Journal of Inclusive Methodology and Technology in Learning and Teaching
ISSN 2785-5104
Anno 6 n. 1sup (2026)

EDIZIONI UNIVERSITARIE ROMANE
editoria scientifica dal 1952

Beyond the Hearing Mind: Deaf Ecosystems, Agentic AI and the Governance of Extended Cognition

Mariagrazia Verbicaro [1] and Patrizia Oliva [2, 1]

[1] PhD student, Department DIMED, University of Messina; mariagrazia.verbicaro@studenti.unime.it
[2] Associate Professor, Department of Health Sciences, "Magna Graecia" University of Catanzaro; poliva@unicz.it

Abstract: The Extended Mind Theory, in its classical formulation by Clark and Chalmers, presupposes an implicitly normotypical and hearing cognitive subject. This paper argues that the Deaf community, as a historically distributed cognitive ecosystem in which body, space, and language are structurally inseparable, does not merely exemplify extended cognition but radically interrogates its residual individualistic assumptions. Drawing on Bateson's cybernetics, Hutchins' distributed cognition, Malafouris' Material Engagement Theory, and 4E cognitive science, the paper develops an analytical proposal articulated across three levels (artefacts, environments, and epistemic governance) for AI-extended educational ecosystems, with particular attention to the structural lacunae of the European Union's Artificial Intelligence Act with respect to signed languages as cognitive systems. The central argument is that the extended mind is institutionally contested and that existing regulatory frameworks reproduce, in algorithmic form, the same structures that have historically impoverished Deaf cognitive ecosystems, perpetuating in algorithmic logic what the Milan Congress of 1880 had sedimented in educational policy.

Keywords: extended mind; Deaf education; sign language; agentic AI; cognitive governance

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The tradition of distributed cognition (Hutchins, 1995), ecological psychology (Gibson, 1979), and extended mind theory (Clark & Chalmers, 1998) has radically destabilised the assumption that the mind is bounded, individual, and sovereign. What remains insufficiently examined is the degree to which this theoretical destabilisation has itself been conducted from within unscrutinised presuppositions: presuppositions about which mind is being described, within which ecological and cultural conditions, and against which normative baseline. The extended mind, in its prevalent formulation, remains implicitly that of a hearing and normotypical subject. The Deaf community does not merely exemplify cognitive extension: it radicalises it, revealing that extension is a constitutive requirement of being cognitive subjects at all. This radicalisation acquires further urgency in the era of agentic artificial

[1] The contribution is the result of the joint work of the authors who share its structure and content.

In particular, paragraphs 2-3-4 are to be attributed to Mariagrazia Verbicaro and paragraphs 1-5 to Patrizia Oliva.

intelligence, whose automatic translation systems, signing avatars, and algorithmic platforms redraw the cognitive boundaries of Deaf linguistic minorities without adequate accountability mechanisms. This paper therefore argues that the Deaf community constitutes a paradigmatic case of extended mind theory, that the Milan Congress of 1880 represents its historical precedent of external cognitive governance, and that the European Union's AI Act of 2024 presents structural lacunae that a Cognitive Governance Framework is called to address.

2. The Extended Mind and Its Unspoken Assumptions

The Extended Mind Thesis rests on the parity principle developed by Clark and Chalmers (1998): any external process that performs the same function as an internal process must be recognised as a cognitive component in its own right. The paradigmatic case of Otto, an Alzheimer's patient who compensates for mnestic deterioration by means of a notebook, illustrates this position elegantly, while simultaneously revealing its implicit presuppositions. The notebook is an artefact designed within an alphabetic writing culture; Otto is a subject whose cognitive architecture was formed within normotypical parameters prior to deterioration. The coupling compensates for a loss. It is on this point that Malafouris (2013) defines this schema as an additive model of cognitive extension, counterposing to it the Material Engagement Theory, in which mind and artefact emerge together through processes of material engagement that precede any distinction between internal and external. The same direction is anticipated by Clark (2003), according to whom human beings are constitutively oriented toward integration with the technologies they inhabit; in the Deaf community, however, this integration becomes structural, since signed language is constitutive of the subject's cognitive architecture.

For the native signer, signed languages are autonomous linguistic systems in which spatial grammar, bodily movement, and facial expression are inseparable from the cognitive act of meaning production (Emmorey, 2002): the signing body is the medium in which the Deaf mind thinks, rendering empirically precise, rather than merely evocative, Merleau-Ponty's phenomenological observation that the body is our very means of having a world. The cognitive system emerges as a systemic property of coordinated activity among persons, tools, and procedures, without residing in any individual mind (Hutchins, 1995), and in the Deaf community this manifests in the intergenerational transmission of the signed language, whose interruption deteriorates individual cognition as a direct effect on the distributed system (Padden & Humphries, 1988). Bateson (1972) radicalises this perspective by shifting the unit of analysis to the ecological relation as such, conceiving the mind as a process organised through information circuits that traverse organism and environment, in which the relevant unit is not the subject but the ecology that constitutes it.

In signed language, the signing space is simultaneously a motor output, a syntactic structure, and a spatial representation of discourse referents (Emmorey, 2002): it follows that an algorithmic system that processes or generates signing intervenes in the circuits that constitute cognition, with implications that no purely technical risk assessment can adequately capture. The 4E cognitive sciences (Varela et al., 1991; Gallagher, 2005) provide the systematic framework for these convergences: signed languages recruit the same neural systems as spoken languages while simultaneously engaging spatial and motor processing systems (MacSweeney

et al., 2008), with motor production and linguistic representation deeply integrated. For the Deaf signer, the relationship between cognitive processes and environment does not merely satisfy the criterion of stable coupling theorised by Clark and Chalmers (1998) but precedes it as an irreducible structural datum. Malafouris (2026) deepens this position by arguing that identity, thought, and language are constituted through engagement with a shared material ecology, in which self-consciousness emerges rather than preexists, rendering any individualistic account inadequate. It follows that any intervention on that ecology is not the provision of a didactic tool but a constitutive act on the formation of the cognitive subject itself, carrying epistemic, ethical, and political responsibilities that research on Deaf education and artificial intelligence has not yet addressed with the analytical depth and theoretical systematicity the stakes require.

3. Cognitive Distribution and History in the Deaf Community

The history of Deaf communities offers extended mind theory a natural, centuries-long, and empirically documented case of constitutively distributed cognition that no laboratory experiment can replicate. Residential schools for the Deaf were not merely sites of instruction but cognitive infrastructures in which signed language was transmitted intergenerationally as a shared cognitive architecture to be inhabited rather than as content to be learned (Hutchins, 1995). Deaf children who arrived at these institutions, often children of hearing parents without competence in signed language, found an ecosystem in which visual-gestural thought was already socially organised in its practices, its narratives, and its cultural artefacts (Padden & Humphries, 1988). The acquisition of signed language was the constitution of a mind within a system: the self-consciousness of the Deaf child does not precede material engagement with that ecosystem but emerges from it, in a process in which subject and linguistic environment determine each other irreducibly (Malafouris, 2026).

In signed languages, the space in front of the signer's body is an active grammatical component in which discourse referents are localised, relations among them expressed through movements that traverse it, and temporal structures disposed along conventionalised spatial axes (Emmorey, 2002). The difference between Otto's notebook and the signing space is not one of degree but of ontological kind: the former functions like memory, while the signing space does not function like grammar but is grammar. Malafouris (2013) places this distinction at the centre of his critique of the additive model of Clark and Chalmers, arguing that the Material Engagement Theory, with its account of cognitive plasticity as the irreversible co-evolution of mind and matter, describes what the parity principle can only approximate. The latter presupposes a preexisting mind that couples with external resources, whereas in signed language mind and space constitute each other simultaneously and are not separable even by analytical abstraction. If the signing space is a constitutive grammatical component rather than an optional support, any institution that determines which languages are legitimate in an educational context performs an act of governance over the cognitive boundaries of the community concerned.

The Second International Congress on the Education of the Deaf, held in Milan in 1880 with the systematic exclusion of Deaf delegates from voting, approved the resolution declaring the superiority of the oralist method and

recommending the prohibition of signed languages in schools (Bauman, 2004). In the following decades signed languages were banned from educational institutions, Deaf teachers removed, and the cognitive infrastructure of intergenerational transmission systematically dismantled. From the perspective of extended mind theory, this was an act through which a predominantly hearing institution determined an entire distributed cognitive system to be illegitimate, replacing it with oralism, which presupposed hearing cognitive architecture as universal norm. Bauman (2004) analysed this phenomenon through the concept of audism, understood as the systematic presupposition of the superiority of hearing cognition, while Skutnabb-Kangas (2000) described structurally analogous dynamics as linguistic genocide in education. Katsui et al. (2024) show how these policies left structurally persistent traces in national educational systems, perpetuating linguistic deprivation that translates into measurable cognitive deficits transmissible across generations.

Deprivation of access to a signed language in early life produces neuropsychological effects comparable to other forms of early sensory deprivation (Hall, 2017), and the language deprivation syndrome that results is the more severe the more prolonged the deprivation (Hall et al., 2017). These data demonstrate empirically that the distributed mind is a functional reality whose interruption produces irreversible effects on individual cognitive capacities. Policies that deprive Deaf children of access to a signed language during the critical period constitute a form of linguistic neglect with precise ethical and legal implications (Humphries et al., 2016), and the presence of interpreters in schools without genuine structural linguistic accessibility creates an illusion of inclusion that masks the continuation of cognitive deprivation (Caselli et al., 2020). The question must be framed in terms of global public health, removing it from the domain of individual and cultural preferences (Murray et al., 2019), since the Milan Congress of 1880 and language deprivation syndrome are the temporal poles of a continuum of external cognitive governance that has maintained its structure of power through diverse institutional forms, and which in the algorithmic era has been transposed with a speed of propagation that the congressional deliberations of the nineteenth century could not have approximated.

4. Agentic AI, the EU AI Act, and Cognitive Governance

AI systems do not merely support human thought but anticipate, orient, and restructure it, participating in decision-making processes that redraw the boundaries of the cognitive system without transparency or adequate accountability, in a measure that exceeds the functional supplementation theorised by Clark and Chalmers (1998). In the context of Deaf education this transformation is concretely operative: deep learning-based systems for automatic sign language recognition have achieved significant technical gains (Adaloglou et al., 2021), while AI-generated signing avatars are deployed as substitutes for human interpretation without critical reflection on their linguistic adequacy keeping pace with their diffusion. A system trained predominantly on data from a single national variety will systematically fail to recognise regional variants and contact signing, while an avatar whose gestural parameters are optimised for algorithmic recognisability will produce output legible to the classifier but linguistically impoverished for the community whose language it purports to render. The distance between genuine support for cognitive autonomy

and its constraint is not determined by the technical sophistication of the system but by the design choices that precede it, each of which incorporates a decision about which forms of Deaf cognition are recognisable and which remain invisible to the system (Alkahtani, 2024).

In Bateson's cybernetic perspective (1972), an AI system that processes signing intervenes in the information circuits that constitute meaning for its participants, rather than managing data at a remove from cognition. AI tools produce cognitive offloading effects that reshape what users attend to and which cognitive operations they continue to perform autonomously (Gerlich, 2024): in contexts where the cognitive system is constitutively distributed, these effects assume a qualitatively different valence from those produced in less deeply distributed cognitive architectures. When an AI avatar mediates a Deaf student's access to educational content, the question of what that avatar renders is not a question about the characteristics of a didactic tool but about the boundaries of the student's cognitive ecosystem.

The Artificial Intelligence Act (European Parliament and Council, 2024) classifies as high-risk those AI systems that determine access to educational institutions or assess learning outcomes, activating obligations of transparency and human oversight. Systems that mediate the linguistic accessibility of educational content for Deaf learners, however, effectively determine whether a Deaf student can access knowledge and in what cognitive form, yet do not receive the same classification. Signed languages are absent from risk classifications and transparency obligations, notwithstanding the fact that the Charter of Fundamental Rights recognises linguistic diversity as a value to be respected and that AI systems affecting signed languages are already widely deployed in educational contexts. This silence reproduces in algorithmic logic the same structural invisibility that characterised the position of signed languages in post-Milan Congress educational policies: then absent from authorised curricula, now absent from the risk classifications governing AI systems deployed in those same institutions, in a continuity whose cognitive consequences the language deprivation literature has documented over more than a century.

The question of cognitive governance in Deaf educational ecosystems is not exhausted by the regulation of AI systems but articulates itself through distinct planes of intervention that mutually implicate each other. The most immediate concerns the technological artefacts that process and generate signed language: training data select which Deaf linguistic productions are sufficiently regular to be recognised, performative standards establish which accuracy is acceptable relative to which reference variety, and output formats determine which signed language is sufficiently standardised to be rendered algorithmically. Each of these choices is an epistemic decision incorporating a norm about which forms of Deaf cognition are recognisable, which is why such systems must be classified as high-risk and subject to mandatory co-design with Deaf communities, since in its absence quality criteria will inevitably reflect hearing epistemic parameters presented as universal. A further plane concerns the physical-digital hybrid spaces in which AI systems and Deaf learners interact: since the environment is a constitutive component of the cognitive system (Hutchins, 1995), its configuration directly determines the possibility of signed language use, and those who hold authority over that configuration exert a direct influence on the processes through which the Deaf subject constructs

knowledge. The deepest and simultaneously most difficult plane to render visible concerns the determination of what counts as legitimate knowledge in Deaf education: it is here that audism (Bauman, 2004) operates through the silent normalisation of hearing epistemic standards incorporated into evaluation metrics and output formats, making it necessary that these norms be rendered explicit, subjected to critical scrutiny, and opened to contestation by the communities that bear their effects. Future research must be empirical, participatory, and capable of returning to Deaf communities the analytical authority over the cognitive processes that concern them, withdrawing the question from the exclusive domain of technical designers and institutional regulators. Until this restitution occurs, any AI system deployed in Deaf education will reproduce the same asymmetry that the Milan Congress of 1880 instantiated in institutional form, perpetuating in algorithmic logic what that congressional logic had sedimented in educational policy. The crux is not technical but epistemological: it concerns the right of Deaf communities to determine the boundaries within which their members think, learn, and recognise themselves as cognitive subjects.

5. Conclusions

In the Deaf community the mind has never resided in the individual but in the system: in the relation between signing body and signing space, in the information circuits of a shared visual-gestural language, in the intergenerational transmission of a constitutively distributed cognitive architecture. Agentic AI introduces into this ecosystem not an absolute novelty but a new actor in the governance of a system that has always already extended beyond any individual boundary. The Extended Mind Theory in the Clark-Chalmers formulation requires, to adequately describe the Deaf case, a constitutive rather than additive account of cognitive extension, and the Material Engagement Theory of Malafouris (2026) provides it with increasing rigour. The Milan Congress of 1880 is the historical precedent that illuminates the structure of the governance relation that agentic AI instantiates today in algorithmic form; the structural silence of the EU AI Act with respect to signed languages is its algorithmic reproduction, a centuries-long invisibility with measurable cognitive consequences. The analytical structure proposed across the levels of artefacts, environments, and epistemic governance does not constitute a definitive solution but an instrument for a future research that must be empirical, participatory, and rooted in Deaf communities themselves.

References

Adaloglou, N., Chatzis, S., Papastratis, I., Stergioulas, A., Papadopoulos, G. T., Zarifis, W., Konstantinidis, D., Daras, P., & Dimitropoulos, K. (2021). A comprehensive study on deep learning-based methods for sign language recognition. IEEE Transactions on Multimedia, 24, 1750–1762. https://doi.org/10.1109/TMM.2021.3070438

Alkahtani, B. N. (2024). The impact of artificial intelligence on quality of life for deaf and hard of hearing students. American Annals of the Deaf, 169(4), 329–347. https://doi.org/10.1353/aad.2024.a947429

Bateson, G. (1972). Steps to an ecology of mind. Chandler Publishing.

Bauman, H.-D. L. (2004). Audism: Exploring the metaphysics of oppression. Journal of Deaf Studies and Deaf Education, 9(2), 239–246. https://doi.org/10.1093/deafed/enh025

Caselli, N. K., Hall, W. C., & Henner, J. (2020). American Sign Language interpreters in public schools: An illusion of inclusion that perpetuates language deprivation. Maternal and Child Health Journal, 24(11), 1323–1329. https://doi.org/10.1007/s10995-020-02975-7

Clark, A. (2003). Natural-born cyborgs: Minds, technologies, and the future of human intelligence. Oxford University Press.

Clark, A., & Chalmers, D. (1998). The extended mind. Analysis, 58(1), 7–19. https://doi.org/10.1093/analys/58.1.7

Emmorey, K. (2002). Language, cognition, and the brain: Insights from sign language research. Lawrence Erlbaum Associates.

European Parliament and Council. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. https://artificialintelligenceact.eu/

Gallagher, S. (2005). How the body shapes the mind. Oxford University Press.

Gerlich, M. (2024). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. Societies, 14(3), 45. https://doi.org/10.3390/soc14030045

Gibson, J. J. (1979). The ecological approach to visual perception. Houghton Mifflin.

Hall, W. C. (2017). What you don't know can hurt you: The risk of language deprivation by impairing sign language development in deaf children. Maternal and Child Health Journal, 21(5), 961–965. https://doi.org/10.1007/s10995-017-2287-y

Hall, W. C., Levin, L. L., & Anderson, M. L. (2017). Language deprivation syndrome: A possible neurodevelopmental disorder with sociocultural origins. Social Psychiatry and Psychiatric Epidemiology, 52(6), 761–776. https://doi.org/10.1007/s00127-017-1351-7

Humphries, T., Kushalnagar, P., Mathur, G., Napoli, D. J., Padden, C., Rathmann, C., & Smith, S. (2016). Avoiding linguistic neglect of deaf children. Social Service Review, 90(4), 589–619. https://doi.org/10.1086/689543

Hutchins, E. (1995). Cognition in the wild. MIT Press.

Katsui, H., Koivisto, M., Meriläinen, N., Rautiainen, P., Tepora-Niemi, S.-M., Tarvainen, M., Rainò, P., & Hiilamo, H. (2024). Deaf people, injustice and reconciliation: Signed memories. Routledge.

MacSweeney, M., Capek, C. M., Campbell, R., & Woll, B. (2008). The signing brain: The neurobiology of sign language. Trends in Cognitive Sciences, 12(11), 432–440. https://doi.org/10.1016/j.tics.2008.07.010

Malafouris, L. (2013). How things shape the mind: A theory of material engagement. MIT Press.

Malafouris, L. (2026). People are STRANGE: Material engagement and the creation of self-consciousness. MIT Press.

Murray, J. J., Hall, W. C., & Snoddon, K. (2019). Education and health of children with hearing loss: The necessity of signed languages. Bulletin of the World Health Organization, 97(10), 711–716. https://doi.org/10.2471/BLT.19.229427

Padden, C., & Humphries, T. (1988). Deaf in America: Voices from a culture. Harvard University Press.

Skutnabb-Kangas, T. (2000). Linguistic genocide in education — or worldwide diversity and human rights? Routledge.

Varela, F., Thompson, E., & Rosch, E. (1991). The embodied mind: Cognitive science and human experience. MIT Press.