ResearchGate

See discussions, stats, and author profiles for this publication at: https://www.researchgate.net/publication/342932474

Dynamic Interactive Artificial Intelligence: Sketches for a Future AI Based on Human-Machine Interaction

Conference Paper · January 2020

DOI: 10.1162/isal_a_00350

CITATION
1

READS
47

2 authors, including:

Dobromir Dotov
McMaster University

16 PUBLICATIONS **46** CITATIONS

SEE PROFILE

All content following this page was uploaded by Dobromir Dotov on 21 October 2021.

The user has requested enhancement of the downloaded file.

Dynamic Interactive Artificial Intelligence: Sketches for a Future AI Based on Human-Machine Interaction

Dobromir Dotov[1] and Tom Froese[2]

[1]McMaster University, Hamilton, ON L8P2S3, Canada
[2] Embodied Cognitive Science Unit, Okinawa Institute of Science and Technology Graduate University, Onna-son, Okinawa 904-0495, Japan
dotovd@mcmaster.ca

Abstract

We propose to designate as dynamic interactive artificial intelligence (dAI) a cross-section of existing work in artificially designed and artificially evolved systems meant for minimal forms of interaction with human users. This approach borrows principles from artificial life and human movement science to avoid pitfalls of traditional AI. Counter to tradition, it prioritizes user-machine inter-dependence over autonomy. It starts small and relies on incremental growth instead of trying to implement advanced complete functionality. It assumes a perceptual ontology founded on movement coordination rather than object classification. Its development process is better described as reverse self-organization rather than reverse engineering. dAI can be viewed as a precursor to or pre-condition for enactive AI and an alternative to traditional frameworks grounded on information representation. We then give examples from our work in human movement science where we have used minimal dynamic interactive agents to induce specific beneficial effects in human participants' movement skills. We also show how dAI can be exploited by both connectionist and symbolic AI.

Introduction

Imagine the following two scenarios for the development of an intelligent vehicle. First, the traditional ideal, the fully self-sufficient car will be a transportation service. You step inside the passenger area and the machine is so independent that you do not even have to see where it is going. The intelligence of the machine substitutes for some of our capacities as perceiving and self-moving animals, to the point that the human drivers only exist to fulfill the intelligent car's purpose. In the second scenario smart cars can serve the role of an adaptive interface between the driver, the road, and the vehicle's complicated physics, increasing and sharpening our abilities to move safely at high speeds on cluttered terrain. Cars exist to enable new forms of human movement. In addition to being somewhat threatening, the former scenario still looks very difficult to achieve, if even slowly becoming possible with the present tools and principles. The latter scenario, however, empowers the human in the loop and is a matter of future development.

Here we suggest that the field of AI can learn from artificial life's interaction-based approach, starting with small functionality and big principles and allowing technology to adapt and evolve. What is needed for fluid integration of biological, artificial, and cognitive systems is to avoid one of AI's original mistakes, namely always aiming for full self-sufficiency and complete mimicry of human skills but failing to achieve them. Before we learn to grow artificial intelligent "organisms", we might have to become better at growing artificial cognitive "organs".

Even though AI is frequently divided into symbolic AI (sAI) and artificial neural nets (ANNs), both traditions try to solve the difficult problems of generativity and generalization. For this reason, it is necessary to define a third flavor of AI that is not concerned with autonomy and as such is less ambitious and more attainable. It can be referred to as dynamic interactive AI (dAI). dAI can lay claims to some of the AI territory because several existing fields of research loosely bound by shared principles have been addressing some of AI's original questions without the explicit label. The incomplete list includes bio-inspired and soft robotics (Pfeifer et al., 2007; Nakajima et al., 2018), embodied intelligence (Pfeifer and Bongard, 2006), social enactive robotics (Fong et al., 2003), enactive AI (Froese and Ziemke, 2009), living technology (Bedau et al., 2010), and physical intelligence (Turvey and Carello, 2012), among others. dAI not so much proposes novel ideas as it captures common principles. The fruitful exchange among such diverse fields requires for shared language and a set of minimal primitives while avoiding the over-generalization of existing concepts such as AI and autonomy. To begin with, these approaches emphasize that for intelligence to develop it first needs a body, where having a body can be understood more generally as dynamically interacting with an environment. Instead of trying to internalize meaning, dAI creates conditions for adaptive behaviors to emerge from dynamic interactions with a machine. This capitalizes on the idea that humans to some extent already are "natural-born cyborgs" (Clark, 2003).

The argument here is not that general AI will somehow emerge from dAI, for example if there were a lot of the latter. Rather, our aim is to show that a minimal definition of dAI

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

applies to ongoing work that does not define itself as AI. Unlike general AI which remains just a promise, the conditions for minimal dAI are already present and deserve more attention. There are also philosophical reasons to be optimistic about dAI. Minimal collective intentionality requires coordinated task performance with shared physical constraints, including technological and AI-driven systems, but not necessarily shared beliefs, theory of mind, symbol-based meta reasoning, etc. (Satne and Salice, 2019). Thus, in contrast to general AI, the aim of dAI is not to communicate with the intellectual mind of the users, but rather to become integrated with their adaptively behaving body. Instead of producing information content to be consumed by the users, the idea is to implicitly reshape their relation with the world (Froese et al., 2012). Importantly, just because dAI examples may rely on minimal internal dynamics, this does not completely limit their functional capacity when multiple agents are interacting. Interaction between simulated mobile robots can enhance internal complexity and the dynamic dimension of the task space beyond the robots' intrinsic complexity (Candadai et al., 2019).

Human movement science, experimental psychology, and sensorimotor neuroscience have also been experimenting with dynamic paradigms for interaction, be it to understand principles for acting in a complex environment or to study human-machine interaction. Assistive movement devices would not typically be categorized as AI but they may qualify as dAI, creating the possibility to learn from human behavioral studies. The purpose of this paper is to show through examples from empirical work how dynamics and interaction already allow for minimal dAI systems to be useful for human users. Such studies are not usually meant to address the field of AI directly and their relevance is not realized. This is the case, we suggest, partly because concepts of AI are prohibitively strong. For this reason, we first need to show where dAI stands relative to ANNs.

dAI and ANNs

It is easier to see the distinction between dAI and sAI than between dAI and ANNs. Aren't neural nets intrinsically dynamic phenomena? Historically there has been an affinity between the embodied/enactive strands of cognitive science and connectionism. Neural nets can serve as structural components of dAI and enable adaptive interaction with an environment (Kadihasanoglu et al., 2017; Beer and Gallagher, 1992). In what follows we point out why ANNs need to become more interactive.

40 years ago, while cognitivism was busy criticizing the radical interaction-centered notions in Gibson's latest book (1979), Geoffrey Hinton - a key figure in today's deep ANNs - pointed out in a commentary (1980) to a BBS article (Ullman, 1980) that there were good ideas in ecological theory and connectionism was poised to exploit them in the context of unsupervised discovery of patterns inherent in the environment, not internally deduced following a rule-based formal logical system. Later Hinton continued to make occasional references to Gibson's work (Hinton and Becker, 1990). Nowadays, machine learning for robotics is re-discovering the importance of affordances (Hsu, 2019) while some theoretical neuroscientists are trying to make sense of deep learning through the lens of direct perception as a process they call direct fit Hasson et al. (2020). Deep learning is also becoming relevant for learning hand-eye coordination (Levine et al., 2018) and it will be interesting to see if this approach will also prove useful in learning agent-agent coordination. The task is not trivial because agent-agent systems do not necessarily have to be synchronizing and/or cooperative. Stable and functional antagonistic patterns in sports are afforded by anisotropic coupling (de Poel, 2016).

Another relevant Gibsonian principle is self-structured information: sensorimotor coordinated interaction generates information that is structured by the agent itself. This could be exploited by ANNs as a signal for self-supervised learning (Pfeifer et al., 2007). A lot of this information is relational, however, implying that for it to be discovered, purposeful interaction with the environment is required, not just passive comparison of the sensory and motor streams. Active inference within the free energy minimization (FEM) framework (Friston et al., 2009) is among the modern attempts to make use of this principle, to the extent that FEM can be applied to AI.

Evident in such efforts is the need to shift emphasis from classification-based to interaction-based training paradigms. ANNs are often placed in a one-directional, sequential pipeline where learning is a blind evolution-like process the goal of which is for perception to achieve full recognition of objects in the environment before action is even to be engaged (Hasson et al., 2020). This classification-based methodology for developing AI implies an object-based ontology: an AI is a developing subject in a world full of pre-existing objects whose identities, classes, and properties are to be learned bit by bit from experience. In this context, generalization from experiencing exemplars to higher-order patterns tends to be the predominant challenge to ANNs. We agree with Hasson et al. that this problem is overemphasized, but for a different reason: ANNs need to face their ontology problem by acknowledging the primacy of interaction. Otherwise ANNs are bound to miss invariants for interaction. Importantly, it is not a new idea that artificial agents driven by ANNs can be evolved to first and foremost interact, rather than detect objects and then interact (Kadihasanoglu et al., 2017; Beer and Gallagher, 1992). dAI emphasizes the importance of such work.

What is missing from an object-based ontology is the discovery of interaction patterns defined at the level of the human-machine system. Where such stable patterns are functional, in theory dAI should enable and exploit their

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

open-ended discovery. This is because closed-loop system-environment coupling affords not only stable patterns but also information about them. In one of the studies described below, participants interacted with a complex system, a chaotic oscillator, affording more than one patterns of stable synchronization. We did not instruct participants exactly how to control the oscillator; they spontaneously settled on a pattern that was feasible for them and that stabilized the chaotic oscillator.

It could be that solutions from an interaction-based ontology scale much better in an increasingly complex environment than solutions from an object-based ontology. Hasson et al. point out that a set of 1000 objects is already a great feat for an ANN and yet does not even account sufficiently for the perceptual world of an animal. This might seem like an insufficient number if one assumes that an animals' perceptual ontology is the same as that of a natural scientist. Affordances greatly simplify the task of the perceptual system inasmuch as getting around in the world is concerned (Dotov et al., 2012). To start with, I just need to perceive two possibilities for action: sturdy surfaces that allow me to maintain my body in a stable relation against them, and vectors of motion directed towards me and fast enough to cancel my ability to escape from them.

Reverse-self-organizing minimal human-machine interaction tasks

Human-machine interaction and human movement science are full of examples of dAI that does not declare itself as such. For example, Raffard et al. (2018) used an embodied, pre-cognitive approach to inter-personal interaction to design an artificial agent that increases synchronization behavior and rapport in individuals with schizophrenia. Iqbal et al. (2016) hard-coded a group anticipation index based on dynamic systems theory in order to facilitate synchronization among a group of a human and several robots.

In developing the stimuli in the following experimental paradigms we assumed an approach that can be referred to as *reverse-self-organizing* instead of reverse-engineering. The ultimate solution was not fully specified to participants and not determined in the task design. Instead, after an in silico search, where we simulated multiple systems and interaction scenarios, we selected a dynamic system and a parameter set allowing for coordination to appear. This approach is needed when investigating the boundary conditions and usefulness of patterns that emerge in the interaction between a human agent and a dynamic system. Note that the opposite approach tends to dominate human experimental psychology where full specification and control of the stimulus is usually required in advance. We also believe that in general 'reverse-self-organizing' is very apt for the sort of work seen in the field of artificial life.

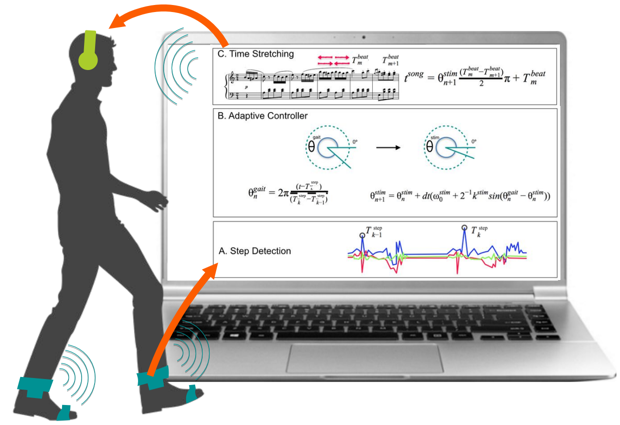

Figure 1: Closed-loop rhythmic auditory stimulation interacting with gait. The aim was to induce, not produce spontaneous synchronization between musical beats and the walker's footsteps by setting the whole task space in a parameter range within the entrainment domain of a two-unit Kuramoto system of coupled oscillators. Figure adapted from (Dotov et al., 2019). (A video clip showing the task dynamics in a simulated trial at https://vimeo.com/297434940.)

Interactive cueing for spontaneous auditory-motor synchronization in PD patients' gait

In the context of improving rhythmic auditory cueing for the gait of patients with PD, we developed a gait-to-hearing feedback system with one degree of internal dynamics (Dotov et al., 2019). The stimulus timing consisted of a phase oscillator with a Kuramoto coupling from gait phase. The intrinsic tempo and coupling gain were manipulated to tune the stimulus to the individual walker's dynamic profile but without going all the way to automatic, forced synchronization. Thus, we did not enforce but provided the conditions for spontaneous entrainment to emerge. Indeed, this cueing strategy proved most useful, compared to non-adaptive or automatically adaptive strategies, as it resulted in effortless synchronization and the most gain in cadence (Dotov et al., 2019). As participants were not told to synchronize with the stimulus and the stimulus was adaptive and tuned to the individual without automatically falling at the precise time of footsteps, see Fig 1, we provided the conditions for spontaneous entrainment to emerge (Dotov et al., 2019).

Interactive learning of chaos control to acquire beneficial variability

Dynamic principles allow implementing natural variability generators that explore and sync to patterns. This work was inspired by robotic studies where intrinsically chaotic con-

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

trollers seen as central pattern generators with adaptive variability discovered stable periodic patterns of interaction with a physical system through the phenomenon of feedback resonance (Pitti et al., 2010). Can a dAI system with certain desirable amount of variability and indeterminacy be used to train human motor variability? In an auditory-motor coordination task we asked participants to try to learn to control a chaotic system with their hand movements, see Fig 2A. Performance was compared against a non-interactive chaotic and a non-interactive periodic condition. As expected, causal interaction measured with the transfer entropy between the human and artificial agents increased along trial but only in the interactive chaotic condition, see Fig 2B-D. Pre-post transfer tests suggested that the interactive task was more beneficial relative to the two other conditions (Dotov and Froese, 2018a,b). Here we performed additional analysis motivated by the dAI ideas of under-determined design. We found a minimal form of open-endedness: without explicit instruction participants discovered that they could perform the task either by locking on a period-1 or a period-2 rhythm, Fig 3. Note that the apparatus obeyed deterministic dynamics but multiple behavioral scenarios were afforded in the system of user and stimulus. Hence, the apparatus supporting dAI can be relatively simple and deterministic but dAI as a phenomenon defined across the human-machine system can possess a small degree of indeterminacy and open-endedness.

dAI is about transparent human-machine interaction, not complete autonomy

The approach advocated here continues ideas advanced in the domain of ecological interface design, with one important difference. Past work applying ecological principles to human-machine interaction has dealt mostly with situations where the machine is a tool towards achieving a certain goal, possibly even spanning a means-end hierarchy (Vicente and Rasmussen, 1990). dAI needs not restrict itself to mediating the user's existing goals. Instead, it can help create novel forms of interaction and even novel tasks that do not make part of the user's typical ecology. For instance, in the first experiment described here participants walked along with a musical auditory stimulus but the task was not to entrain to the music. Walking in synchrony with the beat of a song is not typically a thing that we do in our normal lives. It does happen occasionally, however, and we may even notice that it is interesting and pleasant. More generally, a key aspect of dAI is that the focus of engineering is no longer subsumed only by the question "What is its function?" but also takes seriously the question "How does it feel?"

While dAI allows for open-endedness, it is not its top priority. We assume that there is a long path of multiple open-ended advancements from interactive agents to the full autonomy expected from general AI. In terms of classifying the types and pathways to living technology (Bedau et al.,

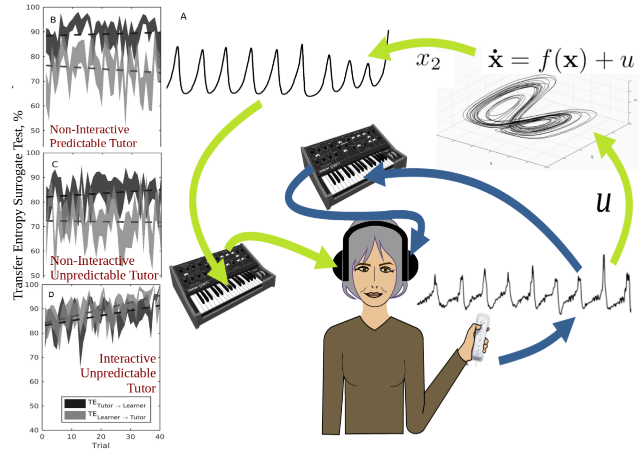

Figure 2: UNpredictable Interactive system with SONnified movement (UNISON). A) The transformed sensor state (accelerometer) of a hand-held device was streamed to a computer and sonified by mapping it to the pitch of a pure tone in the one channel (blue arrows). A state variable (x_2) of the chaotic Chua oscillator was similarly sonified in the left channel (green arrows) and driven by the hand movement signal u scaled by a coupling gain ϵ. $\epsilon = 0$ made for a non-interactive condition. In a third condition that was periodic and non-interactive the tutor was replaced with a sine wave. Sample video clips at https://vimeo.com/267437234. B-D) In each condition, the surrogate test for significant transfer entropy from user to artificial tutor and from artificial tutor to user, averages (SE) across participants shown along practice trials. Figure adapted from (Dotov and Froese, 2018a).

2010), dAI incorporates some pre-conditions for secondary living technology of mixed class as it requires software, hardware, and a human user, but does not cover many criteria such as self-maintenance, self-repair, and agency. For similar reasons dAI is less ambitious than social enactive robotics (Fong et al., 2003) which raises the bar very high. It is a tall task for robots to read off intentional states from humans. We assume, however, that before something is living technology it needs to be interactive. Hence, as a small step within the longer project of reverse-self-organizing living technology, we suggest starting with certain minimal conditions for useful interactive technology.

The apparatus supporting dAI does not do much; it is more of an interface than an agent. Looking at the conditions for something to possess intrinsic meaning like living beings (Froese and Taguchi, 2019), dAI cannot possess much depth because it is supported by deterministic physical phenomena and is not meant to be sensitive and adaptive to its precarious existence. dAI's potential is revealed when seen as an interaction pattern over a human-machine system. In the asymmetric relation between autonomous agent and environment (Froese, 2014), dAI plays a more passive role. Hence, dAI

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

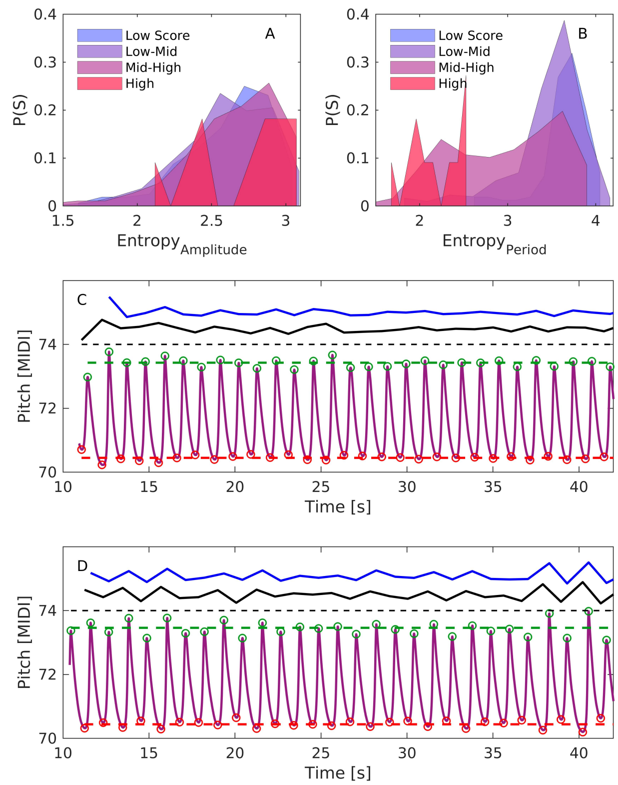

Figure 3: In the interactive condition the participant can discover how to stabilize the stimulus either on a period-1 or period-2 orbit. We analyzed the entropies of the series of periods (inter-peak intervals) and amplitudes of each stimulus cycle. A) The distribution of stimulus amplitude entropies across all trials, divided in four groups by level of performance. A bimodal distribution emerges for the highest scoring trials. B) Same for periods. C) and D) Sample trials exhibiting approximately equal levels of synchronization performance ($C_{max}/RMSE = 1.16$ and $C_{max}/RMSE = 1.17$, respectively) but either a constant (period-1) or an alternating (period-2) stimulus pattern, visualized by the series of half-amplitudes (black) and periods (blue) in the insets, and confirmed by the entropy levels ($S_{Period} = 2.067$ and $S_{Period} = 2.807$, respectively).

avoids (for now) the problem of a clash of autonomies that is to appear if genuinely living technology or strong artificial life is to also serve as technology for human use.

The two empirical studies that we presented here involve minimal forms of dAI with very little capacity for adaptation. dAI is not concerned with autonomy as much as with creating possibilities for users to expand their behavioral domain and embodiment by exploiting principles of animal-environment mutuality. Being artificial, dAI requires that the conditions for its emergence be externally guided. Yet, a given instance of dAI is to be discovered by the user. Importantly, the unspecified boundary between user and machine constitutes a sharp distinction between dAI and other forms of AI research. Here dAI is thought of as something that subsumes both the human user and apparatus, not as the apparatus that the user interacts with.

dAI may contain a seed of the properties associated with this kind of strong artificial life. In our studies the functional coordination patterns were precarious and not fully determined by instruction or design. As such these dynamics point in the direction of an adaptively behaving organ, like a hand autonomously shaping itself appropriately to enable the grasping of a coffee cup. In the future, combining ANNs and dAI promises greater parameter flexibility. Sensorimotor learning with deep nets is starting to gain traction (Levine et al., 2018). It is yet to be determined just how adaptive dAI would become from incorporating deep ANNs to transform sensory signals to motor output.

dAI and sAI

It is useful also to point out where dAI stands in relation to symbolic AI (sAI). To begin with, by definition dAI is easier to incorporate in a human's sensorimotor world than sAI. Furthermore, sAI with its emphasis on functional completeness and symbolic interfacing can become a source of alienation for users. On the other hand, dAI is closer to narrow AI than general AI because dAI is not generative. While tools can provide a range of unexpected affordances to a human user, they are by all means constrained by the user's skills, needs, and capacities.

In sAI, meaning is the object of formalization and internalization in representational structures, accomplished separately in the human and the machine. In dAI, meaning is a feature of human-machine interaction. A dAI agent does not try to internalize such meaning. The pleasure in spontaneously finding yourself moving in synchrony with something in the environment might constitute the meaning emerging while interacting with a given sound-making dAI but this meaning cannot be captured, stored, and transferred.

dAI is not concerned with duplicating particular skills that humans or other animals exhibit. Turing's idea of functional mimicry has become a modus operandi in AI research. The history and theory of technology, however, reveal that truly innovative technology is often not planned (Ihde, 1999). There could also be a pragmatic social virtue in this as it lessens the severity of issues with robot rights and regulation. The aim of empowering humans with dynamic interfaces and augmenting their embodiment is ethically less problematic than designing artificial systems to mimic selected human powers because the mimicry approach threatens to replace those humans whose skills are being reproduced.

The relation between dAI and sAI needs not be adversarial. dAI is not a separate approach to AI but addresses fundamental pre-conditions that can contribute to both sAI

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

and ANNs. Some flavours of AI have been trying to build dynamic, network, and cognitive primitives into their designs from the start, such as Haykin's cognitive radar (2005). A potential future exchange between dAI and sAI would not be totally surprising given that a viable solution to the symbol-grounding problem could be to rely on the sensorimotor system (Barsalou, 1999). The evidence that in humans conceptual processing could rely on the motor system goes beyond mere stimulus-response compatibility, implying that the motor system has a functional role (Vankov and Kokinov, 2013). If AI is to follow the same strategy, it first would need to build a repertoire of task-specific perceptual-motor skill.

Conclusions

We proposed the label dynamic interactive artificial intelligence (dAI) to capture a common thread running through multiple fields of artificial life, robotics, human-machine interaction. dAI is needed in order to clarify the constraints and capacities of a new generation of devices for bi-directional human-machine interaction with a high level of sophistication and dynamic bandwidth. Designing dAI follows a principle of reverse-self-organizing useful patterns of human-machine coordination rather than reverse-engineering specific effects. This means drawing boundary conditions that foster spontaneous patterns to emerge in the human-machine system, rather than designing the full behavior. Unlike other forms of AI that focus on the autonomy of the artificial system, dAI is defined across the human-machine system hence this problem is less relevant. dAI is not ambitious when it gets to its functional range, self-repair, and capacity for self-concern. We believe it can be fruitful to give up on autonomy and generalization in the short term in order to enable narrow but truly spontaneous interactive scenarios. As a form of technology, dAI is to transparently open up new ways of interacting with the world for individuals prior to fulfilling enactive AI's premise of autonomy and pursuit of the agent's own conditions of persistence. A future direction of research is to use ANNs to optimize dynamic interactive agents allowing for co-adaptation between human and artificial agents, in a union of ANNs, information about affordances, and dynamic interactive systems.

References

Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and brain sciences*, 22(4):577–660.

Bedau, M. A., McCaskill, J. S., Packard, N. H., and Rasmussen, S. (2010). Living technology: Exploiting life's principles in technology. *Artificial Life*, 16(1):89–97.

Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive behavior*, 1(1):91–122.

Candadai, M., Setzler, M., Izquierdo, E. J., and Froese, T. (2019). Embodied dyadic interaction increases complexity of neural dynamics: A minimal agent-based simulation model. *Frontiers in psychology*, 10:540.

Clark, A. (2003). *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. OUP.

de Poel, H. J. (2016). Anisotropy and antagonism in the coupling of two oscillators: Concepts and applications for between-person coordination. *Frontiers in psychology*, 7:1947.

Dotov, D. and Froese, T. (2018a). Entraining chaotic dynamics: A novel movement sonification paradigm could promote generalization. *Human Movement Science*, 61.

Dotov, D. and Froese, T. (2018b). Mutual synchronization and control between artificial chaotic system and human. In *Artificial Life Conference Proceedings*, pages 109–110. MIT Press.

Dotov, D. G., Cochen de Cock, V., Geny, C., Ihalainen, P., Moens, B., Leman, M., Bardy, B., and Dalla Bella, S. (2019). The role of interaction and predictability in the spontaneous entrainment of movement. *Journal of Experimental Psychology: General*, 148(6):1041–1057.

Dotov, D. G., Nie, L., and de Wit, M. M. (2012). Understanding affordances: History and contemporary development of Gibson's central concept. *Avant*, 3(2):28–39.

Fong, T., Nourbakhsh, I., and Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and autonomous systems*, 42(3-4):143–166.

Friston, K. J., Daunizeau, J., and Kiebel, S. J. (2009). Reinforcement learning or active inference? *PloS one*, 4(7).

Froese, T. (2014). Bio-machine hybrid technology: A theoretical assessment and some suggestions for improved future design. *Philosophy & Technology*, 27(4):539–560.

Froese, T., Suzuki, K., Ogai, Y., and Ikegami, T. (2012). Using human–computer interfaces to investigate 'mind-as-it-could-be' from the first-person perspective. *Cognitive Computation*, 4(3):365–382.

Froese, T. and Taguchi, S. (2019). The problem of meaning in ai and robotics: Still with us after all these years. *Philosophies*, 4(2):14.

Froese, T. and Ziemke, T. (2009). Enactive artificial intelligence: Investigating the systemic organization of life and mind. *Artificial Intelligence*, 173(3-4):466–500.

Gibson, J. (1979). *The ecological approach to visual perception. Boston, MA, US*. Houghton, Mifflin and Company.

Hasson, U., Nastase, S. A., and Goldstein, A. (2020). Direct fit to nature: An evolutionary perspective on biological and artificial neural networks. *Neuron*, 105(3):416–434.

Haykin, S. (2005). Cognitive radio: brain-empowered wireless communications. *IEEE Journal on Selected Areas in Communications*, 23(2):201–220.

Hinton, G. E. (1980). Inferring the meaning of direct perception. *Behavioral and Brain Sciences*, 3(3):387–388.

Hinton, G. E. and Becker, S. (1990). An unsupervised learning procedure that discovers surfaces in random-dot stereograms. In *Proc. IEEE/INNS International Joint Conference on Neural Networks*, volume 1, pages 218–222. Hillsdale, NJ. Erlbaum.

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

Hsu, J. (2019). Machines on mission possible. *Nat. Mach. Intell*, 1(3):124–127.

Ihde, D. (1999). Technology and prognostic predicaments. *AI & Society*, 13(1-2):44–51.

Iqbal, T., Rack, S., and Riek, L. D. (2016). Movement coordination in human–robot teams: a dynamical systems approach. *IEEE Transactions on Robotics*, 32(4):909–919.

Kadihasanoglu, D., Beer, R. D., and Bingham, G. P. (2017). An evolutionary robotics model of visually-guided braking: Testing optical variables. In *Artificial Life Conference Proceedings 14*, pages 230–236. MIT Press.

Levine, S., Pastor, P., Krizhevsky, A., Ibarz, J., and Quillen, D. (2018). Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 37(4-5):421–436.

Nakajima, K., Hauser, H., Li, T., and Pfeifer, R. (2018). Exploiting the dynamics of soft materials for machine learning. *Soft Robotics*, 5(3):339–347.

Pfeifer, R. and Bongard, J. (2006). *How the body shapes the way we think: a new view of intelligence*. MIT press.

Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318(5853):1088–1093.

Pitti, A., Niiyama, R., and Kuniyoshi, Y. (2010). Creating and modulating rhythms by controlling the physics of the body. *Autonomous Robots*, 28(3):317–329.

Raffard, S., Salesse, R. N., Bortolon, C., Bardy, B. G., Henriques, J., Marin, L., Stricker, D., and Capdevielle, D. (2018). Using mimicry of body movements by a virtual agent to increase synchronization behavior and rapport in individuals with schizophrenia. *Scientific reports*, 8(1):1–10.

Satne, G. and Salice, A. (2019). Shared intentionality and the cooperative evolutionary hypothesis. *Minimal coordination and shared agency. Berlin: Springer*.

Turvey, M. T. and Carello, C. (2012). On intelligence from first principles: Guidelines for inquiry into the hypothesis of physical intelligence (pi). *Ecological Psychology*, 24(1):3–32.

Ullman, S. (1980). Against direct perception. *Behavioral and Brain Sciences*, 3(3):373–381.

Vankov, I. and Kokinov, B. (2013). The role of the motor system in conceptual processing: Effects of object affordances beyond response interference. *Acta psychologica*, 143(1):52–57.

Vicente, K. J. and Rasmussen, J. (1990). The ecology of human-machine systems ii: Mediating'direct perception'in complex work domains. *Ecological psychology*, 2(3):207–249.

Downloaded from http://direct.mit.edu/isal/proceedings-pdf/isal2020/32/139/1908406/isal_a_00350.pdf by guest on 01 July 2021

View publication stats