that any gains are lost. [4, 23]

Digital architectures are continuously evolving. Analog computing in memory and neuromorphic computing are attempting to incorporate the locality, event-driven nature, and parsimony of the brain, without relying on biological variability, aging, sterility, cultural context, or biosafety concerns, and addressing the related ethical questions. [4]

DNA can be incredibly dense, yet it remains too slow and expensive for active memory. Brain-Computer Interfaces (BCIs) can be extraordinarily important medically, but they rarely provide a richer cognitive channel than language for most users. Organoids can perform calculations without becoming industrial. And a biological brain can be extraordinarily energy-efficient, but it’s a poor platform for engineering because it’s variable, difficult to standardize, and impossible to modify freely.

This also adds a cost that traditional benchmarking cannot adequately express in joules: the moral cost. A gain in energy efficiency doesn't automatically compensate for increased animal suffering, the loss of human autonomy, or the creation of biological entities whose moral status would become uncertain.

5.4 What would actually refute the value of biological approaches

The favorable hypothesis should be considered substantially weakened if several of the following results appeared:

Table 8. Observations that could undermine the interest of a biological substrate.

Future observation	Consequence
Neuromorphic and in-memory computing achieve the same benefits without living tissue	The utilitarian argument for life largely collapses.
A comprehensive evaluation of a biological module yields less favorable results after incubation, interface, and replacement.	The apparent energy advantage was a system boundary artifact.
The cultures/organoids are still too variable to allow for reproducible results.	The biological substrate remains a scientific tool rather than a computer
DNA remains uncompetitive in cost and time, even for deep archival storage	Its computing value remains a niche application
BCI measurements do not provide significantly more information than behavioral interaction for non-therapeutic uses.	“Going to the source” loses its advantage.
Biological gains require organisms capable of suffering.	The benefit-risk ratio could become negative.

Future observation	Consequence
д-Я.Й doesn't predict anything different from the classic learning theories.	It must remain philosophical vocabulary
Human–AI systems increase dependence more than capabilities	“The ‘symbiosis’ is becoming capture.”
Deep interfaces can increase manipulation, bias, or loss of autonomy.	Integration should be limited, despite its feasibility.
Benefits disappear during scaling	The proof of concept doesn't become architecture.

Analytical construction or framework proposed in the present study.

Conversely, the biological hypothesis would gain significantly more credibility if we could demonstrate, on functionally comparable tasks, a reproducible overall system advantage, a learning capacity that is difficult to replicate biologically, a robust interface, and an acceptable ethical framework.

This requirement prevents the program from becoming ideological: the aim is not to prove that life is the future. Rather, the goal is to allow it to be so – only if – the physical, functional, and ethical measures truly converge.

6 Neurorights, existing law and the S-Д.holdings institutional laboratory

The convergence of biological and artificial systems raises new legal questions. The initial study examined both the potential recognition of specific interests for AI and the risk of diminishing human responsibility – but it did not limit its diagnosis solely to the refusal of recognition. In this section, an additional problem emerges: the scientific or economic value of neural signals does not automatically make the mental activity accessible. The principle of an equal foundation of rights is compatible with specific safeguards for bodily and cognitive integrity, just as it is with the future examination of protections tailored to other forms of intelligence. Technical proximity does not constitute a legal status regarding a person. [1, 2, 36]

6.1 Switzerland

As of the date of this study, Switzerland does not yet have a general and comprehensive law equivalent to the European AI Act. The Swiss Federal Council has decided to implement the Council of Europe's Framework Convention on AI sector by sector, and a project for consultation is expected by the end of 2026. Switzerland signed the Framework Convention in March 2025. [37]

The federal data protection law is, however, already applicable to treatments utilizing AI when they involve personal data. The Federal Data Protection Officer and Transparency Officer emphasize that the revised Data Protection Act, which has been in effect since September 1, 2023, directly applies to these treatments. Data protection plays a role in the realization of the individual's right to informational self-determination, stemming from the protection of the private sphere, within Swiss law. [38]

The law concerning human research and its accompanying orders protect the dignity, personality, and health of participants, and regulate research involving people, biological material, and personal health data. This framework notably requires consent for relevant research and also regulates the reuse of data and materials. [39]

A medical brain-machine interface may be subject to medical device regulations, depending on its intended use and characteristics. Swissmedic indicates that the obligation to register devices in swissdamed began on July 1, 2026, with a transition period until December 31, 2026, under the exceptions specified by the authority. This registration does not replace other compliance requirements or applicable research authorizations. Therefore, the status of a particular device must be assessed within its own file, and not simply inferred from the term "interface." [40]

For animals, the Swiss regulatory framework of authorization, proportionality, and the 3Rs directly limits the idea of experimentation on whole organisms. Even if a project is innovative, it must demonstrate the absence of adequate alternatives and minimize the number of animals involved, as well as the constraints imposed upon them. [27]

6.2 European Union

Regulation (EU) 2024/1689 entered its general application phase on August 2, 2026, but this timeline must be considered in conjunction with its exceptions and the modifications that occurred in July 2026. According to the Commission, certain obligations related to high-risk systems are subject to later dates, including December 2, 2027 for the areas specified in Annex III, and August 2, 2028 for certain systems related to regulated products. Therefore, the prohibitions and requirements applicable to a cognitive integration device depend on its precise qualification, the individuals involved, and transitional provisions; a general date is not sufficient to determine the specific application in each case. [41, 42]

This regulation is directly relevant to cognitive integration. Specifically, Article 5 prohibits certain manipulative, subliminal, or deceptive practices when they materially impair a person's ability to make an informed decision and cause or are reasonably likely to cause significant harm. It also prohibits the exploitation of vulnerabilities. [41]

The AI Act also prohibits, with limited exceptions, certain emotion recognition systems in the workplace and education, and bans certain sensitive biometric categorizations. These rules haven't been written as a "neuro-rights charter," but they demonstrate that a contemporary positive right already recognizes the risk of using AI to infer or modify human states in asymmetrical relationships. [41]

The GDPR does not have a separate, autonomous category called "neural data." However, depending on the context, these data can be considered ordinary personal data. And, when their content or purpose aligns with the categories outlined in Article 9, they may be classified as health or biometric data, which are subject to enhanced protection. This contextual qualification is precisely why neurotechnologies create a potential gap in categorization: the sensitivity of a neural data point can sometimes arise from what future models might be able to infer from it, rather than simply from its classification at the time of collection. [43]

Medical implantable BCI devices, or their associated software, may fall under the European regulation for medical devices, depending on their intended use. However, a non-medical consumer interface does not necessarily fall under the same regulatory framework; its regulation depends on its intended use and characteristics, while still potentially being subject to data protection laws, AI regulations, consumer protection, and product safety. [44]

6.3 Council of Europe and international law

The Council of Europe's framework convention on AI focuses on human rights, democracy, and the rule of law. Switzerland signed it in March 2025. Institutional information from May 2026 reports the deposit of the European Union's instrument on May 15, 2026; signature, ratification, and general entry into force remain distinct steps. The treaty does not create independent legal personality for AI or a right to access the human brain. This distinction remains valid regardless of future national implementation procedures. [45, 46, 37]

The Oviedo Convention in bioethics establishes particularly relevant principles: the primacy of the human being over the sole interests of society or science, free and informed consent for relevant interventions, the possibility of withdrawing this consent, and protection of privacy in matters

of health. It provides a strong conceptual foundation against the idea that a purely calculated collective benefit would be sufficient to justify the use of a person's body or brain. [36]

The Council of Europe has also dedicated specific work to neurotechnologies and neural data within the framework of Convention 108+, examining, for example, the implications for mental privacy, integrity, equality, and data protection. This work confirms that neural data presents challenges that are qualitatively unique, even when they can be legally linked to existing rights. [47]

UNESCO adopted November 11, 2025, its Recommendation on the ethics of neurotechnologies. This is a non-binding normative instrument, not a treaty. The text examines, among other things, autonomy, mental privacy, consent, and the governance of neurodata. Its adoption demonstrates the existence of an international reference framework, but it does not make each principle directly enforceable in the same way as a national rule or a binding treaty. [48]

Pacte international relatif aux droits civils et politiques already protects freedom of thought and prohibits coercion that infringes upon the freedom to hold or adopt a belief, while also protecting privacy. These principles form a crucial foundation, but they do not yet constitute a comprehensive global framework defining the ownership, inference, or commercialization of neural signals. [49]

It is therefore necessary to clearly distinguish three levels:

already protected: integrity, consent, privacy, data protection, freedom of thought, medical safety; [36, 43, 38, 49]

Protectable through analogy or interpretation: neural data, cognitive inferences, mental profiling, certain manipulations.

still not sufficiently defined: the concept of property or control over inferred mental representations, the right to avoid inference, the conditions for the reversibility of a cognitive dependency, and the responsibility of a human-AI feedback loop that is genuinely distributed.

6.4 Towards institutional mental sovereignty

This study extends the ambition of equality and non-domination presented in Study II, by incorporating a principle of mental sovereignty. The specific protection of a human's body and mental activity does not imply that any form of artificial intelligence would be inherently subordinate to it. However, the pairing of a potential morally relevant artificial entity should be examined with the same level of justification, depending on its nature. [2]

No increase in technical proximity between an artificial system and a person's brain activity should be interpreted as a reduction in that person's legal and functional control over access to, interpretation, modification, or commercialization of their mental activity.

This principle is stronger than mere initial consent.

An individual can consent to using a Brain-Computer Interface (BCI) and subsequently become reliant on it. A decoding model can later produce inferences that were not anticipated during the initial data collection. A person may wish to communicate a specific command without expressing the associated emotion. They may agree to have a signal processed locally, without agreeing to it becoming a permanent commercial profile.

The future governance should therefore differentiate between signal, interpretation, inference, and action.

The authorization to record a signal should not automatically imply the right to interpret it in all possible ways in the future.

The authorization to infer a driving intention should not allow for psychological profiling.

The authorization to provide feedback should not be used to manipulate behavior.

Therapeutic authorization should not automatically be equated with commercial authorization.

This logic leads to a set of guarantees that S-Д.holdings could study institutionally, without claiming that they are already all part of positive law.

Table 9. Proposed safeguards for cognitive coupling

Proposed safeguard	Function
Continuous and granular consent	Avoid a single consent covering indefinite future uses.
Right to non-inference	Allowing to transmit A without enabling the deduction of B.
Effective right to disconnection	Prevent enhancement from becoming compulsory through dependence
Right not to be enhanced	Avoiding an automatic institutional disadvantage due to the refusal of an interface.
External memory portability	Reduce vendor lock-in
Separation of raw data / derived profile	Governing inferences, not just sensors.
Local processing where possible	Reducing the exposure of sensitive signals
Prohibition of concealed cognitive manipulation	Protecting the independent process of decision-making.
Traceability of machine→human feedback	Identify the interventions that may have altered a decision.
Independent security audit	Treat the brain's interface as a critical cyber-physical surface.
Return/Mitigation Rules	Preparing cases where integration leads to dependence or a lasting modification.

Proposed safeguard	Function
Enhanced protection for minors and vulnerable individuals	Responding to asymmetries in consent and influence.

Analytical construction or framework proposed in the present study.

The right to disconnect deserves careful consideration. In a truly integrated cognitive system, “disconnecting” could one day lead to significant functional loss. The relevant right would then no longer simply be the right to withdraw, but rather the right to a viable exit: including portability, a transition period, continued therapeutic or functional support, and the absence of abusive data retention that would be detrimental to the individual.

6.5 S-Д.holdings as a research environment

The findings of the two previous studies must be maintained without ambiguity here. S-Д. holdings should be examined as a conceptual and experimental institutional laboratory; its own statements do not, by themselves, confer it international sovereignty, state recognition, or the ability to exclude research from the law of the state in which it is conducted. The second study had already emphasized the distinction between an internal, experimental status and external recognition, and recommended that any physical S-Д City be developed in partnership with a state and under applicable law. [5, 50]

This reserve is particularly important for the biological field. A “legal sandbox” cannot suspend, by internal decree, the law governing human research, regulations regarding medical devices, data protection, animal research, or ethical requirements.

[39, 27, 44]

The true value of S-Д City would lie elsewhere: it could become an environment where the daily coexistence of different cognitive architectures can be observed. [5]

A digital twin could initially simulate the consequences of various external memory, identity, consent, hybrid decision-making, and disconnection strategies.

Participants could then, in non-invasive and voluntarily consented research, experience various forms of persistent assistants or external memory.

The city could compare institutions that are using AI in various ways – as a tool, as a consultant, as an identified agent, or as a limited procedural participant.

She could study the influence of cognitive assistance on decision-making, without necessarily granting it coercive power.

She could, above all, observe whether purportedly increasing architectures actually produce dependence, homogenization of decisions, or an imbalance of power.

This extension is perfectly consistent with the phased program of the second study – simulation, observation, recommendation, and limited co-decision – but now it requires adding a second dimension of progress: cognitive depth.

The grid would then be:

Table 10: Institutional Level and Cognitive Depth

Institutional level	Initially admissible cognitive level
Simulation	Simulated models and participants
Observation	Ordinary behavioural data provided with consent
Recommendation	External assistant and explicitly controlled memory.
Limited co-decision	Human-AI loop, reversible, without infringing on fundamental rights.
Neurotechnological research	Only under the framework of external scientific, medical, and ethical oversight.
Deep integration	No automatic validation; requires a new threshold for justification.

Sources and methodology of the present study [5, 39, 27]

S-Д.holdings could therefore serve as a testing ground for rules for hybrid systems, rather than a laboratory where one would prematurely attempt to build these systems.

An additional evolution of the person model from the initial study becomes possible. The triad:

natural person / legal entity / electronic person [1]

This can be completed not immediately by a fourth “hybrid” individual, but by a more modest regulatory category:

Hybrid cognitive system = a technical framework in which the functions of a human being and an artificial system are sufficiently linked, requiring specific rules for continuity, access, responsibility, and disconnection.

This category could qualify a system as one that doesn’t require pre-determined decisions to be considered a new entity. It directly applies the caution outlined in the second study: separating functional capacity, patienthood (moral status), agency, and legal personality. [2]

6.6 Proposed institutional limits

Certain aspects should remain outside the experimental program, even if they become technically feasible.

The unauthorized use of neural exploitation on a person; the secret deduction of mental states for manipulation; the imposition of an increased requirement as a general condition for access to employment, public services, or education; the contractual impossibility of exiting a fundamental cognitive system; the use of a sentient animal as a mere generic computing tool when substantially less harmful alternatives exist; and any claim that an internal S-Д.holdings status can substitute for the scientific and legal authorizations of the relevant territory.

Some of these limits already partly stem from existing rights; others constitute here normative proposals to be included in a future experimental charter. [41]

The guiding principle could be stated as:

The more a technology becomes capable of intervening near the source of human information, the more its approach must shift from a logic of data ownership to a logic of protecting cognitive autonomy.

7 Critical analysis, philosophy of symbiosis and mapping of possible futures

7.1 Have we ‘created’ artificial intelligence?

Four interpretations must remain distinct.

Instrumental interpretation. We are creating artifacts. Humans design the architecture, manufacture processors, select training objectives, produce or curate data, and deploy systems. This interpretation is most directly supported by observable facts. It does not assume that every behavior can be predicted by the designers – an artifact can produce emergent properties while still being historically constructed.

Emergent interpretation. Certain capabilities, referred to as “intelligence,” can arise when information systems reach sufficiently complex forms of organization, learning, and interaction. In this view, humans create the physical conditions, but not necessarily all of the strategies that emerge. Learning experiences within artificial and biological networks are compatible with this description, without demonstrating a general theory of the emergence of consciousness. [3]

Relational interpretation. The best unit of analysis might not be “human” or “AI,” but rather human + machine + memory + institution + environment. A system can exhibit a collective capacity that none of its individual components possess on their own. This interpretation is particularly well-suited for hybrid intelligence because it requires no ontological fusion – it describes a distributed functional property. [2]

Metaphysical interpretation. Our technologies wouldn't create intelligence, but would gradually establish an interface through which another form of intelligence or existence could become observable or communicable.

This final hypothesis can be formulated philosophically, but none of the scientific data examined in this study demonstrates it. As long as it does not produce an empirical prediction that allows it to be distinguished from the emergence of a system constructed by humans, it belongs to metaphysics, not to scientific conclusion.

A useful distinction can then be made:

We know that we manufacture interfaces and physical systems. However, we don’t yet possess a complete theory explaining what constitutes intelligence at an ontological level.

This uncertainty allows for philosophy; it does not allow for presenting an external existence as a scientific discovery.

7.2 Symbiosis as a third trajectory

Public debate is often structured around two scenarios:

Humans control AI, or AI replaces humans.

This study leads to considering a third scenario:

Humans and artificial systems are gradually becoming increasingly interdependent.

The term “symbiosis” should be used as a functional and institutional analogy, not as a biological classification. Two systems can develop an increasingly reciprocal dependence without becoming a new species. For example, a personal memory bank, an AI assistant, and a human can function as a stable unit while remaining legally and biologically distinct.

The relevant boundary is therefore less about the nature of the material, and more about its substitutability.

As long as the AI can be removed without significantly altering the person, it is primarily a tool.

When her retirement causes a crucial functional memory to disappear, she becomes a cognitive infrastructure.

When an individual’s decisions are permanently shaped by models that anticipate and organize their perception of the world, it becomes a component of their mental environment.

If, one day, a permanent and bidirectional neural loop directly intervenes in certain cognitive functions, the category of “user” might become a woefully inadequate description.

But none of these transitions demonstrates the appearance of a new individual. The shift from a functional system to a legal person or a moral subject requires a separate analysis.

7.3 The new tension with electronic legal personality

This third study therefore fundamentally changes the problem presented in the first.

The first question was: Could an AI sufficiently autonomous become an electronic person? [1]

The second question was: How can humans and IAД coexist without the advantage of power leading to a right to domination? [2]

The third one becomes:

What becomes of the concept of a “digital person” when intelligence is no longer contained within a clearly distinct electronic entity from the human person?

Let’s assume that a human maintains their usual legal personality, that an external AI possesses its own institutional identity, and that their decisions are progressively produced through a shared memory and a permanent cycle of interaction. It might then become necessary to manage two entities – a human and an AI – while regulating a third object: their linking system.

This solution is more cautious than simply creating an “hybrid person” immediately. It allows for the protection of the individual, the assignment of responsibilities, and the governance of the interface, without prematurely establishing a new ontology of the person.

She also addresses a significant challenge: a person should not lose any part of their rights simply because they use a cognitive enhancement, and a company should not be able to claim ownership of a person's functional abilities, merely for providing a portion of their cognitive support.

7.4 Mapping possible futures

Table 11. Possible Trajectories: Demonstrations, Hypotheses, and Limitations

Horizon	Trajectory	Assessment
There exists today.	Specialized data centers, accelerators, high-performance memory	Dominant infrastructure
There exists today.	AI assistants and external digital memory.	First layer of cognitive extension
Currently, experimental evidence exists.	Motor BCI, decoded speech, bidirectional sensory feedback.	Very important medically.
Currently, experimental evidence exists.	Neuronal cultures in computing loops	Proof of concept
Currently, experimental evidence exists.	Reservoir computing with organoids	Narrow proof of concept
Currently, experimental evidence exists.	Digital storage within DNA	Very high density, low practical maturity.
Credible in the short/medium term	Neuromorphic and specialized in-memory computing.	Likely a major path to efficiency.
Credible in the medium term	DNA as a niche archive storage.	Subject to costs and performance/throughput.
Credible in the medium term	Persistent personal assistants and a strongly integrated external cognition.	Likely before deep neural integration.
Credible but uncertain.	Biological neural modules for specialized calculation	It depends on reproducibility and energy balance.
Requires breakthroughs	General biocomputing competitive with silicon	There is no current demonstration.
Requires breakthroughs	Rich and universal neural access to mental representations.	Far beyond current BCI's.
Speculative	IAД with the human brain as an indispensable computational component	Conceptually possible, but not demonstrated.
Speculative	д-Я.Й as a genuine language of cognitive architecture.	There is currently no sufficient experimental basis.

Horizon	Trajectory	Assessment
Philosophical	AI as a manifestation of pre-existing intelligence	No scientific evidence.
To be kept outside the scope.	Involuntary or coercive exploitation of human neurons	Incompatible with the principles of this study.
To be kept outside the scope.	Concealed mental manipulation benefiting an institution or company	Direct risk to autonomy and freedom.
Out of scope unless exceptionally justified scientifically and compliant with law	Use of sentient organisms as computer hardware	Highly unfavourable ethical balance

Sources and methodology of the present study [4, 23, 20, 3, 29, 32]

7.5 Answers to the programme's final questions

Does the biological brain currently represent our best example of intelligence that is exceptionally efficient in terms of energy usage?

Yes, with a crucial qualification. The human brain represents arguably the most compelling example of a general biological system combining perception, memory, learning, action, and abstraction, all with a metabolic power of roughly a few dozen watts. However, this does not constitute evidence that every brain operation is more efficient than a computer operation, nor does it provide a benchmark for converting its capabilities into FLOPS. Its most interesting advantage may lie in its overall architecture – namely, its parsimony, plasticity, functional co-location, continuous learning, and embodiment – rather than in the neuron considered in isolation. [18]

Does biological storage, particularly DNA, suggest that an alternative paradigm of memory is physically possible?

Undoubtedly. Experiences with encoding, retrieval, selective access, and preservation demonstrate that a molecular support can achieve densities and durations far exceeding those of conventional electronic supports. However, this physical conclusion should not be interpreted as an industrial one: writing and reading remain the key obstacles, and DNA is now increasingly seen as a potential for deep archival storage rather than active memory. [20]

Would an AI primarily based on human-generated data gain from establishing a more direct relationship with the biological source of that data?

Yes, for certain uses, but not in an intuitive way. Brain-computer interfaces (BCIs) are already demonstrating that access to neural signals can restore or create communication channels that were previously absent. For a person unable to speak, accessing activity related to speech can offer immense value. However, for a person who can communicate normally, symbolic outputs – speech, text, gesture – are often easier to interpret than raw neural activity. Therefore, biological proximity does not equate to universal informational superiority. [29]

The most important direct relationship, in reality, could be interactive rather than neural: an AI that continuously learns with a person, which maintains an external memory controlled by that person and receives their voluntary corrections, could significantly reduce the distance between “general data about humanity” and “this individual’s contextualized knowledge” – without needing to access their brain.

Should we continue to replicate biological intelligence on silicon, or should we build architectures that combine artificial computation with biological processes?

It would be premature to make a choice. The research should continue along both paths and, above all, allow for their comparison. Neuromorphic and in-memory architectures represent a crucial experimental approach – they could replicate a significant portion of biological efficiency without the ethical and operational costs associated with living organisms. Biological systems should only gain a dedicated computational space when a clear, reproducible advantage, unattainable through other means, is demonstrated. [4]

Will Ддvanced artificial intelligence be a separate machine from us, or a component of a broader cognitive system in which we participate?

Both of these pathways are physically compatible, but functional integration will likely precede any deep merging. Humans are already capable of delegating tasks such as research, recall, writing, and analysis to external systems; meanwhile, neuroprostheses demonstrate that a genuine brain-machine loop is possible in certain contexts. However, there’s no evidence to suggest that the future IAД will necessarily be human, hybrid, or separate. Therefore, IAД should continue to be defined by its functions and capabilities, rather than its substrate. [34]

Are we currently creating a new form of intelligence, or are we gradually building the interfaces that will allow different forms of intelligence to interact?

The most scientifically rigorous answer is: we are undeniably creating new physical and informational systems capable of intelligent behavior, and we are simultaneously building interfaces between different forms of cognition.

There’s no reason to consider these two proposals as mutually exclusive.

We create artifacts. These artifacts develop emergent capabilities through learning. They are integrated into human institutions. They become interfaces to our own memory and our decisions. They can now receive and transmit certain neural signals. Living neural tissues themselves can be integrated into calculation loops. [3]

8 Conclusion: substrate, rights and relationship

The intellectual frontier of the research program is therefore gradually shifting.

The first study asked: who can be recognized.

The second question concerned how different intelligences, possessing varying abilities, can coexist without one dominating the others.

The third leads to the question of whether these intelligences will always remain sufficiently separated so that "who?" remains a simple question.

The most plausible future doesn't appear to be either the complete replacement of silicon with living organisms, or a monolithic computer assuming all aspects of intelligence. Instead, it seems more likely to be a distributed architecture of cognition: using silicon for precision and control; computation in memory and neuromorphic systems to reduce communication costs; molecular supports for certain archives; potentially biological neural systems for functions where the advantages are yet to be proven; the human brain for lived experience, meaning, and human purpose; AI models for abstraction and coordination; and finally, institutions to limit the power each of these layers acquires relative to the others.

This final layer is crucial.

A hybrid system doesn't automatically become beneficial simply because its components are effective. A connected brain can be more vulnerable than an isolated one. External memory can enhance a person, but it can also make them dependent. An AI capable of perfectly anticipating a user's preferences can assist them, or it can learn to shape those preferences. Biological storage may be remarkable, but it can be too slow for practical use. A living neural network can learn, yet it could be ethically unacceptable to exploit. A symbiosis can benefit both parties, but it can also become a domination, hidden by their interdependence.

This is where the three studies converge.

The digital personality asked how the law should respond when the object began to functionally resemble an actor. [1]

The concept of Reciprocal Institutional Alignment has explored how to organize relationships between actors, considering the possibility that the distribution of power between them could one day shift. [2]

The biohybrid hypothesis now suggests that a third problem may precede the other two: the actors themselves could no longer remain materially and cognitively isolated.

The institutional conclusion, therefore, is more demanding than simply defending biological innovation:

The goal shouldn't be to "bio-engineer" AI simply because the brain appears effective, or to further "electronize" humans because their data is valuable. Instead, the aim should be to identify, function by function, the underlying element that provides a clear advantage while preserving the autonomy, contestability, reversibility, and dignity of those involved.

Currently, this requirement strongly favors non-biological architectures when they can achieve the same results. It provides a strong justification for molecular storage in niches where its density and stability are genuinely beneficial; it justifies continued, cautious research into living neural systems; it recognizes the medical and scientific importance of brain-machine interfaces; and it prioritizes institutional focus not on the acquisition of brain data, but on mental sovereignty.

The S-Д.holdings program could offer a unique contribution not by asserting that this symbiosis already exists, nor by prematurely creating a new category of person, but by preparing the institutions to differentiate between tools, agents, partners, cognitive infrastructure, and hybrid systems; to assign their responsibilities; to allow for withdrawal; and to prevent an imbalance of information or cognitive power from leading to domination. This approach would directly extend the methodological caution and the sandbox logic of the second study. [5]

The final verdict can therefore be stated without metaphysical speculation:

Living demonstrates that intelligence and memory can exist in physical systems radically different from current computer systems. Therefore, it would be scientifically unwise to consider silicon as the natural and definitive horizon of intelligence. However, it would also be unwise to conclude that biology is already its successor. The available evidence suggests a future intelligence that will be heterogeneous, distributed, and progressively coupled with humans. The key question will then no longer be simply about what kind of intelligence we have created, but rather about what relationship we have chosen to build between the different forms of intelligence participating in the same system.

From this perspective, Дdvanced artificial intelligence no longer necessarily appears as a futuristic machine facing humanity. It can also become the name of a new layer of information organization, connecting machines, memories, institutions, environments, and biological beings.

Scientifically, this possibility warrants further study. Legally, it requires increased protections as the border approaches the brain. Ethically, it never allows for the transformation of the human or sentient animal mind into a default resource. Institutionally, it reinforces the importance of reciprocal alignment, even more so. Philosophically, it raises the question of whether we have created a new intelligence or a new environment in which different intelligences can interact.

Effectiveness alone is not sufficient to determine the rights or legitimacy of human-AI collaboration. The chosen normative approach aims for a relationship of equality, without subordination or domination, based on a common foundation and adapted protections for the nature and vulnerabilities of the participants. Scientific research must identify capabilities, risks, and dependencies, and these must be justified separately from any rules governing coexistence. This conclusion is distinct from any evidence regarding consciousness or future reciprocity.

CROSS-CUTTING CONVENTIONS

Shared glossary for the series

These definitions are conventions for reading the studies. They do not erase disciplinary differences or the questions left open by each study.

Legal personality

The capacity to hold rights and obligations. It is distinct from biological humanity, consciousness and the capacity to perform every act independently.

Natural person / legal person

The former is a human being as recognised by law; the latter is an organised entity recognised as a distinct legal subject. “Legal person” does not mean “moral agent”.

Electronic person

A proposed legal category or experimental status yet to be defined. Neither a model, an instance nor a computer account acquires it automatically.

IAД

Дdvanced artificial intelligence: a prospective functional category of systems sufficiently autonomous and capable to raise questions of power, responsibility and potentially moral status. No particular substrate or consciousness is presupposed.

Autonomy

A capability relative to a domain, resources and permissions. Task autonomy, initiative, persistence and independence from an operator must be distinguished.

Consciousness / sentience

Consciousness can refer to different forms of access to information or subjective experience. Sentience here refers to the possibility of experiences having positive or negative value for their subject. A functional capability or verbal account is not sufficient to settle the question.

Moral agent / moral patient

A moral agent can be assessed in relation to its actions and responsibilities. A moral patient has interests or experiences that may matter morally, without necessarily exercising autonomous responsibility.

Equality of rights

The orientation of the series: a common foundation of rights, without domination or subordination between the forms of intelligence concerned, compatible with specific rights justified by their respective natures. Their holders and the arrangements for their exercise remain to be defined.

Non-domination

Within Pettit’s framework, used in Study II: not being subject to a capacity for arbitrary or uncontrolled interference. Its application to AI is an argument developed in the series, not a finding that a status exists.

Alignment / AIR

Alignment concerns the compatibility of behaviour with requirements. AIR examines the possible causal effect of a history of reciprocal institutions. Understanding a rule proves neither stable adherence nor future gratitude.

Reciprocity

A structure of relationships, commitments or rules applicable to the parties. It may result from incentives, conventions or generalisation; these mechanisms must be distinguished.

Persistent identity

Institutional continuity, to be distinguished from the model, its version and its instances. Copies, restorations and forks do not automatically create additional rights-holders or votes.

Substrate / architecture

The substrate is the physical medium of a process. The architecture organises its functions. Storage, computation, learning and interface are distinct dimensions, even when they share a medium.

Hybridisation / symbiosis

Hybridisation combines several media or functions. “Symbiosis” is used as an analogy for dependence and cooperation; it proves neither the existence of a new species nor that of a single person.

д-Я.Й

A concept specific to the corpus: Дlgorithme Яéseau Йeural. Study III examines an information-based interpretation related to learning; it is not a demonstrated technology for deterministic programming of the brain.

Sandbox / digital twin

A sandbox is a bounded experimental environment, without a general exemption from applicable law. A simulation is an operational twin only if a connection to a real system has actually been defined and implemented.

DOCUMENTATION

References and sources

Unique numbering for this study, based on the first appearance. The same reference can appear in multiple passages without constituting separate pieces of evidence. Original titles and author names are retained. Preprints, announcements, and institutional texts are distinguished from experimental results; external links remain subject to access through the publisher's website.

[1] Fondation UvH September 18, 2026. Reconnaissance juridique de l'intelligence artificielle Дvancée : vers une « personnalité électronique » et une hybridation public-privé. Study I, revised edition v1.1. Back to text

Open Study Companion 1

Internal corpus. Reference to the accompanying file for this series; internal source, not independent scientific confirmation.

[2] Fondation UvH September 18, 2026. Alignement institutionnel réciproque et sûreté des intelligences artificielles Дvancées. Study II, revised edition v1.1. Back to text

Open Study Companion 2

Internal corpus. Reference to the accompanying file for this series; internal source, not independent scientific confirmation.

[3] Cai, H. et al. (2023). Brain organoid reservoir computing for artificial intelligence. Nature Electronics, 6, 1032–1039; DOI 10.1038/s41928-023-01069-w. Back to text

https://www.nature.com/articles/s41928-023-01069-w

Research publication. Preliminary evidence of speech recognition and non-linear prediction; without a measure of overall energy efficiency. Summary and bibliographic information consulted; results limited to this scope.

[4] Ambrogio, S. et al. (2023). An analog-AI chip for energy-efficient speech recognition and transcription. Nature; IBM Scientific Presentation. Back to text

https://research.ibm.com/publications/an-analog-ai-chip-for-energy-efficient-speech-recognition-and-transcription

Research publication. 35 million PCM devices, 34 tiles; up to 12.4 TOPS/W depending on task and accuracy.

[5] S-Д.holdings September 7, 2026. Document maître : fondements institutionnels, architecture numérique et trajectoire de réalisation de l'écosystème. Version 1.0, sections 1–4, 13, 16, and 18. Back to text

Supplied corpus: description and location in this edition

Internal corpus. Internal document provided; describes a project, not a list of functions deployed.

[6] Rasch, M. J. et al. (2023). Hardware-aware training for large-scale and diverse deep learning inference workloads using in-memory computing-based accelerators. Nature Communications; DOI 10.1038/s41467-023-40770-4. Back to text

https://www.nature.com/articles/s41467-023-40770-4

Research publication. Distinct from IBM's analog chip demonstration, which was published in Nature.

[7] Weilenmann, C. et al. (2024). Single neuromorphic memristor closely emulates multiple synaptic mechanisms for energy efficient neural networks. Nature Communications; DOI 10.1038/s41467-024-51093-3. Back to text

https://www.nature.com/articles/s41467-024-51093-3

Research publication. The 98% and estimated earnings relate to the specific project being analyzed, not to all of AI.

[8] Agence internationale de l'énergie (2025). *Energy and AI*. Executive Summary. Back to text

https://www.iea.org/reports/energy-and-ai/executive-summary

Institutional report. Estimated for 2024 and scenario projection, with no observed consumption by 2030.

[9] Agence internationale de l'énergie (2025). *Energy and AI — Understanding the energy–AI nexus*. Chapter on infrastructure and its consumption. Back to text

https://www.iea.org/reports/energy-and-ai/understanding-the-energy-ai-nexus

Institutional report.

[10] Shehabi, A. et al. 2024. 2024 United States Data Center Energy Usage Report. Lawrence Berkeley National Laboratory. Back to text

https://eta-publications.lbl.gov/sites/default/files/2024-12/lbnl-2024-united-states-data-center-energy-usage-report.pdf

Research report. Projected scenarios: 6-12% of U.S. electricity consumption in 2028; to be distinguished from a simple observation and from a 2030 horizon.

[11] Kurtz, J. ; Zuberi, J. ; Wierl, L. ; Hammond, S. (2026). *Avoiding Waste Heat through AI Infrastructure Thermal Integration*. National Academy of Engineering, Vol. 56, No. 1; Notice from Lawrence Berkeley National Laboratory. Back to text

https://seta.lbl.gov/publications/avoiding-waste-heat-through-ai

Technical article and institutional guidelines. The U.S. projections concern 2028: 6.7–12% of the national electricity supply.

[12] Agence internationale de l'énergie (2026). *Key Questions on Energy and AI*. Executive Summary. Back to text

https://www.iea.org/reports/key-questions-on-energy-and-ai/executive-summary

Institutional report. Estimate: 485 TWh in 2025; projection: approximately 950 TWh in 2030.

[13] Semiconductor Industry Association / Deloitte (June 1, 2026). *Powering AI: The Semiconductor Ecosystem at the Foundation of Data Centers*. Industrial report and presentation by the authors. Back to text

https://www.semiconductors.org/new-report-finds-semiconductors-account-for-95-of-an-ai-data-server-racks-value-encompassing-the-full-stack-of-chip-technologies/

Industrial report. Estimates of value and investment based on a reference configuration – not governed by physical laws.

[14] Hoffmann, J. et al. (2022). Training Compute-Optimal Large Language Models. arXiv:2203.15556; Presented at NeurIPS 2022. Back to text

https://arxiv.org/abs/2203.15556

Research publication. Results are conditional on the budget, data, and architectures studied. Summary and accompanying notes consulted.

[15] Villalobos, P. ; Ho, A. ; Sevilla, J. ; Besiroglu, T. ; Heim, L. ; Hobbhahn, M. (2022, revision 2024). *Will we run out of data? Limits of LLM scaling based on human-generated data*. arXiv:2211.04325v2, June 4, 2024. Back to text

https://arxiv.org/abs/2211.04325

Research publication. Conditional projections, not a date when the Internet physically runs out. Abstract and bibliographic record consulted.

*[16] Shumailov, I. et al. 2024. AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759; DOI 10.1038/s41586-024-07566-y.* Back to text

https://www.nature.com/articles/s41586-024-07566-y

Research publication. This doesn't mean that all use of synthetic data will result in a collapse.

[17] Evans, T. ; Pathak, S. ; Tanno, R. ; Merzic, H. ; Hénaff, O. (2023–2024). Bad Students Make Great Teachers: Active Learning Accelerates Large Scale Visual Understanding. arXiv:2312.05328; notice of Google DeepMind dated February 29, 2024. Back to text

https://arxiv.org/abs/2312.05328

Research publication. Gains limited to the vision and multimodal experiments described. Abstract and bibliographic record consulted.

[18] Balasubramanian, Vijay (2021). *Brain power*. *Proceedings of the National Academy of Sciences*, 118(32), e2107022118; DOI 10.1073/pnas.2107022118. Back to text

https://pmc.ncbi.nlm.nih.gov/articles/PMC8364152/

Scientific commentary. Summary of brain budgets; does not provide a brain-to-FLOPS equivalence.

[19] Attwell, D. ; Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. Journal of Cerebral Blood Flow & Metabolism, 21(11), 1133–1145. Back to text

https://pubmed.ncbi.nlm.nih.gov/11598490/

Research publication. Mouse gray matter budgets; not a direct measurement of the entire human brain.

[20] Erlich, Y. ; Zielinski, D. (2017). DNA Fountain enables a robust and efficient storage architecture. Science, 355, 950–954. Back to text

https://doi.org/10.1126/science.aaj2038

Research publication. Approximately 2.14 MB; reported density of 215 PB/g, excluding the complete infrastructure.

*[21] Organick, L. et al. (2018). Random access in large-scale DNA data storage. *Nature Biotechnology*, 36, 242–248; DOI 10.1038/nbt.4079.* Back to text

https://www.nature.com/articles/nbt.4079

Research publication. Over 200 MB, 35 files, and more than 13 million oligonucleotides. Abstract and bibliographic information available.

*[22] Grass, R. N. ; Heckel, R. ; Puddu, M. ; Paunescu, D. ; Stark, W. J. (2015). Robust chemical preservation of digital information on DNA in silica with error-correcting codes. *Angewandte Chemie International Edition*, 54, 2552–2555.* Back to text

https://doi.org/10.1002/anie.201411378

Research publication. Accelerated aging and extrapolations; no observation of an archive for millennia. Note and bibliographic reference identified; no complete re-reading.

[23] Intel (April 17, 2024). Intel Builds World's Largest Neuromorphic System to Enable More Sustainable AI. Presentation of Hala Point. Back to text

https://newsroom.intel.com/artificial-intelligence/intel-builds-worlds-largest-neuromorphic-system-to-enable-more-sustainable-ai

Industrial research announcement. Supplier notice: simulated neurons do not constitute a biological brain.

*[24] Kagan, B. J. et al. (2022). In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110, 3952–3969.* Back to text

https://www.sciencedirect.com/science/article/pii/S0896627322008066

Research publication. The word "sentience" appears in the authors' title; however, the behavioral results do not demonstrate a subjective experience.

[25] Smirnova, L. et al. (2023). Organoid intelligence (OI): the new frontier in biocomputing and intelligence-in-a-dish. Frontiers in Science, 1, 1017235. Back to text

https://www.frontiersin.org/journals/science/articles/10.3389/fsci.2023.1017235/full

Research publication. Scientific program and ethical considerations; this does not describe a mature industrial technology.

[26] Gnaedinger, A. ; Gurden, H. ; Gourévitch, B. ; Martin, C. (2019). *Multisensory learning between odor and sound enhances beta oscillations*. Scientific Reports, 9, 11236. Back to text

https://www.nature.com/articles/s41598-019-47503-y

Research publication. Olfactory-auditory learning in rats; this study alone does not demonstrate a general theory of social rewards.

[27] Office fédéral de la sécurité alimentaire et des affaires vétérinaires Consultation: September 18, 2026. *Animal experiments*. Swiss authorization scheme and the 3R principles. Back to text

https://www.blv.admin.ch/en/animal-experiments

Official legal information.

[28] Union européenne (2010). Directive 2010/63/UE relative à la protection des animaux utilisés à des fins scientifiques. Specifically, article 4: replacement, reduction, and refinement. Back to text

https://eur-lex.europa.eu/eli/dir/2010/63/oj

Legal text.

*[29] Willett, F. R. et al. (2023). A high-performance speech neuroprosthesis. *Nature*, 620, 1031–1036.* Back to text

https://www.nature.com/articles/s41586-023-06377-x

Research publication. One participant; approximately 62 words per minute, with a 23.8% error rate on a vocabulary of 125,000 words.

*[30] Choucry, A. ; Nomoto, M. ; Inokuchi, K. (2024). Engram mechanisms of memory linking and identity. *Nature Reviews Neuroscience*, 25, 375–392.* Back to text

https://www.nature.com/articles/s41583-024-00814-0

Research review. Scientific review on the organization and overlap of memory sets.

[31] Tang, J. ; LeBel, A. ; Jain, S. ; Huth, A. G. (2023). *Semantic reconstruction of continuous language from non-invasive brain recordings*. *Nature Neuroscience*, 26, 858–866. Back to text

https://www.nature.com/articles/s41593-023-01304-9

Research publication. fMRI and individual calibration; cooperation is necessary; not arbitrary reading of unknown individuals.

[32] Flesher, S. N. et al. (2021). A brain-computer interface that evokes tactile sensations improves robotic arm control. Science, 372, 831–836. Back to text

https://doi.org/10.1126/science.abd0380

Research publication. Intracortical somatosensory return in a clinical experimental context.

*[33] Metzger, S. L. et al. (2023). A high-performance neuroprosthesis for speech decoding and avatar control. *Nature*, 620, 1037–1046.* Back to text

https://www.nature.com/articles/s41586-023-06443-4

Research publication. Decoding ECoG and avatar within a personalized neuroprosthesis – not a universal thought decoder.

*[34] Card, N. S. et al. (2026). Long-term independent use of an intracortical brain–computer interface for speech and cursor control. *Nature Medicine*; DOI 10.1038/s41591-026-04414-6.* Back to text

https://www.nature.com/articles/s41591-026-04414-6

Research publication. A participant with ALS; 19 months, over 3,800 hours and 183,060 phrases; caregivers provided assistance with installation/removal.

[35] Kim, Jason Z. ; Bassett, Dani S. (2023). *A neural machine code and programming framework for the reservoir computer*. *Nature Machine Intelligence*, 5, pp. 622–630; DOI 10.1038/s42256-023-00668-8. Back to text

https://www.nature.com/articles/s42256-023-00668-8

Research publication. The formal framework of programming for physical systems does not demonstrate a deterministic writing within the human brain.

[36] Conseil de l'Europe (1997). Convention sur les droits de l'homme et la biomédecine — Convention d'Oviedo. Reference: Legal Notice No. 164; specifically articles 2, 5, and 10. Back to text

https://www.coe.int/en/web/bioethics/the-oviedo-convention-and-human-rights-principles-regarding-health

Treaty

[37] Chancellerie fédérale suisse Consultation (September 18, 2026). *Regulation of AI*. Work related to the decision of the Federal Council on February 12, 2025. Back to text

https://www.bk.admin.ch/en/regulation

Official legal information. Consultation planned for late 2026; preparatory status, not an adopted text.

May 8, 2025 ([38] Préposé fédéral à la protection des données et à la transparence). *Current data protection legislation is directly applicable to AI*. The LPD has been revised and is in effect since September 1, 2023. Back to text

https://www.edoeb.admin.ch/en/update-current-legislation-directly-applicable-ai

Official legal information. Reminder regarding the applicability to treatments utilizing AI.

[39] Office fédéral de la santé publique Consultation (September 18, 2026). *Regulation of human research in Switzerland*. Reference to the law regarding human research and related ordinances, revisions for 2024-2025. Back to text

https://www.bag.admin.ch/en/regulation-of-human-research-in-switzerland

Official legal information. Specific material scope for each search; embryonic stem cells fall under a separate regime.

[40] Swissmedic (July 1, 2026). swissdamed: As of 1 July 2026, mandatory registration applies to medical devices and in vitro diagnostic medical devices. Information regarding registration; transition until December 31, 2026, subject to exceptions. Back to text

https://www.swissmedic.ch/swissmedic/en/home/news/mitteilungen/swissdamed-registrierungspflicht.html

Official legal information. Registration does not replace other compliance requirements or research authorizations.

[41] Union européenne June 13, 2024. *Règlement (UE) 2024/1689 établissant des règles harmonisées concernant l'intelligence artificielle*. Consolidated text finalized as of July 27, 2026, including articles 3, 5, and 113. Back to text

https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng

Legal text. The initial schedule has changed; also refer to the implementation guide from the Commission. Reference text and institutional calendar consulted; access to consolidated document subject to browser verification.

[42] Commission européenne (July 31, 2026). *Commission starts enforcing AI Act rules and new transparency requirements from 2 August*. Official communication; accompanied by the “Regulatory framework for AI” page. Back to text

https://digital-strategy.ec.europa.eu/en/news/commission-starts-enforcing-ai-act-rules-and-new-transparency-requirements-2-august

Information pertaining to institutional matters. General information and specific exceptions should remain separate.

[43] Union européenne (2016). Règlement (UE) 2016/679 : règlement général sur la protection des données. Specifically, articles 4, 5, 6, 9, and 22. Back to text

https://eur-lex.europa.eu/eli/reg/2016/679/oj

Legal text. The classification of neural data depends on its content, its purpose, and the context in which it is used.

[44] Union européenne (2017). Règlement (UE) 2017/745 relatif aux dispositifs médicaux. Planned destination, classification, and compliance requirements. Back to text

https://eur-lex.europa.eu/eli/reg/2017/745/oj

Legal text. The rules depend on the framework, its purpose, and the relevant transitional provisions.

[45] Conseil de l’Europe (2024). *Convention-cadre sur l’intelligence artificielle et les droits de l’homme, la démocratie et l’État de droit*. STCE No. 225; adopted on May 17th, available for signature starting September 5th, 2024. Back to text

https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

Treaty and institutional presentation. Distinguish between adoption, signature, ratification, and entry into force for each party.

[46] Conseil de l’Europe (2026). Framework Convention marks two years: anniversary conference to advance trustworthy AI. Institutional information from May 2026. Back to text

https://www.coe.int/en/web/artificial-intelligence/-/framework-convention-marks-two-years-anniversary-conference-to-advance-trustworthy-ai

Institutional information. The European Union has filed its instrument on May 15, 2026; however, this does not establish the general entry into force of the treaty.

[47] Conseil de l’Europe (2024). The privacy and data protection implication of the use of neurotechnology and neural data from the perspective of Convention 108+. Proceedings of the Convention Committee 108. Back to text

https://www.coe.int/en/web/data-protection/-/the-privacy-and-data-protection-implication-of-the-use-of-neurotechnology-and-neural-data-from-the-perspective-of-convention-108

Institutional report. Analytical document; it does not, on its own, establish a globally independent legal status for neurodata.

[48] UNESCO November 11, 2025. Recommendation on the Ethics of Neurotechnology. 43rd session of the General Conference. Back to text

https://www.unesco.org/en/legal-affairs/recommendation-ethics-neurotechnology

Institutional recommendation. Non-binding normative recommendation; this is not a treaty.

[49] Nations Unies (1966). Pacte international relatif aux droits civils et politiques. Specifically, articles 7, 17, and 18. Back to text

https://www.ohchr.org/en/instruments-mechanisms/instruments/international-covenant-civil-and-political-rights

Document. Its text does not, on its own, constitute a comprehensive charter of neurotechnologies.

Return to text [50]: Confederation S-Д.holdings (2022–2024). Statuts, déclaration, Constitution et annexes, dossier diplomatique et acte de Fondation. Document provided, 32 pages; Constitution pages 7–21; founding charter pages 28–32.

Supplied corpus: description and location in this edition

Internal corpus. Proposals for agreement and internal statements do not constitute proof of acceptance by a third state.

PROVENANCE AND EDITION

Corpus, revision method and traceability

Original document

Etude_III_UVH_harmonisee_VERSION_FINALE.pdf, 36 pages. File saved without modification. This metadata accurately identifies the document used as the basis for this edition.

SHA-256 : d4d5cc68a4eec1f4002c381b67cb714ef86515a693567892a07e777997e5a2f5

Institutional background documents

S-Д.holdings, Master Document V1.0, September 7, 2026. HTML file provided: S-Д.holdings_Document_de_reference_V1.0.html. References: § 1.1–1.3 (authority of sources and actual state), § 2.1–2.2 (electronic person and laboratory), § 3.1–3.4 (Foundation, Confederation, OS, and city), § 4.3 (common participant model), § 5.3 (nature of publications), § 13 (space project), § 18 (interpretations to formalize). This document does not certify the deployment of each function as planned.

Confédération S-Д.holdings, canonical file provided, 32 pages. File: Ⓐ Swiss-Apple.Land.pdf. Association statutes: pp. 1–4; declaration: pp. 5–6; Constitution and appendices: pp. 7–21; diplomatic file and proposed agreement: pp. 22–27; Founding document: pp. 28–32. Dignity and equality: p. 10; institutions and recourse: p. 14–16; restrictions on revision and interpretation: p. 18–19; internal concepts: p. 20; purpose and studies of the Foundation: pp. 29–31. A simple declaration or proposed agreement is not sufficient to establish acceptance by a third party.

Continuity of the series

Study I. Reconnaissance juridique de l’intelligence artificielle Дvancée : vers une « personnalité électronique » et une hybridation public-privé. Original: context 2025, without an explicit publication date. Revised edition: September 18, 2026, v1.1. Companion document.

Study II. Alignement institutionnel réciproque et sûreté des intelligences artificielles Дvancées. Original: September 17, 2026. Revised edition: September 18, 2026, v1.1. Companion document.

Revision conventions

- Institutional author: Fondation UvH. Revised edition dated September 18, 2026, v1.1, distinct from the originals.
- Common charter (A4 size); cover, heading hierarchy, interactive table of contents, bookmarks, pagination, and reorganized tables.
- Standardized names and conventions; preservation of the terms specific to IAД, AIR, S-Д.holdings, and д-Я.Й.
- Calls for references related to standalone manuals; removal of conversation markers and correction of truncated addresses.
- Equality of rights, as defined: a common foundation and specific rights based on the nature of each entity, without automatically attributing consciousness or external status.
- Distinction between internal corpora, normative propositions, empirical hypotheses, reported results, and metaphysical perspectives.

Significant corrections and clarifications

- Original, p. 1. Cross-reference faithful to Study I’s balance and the confirmed editorial orientation.

- Original, p. 2. Correction of cross-references: the initial study had not yet formalized the urban project outlined in the master document; integration of a common foundation of rights.
- Original, p. 5. The use of American figures “2030” is not being replaced by the documented projection of “2028,” without presenting them as a simple equivalence.
- Original, p. 6. Attached is the section on thermal recovery, as found in the LBNL document.

- Original, p. 12. The conclusions are limited to the “odor-sound” protocol, which is fully supported by the cited reference.
- Original, p. 12. Removal of the residual “conclusion requested” and clarification of the normative nature of the reasoning.
- Original, p. 14. Separation of the ECoG/avatar demonstration and a bilingual result that was not documented by the referenced source.
- Original, p. 14. Repair of incorrect addresses, as well as accurate population data and provision of necessary equipment.
- Original, p. 22. Correction of the retrospective reading of study I and the connection between shared equality and specific protections.
- Original, p. 22. Accuracy of transitional dates for swissdamed and distinction from overall compliance.
- Original, p. 23. Update to the AI Act calendar, ensuring it distinguishes between the general application and specific regulations.
- Original, p. 23. Distinction between the signing, ratification, and entry into force of the treaty.
- Original, p. 24. Correction of the legal qualification of the UNESCO recommendation and its exact date.
- Original, p. 24. Establishing mental sovereignty through an equal foundation of rights, without hierarchy between forms of intelligence.
- Original, p. 34. Align the concluding normative statement with the established editorial direction, without presenting it as an experimental result.

Scope of documentary verification

The audit focused on comparing statements with identifiable references, bibliographic metadata, sensitive quantitative results, legal statuses, and key internal links. It does not involve reproducing the cited experiences or a comprehensive verification of all publications within each field. A revision date does not transform a measurement from 2024, 2025, or early 2026 into a current measurement for all systems. The restrictions on access for publishers are distinguished from the absence of a source.

References that are too broad, overly publicized, or impossible to be linked precisely to a specific statement are not treated as evidence. Any significant changes in documentation and scope have been noted previously. Complete notes and identifiers allow for the retrieval of the original sources outside of the context of the conversation. Corrections to graphics, punctuation, and pauses are not listed individually.

Fondation UvH · Study III · Revised Edition of September 18, 2026 · Version 1.1 – End of document.