Available online at www.sciencedirect.com

ScienceDirect

Procedia Manufacturing 00 (2015) 000–000

Procedia MANUFACTURING

www.elsevier.com/locate/procedia

6th International Conference on Applied Human Factors and Ergonomics (AHFE 2015) and the Affiliated Conferences, AHFE 2015

Situating Cognition within the Virtual World

Paul R. Smart[a]* and Katia Sycara[b]

[a]*University of Southampton, Southampton, SO17 1BJ, UK*
[b]*Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*

Abstract

Cognitive architectures and virtual environments have a long history of use within the cognitive science community; however, few studies have sought to combine the use of these technologies to support computational studies into embodied, extended, situated and distributed (EESD) cognition. Here, we explore the extent to which the ACT-R cognitive architecture and the Unity game engine can be used for these purposes. A range of issues are discussed including the respective responsibilities that the cognitive architecture and game engine have for the implementation of specific processes, the extent to which the representational and computational capabilities of cognitive architectures are suited to the modeling of EESD cognitive processes, and the extent to which the kind of embodiment seen in the case of so-called embodied virtual agents resembles that seen in the case of real-world bio-cognitive systems. These issues are likely to inform the focus of future research efforts concerning the integrative use of virtual environments and cognitive architectures for the computational modeling and simulation of EESD cognitive processes.
© 2015 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of AHFE Conference.

Keywords: virtual environment; cognitive architecture; embodied cognition; distributed cognition; situated cognition

1. Introduction

In the attempt to develop computational models of cognitive processes, researchers have tended to adopt conceptions of cognition that are both individualistic and heavily neurocentric. Cognitive processes are thus seen as the properties of individual human agents, and the materialistic underpinning of such processes is seen as firmly

* Corresponding author. Tel.: +44(0)23-8059-4474; fax: +44(0)23-8059-2783.
E-mail address: ps02v@ecs.soton.ac.uk

2351-9789 © 2015 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of AHFE Conference.

rooted in the operations of the biological brain. Challenging this view (or at least tempering it) is research that goes under the headings of embodied [1], extended [2], situated [3] and distributed [4] (EESD) cognition. Research in these areas tends to emphasize the role that material embodiment, environmental embedding and social interaction play in the molding (and perhaps even the mechanistic realization) of cognitive processes. Given the recent popularity of EESD approaches within the cognitive scientific community, there is likely to be considerable interest in cognitive computational models that approach cognition from these perspectives. Such models should aim to represent at least some of the forces and factors that are believed to shape the profile of cognition in ecologically-realistic situations. In particular, they should aim to model the closed-loop interactions that agents often have with their physical and social environments, capturing the way that agents (at least sometimes) actively structure or manipulate their environments in order to simplify, sequence or otherwise support cognitive processes. In addition to their role in evaluating the ideas and assumptions of EESD cognitive science, such models are also likely to be useful in a human factors context, enabling us to evaluate the potential impact of social and technological interventions on cognitive capabilities at both an individual and collective (e.g., team) level. Of particular interest, in this respect, is the attempt to understand the transformational changes wrought by emerging digital technologies on our cognitive and epistemic profiles. The advent of Web-based technologies, for example, presents new opportunities for the exploitation of cognitively- and epistemically-significant information resources, as well as new opportunities for social interaction and engagement. While such technologies have sometimes been denigrated on the grounds that they undermine our ability to think, read and remember [5], others have adopted a more positive view, suggesting that the new technologies form part of an emerging cognitive ecology [see 6] whose cognitive impacts are best understood through the conceptual lens of EESD cognitive science [7, 8].

Approaching cognitive modeling from an EESD perspective requires technologies that can adequately represent at least some of the agent-world and agent-agent interactions that occur in real-world situations. In this respect, the use of virtual environments, such as those built using contemporary game engines, may be of significant value. Virtual environments have been used in the context of artificial intelligence and artificial life research to simulate the behavior of embodied and environmentally-embedded agents [9]. They have also featured as part of research into intelligent virtual agents where issues of inter-agent interaction are the primary focus of interest [10]. Importantly, the use of virtual environments enables us to simulate at least some of the agent-world interactions that occur in cases of real-world cognition: virtual agents can thus be made to act on their environments according to the constraints of a virtual body, and every action implemented by a virtual agent results in a change in the environment, or the agent's position within it. The increasing sophistication of 3D modeling, graphic rendering, character animation and physics simulation technologies adds to the appeal of using virtual environments for computational studies into EESD cognition. Most notably, technical advances in the games industry tend to be driven by the quest for ever-greater levels of realism in virtual environments. This obviously expands the kind of opportunities that researchers have to enhance the richness and fidelity of computer simulations.

In addition to the use of virtual environments, it is important to consider the kind of systems that are used to control agent behavior. These could assume a variety of forms, such as state machines, neural networks and production systems. Of particular interest, in the current context, are a class of systems that have been referred to as cognitive architectures. These are computational frameworks that support the modeling and simulation of cognitive processes, often by providing a set of prefabricated representational and computational elements that can be adapted to specific task contexts. Two of the most well-known cognitive architectures – Soar [11] and ACT-R [12] – have a long history of use within the cognitive science community. They have been used to model many aspects of human cognition, and they have also been used to implement the cognitive capabilities of a variety of intelligent systems, such as real-world robots. In the context of research using virtual environments, there is considerable interest in using cognitive architectures for the purposes of virtual character design [13]. Fuelling this interest is the possibility that cognitive architectures could be used to improve the behavioral sophistication of non-player characters in video games. There is also the possibility of using cognitive architectures to support the modeling of human game-player responses and thus improve the level of support that game developers have for the automated testing of video games.

Although there are a number of instances in the scientific literature where cognitive architectures have been combined with virtual environments [10, 14], few studies have specifically sought to apply the two technologies to investigate issues in EESD cognition. One of the shortcomings of the scientific literature in this area concerns the availability of generic and reusable integration solutions: few studies have, to date, sought to articulate a standard

approach to the integration of cognitive architectures with virtual environments. A second shortcoming relates to an awareness of the technical challenges and conceptual issues associated with the use of cognitive architectures and virtual environments in EESD cognitive science research. In particular, a program of scientific research into the computational modeling of EESD cognition requires careful attention to the mechanisms used to situate cognitive agents within virtual environments. This throws up a range of issues that are seldom given adequate attention in the scientific and engineering literature. For example, when it comes to the implementation of perceptuo-motor capabilities, what level of control should the cognitive architecture have for the implementation of agent actions? Should the movements of agents be subject to detailed control by the cognitive architecture, or is it sufficient for the cognitive architecture to simply emit signals that trigger the execution of predefined animation clips? When it comes to perception, should perceptual processing assume the form of operations over low-level sensory data obtained from simulated sensors (e.g., the raw output of a scene camera), or is it sufficient to simply assert the existence of high-level perceptual information based on the kind of access the virtual environment provides to the geometric and photometric properties of objects within an agent's field of view. Many of these issues concern the balance of responsibility that the cognitive architecture and game engine have for the implementation of specific processes, as well as the level of functional abstraction that should be adopted as part of the design and development process.

The aim of the current paper is to raise community awareness of at least some of the issues surrounding the use of cognitive architectures and virtual environments in EESD cognitive science research. We first describe an approach to integrating a particular cognitive architecture (i.e., ACT-R) with a particular game engine (i.e., Unity). This integration effort has yielded a framework – the ACT-R Unity Interface (UI) Framework – which can be used to integrate multiple ACT-R models (each representing a distinct cognitive agent) into Unity-based virtual environments [15]. The functionality of the ACT-R UI Framework is described, albeit briefly, in Section **Error! Reference source not found.**. Subsequent sections provide an overview of issues arising from our experience of using the ACT-R UI Framework to undertake studies into EESD cognition. In particular, we focus on issues relating to perceptual processing, motor control and cognitive processing. Although these issues have surfaced in a particular technical context (one involving ACT-R and Unity), they are likely to be relevant to any program of research that seeks to combine the use of cognitive architectures and virtual environments for the purpose of undertaking computational studies into EESD cognition.

2. ACT-R/Unity Integration

The ACT-R UI Framework is an integration solution that enables multiple ACT-R models to inter-operate with applications built on top of the Unity game engine [15]. The solution is grounded in the use of networking technologies, exploiting the native support that Unity provides for TCP/IP socket-based communications. It also relies on the use of the JavaScript Object Notation (JSON) data interchange format to support platform-neutral and application-agnostic modes of information exchange. The result is a generic and reusable integration framework that provides the foundation for cognitive science simulations involving the use of the ACT-R architecture and the Unity game engine. The use of a network-based approach to integration means that ACT-R models need not be hosted on the same machine that runs the Unity environment. This helps to address some of the performance issues that inevitably arise in the context of complex simulations.

Previous studies involving the integrative use of cognitive architectures and virtual environments have tended to focus on one particular game engine, namely the Unreal Tournament Engine. The decision to focus on the Unity game engine in the context of the current development effort is based on a number of factors. These include the popularity of the game engine within the developer community, the level of support provided to Unity developers, the availability of community-supplied game assets (models, animations, textures, etc.), and the extensive support that Unity provides for multiplatform deployment. ACT-R was chosen as the preferred cognitive architecture based on its popularity and its grounding in the experimental psychology literature. In addition, the design of ACT-R is inspired by the functional organization of cognitive processes in the human brain: specific modules in ACT-R are deemed to be functionally-equivalent to the processes carried out by specific brain regions (e.g., ACT-R's `goal` module maps onto the anterior cingulate cortex), and this enables cognitive modelers to make explicit predictions concerning the recruitment of specific neuroanatomical structures at particular junctures in a cognitive process.

Using the ACT-R UI Framework, we have developed a range of simulations in which one or more ACT-R models control the behavior of virtual characters that are embedded in a virtual environment [15, 16]. For the most part, these simulations have focused on the exploration of a simple maze environment by non-humanoid robotic characters. In a more recent simulation, codenamed THESEUS, we have combined the use of ACT-R models with humanoid character models in order to study the exploration of a more complex spatial environment. Unlike previous studies, THESEUS emphasizes a collaborative approach to exploration: multiple ACT-R agents (four in total) are required to coordinate their exploratory efforts and create a shared cognitive representation (a collective cognitive map) of the spatial environment as quickly as possible. As with previous simulations using the ACT-R UI Framework, agents in the virtual environment orient themselves by integrating a variety of forms of sensory information, and they are also able to respond to motor-related instructions from ACT-R in order to effect specific movements. In contrast to previous simulations, however, agents in the THESEUS simulation are able to engage in more advanced forms of sensorimotor processing: visual information processing, for example, relies on the use of image processing algorithms, while motor control processes combine the use of character animation clips with finite state machines.

3. Perceptual Processing

The behavior of ACT-R-controlled agents in the THESEUS simulation is influenced by a variety of kinds of sensory information. In fact, the virtual characters are equipped with 'sensors' that enable them to process visual, auditory, kinesthetic, directional and tactile information. These sensory processing capabilities are implemented in a variety of ways; however, they all function to yield information about the local (virtual) environment of an agent, and they all contribute to an agent's ability to recognize specific locations and make navigation-related decisions.

Perhaps the most complex form of processing occurs in the context of the visual modality. ACT-R agents in the THESEUS simulation are required to detect the presence of colored shapes and numbers within the visual field at particular locations in the environment. These objects function as visual landmarks, enabling agents to recognize specific locations and orient themselves within the environment. Each agent has a Unity `Camera` component that provides a view of the scene from the agent's perspective. During the course of the simulation, the `Camera` renders to a `RenderTexture` asset, which is a specialized form of 2D image asset. The pixel data associated with this asset is analyzed using a combination of image processing techniques in order detect target features (i.e., colors and shapes) within the character's visual field. A colored shape, for example, is detected using a combination of color filters and blob processing algorithms.

For the most part, the processing of sensory information occurs in the context of the Unity environment and the outputs are periodically posted to ACT-R as part of a 'sensor processing cycle'. By default, this cycle executes at a rate of 2Hz, although the rate can be increased or decreased as desired. Perceptual information is ultimately represented in the form of an ACT-R chunk that is automatically presented to an ACT-R model via a custom ACT-R module, called the `sensorimotor` module. One important thing to note here is that all the perception-related information is integrated into a single chunk rather than being split according to perceptual modality and presented via modality-specific ACT-R modules and buffers. One of the factors motivating this design choice concerns the limited support that the core ACT-R architecture provides for the representation of perceptual information (by default ACT-R only supports the representation of visual and auditory information). This derives from the fact that ACT-R is primarily intended for situations in which a human subject is sat in front of a computer display screen and participates in conventional cognitive psychology experiments. This makes the core architecture ill-suited for the kind of simulations that are likely to be encountered in the context of EESD computer simulations. In particular, computational studies of embodied cognition are likely to require modules that can process a variety of forms of interoceptive information, such as kinesthetic, tactile and vestibular information. These modules will obviously need to be the focus of future development efforts. The core ACT-R architecture *does* provide support for the representation of visual and auditory information via the `vision` and `audio` modules; however, the functionality of these modules may need to be extended to support studies with the kind of virtual environments that can be engineered using contemporary game engines.

Aside from the addition/extension of perceptual modules, attention needs to paid to the structure of the sensory-information that is passed to ACT-R from the Unity environment, as well as the location of sensory processing

routines. In the case of THESEUS, we opted to perform all low-level sensory processing within the Unity environment, and we limited the content of network communications to high-level symbolic representations of perceptual features. In the case of visual processing, for example, a complex array of image pixel data was reduced to a symbolic representation indicating the presence of colored shapes within the agent's field of view. This strategy was selected for performance reasons. In particular, we aimed to reduce the amount of information communicated over the network, and we also aimed to reduce the perceptual processing burden placed on ACT-R models. There is, of course, a tradeoff here between the complexity of the sensory processing that needs to be performed, the number of agents that feature in a simulation, the frequency of execution of the sensor processing cycle, and the extent to which simulation capabilities are distributed across multiple machines.

One way of reducing the computational overhead associated with sensory processing is to eliminate the need for the low-level processing of sensor information and simply rely on the information provided by the virtual environment. The approach to the processing of visual information in the context of the THESEUS simulation resembles the approach adopted by some artificial life researchers who have resorted to a combination of simulated retinas and computer vision algorithms to equip artificial animals with biomimetic visuo-motor capabilities [17]. An alternative approach, here, is to rely on the information that is already available in the virtual environment and use this to predict the outcome of the low-level sensory processing routines. Thus, rather than resort to image processing techniques to detect colored shapes in the case of the THESEUS simulation, an alternative approach is to simply ensure that the game object hosting the relevant texture asset is located within the view frustum of the virtual agent's `Camera` component. Providing this condition is met, it is reasonable to assume that the agent can 'see' the shape in question without the need for further processing. Something akin to this approach is seen in the case of auditory processing for the THESEUS simulation. Thus, rather than resort to the processing of low-level representations of audio signals, the volume of auditory stimuli can be determined based on the relative locations of ACT-R agents and `AudioSource` objects (these are responsible for playing audio files at particular locations in the environment). Based on this information it is possible to predict whether an agent can hear a particular auditory stimulus (by determining whether the volume exceeds a cutoff threshold) and what the frequency of the stimulus is (based on the particular audio file played by the `AudioSource`). Clearly, the kind of information that is available in a virtual environment can be put to good use in terms of simplifying perceptual processing routines, and this is one of the obvious advantages of using virtual environments for simulation purposes. There may, however, be situations where the use of low-level processing techniques is still important in EESD cognitive science research – the zoomimetic simulations encountered in the case of artificial life experiments [e.g., 17] are a case in point.

4. Motor Control

The ACT-R UI Framework does more than support the communication of sensory information to ACT-R models, it also supports the processing of commands related to an agent's behaviour within the virtual environment. In the THESEUS simulation, for example, ACT-R models control the movements of characters by issuing simple motor commands, such as 'turn left 90 degrees' or 'move forward'. These are represented in the ACT-R model as chunks that are created following a request to a specialized motor buffer (which forms part of the aforementioned `sensorimotor` module). Essentially, ACT-R models control their characters by creating chunks to represent specific actions based on available sensory information, as well as the contents of declarative memory. These chunks are subsequently serialized as JSON messages and posted to the Unity environment. Once received by Unity, the messages are used to trigger the execution of predefined animation clips for a particular character (e.g., an instruction to move forward will trigger the execution of a walking animation that moves the character forward). In general, the receipt of motor commands affects the state of parameters associated with a finite state machine (implemented using Unity's Mecanim system). This determines which animation clip is active at any given time.

One of the issues that arises in the context of motor control mechanisms has already been encountered. It concerns the suitability of existing ACT-R modules for simulations that specifically target EESD cognitive processes. Thus although ACT-R already comes with a module that is intended to handle motor information (i.e., the `motor` module), this module seems ill-suited for the kinds of motor control processes demanded by at least some forms of EESD cognitive simulation. In particular, the ACT-R `motor` module is designed for situations in which

motor movements are directed to a computer keyboard and mouse device. As was the case with perceptual processing, additional modules may need to be developed to adapt ACT-R for EESD cognitive simulations.

A second issue that arises in a motor control context concerns the complexity of the commands created by ACT-R models. In the case of perceptual processing, we saw that there was a tension between the way in which high-level perceptual abstractions were computed and where these computations were actually located (either at the level of the Unity environment or at the level of individual ACT-R models). A similar concern arises in the case of motor control. Firstly, there is the issue of the structural complexity of the messages issued by an ACT-R model. In the context of the THESEUS simulation, we have assumed that ACT-R motor commands code for specific actions or action sequences that are implemented as animation clips. Clearly, however, this is only one way in which motor control processes could operate. A more complex approach would be to abandon the use of predefined animation clips and directly control the individual movements of a virtual character's body parts. The advantage of this approach is that it supports the adaptive fine tuning of character behaviour in response to specific situations and sensory contingencies. This is likely to be the preferred approach for computer simulations of EESD cognitive processes, although the need for detailed motor control mechanisms should ultimately be judged relative to the aims of a particular simulation. In the case of the THESEUS simulation, for example, it was not clear that fine-grained motor control solutions added anything to the ability of agents to coordinate their exploration-related activities.

Even in situations where complex motor control schemes do seem to be warranted, there remains the issue of how to go about implementing these schemes. One possibility is that all relevant information is computed by an ACT-R model, and this information is then used to update the position and rotation of the relevant body part transforms of the virtual character. The problem with this approach is that it places an excessive computational burden on the cognitive architecture (and it also inflates the amount of information that needs to be exchanged with the virtual environment). A better solution, in our view, is for the cognitive architecture to simply decide what actions need to be performed and delegate the actual implementation of these actions to the virtual environment.

Whatever the nature of the motor control information communicated to the virtual environment, character movements ultimately entail the modification of body part transforms. Computational approaches to this problem typically involve the use of forward and inverse kinematic solutions. As an extension of such techniques, it might be worth considering some of the control mechanisms that have been studied in the context of the embodied cognition literature, particularly those that deal with issues of morphological computation and ecological control systems [2, 18].

A final point to be made about motor control relates to the notion of embodiment itself. We have suggested that virtual environments and cognitive architectures can be used to undertake studies into computational embodied cognition. However, different levels or grades of embodiment have been identified in the philosophy of mind literature. Clark [2] thus distinguishes between 'mere embodiment', 'basic embodiment' and 'profound embodiment'. He suggests that profound embodiment is primarily a feature of bio-cognitive systems, and that it is unlike the kind of embodiment seen in the case of virtual or real-world synthetic agents. The details of what makes something a profoundly embodied system as opposed to a merely or basically embodied system need not concern us here. What matters is simply the fact that it is possible to distinguish between different forms of embodiment. Attempts to investigate embodied cognition using computer simulation techniques need to be sensitive to these conceptual distinctions. In particular, it is important to recognize that simply referring to synthetic agents (either physical or virtual) as 'embodied' does not necessarily mean that they are exhibiting the *same* form of embodiment as that seen in the case of real-world biological agents.

5. Cognitive Processing

The aim of agents in the THESEUS simulation is to collaboratively explore a complex spatial environment in the most efficient manner. This requires agents to exhibit a range of spatial cognitive capabilities, such as recognizing previously visited locations and updating the structure of cognitive map representations (it should also be clear that agents are required to communicate information as they explore the environment; however, due to space limitations, we do not attempt to cover communication issues in the current paper).

The spatial cognitive capabilities of ACT-R agents are realized by ACT-R models that enable agents to form a topologically-structured cognitive map of the environment. Previous studies have shown that this approach can be

used to good effect in terms of supporting navigation-related decisions [15]. One of the potential problems with the approach, however, is that it can be seen to embrace precisely that kind of view of cognition that EESD cognitive simulations are supposed to challenge. In other words, by situating cognitive processing within the cognitive architecture, or by labeling whatever the cognitive architecture does as 'cognitive', we risk endorsing the internalistic and neurocentric prejudices that EESD approaches have attempted to undermine. In particular, there is a tendency to see the cognitive architecture as the locus of all the cognitive processing capabilities of an agent, and then see the virtual environment as simply a space for the expression of motor outputs and the receipt of sensory inputs. This kind of strict separation between perception, cognition and action is exactly the kind of view that has troubled EESD theorists, but it is a view that is difficult to resist when a cognitive architecture is used in conjunction with a game engine, especially when the two environments are already subject to a degree of logical/physical separation and functional specialization. The temptation is perhaps all the greater when the cognitive architecture in question is ACT-R. In particular, ACT-R is deemed by its proponents to provide a functional abstraction of brain-based processes [12], with core ACT-R modules being mapped to specific regions of the human brain. Although this neuroanatomical grounding has been highlighted as one of the reasons to choose ACT-R over other cognitive architectures (the current paper is no exception – see Section 2), there is an obvious concern that by electing to use the ACT-R architecture we are already committing ourselves to a subtle neurocentrism (and individualism) regarding the physical organization of cognitive processes.

It is not clear, at the present time, to what extent these sorts of issues undermine the use of cognitive architectures in EESD cognitive science research. Our own view is that ACT-R is sufficiently flexible to support its use in the modeling of EESD cognitive processes. As we have seen in the context of this paper, the addition of new modules can support forms of perceptuo-motor integration that are not possible with the core ACT-R architecture. In addition, we see no reason why the cognitive mechanisms targeted by ACT-R cannot form part of a larger information processing ensemble that transcends the boundaries of individual cognitive models. When it comes to memory processes, for example, it is possible to see social interactions as contributing to forms of collaborative recall and rehearsal, both of which arguably serve to enhance the mnemonic capabilities of a cognitive agent [19]. Since these mechanisms are believed to work in concert with brain-based modes of information recall and encoding, there is no reason to assume that ACT-R cannot be used to study these sort of socially-situated cognitive processes.

6. Conclusion

Our aim in this paper has been to discuss at least some of the issues that arise in the context of research that uses cognitive architectures and virtual environments for the purposes of research into EESD cognition. There are a number of factors that motivate the use of cognitive architectures and virtual environments for this sort of research. Firstly, cognitive architectures provide a framework that supports cognitive modeling and cognitive simulation. The use of these architectures in a wide variety of cognitive science and artificial intelligence contexts over the past few decades has yielded a wealth of insights and resources that should not be overlooked in the context of contemporary cognitive system engineering efforts. Secondly, virtual environments provide a means to simulate agent-world interactions in an increasingly realistic fashion. The current state-of-the-art in graphic rendering techniques, character animation and physics-based simulation makes these environments an attractive alternative to experiments that rely on real-world robotic platforms.

The attempt to combine cognitive architectures with virtual environments for the purposes of modeling EESD cognitive processes has been studied using a integration framework featuring the ACT-R architecture and the Unity game engine [15]. We have used this framework to develop simulations in which one or more virtual agents are tasked with the exploration of a spatial environment. In our latest simulation, codenamed THESEUS, multiple humanoid characters are required to collaboratively explore a complex spatial environment using a combination of cognitive, communicative and perceptuo-motor capabilities. While THESEUS is currently being used to undertake a number of experiments exploring the relationship between agent characteristics and collective performance, its development has also proved instrumental in terms of focusing attention on a range of issues that sit at the heart of scientific attempts to apply cognitive architectures and virtual environments to the modeling and simulation of EESD cognitive processes. A number of these issues have been discussed in the current paper. They include the

balance of responsibility that the cognitive architecture and game engine have for the implementation of specific processes, the extent to which the representational and computational capabilities of cognitive architectures are suited to the modeling of EESD cognitive processes, and the extent to which the kind of embodiment seen in the case of so-called embodied virtual agents resembles that seen in the case of real-world biological systems (recall the distinction between mere, basic and profound forms of embodiment). Of greatest concern, perhaps, is the extent to which a cognitive architecture, such as ACT-R, is suited to the whole enterprise of EESD cognitive simulation. There is a risk that cognitive architectures in general, and perhaps ACT-R in particular, both reflect and reinforce long-standing intuitions about the physical organization and mechanistic realization of cognitive processes: intuitions that EESD cognitive scientists have long attempted to undermine and overthrow. There is thus a question to be addressed in future theoretical and empirical work concerning the extent to which cognitive architectures are genuinely useful in furthering the interests of the EESD cognitive science research agenda. Our own view on this matter is that cognitive architectures can be adapted to support research into at least some forms of EESD cognition. We recognize, however, that further work is required to improve our understanding of the conceptual issues and technical challenges associated with this complex but scientifically-significant endeavor.

Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

[1] Shapiro, L. A. Embodied Cognition, Routledge, Abingdon, Oxfordshire, UK, 2011.
[2] Clark, A. Supersizing the Mind: Embodiment, Action, and Cognitive Extension, Oxford University Press, New York, New York, USA, 2008.
[3] Robbins, P., Aydede, M. (Eds.) The Cambridge Handbook of Situated Cognition, Cambridge University Press, New York, USA, 2009.
[4] Hutchins, E. Cognition in the Wild, MIT Press, Cambridge, Massachusetts, USA, 1995.
[5] Carr, N. The Shallows: How the Internet Is Changing the Way We Think, Read and Remember, Atlantic Books, London, UK, 2010.
[6] Hutchins, E. Topics in Cognitive Science 2 (2010) 705-715.
[7] Smart, P. R. Metaphilosophy 43 (2012) 426-445.
[8] Smart, P. R. in: Shapiro, L. A. (Eds.) The Routledge Handbook of Embodied Cognition, Routledge, New York, New York, USA, 2014.
[9] Luck, M., Aylett, R. Applied Artificial Intelligence 14 (2000) 3-32.
[10] Rickel, J., Johnson, L. W. in: Cassell, J., Sullivan, J., Prevost, S. (Eds.) Embodied Conversational Agents, MIT Press, Cambridge, Massachusetts, USA, 2000.
[11] Laird, J. E., Newell, A., Rosenbloom, P. S. Artificial Intelligence 33 (1987) 1-64.
[12] Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., Qin, Y. Psychological Review 111 (2004) 1036-1060.
[13] Turner, J. O., Nixon, M., Bernardet, U., DiPaola, S. (Eds.) Integrating Cognitive Architectures into Virtual Character Design, IGI Global, Hershey, Pennsylvania, USA, in press.
[14] Best, B. J., Lebiere, C. in: Sun, R. (Eds.) Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Interaction, Cambridge University Press, New York, New York, USA, 2006.
[15] Smart, P. R., Scutt, T., Sycara, K., Shadbolt, N. R. in: Turner, J. O., Nixon, M., Bernardet, U., DiPaola, S. (Eds.) Integrating Cognitive Architectures into Virtual Character Design, IGI Global, Hershey, Pennsylvania, USA, in press.
[16] Smart, P. R., Stone, P., Braines, D., Sycara, K., Shadbolt, N. R. 15th International Conference on Intelligent Virtual Agents, Delft, The Netherlands, in press.
[17] Terzopoulos, D., Rabie, T., Grzeszczuk, R. in: Langton, C. G., Shimohara, K. (Eds.) Artificial Life V: Proceedings of the 5th International Workshop on the Synthesis and Simulation of Living Systems, MIT Press, Nara, Japan, 1997, pp. 346-353.
[18] Pfeifer, R., Bongard, J. How the Body Shapes the Way We Think: A New View of Intelligence, MIT Press, Cambridge, Massachusetts, USA, 2007.
[19] Smart, P. R. 1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition, Maryland, USA, 2010.