Dialogue

One Developmental Cognitive Architecture to Rule Them All?

Matthias Rolf

Dept. of Computing and Communication Technologies
Oxford Brookes University, UK
mrolf@brookes.ac.uk

Lorijn Zaadnoordijk

Donders Institute for Brain, Cognition and Behaviour
Radboud University
Nijmegen, Netherlands
l.zaadnoordijk@donders.ru.nl

Johan Kwisthout

Donders Institute for Brain, Cognition and Behaviour
Radboud University
Nijmegen, Netherlands
j.kwisthout@donders.ru.nl

A (cognitive) architecture describes the structure of an intelligent agent's mind, which may include emergent or even purely reactive approaches. Classical cognitive architectures typically describe grown behavioral or reasoning skills and are typically not embodied and structurally static, which makes their transfer to developmental problems problematic (Vernon et al., 2007). There have been some efforts to dedicatedly create architecture of learning and development, e.g. (Morse et al., 2010, Bellas et al., 2010). Many studies in developmental science describe or investigate the interplay of action and perception, motivation, and other aspects in closed and often embodied loops, thereby inevitably describing architectural aspects, even though not comprehensive ones that can span entire skill sets.

There clearly is not any cognitive architecture or general structural description that could "rule" developmental science (psychology/robotics), yet. The real questions of this dialogue initiation are therefore what purpose a single standard model and architecture could serve, and in how far the process of searching for one could be useful along the way.

What is the purpose of a developmental cognitive architecture?

The answer likely depends on whether one specifically looks at the scientific understanding of (human) intelligence, or at the engineering capability to build intelligence (that is, besides generally providing a potentially common language for researchers). Architectures potentially do more for science than "this is how it could work" descriptions. Unlike purely behavioral or descriptive models (e.g. sheer statistics of behavior), architectures describe hidden structure that is meant to explain the "how" and that might be experimentally verified. At the engineering end we might, in fact, find ourselves developing toolkit like solutions that also practically aid the creation of a developing intelligence.

Within either science or engineering, where would we find benefits from striving for unifying architectures?

Complexity monster or shackle?

Architectures naturally aim to address more than a single skill or a single scenario. If any single skill is investigated at a time (which is the practical norm), using a whole architecture involves complexity that is not strictly necessary for the task at hand, potentially violating Occam's razor. On the other hand, it has been argued that architectures actually constrain (instead of inducing unnecessary complexity) by confining models to a fixed formal language (Jones et al., 2000).

What areas or research could currently benefit from architectural efforts without being over-constraint by such a fixed language? How can practically good scientific experiments be conducted with such architectures?

Acknowledgements

We would like to thank all participants of the Lorentz-NIAS workshop "Perspectives on Developmental Robotics" (May 2017), whose discussions shaped this dialogue initiation.

Morse, A. F., De Greeff, J., Belpeame, T., & Cangelosi, A. (2010). Epigenetic robotics architecture (ERA). IEEE Transactions on Autonomous Mental Development, 2(4), 325-339

Vernon, D., Metta, G., & Sandini, G. (2007). A survey of artificial cognitive systems: Implications for the autonomous development of mental capabilities in computational agents. IEEE transactions on evolutionary computation, 11(2), 151-180.

Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. Psychological Science, 11(2), 93-100.

Bellas, F., Faiña, A., Varela, G., & Duro, R. J. (2010). A cognitive developmental robotics architecture for lifelong learning by evolution in real robots. In Neural Networks (IJCNN), The 2010 International Joint Conference on (pp. 1-8). IEEE.

Cognitive Architectures Should Not Be Computer Architectures

Niels Taatgen

Artificial Intelligence,
University of Groningen,
Netherlands
n.a.taatgen@rug.nl

The goal of cognitive architectures, as defined by Newell (1990) and Anderson (1983), is to provide a single theory that can explain and predict all aspects of the human mind. An architecture is not just a theory, it is also a simulation platform to build models of human task performance. A simulation platform means that certain representations have to be chosen and primitives have to be defined that serve as the building blocks of human cognition. Therefore, current cognitive architectures operate at a certain level of abstraction: Newell's Soar has chosen a purely symbolic level of abstraction, whereas Anderson's ACT-R has symbolic representations augmented with subsymbolic parameters. This choice of representational level is what Newell calls "carving nature at its joints". Despite the poetic nature of this claim, I would like to put forward the idea that this is a mistake.

Cognitive architectures provide the innate capabilities of the mind, which means that anything that is learned is not part of the architecture, even though the learning mechanisms themselves are. This means, assuming cognitive architectures operate on a certain level of abstraction, that anything below the level of abstraction of the architecture is implementation, and anything above that level has to be explained by knowledge that is accumulated through the learning mechanisms of the architecture. This reflects how architectures are designed in computer science, where each level of abstraction (e.g., logical circuits, microprogramming, machine language, higher level programming) is self-contained and virtually independent of constraints of the lower level.

Although this approach has been very successful in modeling many aspects of cognition, it fails if the phenomena that it wants to model are too far removed from the symbolic (or subsymbolic) level. For example, mechanisms for perceptual and motor learning are considered to be part of the implementation, and are therefore not covered. On the other hand, understanding natural language instruction requires so much knowledge and skills (all learned) that modeling that process in a cognitive architecture becomes programming in an awkward programming language. Instead, we have to acknowledge that learning and processing happens at many different levels of abstraction, and that we need a cognitive architecture with multiple levels of abstraction to be able to "rule them all": more similar to levels of abstraction in physics, chemistry and biology.

A multi-level cognitive architecture (Taatgen, 2017) provides representations and learning mechanisms for different levels of abstraction, from the neural level to higher-level reasoning. Each level has its own representations, which are composed of units from the level below. The composition or learning processes differs by level. At the lower levels of abstraction, where processing is typically fast, learning is often slow, for example attenuation of cells in the visual cortex to particular line orientations in the visual field. Learning at that level is a form of unsupervised learning. At the highest level of abstraction people can interpret natural language instructions for a new task, translate these into an instantiation of the necessary skills, and carry out that new task right away. Therefore, learning is fast ("one-shot") but processing is slow relative to processing at the lowest level.

In between we probably need several levels of abstraction: one in which we learn new skills to carry out tasks, such as the ability to count, or interpret language. Below that is traditional level of cognitive architectures, where units of representation are single "thinking steps" of in the order of 100ms each. Although there are many proposals for many levels of abstraction, tying them all together into a stable system will still be a big puzzle, which will hopefully not require a dark overlord.

Anderson, J. R. (1983). The architecture of cognition. book, Cambridge, MA: Harvard university press.
Newell, A. (1990). Unified theories of cognition. book, Cambridge, MA: Harvard university press.
Taatgen, N. A. (2017). Cognitive architectures: innate or learned? In The 2017 AAAI Fall Symposium Series (pp. 476–480). Palo Alto, CA: AAAI Press.

One (Dynamic Field) Theory to Rule Them All

John P. Spencer

School of Psychology, University of East Anglia, Norwich, UK

j.spencer@uea.ac.uk

Background

In a 2015 book, my colleagues and I proposed a large-scale neural process model of executive function (EF) and word learning (Schöner, Spencer, & The DFT Research Group, 2015). The model was constructed by integrating smaller models of attention and visual exploration (Perone & Spencer, 2012; Ross-Sheehy, Schneegans, & Spencer, 2015), visual working memory (Johnson, Spencer, Luck, & Schöner, 2009; Simmering, 2016), object representation and binding (Johnson, Spencer, & Schöner, 2008), and a variety of word learning phenomena (Samuelson, Smith, Perry, & Spencer, 2011). Thus, this line of work demonstrated the scalability of dynamic field theory—to move from smaller scale models and integrate them into a larger scale neural architecture that explains key aspects of higher-level cognition. And in each case—at small and large scales—the theory has led to testable predictions about how infants and children develop (Buss & Spencer, 2014; Perone, Molitor, Buss, Spencer, & Samuelson, 2015).

Here, I reflect on where this work might go: are we systematically moving toward one (dynamic field) theory to rule them all? If that is the aim, then what challenges do I see at present and what benefits might there be for pursuing this future? I discuss each in turn.

Challenges of large-scale modeling

• The integrated word learning/EF model is really complicated: The large scale architecture we are currently working with has 24 coupled cortical fields representing the neural activation of roughly 24,000 neurons (i.e., about 1000 neurons per population). How do we communicate the details of a model like this to a non-expert? Our first presentation of this model was in a book. That was reasonable because we could build the model across chapters. In a journal format, it is not possible to review all of the 'smaller' models that were integrated into the larger scale architecture in detail. This presents a unique communication challenge.

• Fitting the data must be done by hand. We have tried multiple model-fitting approaches. The computational load is too great—years of computation would be required to search the parameter space, even using the most advanced approaches. Indeed, we are not yet confident that Markov Chain Monte Carlo (MCMC) methods would even converge with our smaller scale models—we are exploring that now. This means that there is an entire literature in mathematical psychology dealing with model-fitting and 'free' parameters that we have to step outside of. It is not that we disagree with the arguments there—they just are not applicable to large scale neural models.

Benefits of large-scale modeling

• The integrated model should have massive generalizability. Our word learning model is the same as our executive function model—this is striking. Few models of EF even mention word learning; the fact that we have developed a large-scale model in two literatures in parallel suggests a deep theoretical link between these domains. Moreover, because our large-scale model has been built on the back of smaller architectures, this means we have integrated a large set of phenomena ranging from early attention to dimensional category labelling.

• The integrated model gives us a glimpse of neural reality. We have developed ways to generate hemodynamic predictions from dynamic field models (Buss, Wifall, Hazeltine, & Spencer, 2014). Here, the neural complexity of the model can be an advantage. We often simplify neural functions in smaller scale models (often to create fewer 'free' parameters). When we fit a brain model, however, those details are useful—each piece of the architecture hopefully leaves a neural signature that can be detected in the brain. Interestingly, when we use a general linear modelling approach (Wijeakumar, Ambrose, Spencer, & Curtu, 2017), we don't have to create separate regressors for each function or event in the task. Instead, the model specifies the entire neural pattern through time as the task unfolds. This approach could be very powerful.

• Large scale neural models might help us understanding developmental cascades. Embedded in our large scale model is a potential cascade of developmental changes moving from early perceptual and motor systems to attention and working memory to higher-level word learning and EF. Embedding these different functionalities into a single system allows us to look for signatures of cascade effects—as one system wires itself up, how does that affect the other systems to which it is coupled? This could shed new light on how development constructs itself step-by-step.

Conclusions: Are we hobbits or wizards?

• In my view, cognitive and developmental science have a Tolkien problem. As in Tolkien's trilogy, there is an intrinsic fear about 'one theory to rule them all'. Often, the idea is dismissed out of hand—the idea that someone might propose a large-scale model of the brain that is accurate is deemed to be a pipe dream. But if you dig deeper, I think there is something threatening about the idea. Let's say our model of word learning and executive function does a good job of explaining behavioural and neural data with kids and adults. This would be threatening in that the theory does something no other theory currently does—and if we are good scientists, that should matter! Now, everyone working in these domains should be required to learn about that theory. This places constraints on future work in this area. That's often uncomfortable for scientists—it is much easier to ignore the theory; to live the life of a hobbit.

• Other people might want to rise above—to become wizards and master this new theory. That takes time and energy, and should be encouraged. Physics has chosen this route. Does that mean there can only be a small handful of wizards? Here it depends on your definition of a wizard. There are certainly only a handful of theoreticians who understand Einstein, string theory, and the like. But there are wizards in experimental physics who are absolutely central to modern progress in physics. It is just that the wizarding world has bifurcated into theoretical physics and experimental physics. We think this is healthy.

• Thus, in the near future, we think cognitive and developmental science will have to decide whether to pursue the route of the hobbit or the wizard. It will soon become unreasonable to think that scientists could master empirical work with children, computational modelling, neuroscience, and robotics—something has to give. And we think that large scale models of brain function might be the tipping point that forces a sea change. So don't fear the ring. Ultimately, science is about discovering truth. What we do with that knowledge is a completely different dialog.

Buss, A. T., & Spencer, J. P. (2014). The emergent executive: a dynamic field theory of the development of executive function. Monographs of the Society for Research in Child Development. http://doi.org/10.1002/mono.12096

Buss, A. T., Wifall, T., Hazeltine, E., & Spencer, J. P. (2014). Integrating the behavioral and neural dynamics of reponse selection in a dual-task paradigm: A dynamic neural field model of Dux et al. (2009). Journal of Cognitive Neuroscience, 26(2), 334–351.

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. Psychological Science, 20(5), 568–577. http://doi.org/10.1111/j.1467-9280.2009.02329.x

Johnson, J. S., Spencer, J. P., & Schöner, G. (2008). Moving to a higher ground: the dynamic field theory and the dynamics of visual cognition. New Ideas in Psychology, 26, 227–251.

Perone, S., Molitor, S., Buss, A. T., Spencer, J. P., & Samuelson, L. K. (2015). Enhancing the executive functions of 3-year-olds in the dimensional change card sort task. Child Development, 86, 812–827.

Perone, S., & Spencer, J. P. (2012). Autonomy in action: Linking the act of looking to memory formation in infancy in infancy via dynamic neural fields. Cognitive Science, 1–59.

Ross-Sheehy, S., Schneegans, S., & Spencer, J. P. (2015). The Infant Orienting With Attention Task: Assessing the Neural Basis of Spatial Attention in Infancy. Infancy, 20(5), 467–506. http://doi.org/10.1111/infa.12087

Samuelson, L. K., Smith, L. B., Perry, L. K., & Spencer, J. P. (2011). Grounding word learning in space. PloS One, 6, E28095.

Schöner, G., Spencer, J. P., & The DFT Research Group. (2015). Dynamic Thinking: A Primer on Dynamic Field Theory. New York, NY: Oxford University Press.

Simmering, V. R. (2016). Working memory capacity in context: Modeling dynamic processes of behavior, memory, and development. Monographs of the Society for Research in Child Development, 81(3), 7–148.

Wijeakumar, S., Ambrose, J. P., Spencer, J. P., & Curtu, R. (2017). Model-based functional neuroimaging using dynamic neural fields: An integrative cognitive neuroscience approach. Journal of Mathematical Psychology, 76, 212–235. http://doi.org/10.1016/j.jmp.2016.11.002

Considering Simple Developmental Cognitive Architectures Before Complex Ones

Gary Jones

Department of Psychology, Nottingham Trent University, Nottingham, UK

gary.jones@ntu.ac.uk

What is the starting point of a developmental cognitive architecture? The minimum would presumably involve a mechanism that enabled one to learn things and another that constrained what could be learned. That way, we have something that is able to change its behaviour based on interacting with the world, while also "developing" i.e. not learning too much in a short space of time. Such a basic cognitive architecture is clearly not going to be able to explain the multi-faceted ways in which children develop so why not introduce complexity? At least two reasons spring to mind: first, how confident are we that the extra processes we include actually exist in the developing child? Second, hadn't we better find out how far we can get with the most basic architecture, because it might tell us which of the more complex processes aren't needed?

What can be taken as a given is that the child experiences the world and learns something from that experience. What does this experience involve? Unfortunately this is the magic question yet we are beginning to be able to answer it. Work by Linda Smith and colleagues (Jayaraman, Fausey & Smith, 2015; Yu & Smith, 2013) begins to illustrate infant experience of the visual world while linguistic experience in infancy can be estimated from large-scale transcripts and videos of mother-child interactions.

My own work involves the latter case of language where one estimates the child's experience based on the large-scale maternal speech that they hear. It would seem that gradual associative learning together with some constraint on the information processed from linguistic input can tell us a great deal about the language learning process and what may be involved. For example, measures of what can be referred to as verbal short-term memory appear to assess the child's current linguistic knowledge, be it sublexical information (where tests of nonword repetition can be simulated on the basis of the child's current linguistic knowledge, Jones, 2016) or lexical information (where tests of digit span can be explained from exposure to seemingly random digit sequences that appear in natural language, Jones & Macken, 2015). One even sees effects of syntax on the basis of associative learning (Kidd, 2012). Given that language is often perceived to be an indicator of intelligence, it would seem that one can get reasonably far in explaining linguistic phenomena purely from a straightforward associative learner and a good estimate of linguistic input.

Clearly, such a simplistic view is not going to account for all of child development, and that is where adding complexity comes in. But it is important to have a good idea of the child's experience of the particular tasks being examined because it seems that a lot of what appears as 'complexity' may reflect the environmental stimulus. It seems to me then that the role of the cognitive architecture is to encapsulate what the child learns from experience while also capturing higher-level cognition that can use learned knowledge to create new knowledge. In addition, this needs to go across numerous domains. That, I believe, is the challenge for a developmental cognitive architecture and whether we are yet at the stage where one can achieve this appears debatable.

Jayaraman, S., Fausey, C. M., & Smith, L. B. (2015). The faces in infant-perspective scenes change over the first year of life. PLoS ONE, 10: e0123780.
Jones, G. (2016). The influence of children's exposure to language from two to six years: The case of nonword repetition. Cognition, 153, 79-88.
Jones, G. & Macken, B. (2015). Questioning short-term memory and its measurement: Why digit span measures long-term associative learning. Cognition, 144, 1-13.
Kidd, E. (2012). Implicit statistical learning is directly associated with the acquisition of syntax. Developmental Psychology, 48, 171.
Yu, C. & Smith, L. B. (2013). Joint attention without gaze following: Human infants and their parents coordinate visual attention to objects through eye- hand coordination. PLoS ONE, 8: e79659.

Yes, One Developmental Cognitive Architecture Is Necessary and Feasible

J. Gerard Wolff

CognitionResearch.org, Menai Bridge, UK

jgw@cognitionresearch.org

In their dialog initiation, Matthias Rolf, Lorijn Zaadnoordijk, and Johan Kwisthout ask:

> '… whether and how it would be useful both epistemologically and in practice to aim towards the development of a "standard integrated cognitive architecture", akin to "standard models" in physics. In particular, [the authors] ask this question in the context of understanding development in infants, and of building developmental architectures, thus addressing the issue of architectures that not only learn, but that are adaptive themselves.'

In brief, my answer is "yes", it is essential to aim for such an architecture, and to abandon the quest only when there is overwhelming evidence that it cannot be done. The main reasons are: 1) That, while the gathering of empirical evidence is an important part of any science, it is at least as important to try to develop parsimonious theories to make sense of empirical evidence; 2) In a 20-year programme of research, derived from earlier research on language learning (Wolff, 1988), I have developed a cognitive architecture which already has a lot to say about the nature of cognition, including learning, adaptation, and cognitive development. This research demonstrates what can be achieved, suggesting that it will indeed be possible to develop a "standard integrated cognitive architecture".

The SP theory of intelligence and its realisation in the SP computer model is a unique attempt to simplify and integrate observations and concepts across artificial intelligence, mainstream computing, mathematics, and human learning, perception, and cognition, with information compression as a unifying theme (See Wolff, (2006, 2013, 2016) and other papers on www.cognitionresearch.org/sp.htm).

A central idea in the SP system is the powerful concept of SP-multiple-alignment, borrowed and adapted from bioinformatics. This yields:

• Versatility in aspects of intelligence. The SP system has strengths in several aspects of intelligence including: 'unsupervised' learning (which has the potential to be the foundation of other kinds of learning); the analysis and production of natural language; pattern recognition that is robust in the face of errors in data; pattern recognition at multiple levels of abstraction; computer vision; best-match and semantic kinds of information retrieval; planning; problem solving; and:

  – Versatility in the representation of knowledge. The SP system has strengths in the representation of diverse kinds of knowledge including: the syntax of natural languages; class-inclusion hierarchies (with or without cross classification); part-whole hierarchies; discrimination networks and trees; if-then rules; entity-relationship structures; relational tuples; and concepts in mathematics, logic, and computing.

– Versatility in reasoning. Strengths of the SP system in reasoning include: one-step 'deductive' reasoning; chains of reasoning; abductive reasoning; reasoning with probabilistic networks and trees; reasoning with 'rules'; nonmonotonic reasoning and reasoning with default values; Bayesian reasoning with 'explaining away'; causal reasoning; reasoning that is not supported by evidence; the inheritance of attributes in class hierarchies; and inheritance of contexts in part-whole hierarchies. There is also potential for spatial reasoning, and for what-if reasoning.

• Seamless integration of diverse kinds of knowledge and diverse aspects of intelligence. Because diverse kinds of knowledge and diverse aspects of intelligence all flow from a single coherent and relatively simple source—the SP-multiple-alignment framework—there is clear potential for the SP system to provide seamless integration of diverse kinds of knowledge and diverse aspects of intelligence, in any combination. It appears that that kind of seamless integration is essential for human levels of fluidity, versatility and adaptability in intelligence.

There is more detail in Wolff (2017, Appendix B).

J. G. Wolff. Learning syntax and meanings through optimization and distributional analysis. In Y. Levy, I. M. Schlesinger, and M. D. S. Braine, editors, Categories and Processes in Language Acquisition, pages 179–215. Lawrence Erlbaum, Hillsdale, NJ, 1988.
J. G. Wolff. Unifying Computing and Cognition: the SP Theory and Its Applications. CognitionResearch.org, Menai Bridge, 2006. ISBNs: 0- 9550726-0-3 (ebook edition), 0-9550726-1-1 (print edition).
J. G. Wolff. The SP theory of intelligence: an overview. Information, 4(3):283–341, 2013. arXiv:1306.3888.
J. G. Wolff. The SP theory of intelligence: its distinctive features and advantages. IEEE Access, 4:216–246, 2016. arXiv:1508.04087.
J. G. Wolff. Software engineering and the SP theory of intelligence. Tech- nical report, CognitionResearch.org, 2017. Submitted for publication. arXiv:1708.06665.

Clément Moulin-Frier
Synthetic Perceptive, Emotive and Cognitive Systems (SPECS), Universitat Pompeu Fabra, Barcelona, Spain
clement.moulinfrier@gmail.com

Paul F.M.J. Verschure
Director of SPECS Lab (IBEC, BIST, ICREA), Barcelona, Spain
pverschure@ibecbarcelona.eu

Distributed Adaptive Control As an Integration Framework for Cognition

As pointed out by Matthias Rolf, Lorijin Zaadnoordijk and Johan Kwisthout, there is a gap between the description of the structure of an intelligent agent's mind, which may include emergent or even purely reactive approaches and classical cognitive architectures that describe advanced behavioral or reasoning skills. The latter are typically not embodied and structurally agnostic. This contrast reflects an old debate in Cognitive Science and Artificial Intelligence, where two opposing approaches have been advanced to explain how cognitive functions can arise. Top-down approaches rely on a priori symbolic representations of a task, which have to be recursively decomposed into simpler ones to be executed by the agent. These depend principally on methods from symbolic artificial intelligence, as, e.g., in Soar (Laird, Newell, & Rosenbloom, 1987). The alternative, bottom-up approaches instead implement behavior without relying on advanced knowledge representation and reasoning. This is typically the case in behavior-based robotics, where low-level sensory-motor control loops form the starting point of emergent behavioral complexity as Simon's "ant on the beach" example (Simon, 1969), Braitenberg's Vehicles (Braitenberg, 1986) and implemented in the Subsumption architecture (Brooks, 1986).

Interestingly, a machine-learning-oriented version of this old debate has emerged from recent advances in Artificial Intelligence. On the one hand, strong emphasis is placed on so-called Deep Learning frameworks, where large feed-forward networks are trained end to end with an extremely large amount of training data. On the other hand, a drastically different approach has also received considerable attention, arguing that Deep Learning is not able to solve key aspects of human cognition without having access to advanced prior knowledge (Lake, Ullman, Tenenbaum, & Gershman, 2017). This approach states that human cognition relies on causal models of the world built through combinatorial processes to rapidly acquire knowledge and generalize it to new tasks and situations. This solution, however, comes at a cost: the underlying algorithms require non-trivial a priori knowledge, and an assumption of such models is that learning should be grounded in intuitive theories of physics and psychology.

This illustrates that despite the recent advances, we still face the old debate between bottom-up and top-down models of cognition. It is, therefore, a major challenge to structure these heterogeneous aspects of cognition in one unified theory. The Distributed Adaptive Control (DAC) theory of the mind and brain (see Verschure, 2016, for a recent review) provides a principled framework for realizing this structuring and integration effort by grounding it into biology, neuroscience, and ecology. DAC proposes that cognition is based on the interaction of four interconnected control layers operating at different levels of abstraction (see Figure 1). The first level, the somatic layer, corresponds to the embodiment of the agent within its environment, with its sensors and actuators as well as the physiological needs (e.g. exploration or safety). Extending bottom-up approaches with drive reduction mechanisms, complex behavior is bootstrapped in DAC from the self-regulation of an agent's physiological needs when combined with reactive behaviors (the reactive layer). This reactive interaction with the environment drives the joint acquisition of both perceptual and action hierarchical representations modulated by value signals in the adaptive layer. These compressed (hierarchical) and informative (modulated by value) representations support the acquisition of causal models of the world for goal creation and planning at the fourth contextual layer, which comprises systems for episodic, procedural and working memory and an autobiographical memory supporting life-long learning. These high-level processes, in turn, modulate the activity at the lower levels via top-down pathways shaped by behavioral feedback, i.e. acting through the environment itself rather than through direct internal control signals. The control flow in DAC is therefore distributed, both from bottom-up and top-down interactions between layers, as well as from lateral information processing within each layer.

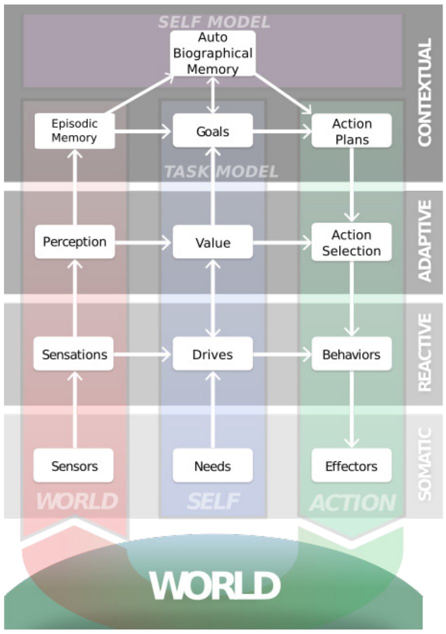

Figure 1: Abstract representation of The DAC architecture, see text and (Verschure, 2016) for details.

Braitenberg, V. (1986). Vehicles. Cambridge, MA, USA: MIT Press.
Brooks, R. (1986). A robust layered control system for a mobile robot. IEEE J. Robotics and Automation, RA-2, 14–23.
Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. Artificial Intelligence. https://doi.org/10.1016/0004-3702(87)90050-6
Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building Machines That Learn and Think Like People. Behavioral and Brain Sciences. https://doi.org/10.1017/S0140525X16001837
Moulin-Frier, C., Puigbò, J.-Y., Arsiwalla, X. D., Sanchez-Fibla, M., & Verschure, P. F. M. J. (2017). Embodied Artificial Intelligence through Distributed Adaptive Control: An Integrated Framework.
Simon, H. A. (1969). Sciences of the artificial. Cambridge: Cambridge, Ma.: MIT Press.
Verschure, P. F. M. J. (2016). Synthetic consciousness: the distributed adaptive control perspective. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 371(1701), 263–275. https://doi.org/10.1098/rstb.2015.0448
Verschure, P. F. M. J. P., & Althaus, P. (2003). A real-world rational agent: unifying old and new AI. Cognitive Science, 27(4), 561–590. Retrieved from http://dx.doi.org/10.1016/S0364-0213(03)00034-X

One Developmental Cognitive Architecture to Rule Them All? Responses to Commentaries

Matthias Rolf

Dept. of Computing and Communication Technologies
Oxford Brookes University, UK
mrolf@brookes.ac.uk

Lorijn Zaadnoordijk

Donders Institute for Brain, Cognition and Behaviour
Radboud University Nijmegen, Netherlands
l.zaadnoordijk@donders.ru.nl

Johan Kwisthout

Donders Institute for Brain, Cognition and Behaviour
Radboud University Nijmegen, Netherlands
j.kwisthout@donders.ru.nl

During the NIAS-Lorentz workshop on 'perspectives on developmental robotics', that we organized in May 2017, the participants identified the need for a *developmental* cognitive architecture as a means for modeling development. In our dialogue initiation we focused on two questions: What is the research goal of a developmental cognitive architecture? And is the complexity inherent in a cognitive architecture a hindrance or a blessing?

Jones responds negatively to our question whether it is feasible and timely to focus on developmental cognitive architectures. He suggests that it might be too early to aim for such a general architecture, and that it is probably better to focus first on simple (associative) models before postulating more elaborate mechanisms and processes, as we yet do not really understand what actually constitutes the experiences that infants use in their learning. *Spencer*, in contrast, is more positive towards a general architecture; his approach is to focus on the low-level neural architecture (dynamic field theory) and scale up to simulations with thousands of neurons. Here, explaining and fitting the model is a challenge, but might hint at the neural reality of developmental cascades.

Moulin-Frier and Verschure reflect on the historical debate between symbolic and connectionist AI, that recently re-emerged in the form of structure-rich (Bayesian) causal models on the one hand and model-free (deep-learning) approaches on the other hand. Their DAC (Distributed Adaptive Control) approach aims to connect higher cognitive causal models with somative, reactive, and adaptive layers. *Wolff* proposes the SP theory of intelligence (and corresponding computer model) which is rooted in information compression. It still remains to be seen how either architectural approaches allow for studying development, i.e., the effect of physical change on cognition. *Taatgen*, in addition, points at the difficult problem of uniting in a single multi-level architecture both low-level learning (often slow and model free) and higher-level (instructed and/or planned) learning that is relative fast.

Our preliminary conclusion of this discussion is that there is (obviously) no 'simple solution' towards the developmental problem, where 'development' transcends 'learning' in the sense that the architectural cognitive features *themselves* are changing during development. Important questions are still open: For example whether we can define development in and through architectures, and whether the architectural level is necessary to define "development", as opposed to "learning". A dedicated workshop on these questions—for example at a future ICDL/EpiRob conference—might be timely to address these vital research questions.