Allostasis Machines: a model for understanding internal states and technological environments

Bradly Alicea^{1,2}, Daniela Cialfi^{1,3}, Anson Lim¹, Jesse Parent¹

Abstract

In the present paper we will approach enactivism from the perspective of internal regulation: while the environment shapes the organism, it is also true that organisms have complex internal states with regulatory machinery with a set of continuous phenotype-environment interactions. The aim of the present paper is to provide a visual means to analyze these interactions in individuals and computational agents alike. An essential component of our approach is the representation of continuous internal states through the usage of the single continuous indicator we call an Allostasis Machine (AM). Consequently, we consider potential perturbation regimes for both naturalistic and virtual environments: within the naturalistic cases, it is possible to observe the effects of perturbations in isolation, or as overlapping, multiplicative events. In virtual cases, we can observe perturbations as the outcome of both realistic and fantastical environments. To conclude, we discuss how AMs can be utilized to improve our understanding of both the theoretical basis of embodied interaction and the dynamic regulation of complex psychophysiological states.

Introduction

There are many exciting opportunities to unify Human-Computer Interaction (HCI) concepts with Neuroscience. Some of these involve the implementation of brain-computer interfaces [1], which results in a number of intentional and unintentional feedback mechanisms [2]. Contrary to that approach, we propose that both the Neuro-HCI interdisciplinary relationship and the associated feedback mechanism is much richer: feedback mechanisms result in dynamic regulation, and can be understood as a relationship between the body, the internal model, and the environment.

In this context, we propose that the brain is not a computer, or extension of a computer, but an internal model that provides output to an Allostasis Machine (AM). AMs are composed of two parts, an internal model and a dynamical representation of internal model output, the latter of which produced a complex and highly nonlinear internal state trajectory over time. We derive our definition of a machine from Bongard and Levin [3]: AMs are an integrative mechanism that enable prediction of behavior at multiple scales. AMs are a representation of internal state as it interacts with the environment. In this way AMs are governed by the constraints of both environmental physics and cell/molecular biology of the nervous system. To understand the behavior of this organic machine, we will discuss the role of allostasis in physiological

¹ Orthogonal Research and Education Laboratory, Champaign, IL USA bradly.alicea@outlook.com

² OpenWorm Foundation, Boston, MA USA

³ University of Chieti-Pescara, Pescara Italy danielacialfi@gmail.com

regulation, which is influenced by physiological adaptation, physiological constraints, learning, and the fluctuations in the environment.

Role of AMs in Neuro-HCI Systems

We can begin to understand the role of AMs in Neuro-HCI systems by understanding how the interactions between brain, behavior, and environment are characterized by a form of self-organizing stability called autopoiesis. Autopoietic systems [4] are defined by fluctuating interactions between parts of a system that continually change as they interact with the environment. This not only characterizes the interactivity of a Neuro-HCI system, but directly supports the notion that the stability of a given AMs is contingent upon both its short- and long-term history. When autopoiesis is disrupted through anti-homeostatic [5] conditions, a system's organization disintegrates and in a Neuro-HCI leads to poor performance. On the other hand, autopoietic systems can demonstrate learning when the system's configurations are transformed in a beneficial manner [6]. The effect of perturbations on such a system can be demonstrated by the role of capacity in attention and memory. As a perturbation is experienced, the current configuration of the internal model can overcome low-magnitude perturbations, and capacity is maintained. At the same time, the internal model changes its configuration so that adaptive (or maladaptive) changes become possible. In an example where perturbation strength is held constant, the internal model will adapt and become more robust to future perturbation. In walking through our examples of AMs, the angle of perturbation and recovery characterize the strength of the perturbation and the recovery capacity of the internal model, respectively.

The relationship between information processing capacity and noise is also a very important component of internal regulation. It has been proposed [7] that so-called pink ($1/f$) noise characterizes interaction patterns (or cognitive coupling) between individuals and the technological environment. As noise is an integral feature of interaction, most of the noise is tolerable to the regulatory capacity of the internal model. However, the individual can experience large bursts of nonlinear turbulence that occur at multiple scales of complexity [8]. Multi-scale noise can occur in many different types of Neuro-HCI systems, including sensorimotor control [9]. Experiencing noise in this unpredictable way can overwhelm the individual's capacity to recover. Understanding the strategies used by an individual in the situations requires both a more formal representation of internal regulation and more explicit internal model mechanisms for the triggering/alleviation of breakdowns in cognitive coupling. The AM approach captures this quite well: as we will see, the combination of a self-regulating internal model and a dynamical representation allow us to observe the effects of these potential instabilities on performance in a naturalistic content.

Homeostasis as Internal Regulation

From our perspective, the nervous system is an embodied, autopoietic system in which stability is determined by dynamic regulation. While we propose that this form of regulation

involves allostasis, and particularly allostatic drive, it is necessary to revisit Cannon's [10] original concept of homeostasis. Homeostasis refers to the regulation of internal state relative to a set point or equilibrium, which can be derived from properties of the organism such as stress, psychophysiological systems, or emotional state. One demonstration of homeostasis in intelligent systems is W.R. Ashby's Homeostat [11], which was the first machine to maintain a steady state in the face of varying environmental input. Furthermore, homeostasis allows us to think about regulation in terms of feedback loops. Von Bertalanffy [12] describes this diagrammatic arrangement as circular, and has been used to describe the nonlinear effects resulting from interactions between environmental inputs and internal feedback. Homeostasis is often said to be regulation of the internal "milieu", which implies that a very broad conceptual variable is responsible for our measurement of this state. An internal state governed by homeostasis is an autonomous adaptive system that acquires regulatory capacity either actively through encoding memories or passively through an ability to recall past events [13]. Consequently, homeostasis, according to Puglisi et al [14] is an equilibrium model, which does not allow us to represent many far-from-equilibrium phenomena.

While we can think of homeostatic regulation as a vectorized dynamical system (not unlike time-series measurements of a thermostat), applying homeostasis as a model for characterizing change in the internal system forces us to reevaluate the notion of the physiological milieu. Bernard initially conceived of the internal milieu as an inherently fixed system that is able to perfectly match any challenge imposed by its environment [15]. The modern understanding of homeostasis predicts that internal systems make two types of adaptations: a regulatory adaptation to match the set point, and a dynamic adaptation that allows the internal system to expand the range of conditions under which the set point can be maintained. This latter type of adaptation allows the internal system to expand (and potentially contract) its homeostatic range [16], which ultimately requires us to think more broadly regarding the existence of stable states in our internal system.

Allostasis as Dynamic Homeostasis

While homeostasis allows us to think about internal regulation over time, it does not lend itself to the application of a more formal dynamical systems approach. Sterling [17, 18] and Sterling and Eyer [19] have proposed allostasis as an alternative to homeostasis. In their formulation, allostasis provides a set of regulatory principles with respect to homeostasis, particularly by introducing a formal mechanism for change over time (allostatic drive). Allostatic drive occurs when a system is perturbed in some way by the environment. Traditionally, these perturbations take the form of psychophysiological stressors, but can also be represented by phenomena such as multitasking distractions in a model of attentional capacity. Other authors [20, 21] have made the claim that allostasis is a predictive form of inference, and allows our internal model to respond to future challenges in an anticipatory fashion. The mechanism of

allostatic drive, along with the cognitive implications of allostatic control enable an internal model's trajectory to move away from the original setpoint and find new dynamic equilibria.

Allostasis Machines

AMs are vectors that capture the perturbations and stable states of a dynamical system. Given an initial condition (baseline), this vector can either suffer negative deflections due to environmental perturbation or positive recoveries due to the adaptability of an individual's internal model. The net movement of the vector over time captures allostatic drive, which leads to a regulated state. At any moment during this process, a perturbation can be delivered to the internal state: the perturbation can either temporarily distort the vertical (y-axis) component of the vector. Consequently, the process either returns to normal or creates a hysteretic distortion where the vector does not return to baseline. A generic example of an AMs is shown in Figure 1.

Figure 1. A non-technical example of an AMs internal state output trajectory, from initial condition to regulated state. Inset: demonstration of the flow of time, from initial condition to regulated state.

Complexity of the Internal State

According to Bernard's naive assumptions about the internal milieu [11], the selective representations of most cognitive models need to be considered as a complex, dynamical system. Yet internal states are much more complicated than the output of a simple dynamical system. Based on the variety of internal models featured in the neuroscience literature, such as those representing the cerebellum [22] and attention schema for subjective processing [23], internal models generate representations of cognitive states that are directly relevant to behavioral outputs. As such, they could be defined as invariant to psychophysiological states or innate mechanisms that might influence this state. This becomes important when we take into

consideration regulatory mechanisms that determine an individual's internal state as it actively explores its environment.

These innate mechanisms can be represented in a number of ways. Computational representations of the genome [24, 25], particularly gene expression networks (GRNs, see [26]), can approximate both the constraints and adaptive capacity of an organism's internal biology on processing environmental inputs. Particularly interesting in the phenomenon of synthetic effects that occur when multiple genes interact with each other [27]. This level of description is utilized in the internal model to facilitate robust and vulnerable responses to environmental perturbations. Nonlinear effects also emerge from the operation of the internal model, as gene expression interacts richly with the multiple sensory channels of naturalistic interaction with the environment.

Trajectories of the Internal State

Although Figure 1 is meant to capture event-oriented signals related to naturalistic behavior, the evolution of an internal model is open-ended [28] in nature. In the context of AMs, open-ended evolution is the continual reconfiguration of the internal model with no explicit goal or objective. This means AMs can sojourn to new regulatory regimes (and ultimately stable states) that correspond to adaptive changes due to continual interaction with the environment. The representation of internal model dynamics captures this dynamic balance between environmental influence and biological robustness as a generic dynamical system. Critically, our representation of internal model output (Figure 1) captures the strength of perturbations and recovery in the form of relative magnitudes and phase angles (Figure 2).

Yet how do we understand these trajectories in terms of specific interactions with tasks and the environment? One possible definition, rooted in Psychophysics, is a sensory signal that changes its magnitude with respect to time. Moving your fingertips over a ridged surface results in environmental perturbations of a given frequency and duration because perturbation frequency depends upon the spatial location of each ridge as well as the temporal interval between each ridge. This is further dependent on the speed of the touch effector, which is modulated by feedback from the nervous system. By contrast, perturbation duration involves both speed of the touch effector and the overall capacity of the internal state to withstand perturbations. In essence, this can be reduced to a simple oppositional relationship: robustness of the internal state and its trajectory versus the magnitude/duration of perturbation.

From this definition we can understand several components that influence the dynamical trajectory of our AM. Figure 2 demonstrates these steps in graphical form. The first step is to establish a baseline state (Figure 2a). Then a single perturbation of a given strength is delivered to the internal state (Figure 2b). This magnitude has an effect on the strength of a single perturbation (Figure 2c), which drives the state away from the baseline in a negative direction.

Perturbation depth is generally at a 45 degree angle to the baseline, but can be characterized by variable phase angles depending on whether the perturbation is encountered in a gradual or sudden manner. In our haptic ridges example, perturbations are rather sharp. By contrast, haptic exploration of a thermal gradient would be more gradual.

This leads us to a refractory period, during which the state trajectory generally recovers at a symmetrical phase angle to the perturbation (Figure 2d). Once again, this can depend on the application context. However, the recovery of an internal state trajectory does not always result in a return to the baseline state. When recovery is not equivalent to perturbation strength, then the trajectory is said to be hysteretic. Hysteresis might result from multiple perturbations in rapid succession, or from an insufficient memory of the previous state. The latter will be common for out-of-distribution magnitudes. Finally, the recovered internal state trajectory (Figure 2e) is dependent upon the timing of additional perturbations, which can determine the likelihood of transition to a new state. In material science, hysteresis is often referred to as a form of memory [29]. While a memory system is not made explicit in AMs, they nevertheless serve an integral role of their function.

Figure 2. A diagram of the external perturbation/internal state trajectory relationship. Lack of recovery capacity (e.g. hysteresis) of our internal state trajectory is equivalent to allostatic drive.

One area for further understanding is how to characterize the relative strength of a perturbation. Even when perturbations are described in terms of single sensations (quantities such as Newtons or Lumens), it is hard to describe their effect on the internal state. One way to do this is to define the baseline state as a mean background level of some sensory channel, with local deviations from this value as a perturbation of magnitude m . Corresponding to this is a

distribution function that can be estimated for the organism or agent that describes a general degree of robustness to outliers, rare events, and out-of-distribution phenomena.

Examples from Specific Environmental Representations

AMs are broadly applicable to a wide range of Neuro-HCI systems. Our approach is particularly amenable to naturalistic experiments, where organisms or agents are allowed to freely explore their environment. Future work will also include the study of cognitive phenomena such as visual flow and spatial cognition. In these cases, perturbations would be introduced through controlled misalignments of visual cues and affordances. This will require the integration of AMs with agent-based models such as developmental Braitenberg Vehicles [30] or serve as predictive models for Neuroimaging and Psychophysiological data.

As the AM approach might also serve as the basis for computational models of human-computer interaction, one highly relevant area is the study of virtual environments. With respect to hyperrealistic virtual environments, perturbations might take the form of latencies and stimuli of differential resolution. More impressionistic and fantastical virtual representations are less dependent on being as faithful to realism. However, affordances are still required to navigate the environment, and disruptions to their recognizability and relative location may serve to introduce perturbations of different frequencies and magnitudes.

AMs also allow us to model adaptation to virtual environments that occur both as a consequence of sensory accommodation [31] and longer-term biological plasticity [32]. For example, we might expect the ability to recover from a perturbation to be weakened upon first entering the virtual environment, while accommodation allows for the response to perturbation to be more robust over time [31]. Similarly, walking on a treadmill at degrees of incline while interacting with a CAVE environment [32] can have the same effect on psychophysiological indicators of performance.

Discussion

AMs are decoupled from measurements of neuronal or network output state, which makes them controversial as a model of intelligent behavior. Although they are simplistic representations for systems with many interacting components, they nevertheless serve as a heuristic to understand the mapping between internal state and the effects of a naturalistic environment. We can also use AMs to understand the effects of feedback at multiple temporal scales. This might be useful in the approximation of phenomena such as predictive processing [33]. Indeed, AMs allow for both minimal representational capacity and evaluation of error minimization [34]. It is important to stress that “evaluation” does not imply optimization: in the case of virtual environments, selective mismatch might be a desirable feature of the interaction.

Furthermore, considering AMs from the perspective of existing theoretical work on embodied interaction is also important. While we define a simple duality between an internal model and the external environment, affordances encountered in the environment present three-way interactions between an internal model, an external environment, and the body (morphology) itself [35]. While morphology mediates the interaction between internal and external worlds in this relationship, representations of internal state serve to mediate between innate mechanisms and environmental stresses. AMs might also help us understand how technology usage interplays with the internal state. The idea that technology becomes incorporated into the body over time [36] can be demonstrated quantitatively using the AM approach. This can be accomplished by using a range of model systems as a quantitative basis for a given trajectory. Examples of these include directly measuring the magnitude of perturbations such as attentional distractions or inertial forces, employing dimensionality reduction techniques on neurophysiological measurements, and using reinforcement learning simulations to derive quantitative learning curves.

As perturbations are disruptions to performance, they result not only from purely environmental factors, but from interactions with affordances. Under normal circumstances, interactions with affordances facilitate smooth interactions [37]. For example, doorknobs can afford twisting movements, while handles afford pulling movements. When these affordances are unclear or purposefully decoupled from environmental function, a specific perturbation is introduced. In the language of interaction, the information transfer rate between environment and individual is decreased due to an increase in noise [38]. This can be understood in the context of task performance and tool use. Tools that are compatible with the individual, particularly those that require a set of familiar skills, serve to minimize noise and maintain an allostasis machine in a stable state with little to no hysteretic response.

Tools that are well-practised also serve as extensions of the body schema [39], which has a basis in Primate physiology [40]. The idea of tool embodiment has been developed by [41] to demonstrate how the internal model, tools, and the environment interact to form a functional unit. As a given tool becomes familiar to the individual, the tool becomes an extension of the body schema. This involves attentional shifts to the task rather than the tool, thus improving performance on the tasks themselves. This is particularly advantageous in virtual environments, where tools are not as familiar as those found in real-world settings. This allows for technologies (tools in the environment) to become transparent with respect to the individual [42], thus acting as a type of affordance that assists in task performance. We can also interpret the incorporation of tools and other components of the environment into the body schema as a form of allostasis. The notion of allostatic control is a form of internal model regulation that acts as a form of selection for configurations that minimise expected free energy during future interactions [43]. Not only does this imply that prediction is a key component in regulation, but also that internal regulation is constrained by energetic processes.

There are several problem domains that might provide potential use cases in the future. The first of these involve the use of sensorimotor affordances that require a smooth action-perception loop and minimize the number and magnitude of perturbations during performance. This could include something like a handwriting-assistance technology for people with movement disorders. Musical interfaces that enable people to keep rhythm with a task is another potential application domain. The regulation and maintenance of spatial context is yet another potential use case, whether that be in terms of incorporating landmarks into the body schema or through eliminating spatial perturbations due to mental rotations. Finally, these domains can be combined in the composition of virtual environments, where perturbations are introduced due to a host of potential cognitive mismatches.

Moreover, AMs can also be applied to a broad range of agent-based settings, from social interaction to tests of biopsychological variation across a common set of interactions. In the case of social interaction, the AMs of different individuals can be compared during cooperation and competition. AMs also lend themselves to understanding the psychophysiology of cognitive states, particularly during continuous tasks like driving [44]. We can even use AMs as the basis for game-theoretic analysis, as switching between various internal states can be modeled as a strategic pursuit. In conclusion, it is possible to use agent-based modeling approaches to better understand variation in AMs with respect to the potential difficulty of an environment given a large number of possible internal representations.

References

- [1] Zander, T.O. and Kothel, C. (2011). Towards passive brain–computer interfaces: applying brain–computer interface technology to human–machine systems in general. *Journal of Neural Engineering*, 8, 025005.
- [2] Ritter, W. (2011). Benefits of subliminal feedback loops in Human-Computer Interaction. *Advances in Human-Computer Interaction*, 346492.
- [3] Bongard, J. and Levin, M. (2021). Living Things Are Not (20th Century) Machines: updating mechanism metaphors in light of the modern science of machine behavior. *Frontiers in Ecology and Evolution*, 9, 650726.
- [4] Maturana, H. and Varela, F. (1980). Autopoiesis and Cognition: the realization of the living. D. Reidel Publishing, Dordrecht, Holland.
- [5] Valentinov, V. (2017). Wiener and Luhmann on feedback: from complexity to sustainability. *Kybernetes*, 46(3), 386-399.

- [6] Winograd, T. and Flores, F. (1987). Understanding Computers and Cognition: A New Foundation for Design. Addison-Wesley, Boston, MA.
- [7] Meacham, D. and Casanova, M.P. (2018). The Over-Extended Mind? Pink Noise and the Ethics of Interaction-Dominant Systems. *Nanoethics*, 12, 269-281.
- [8] Holden, J.G., Van Orden, G.C., and Turvey, M.T. (2009) Dispersion of response times reveals cognitive dynamics. *Psychological Review*, 116, 318–342.
- [9] Bennett, D. (2020). Multifractal Mice: measuring task engagement and readiness-to-hand via hand movement. *PsyArXiv*, doi:10.31234/osf.io/tq53k
- [10] Cannon, W.B. (1930). The Wisdom of the Body. W.W. Norton, New York.
- [11] Ashby, W.R. (1960). Design for a Brain: the origin of adaptive behavior. Chapman and Hall, New York.
- [12] Von Bertalanffy, L. (1968). General System Theory. George Braziller, New York.
- [13] Dworkin, B.R. (1993). Learning and physiological regulation. University of Chicago Press, Chicago, IL.
- [14] Puglisi, A., Sarracino, A., and Vulpiani, A. (2017). Temperature in and out of equilibrium: A review of concepts, tools and attempts. *Physics Reports*, 709-710, 1-60.
- [15] Gross C.G. (1988). Claude Bernard and the constancy of the internal environment. *The Neuroscientist*, 4, 380–385.
- [16] Davies. K.J.A. (2016). Adaptive Homeostasis. *Molecular Aspects of Medicine*, 49, 1–7.
- [17] Sterling, P. (2004). Principles of allostasis: optimal design, predictive regulation, pathophysiology and rational therapeutics. In: Allostasis, Homeostasis, and the Costs of Adaptation. J. Schulkin (ed.) Cambridge University Press: Cambridge, UK.
- [18] Sterling, P. (2011). Allostasis: A Model of Predictive Regulation. *Physiology & Behavior*, 106(1), 5-15.
- [19] Sterling, P., and Eyer, J. (1988). Allostasis: a new paradigm to explain arousal pathology. In *Handbook of Life Stress, Cognition and Health*, S. Fisher & J. Reason (eds.), 629–649. Wiley and Sons, New York..

- [20] Corcoran, A.W., Pezzulo, G., and Howhy, J. (2020). From allostatic agents to counterfactual cognisers: active inference, biological regulation, and the origins of cognition. *Biology and Philosophy*, 35. doi:10.1007/s10539-020-09746-2.
- [21] Pezzulo, G., Rigoli, F., and Friston, K. J. (2015). Active inference, homeostatic regulation and adaptive behavioural control. *Progress in Neurobiology*, 134, 17–35.
- [22] Wolpert, D.M., Miall, C., and Kawato, M. (1998). Internal models in the cerebellum. *Trends in Cognitive Science*, 2(9), 338-347.
- [23] Graziano, M.S.A. and Webb, T.W. (2015). The attention schema theory: a mechanistic account of subjective awareness. *Frontiers in Psychology*, 6, 500.
- [24] Albert, J. (1999). Computational modeling of an early evolutionary stage of the nervous system. *Biosystems*, 54(1-2), 77-90.
- [25] Aguilar, W., Santamaria-Bonfil, G., Froese, T., and Gershenson, C. (2014). The past, present, and future of artificial life. *Frontiers in Robotics and AI*, 1, 8.
- [26] Cussat-Blanc, S., Harrington, K., and Banzhaf, W. (2019). Artificial Gene Regulatory Networks: a review. *Artificial Life*, 24(4), 296-328.
- [27] Sanford, E.M., Emert, B.L., Cote, A., and Raj, A. (2020). Gene regulation gravitates toward either addition or multiplication when combining the effects of two signals. *eLife*, 9, e59388.
- [28] Adams, A., Zenil, H., Davies, P.C.W., and Walker, S.I. (2017). Formal definitions of unbounded evolution and innovation reveal universal mechanisms for open-ended evolution in dynamical systems. *Scientific Reports*, 7, 997.
- [29] Ortin, J. and Delaey, L. (2002). Hysteresis in shape-memory alloys. *International Journal of Non-Linear Mechanics*, 37(8), 1275-1281.
- [30] Dvoretiskii, S., Gong, Z., Gupta, A., Parent, J., and Alicea, B. Braitenberg Vehicles as developmental neurosimulation. *arXiv*, 2003.07689.
- [31] Redding, G.M., Rossetti, Y., and Wallace, B. (2005). Applications of prism adaptation: a tutorial in theory and method. *Neuroscience and Biobehavioral Reviews*, 29(3), 431-444.

- [32] Kim, A., Darakjian, N., and Finley, J.M. (2017). Walking in fully immersive virtual environments: an evaluation of potential adverse effects in older adults and individuals with Parkinson’s disease. *Journal of NeuroEngineering and Rehabilitation*, 14(16), doi:10.1186/s12984-017-0225-2.
- [33] Nave, K., Deane, G., Miller, M., and Clark, A. (2020). Wilding the predictive brain. *WIREs Cognitive Science*, doi:10-1002/wcs.1542.
- [34] Williams, D. (2017). Predictive Processing and the representation wars. *Minds and Machines*, doi:10.1007/s11023-017-9441-6.
- [35] Dourish, P. (2001). *Where the Action Is: the foundations of embodied interaction*. MIT Press, Cambridge, MA.
- [36] Clark, A. and Chalmers, D.J. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- [37] Chemero, A. (2003) An Outline of a Theory of Affordances. *Ecological Psychology*, 15(2), 181-195.
- [38] Hornback, K. and Oulasvirta, A. (2017). What is Interaction? *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 5040-5052.
- [39] Bergstrom, J., Mottelson, A., Muresan, A., and Hornback, K. (2019). Tool Extension in Human–Computer Interaction. *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1-11.
- [40] Maravita, A. and Iriki, A. (2004). Tools for the body (schema). *Trends in Cognitive Science*, 8(2), 79-86.
- [41] Alzayat, A., Hancock, M., and Nacenta, M.A. (2017). Measuring Readiness-to-Hand through Differences in Attention to the Task vs. Attention to the Tool. *Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces*, 42–51.
- [42] Wheeler, M. (2019). The reappearing tool: transparency, smart technology, and the extended mind. *AI and Society*, 34, 857-866.
- [43] Kiverstein, J. and Sims, M. (2021). Is free-energy minimisation the mark of the cognitive? *Biology and Philosophy*, 36, 25. doi: 10.1007/s10539-021-09788-0.

[44] Lohani, M., Payne, B.R., and Strayer, D.L. (2019). A review of Psychophysiological measures to assess cognitive states in real-world driving. *Frontiers in Human Neuroscience*, 13, 57. doi:10.3389/fnhum.2019.00057.