IEEE CIS

CDS NEWSLETTER

The Newsletter of the Technical Committee on Cognitive and Developmental Systems

Volume 14, number 1
Spring 2017

Developmental Robotics
Machine Intelligence
Neuroscience
Psychology

Editorial

Hierarchical Predictive Coding in Autonomous AI and Development

Pierre-Yves Oudeyer

Inria and Ensta
ParisTech, France
Editor

pierre-yves.oudeyer@inria.fr

The idea that the brain is pro-actively making predictions of the future at multiple levels of hierarchy has become a central topic to explain human intelligence and to design general artificial intelligence systems. In this issue, Jun Tani, who has been studying recurrent neural networks models of sensorimotor development for the last 20 years, introduces a dialog to ask whether hierarchical predictive coding enables a paradigm shift in development robotics and AI. Andy Clark, Doug Blank, James Marshall, Lisa Meeden, Stephane Doncieux, Giovanni Pezzulo, Martin Butz, Ezgi Kayhan, Johan Kwisthout and Karl Friston give their perspectives on this topic. In particular, they discuss the importance of various complementary mechanisms to predictive coding, which happen to be right now very actively researched in artificial intelligence: intrinsic motivation and curiosity, multi-goal learning, developmental stages (also called curriculum learning in machine learning), and the role of self-organization. They also underline several major challenges that need to be addressed for general artificial intelligence in autonomous robots, and that current research in deep learning fails to address: 1) the problem of the poverty of stimulus: autonomous robots, like humans, have access to only little data as they have to collect it themselves with severe time and space constraints; 2) the problem of information sampling: which experiments/observations to make to improve one's world model. Finally, they also discuss the issue of how these mechanisms arise in infants and participate to their development.

In a new dialog initiation, Matthias Rolf, Lorijn Zaadnoordijk and Johan Kwisthout extend this discussion by asking 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, they 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. Those of you interested in reacting to this dialog initiation are welcome to submit a response by November 30th, 2017. The length of each response must be between 600 and 800 words including references (contact pierre-yves.oudeyer@inria.fr).

Message From the CDS TC Chair

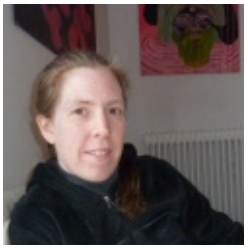

Kathryn Merrick

School of Engineering and Information Technology, University of New South Wales, Canberra, Australia

Chair of the Technical Committee on Cognitive and Developmental Systems

k.merrick@adfa.edu.au

As we reach the mid-point of 2017 it seems a good time to both reflect on our work this year and plan ahead.

So far this year we have discussed the addition of new goals for our technical committee and added two goals that consider machine recognition of cognitive characteristics in humans. Our technical committee goals now include:

- building machines capable of life-long adaptation and interaction with the physical and social world (existing goal)
- building machines that can model and recognise cognitive characteristics relevant to development in their human collaborators, and act accordingly to assist human activities (new goal)
- using machines as tools to better understand human and animal development and cognition (existing goal)
- using machines to support human learning and development (new goal).

In June this year, I attended the IEEE Computational Intelligence Society Technical Activities meeting in San Sebastian, Spain. News relevant to researchers in cognitive and developmental systems includes the formation of two new technical committees: the first tasked with identifying the research challenges of the future that span across the existing Computational Intelligence Society technical committees; and the second tasked with considering the social and ethical challenges that may accompany future advances in computational intelligence. Both of these technical committees will benefit from input from CDS researchers and interested members of our community are encouraged to contact the Technical Activities vice president.

I am also aware that members of our community have contributed to recent events in human-robot interaction (at HRI 2017 Vienna, Austria) and designing for curiosity (at CHI 2017).

After the Second Workshop on Evolution in Cognition to be held at GECCO in July at GECCO (http://gecco-2017.sigevo.org/index.html/Workshops), we look forward to ICDL-Epriob (http://www.icdl-epirob.org/) in Portugal in September as well as a host of workshops including and The Third International Workshop on Intrinsically Motivated Open-ended Learning (http://www.imol-conf.org/) in Rome in October.

It is exciting to see the ongoing efforts of members of our community and I look forward to further expanding our technical committee and task force members in the second part of this year as we seek to develop new task forces in the area of cognitive modelling.

Previous open-access editions of the newsletter can be found at: http://icdl-epirob.org/cdsnl
Web site of the IEEE TC on Cognitive and Developmental Systems: http://icdl-epirob.org/cdstc
IEEE ICDL-Epirob conference: http://www.icdl-epirob.org

Table of Contents

Editorial

Pierre-Yves Oudeyer
Hierarchical Predictive Coding in Autonomous AI and Development 1

Kathryn Merrick
Message From the CDS TC Chair 2

Dialogue

Jun Tani
Exploring Robotic Minds by Predictive Coding Principle 4

Andy Clark
Precisions, Slopes, and Representational Re-description 5

Douglas Blank, James Marshall, Lisa Meeden
A Developmental Robotics Manifesto 6

Stephane Doncieux
The Challenges and Pitfalls of Emergence in Developmental Robotics 7

Giovanni Pezzulo, Martin Butz
Predictive Processing in Developmental Robotics: Three Challenges 8

Ezgi Kayhan, Johan Kwisthout
Predictive Processing in Development 10

Karl J. Friston
Predictive Coding, Active Inference and Sentient Robots 11

Jun Tani
Response to Commentaries on "Exploring Robotic Minds by Predictive Coding Principle" 12

New Dialogue Initiation

Matthias Rolf, Lorijn Zaadnoordijk, Johan Kwisthout
One Developmental Cognitive Architecture to Rule Them All? 14

IEEE TCDS Table of Contents

Volume 9, Issue 1, March 2017 15

Volume 9, Issue 2, June 2017 17

Dialogue

Exploring Robotic Minds by Predictive Coding Principle

Jun Tani

Cognitive Neurorobotics Research Unit, Okinawa Institute of Science and Technology, Okinawa, Japan

tani1216jp@gmail.com

This dialogue discusses the topic of predictive coding in developmental robotics, highlighted from my newly published book (Tani, 2016).

The book proposes that the mind is comprised of emergent phenomena, which appear via intricate and often conflictive interactions between the top-down intention for acting on the external world and the bottom-up recognition of the resultant perceptual reality. It is presumed that the skills for generating complex actions as well as knowledge and concepts for representing the world naturally develop through entangled interactions between these two processes. This hypothesis has been evaluated by conducting nearly two decades of neurorobotics experiments using various recurrent neural network models based on the principle of predictive coding.

Is predictive coding a paradigm shift in developmental or learning robots?
The idea of sensory-motor mapping has dominated for a long period in the study of behavior-based robotics. However, robots based on just sensory-motor mapping schemes cannot achieve human-level thinking and acting because they should be much more proactive toward the future as well as reflective of the past. In predictive coding, the intention for an action is generated with prediction of the action's consequence. Likewise, the recognition of the actual consequence in the open environment reflects on the current intention by means of the error regression with the prediction error.

Is implementation by RNN using error back-propagation through time (BPTT) essential?
A notable advantage of RNN models is that they are differentiable. If the whole network is built on a set of modular RNNs—for instance one RNN for each sensory modality of a robot, one to learn multi-modality associations, and one for executive control—the whole also becomes differentiable. In this situation, a prediction error appearing at a particular spatio-temporal point in the perceptual flow can be distributed into the whole network retrospectively using error backpropagation through time. If the whole network activity is imposed with particular macroscopic constraints such as multiple timescales (for instance, different local subnetworks functioning at different timescales) or multiple spatial scales (for instance, different local connectivity distribution among subnetworks), some meaningful structures such as spatio-temporal hierarchy can self-organize as the result of end to end learning on this differentiable network. This type of development by means of the downward causation cannot be expected if the whole system is composed of patchy assemblies of different computational schemes.

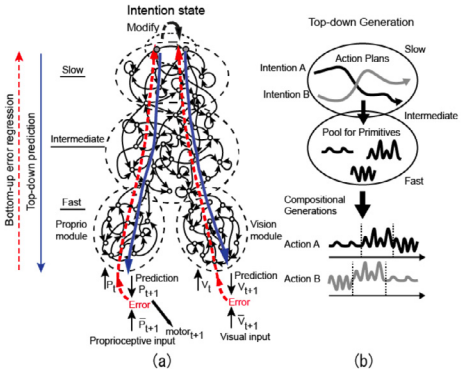

(a) Predictive coding implemented by multiple timescales RNN and (b) self-organization of functional hierarchy for action generation.

Is staged development essential?
It is fair to say that the recent success of deep learning is owed to a few researchers who have strongly believed for decades that the error backpropagation applied to differentiable networks is the most effective machine learning scheme. Now, we witness that convolutional neural networks, long-term and short-term memory as well as neural Turing machine built on this idea show significant learning performance by using millions of training data available on the internet.

However, this deep learning approach supported by usage of huge amount of data cannot be applied directly to developmental robots because they are constrained by the so-called poverty of stimulus, just like human infants. For both robots and infants the amount of experience in the real world is quite limited. Still at least for infants, skills and knowledge can be developed adequately with generalization even under such conditions. As pointed out by many others, it is expected that learning in one developmental stage can provide a "prior" for the one in the next stage thus drastically reducing freedom of learning. By this means, generalization with less amount of tutoring experience becomes possible. Based on this conception, developmental stage would proceed from physical embodiment levels to more symbolic ones. Tutoring should require a lengthy period wherein physical interactions between robots and tutors involve "scaffolding": guiding support provided by tutors that enables the bootstrapping of cognitive and social skills

required in the next stage.

Can robots attain free will and consciousness? For robots built on predictive coding, action and thoughts are generated as emergent phenomena when dense interactions between the top-down and the bottom-up process are developed in circular causality. It has been shown that chaos developed in the higher cognitive levels drives the spontaneous generation of the next intentional action, which will then be modified by means of minimizing the resultant conflictive error with the outer world (Tani, 2016). It is speculated that the spontaneity in generating the next intention by chaos might account for the unconscious generation of free will reported by Benjamin Libet whereas effortful process of minimizing the conflictive error does the same for the post-dictive conscious awareness of it. When robotic minds are built on such emergent phenomena, those robots could have subjective experiences, just like us.

Tani, J. (2016). Exploring Robotic Minds: Actions, Symbols, and Consciousness as Self-Organizing Dynamic Phenomena. Oxford University Press.

Precisions, Slopes, and Representational Re-description

Andy Clark

School of Philosophy, Psychology and Language Sciences, Edinburgh University, Scotland. UK

andy.clark@ed.ac.uk

Jun Tani's robotic explorations reveal the power and promise of hierarchical predictive coding as a bridge linking basic forms of sensorimotor engagement with the emergence of higher and higher forms of abstraction and control. Prediction-based learning yields representational forms, at higher processing levels, that act to summarize, compress, and control, activity at lower levels. Staged development with increasing flexibility results, since the process of level-by-level re-coding make lower-level knowledge available as 'chunks' for higher-levels to 'program' (re-purpose and re-organize).

These architectures give concrete computational form to 'representational re-description'—an endogenously-driven process in which sensory information is repeatedly re-coded ('re-described') in ways that support wider and more flexible kinds of use (Karmiloff-Smith (1992)—see also Clark and Karmiloff-Smith (1993), Cleeremans (2014), and Doncieux (2015). Prediction-driven hierarchical learning results in just such a process of staged development—one in which each higher level seeks to separate out causes and regularities that govern or explain patterns extracted at the level below. This whole process—just as Karmiloff-Smith suggested—is constrained by powerful endogenous forces favoring elegance and simplicity. This is because the learning routine (see Pezzulo, Rigoli, & Friston (2015)) favors the fewest-parameter model able to deliver (across a wide variety of contexts) apt action and choice. Complexity-reducing re-descriptions will thus continue to be sought even after behavioral success has been achieved. Such systems continually work on themselves to generate better and better (more powerful, less complex) models.

It is interesting to consider the potential (and potentially synergistic) influence of some potent additional elements prominent elsewhere in the literature on the 'predictive brain' (for a review, see Clark (2016)). One such is the variable 'precision-weighting' of the prediction error signal. Precision-weighting reflects the self-estimated reliability, for a given task in a given context, of specific prediction error signals. Increasing precision means increasing the post-synaptic gain or 'volume' on select prediction error signals, thus temporarily accentuating their influence. On a foggy day (to take a common example) this would enable the system to increase the influence of auditory information and to reduce the impact of incoming visual evidence, allowing a greater-than-usual role for top-down visual prediction.

Estimated precision also helps determine the nature and locus of control (Pezzulo et al. (2015)). 'Habitual' control occurs when reliable (precise) sensory prediction error is rapidly resolved at lower levels of the processing hierarchy. More reflective means of control occur when precise (salient, reliable) prediction error arises and is resolved at higher levels of processing. Variable precision-weighting would thus enable the selection of which 'representational re-description' should control behavior at a given moment. An important research horizon is to better understand forms of control (realized as top-down predictions) that entrain temporally extended sequences of inputs, so as to sustain long-term plans and projects of the kind we associate with distinct human agents. Distinctively human forms of conscious experience may emerge only when we ourselves turn up as 'control elements' in long-term predictive models governing our own future

actions (see our ongoing project at www.x-spect.org).

Another potent additional element may be the slope of prediction—error minimization itself. An emerging proposal is that an adaptively valuable strategy is to seek out situations in which the slope of minimization of prediction error is itself maximized (Oudeyer and Smith (2016), Joffily & Coricelli (2013), Miller and Clark (forthcoming)). This may help bring valence and emotion into the picture. The idea is that these track the rate at which prediction errors are being minimized relative to expectations. When error is minimized at a greater rate than expected, positive valence results. Such agents will actively seek out good learning situations—'sweet spot' learning environments, where they can significantly improve their predictive model of some salient aspect of the world.

Finally, perhaps it is not just the slope but the location (within the predictive hierarchy) of 'better-than-expected' prediction error minimization that matters. In a re-descriptive hierarchy, unexpectedly resolving prediction errors occurring at the higher levels will often signal a kind of 'falling into place' in which multiple tensions and inconsistencies are resolved at a single stroke—as when we suddenly succeed in seeing the hidden image in a 'magic eye' (autostereogram) display, or spot a mathematical derivation linking one body of results to another. Positive valence would then track not merely the rate, or the quantity, of prediction error minimization (relative to expectations) but also the quality.

This work was supported by ERC Advanced Grant XSPECT - DLV-692739

Clark, A. (2016) Surfing Uncertainty: Prediction, Action, and the Embodied Mind (Oxford University Press, NY)
Clark, A. and Karmiloff-Smith A. (1993) The Cognizers Innards Mind And Language 8: 4: 487-519
Cleeremans, A. (2014). Connecting conscious and unconscious processing. Cognitive Science, 38(6), 1286–1315.
Doncieux, S. (2015) "Representational redescription: the next challenge?" CDS TC Newsletter 12:
Joffily M., and Coricelli G. (2013) Emotional Valence and the Free-Energy Principle. PLoS Comput Biol 9(6): e1003094.
Karmiloff-Smith, A. (1992) Beyond Modularity Cambridge, MA: MIT Press
Millar, M. and Clark, A. (Forthcoming) Happily Entangled: Prediction, Emotion, and the Embodied Mind Synthese
Oudeyer, P-Y. and Smith. L. (2016) How Evolution may work through Curiosity-driven Developmental Process Topics in Cognitive Science, 1-11.
Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. Progress in Neurobiology, 134, 17–35.

A Developmental Robotics Manifesto

Douglas Blank, James Marshall, Lisa Meeden

Douglas Blank

Computer Science Dept.
Bryn Mawr College,
Pennsylvania, USA
dblank@cs.brynmawr.edu

We are largely in agreement with Tani's approach to developmental robotics as elucidated in this dialog and his recent book. The basic assumptions inherent in his approach, such as that agents are embodied in the world and that neural systems are capable of complex learning, are now established wisdom. Although this has been a relatively recent shift in AI and Cognitive Science, we consider these underlying assumptions to be a given and thus do not address them further. Here we expand on Tani's questions and offer a broader set of principles for guiding developmental robotics research.

The learning system should be capable of taking advantage of the world's structure and continuity.
For any specific learning algorithm, there will always be some types of problems that cannot be easily learned. Thus one should choose a learning method that best matches the domain. Our world is a highly structured and largely continuous environment through time and space. Because of these regularities, we should use those learning systems that can exploit the gradients between similar situations. Gradient descent procedures, such as backpropagation, are well suited for domains of this type.

In addition, an embodied agent moving through the real world typically needs to execute a sequence of actions in order to achieve its current goal or plan. Recurrent neural networks are able to construct representations of sequential memory using gradient descent procedures. Thus, Tani's approach of using BPTT applied to a recurrent neural system clearly meets our guidelines. However, we do not want to commit to any particular architecture or technique, even though we use very similar implementations in our own research.

Intrinsic motivation and the ability to make predictions and abstractions should be innate.
Developmental robotic agents are confronted by a dynamically changing and immensely complex world, which is only partially predictable. An agent's sensory systems provide a ceaseless flood of multimodal information about the surrounding environment. Initially, a robotic agent has no understanding of the relationship between its sensors and the world, nor how its actions affect its sensors. Intrinsic motivation imbues the robot with curiosity and a desire to learn, which

James Marshall

Computer Science Dept.
Sarah Lawrence College,
New York, USA
jmarshall@sarahlawrence.edu

Lisa Meeden

Computer Science Dept.
Swarthmore College,
Pennsylvania, USA
meeden@cs.swarthmore.edu

guides the robot in seeking out a low-level understanding of itself. At this stage, moment-to-moment sensory predictions about the outcomes of actions drive the robot to make low-level abstractions. Even though the system will never be able to perfectly predict the world, attempting to predict it will generate an error signal that gives the robot useful information. This information can be exploited in a variety of ways, for example to measure learning progress, to trigger attention, or to recognize sources of variability, which could include other agents.

Intelligence emerges through bottom-up and top-down interactions.
Once the robot has developed a bottom-up understanding of how its sensors and actions interact, and has created low-level abstractions based on this understanding, higher-order predictions and chunked abstractions can emerge. To be useful, meaning must be extracted from the sensory stream, in a continuous process that filters out enormous amounts of noisy, extraneous, redundant, or irrelevant information, depending on the situation at hand, and a coherent, abstract interpretation of the situation must be constructed.

This abstract interpretation can be bootstrapped from the knowledge gained at the lower levels. First, a higher level would learn to predict the sequence of states that occur at a lower level, leading to the development of its own higher-level abstractions. We note that these higher-level abstractions can be much more sophisticated than merely predicting one's own sensor readings. For example, a system could compare possible future action sequences in a type of counterfactual exploration. A higher level could then manipulate the sequence of states at the lower lever in order to achieve a chosen goal.

This process proceeds simultaneously on many levels of abstraction, and gradually, through development, becomes ever more efficient over the lifetime of the agent, as its knowledge of the world increases and it learns to better exploit that knowledge in pursuit of its goals.

We believe that continual, sustained, and interacting pressures will be necessary to create a system of lifelong learning. Over time, such an emergent, self-ratcheting system will have the potential to achieve robust levels of intelligent behavior in dynamic, unpredictable environments.

Conclusions
Returning to Tani's primary questions: we see predictive coding as essential; BPTT as a promising approach, though not essential; staged development as emergent; and foresee that by following the philosophy outlined above, robots could one day be conscious. We believe that Tani's work is a valuable contribution to better understanding the potential of developmental robotics, and, if combined with self-motivation and self-ratcheting pressures, is firmly in line with our manifesto.

The Challenges and Pitfalls of Emergence in Developmental Robotics

Stephane Doncieux

ISIR, UPMC,
Paris, France

doncieux@isir.upmc.fr

Cognitive development can be studied from different perspectives, may it be, for instance, dynamic systems or Bayesian learning (Newcombe, 2013). Connectionism is one of them and follows the empiricist perspective of Locke in which the baby starts its development with very little knowledge ('tabula rasa') and builds itself through his interaction with the environment. It leads to a notion that is fundamental for such approaches: emergence. After interaction, new features, may they be functions or representations, are expected to appear in the system whereas they were not built in it initially. Emergence is a structuring principle of the connectionist view of development. It frames the discussions and has a significant impact on the approaches that are considered and those that are avoided. Tani's work fits in the connectionnist view of cognition and his questions need to be placed in this particular context.

Predictive coding is a theory proposed to reconcile bottom-up and top-down approaches Tani (2016). In this framework, predictions are made to give a meaning to the complexity of the perceptions. The discrepancies between the two can drive a learning process towards a better matching. It is clearly not a paradigm shift in general, as many approaches rely on models to predict the effect of actions, model-based reinforcement learning, for instance Kaelbling et al. (1996). From a neuroscience point of view, the brain itself is known for long to be a prediction machine (Bubic et al. , 2010). If predictive coding is a paradigm shift, it is then from the perspective of the emergent paradigm of connectionism as used in robotics.

Staged development allows to progressively acquire information and bootstrap new capabilities when the required knowledge has been built. According to Piaget and to

the other constructivists, this is an important feature of human development. As Tani asks, is it essential for a robot to develop? It is hard to demonstrate, but this approach has an important methodological advantage: it allows to decompose the developmental process and study it piece by piece. It raises anyway some challenges for a purely connectionnist approach focused on emergence. How would the different stages emerge and structure themselves in a large neural network? These questions actually raise a critical challenge with respect to a connectionist view of development in robotics: the challenge of the methodology to follow to answer such questions. Contrary to what happens in nature, a scientist working in this field has motivations. They are of two different kinds: helping understand how the brain works or building robots with new capabilities. In both cases, researchers are expecting their robot, may it be real or simulated, to behave in a certain way after a certain amount of computation that is bounded by their computational resources and the time they spend on the study. These expectations are important and required to get a work worth to be published. It is hard to avoid as these expectations will drive the choice of the tools to use to analyse the system. These tools will drive the design, allow to fix bugs and compare robot's behavior to biological data for a neuroscience work or to alternative approaches for a robotics work. These methods make it hard to deal with emergence of unexpected features, as the researcher would not know what to measure or look at. Routine work in this field deals then with expected emergence, i.e. emergence of features that are explicitly looked for by the researcher. If the features are not known, the only possible method is serendipity. Beyond serendipity, can emergence of development be intentionnally studied? What are the intermediate steps? How to build a neural network that would make them emerge? What theory can drive this work? What methodology can be used? How to be sure that the intermediate steps chosen are the right ones? Does the choice of the intermediate steps not go against the principles of emergence?

If the recent progresses of deep learning show that neural networks are powerful machine learning tools, they are still used in a single and well identified learning process. Going one step further and developing a connectionnist approach to development requires either to, at least partially, abandon emergence and turn to hybrid approaches, in which neural networks are black box modules used in a modular, Fodorian architecture or to develop a methodology that would reconcile emergence with researchers' work. The question is then do we really want to keep emergence and, if the answer is yes, how?

Andreja Bubic, D. Yves Von Cramon, and Ricarda Schubotz. Prediction, cognition and the brain. Frontiers in Human Neuroscience, 4:25, 2010. ISSN 1662-5161. doi: 10.3389/fnhum.2010.00025. URL http://journal.frontiersin.org/article/10.3389/fnhum.2010.00025.
Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. Journal of artificial intelligence research, 4:237–285, 1996.
Nora S Newcombe. Cognitive development: changing views of cognitive change. Wiley Interdisciplinary Reviews: Cognitive Science, 4(5):479–491, 2013.
Jun Tani. Exploring robotic minds: actions, symbols, and consciousness as self-organizing dynamic phenomena. Oxford University Press, 2016.

Predictive Processing in Developmental Robotics: Three Challenges

Giovanni Pezzulo, Martin Butz

Predictive Processing (PP) and the closely related Free Energy Principle (FEP) foster an increasingly popular perspective on the mind, promising to integrate various theories from neuroscience, cognitive science, and philosophy (Butz et al., 2003; Butz, 2008; Clark, 2016; Friston, 2010; Friston et al., 2016; Pezzulo et al., 2015, 2008). In this respect, Tani's book is timely and intriguing: it reports the results of an ambitious research program, which applied a dynamical systems approach implemented in recurrent neural networks (RNNs) to robotics for 20 years. From our research we would like to raise three points that seem to be critical to succeed in the open-ended development of truly autonomous, artificial systems.

Balancing exploration and exploitation. FEP suggests that both epistemic drives (active information gathering) and goal achievement may stem from a unique imperative, i.e., reducing (anticipated) free energy. FEP was shown to enable balancing these generally competitive drives, for example, by noticing that in conditions of uncertainty it is better to first pursue epistemic drives (reducing uncertainty in the relevant task dimensions) before extrinsic (utilitarian) goals can be pursued (Butz and Kutter, 2017; Friston et al., 2015), although successful applications to complex scenarios are still pending. Tani's approaches either rely on chance induced by chaotic dynamics or on teacher-based demonstrations. Thus, the open challenge remains to build scalable, autonomous systems that are able to properly balance exploration and exploitation across development, possibly supported by (genetically)

Giovanni Pezzulo

Institute of Cognitive Sciences and Technologies, Rome, Italy

giovanni.pezzulo@istc.cnr.it

Martin Butz

Computer Science & Psychology, University of Tübingen, Tübingen, Germany

martin.butz@uni-tuebingen.de

pre-determined developmental pathways and tendencies towards curiosity and epistemic actions (Baldassarre and Mirolli, 2013; Butz and Kutter, 2017; Donnarumma et al., 2017; Oudeyer et al., 2007; Pezzulo et al., 2016; Schmidhuber, 1991).

Inductive biases for learning generative models. From a developmental perspective, a further open problem is how to guide the construction of increasingly more sophisticated, abstract generative models, such as object models or object concepts (e.g. a "container"), that build upon sensory and motor signals. At the moment, Tani's recurrent neural networks have a predetermined hierarchical structure, partially including different temporal resolutions; but mechanisms for inferring structure automatically during development would be desirable. One example method may be an event-segmentation bias, which can be based on lasting, significant changes in the active predictive encodings. This bias may be the key to foster progressive abstractions of generative models into suitable, behavior-oriented, hierarchical event taxonomies (Butz, 2016, 2017). Additionally, factorization approaches (which are also employed by Tani's ANNs with parametric biases) may allow the "splitting" of generative models into manageable, meaningful encodings (e.g. where, what, and when) (Butz, 2016, 2017; Verschure, et al., 2014). The involvement of inductive biases into dynamical ANNs or FEP-related systems seems to be essential to enable scalable system development in real-world, open-ended environments.

Exploiting embodiment. Tani's models are embodied in the sense that they were applied and developed in real robots. The robot bodies and the addressed tasks were selected to be compatible. Thus, the applicability of the chosen techniques in open-ended developmental system remains as a critical challenge. Seeing that a manifold of examples exists, showing that embodiment can significantly facilitate and bootstrap cognitive development, including inferential abstraction (Butz and Kutter, 2017), an important challenge for the future is to extend Tani's approach and FEP to combine "the best of two words", that is, embodied and hierarchical cognitive inference.

Thus, while Tani's work sets a milestone in the development of truly autonomous systems, there is still a long way to go. We believe that the integration of considerations of embodiment, inductive biases, and balancing exploration and exploitation within the general framework of PP-based robotics will be critical for success.

Baldassarre, G., Mirolli, M. (Eds.), 2013. Intrinsically Motivated Learning in Natural and Artificial Systems. Springer Berlin Heidelberg, Berlin, Heidelberg.
Butz, M.V., 2008. How and Why the Brain Lays the Foundations for a Conscious Self. Constr. Found. 4, 1–14.
Butz, M.V., 2016. Toward a Unified Sub-symbolic Computational Theory of Cognition. Front. Psychol. 7.
Butz, M.V., 2017. Which Structures Are Out There, in: Metzinger, T.K., Wiese, W. (Eds.), Philosophy and Predictive Processing. MIND Group, Frankfurt am Main. doi:10.15502/9783958573093
Butz, M.V., Kutter, E.F., 2017. How the Mind Comes into Being: An introduction to cognitive science from a functional and computational perspective. Oxford University Press.
Butz, M.V., Sigaud, O., Gerard, P. (Eds.), 2003. Anticipatory Behavior in Adaptive Learning Systems: Foundations, Theories, and Systems, Springer. ed.
Clark, A., 2016. Surfing Uncertainty: Prediction, Action, and the Embodied Mind. Oxford University Press.
Donnarumma, F., Costantini, M., Ambrosini, E., Friston, K., Pezzulo, G., 2017. Action perception as hypothesis testing. Cortex. doi:10.1016/j.cortex.2017.01.016
Friston, K., 2010. The free-energy principle: a unified brain theory? Nat Rev Neurosci 11, 127–138. doi:10.1038/nrn2787
Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., Pezzulo, G., 2016. Active Inference: A Process Theory. Neural Comput. 29, 1–49. doi:10.1162/NECO_a_00912
Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., Pezzulo, G., 2015. Active inference and epistemic value. Cogn. Neurosci. 0, 1–28. doi:10.1080/17588928.2015.1020053
Oudeyer, P.-Y., Kaplan, F., Hafner, V., 2007. Intrinsic Motivation Systems for Autonomous Mental Development. IEEE Trans. Evol. Comput. 11, 265–286.
Pezzulo, G., Butz, M.V., Castelfranchi, C., Falcone, R. (Eds.), 2008. The Challenge of Anticipation: A Unifying Framework for the Analysis and Design of Artificial Cognitive Systems, LNAI 5225. Springer.
Pezzulo, G., Rigoli, F., Friston, K.J., 2015. Active Inference, homeostatic regulation and adaptive behavioural control. Prog. Neurobiol.
Pezzulo, G., Vosgerau, G., Frith, U., Hamilton, A., Heyes, C., Iriki, A., Jörntell, H., König, P., Nagel, S., Oudeyer, P.-Y., Rupert, R., Tramacere, A., 2016. Acting up: An approach to the study of cognitive development 18.
Schmidhuber, J., 1991. Adaptive Confidence and Adaptive Curiosity (No. FKI-149-91). Institut fuer Informatik, Technische Universitaet Muenchen.
Verschure, P.F.M.J., Pennartz, C.M.A., Pezzulo, G., 2014. The why, what, where, when and how of goal-directed choice: neuronal and computational principles, Philos Trans R Soc Lond B Biol Sci, 369: 20130483. doi: 10.1098/rstb.2013.0483

Predictive Processing in Development

Ezgi Kayhan

Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany

kayhan@cbs.mpg.de

Johan Kwisthout

Donders Institute for Brain, Cognition and Behaviour
Radboud University Nijmegen, Netherlands

j.kwisthout@donders.ru.nl

A novel analogy is taking hold in theoretical neuroscience: "Prediction is to brain as digestion is to stomach." This analogy, provocative as it is, expresses the essence of what has become known as the Predictive Processing account (Clark, 2015). According to this account, the brain is essentially "a sophisticated hypothesis-testing mechanism, which is constantly involved in minimizing the error of its prediction of the sensory input it receives of the world" (Hohwy, 2013, p.1). Using generative internal models the brain predicts its own inputs in a cascading hierarchy of increasingly complex hypotheses about hidden states of the world. The part of the inputs that could not be correctly predicted (viz., the prediction error) is used to update the hypotheses to eventually maximize the accuracy of the internal models (Friston, 2010). Notwithstanding its empirical and theoretical successes for explaining the adult brain (Brown et al., 2011; Seth, 2013; van Pelt et al., 2016) the Predictive Processing account is lacking one key ingredient: A coherent and consistent explanation of how generative models that allow for making predictions are formed and improved in development.

Although some evidence point at the early predictive architecture of the human brain (Emberson et al., 2015; Kouider et al., 2015), there are still open issues when considering whether Predictive Processing account can explain development. How do infants use prediction errors to generate and refine models? How do prediction error minimization account for the innovative, creative part of learning: forming new concepts and associations and enriching existing models with contextual dependencies? Can individual differences in infant learning be explained in terms of different parameters or strategies in prediction error minimization?

Despite big questions, some concepts in the framework might be tailored to explain infant development. For example, although, mathematically, minimizing free energy is equivalent to maximizing the accuracy of the models (Friston et al., 2016), one would argue that the latter better describes infant behavior. Observing natural infant locomotion would simply speak for this argument. Infants around 12 to 19-months take 2367.6 steps and fall 17.4 times per hour (Adolph et al., 2012). One would wonder why infants would repeatedly try to take steps, as each try would presumably elicit prediction errors, perhaps, until they master the skill. However, if the behavior were driven by the goal of maximizing the accuracy of the internal models, this would potentially better explain what drives infants' behavior. Relatedly, exploration and curiosity, which are known to be crucial to infant learning and development (Oudeyer & Smith, 2016), might also be addressed by the aim of maximizing the accuracy of the world models.

Future directions

Providing theoretical, empirical, and computational evidence on whether and how the Predictive Processing framework could explain infant learning and development would pave the way to a novel and interdisciplinary research genre, drawing upon the joint experience of theoretical neuroscientists, developmental roboticists, and developmental researchers. Not only would such research inform developmental scientists to understand infant behavior and brain function, but also it will enrich the Predictive Processing framework to explain how generative models are developed, which is currently underspecified in the framework. Among others, these important questions are awaiting answers in the future.

Adolph, K. E., Cole, W. G., Komati, M., Garciaguirre, J. S., Badaly, D., Lingeman, J. M., Chan, G.L.Y. & Sotsky, R. B. (2012). How do you learn to walk? Thousands of steps and dozens of falls per day. Psychological Science, 23(11), 1387-1394. doi: 10.1177/0956797612446346.

Brown, H., Friston, K. J., & Bestmann, S. (2011). Active inference, attention, and motor preparation. Frontiers in Psychology, 2, 218. doi: 10.3389/fpsyg.2011.00218.

Clark, A. (2015). Surfing uncertainty. Oxford, UK: Oxford University Press.

Emberson, L.L., Richards, J.E., & Aslin, R.N. (2015). Top-down modulation in the infant brain: Learning-induced expectations rapidly affect the sensory cortex at 6 months. PNAS, 112(31), 9585–9590.

Friston, K. (2010). The free-energy principle: a unified brain theory?. Nature Reviews Neuroscience, 11(2), 127-138. doi:10.1038/nrn2787.

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2016). Active inference and learning. Neuroscience & Biobehavioral Reviews, 68, 862-879. doi: 10.1016/j.neubiorev.2016.06.022.

Hohwy, J. (2013). The predictive mind. Oxford, UK: Oxford University Press.

Kouider, S., Long, B., Le Stanc, L., Charron, S., Fievet, A. C., Barbosa, L. S., & Gelskov, S. V. (2015). Neural dynamics of prediction and surprise in infants. Nature Communications, 6. doi:10.1038/ncomms9537.

Oudeyer, P. Y., & Smith, L. B. (2016). How Evolution May Work Through Curiosity-Driven Developmental Process. Topics in Cognitive Science, 8(2), 492-502. doi: 10.1111/tops.12196.

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. Trends in Cognitive Sciences, 17(11), 656–663.

van Pelt, S., Heil, L., Kwisthout, J., Ondobaka, S., Van Rooij, I, and Bekkering, H. (2016). Beta- and gamma activity reflect predictive coding in the processing of causal events. Social Cognitive and Affective Neuroscience, 11(6), 973-980.

Predictive Coding, Active Inference and Sentient Robots

Karl J. Friston

The Wellcome Trust Centre for Neuroimaging Institute of Neurology, London, UK

k.friston@ucl.ac.uk

Jun Tani touches on many intriguing points. I will focus on three of my favourite; namely: *generative models of the future, hierarchical inference* and the *poverty of stimulus* challenge. Before expanding on his observations, I want to set the scene for this focus:

I fully concur with Jun that hierarchical inference under generative models—implicit in predictive coding—is the way forward in developmental neurorobotics, and perhaps generalised artificial intelligence. However, predictive coding, in and of itself, is only part of the story. One could argue that any recognition scheme that uses back propagation of prediction errors falls under the rubric of predictive coding (e.g., hierarchal or deep convolution networks). However, the recognition problem is almost trivial in relation to the problem faced by neurorobotics. The real problem is not how to recognise the causes of sensor data but *how to select the data that best discloses its causes*. In my world, this is referred to as active inference (Friston et al., 2015); namely, the bilateral use of action and perception to navigate an uncertain world. In short, active (hierarchical Bayesian) inference may entail predictive coding but not vice versa.

Predicting the future

If one puts action into the mix, a whole world of 'planning as inference' emerges (Attias, 2003; Botvinick and Toussaint, 2012; Mirza et al., 2016); which begs the question; "how do we do predictive coding of the future"? If one subscribes to generative modelling, the answer is clear: one has to have generative models of the future. This is becoming increasingly clear in theoretical neuroscience, where epistemic behaviour is a natural consequence of Bayesian inference, under prior beliefs about the consequences of action (Friston et al., 2015). Furthermore, several nice devices present themselves for use in robotics. For example, one can cast a policy selection as Bayesian model selection (based upon the marginal likelihoods of policies that treat future outcomes as hidden states). This has the fundamental advantage of covering epistemics and intrinsic motivation (Oudeyer and Kaplan, 2007); namely, behaving in a way that reduces uncertainty through sampling salient sensory cues—or engaging in novel behaviours to discover "what happens if I do this" (Schmidhuber, 2006). In brief, the minimisation of expected free energy or maximisation of expected model evidence leads naturally to self-organisation and self-evidencing (Hohwy, 2016). In short:

> *"Robots based on just sensorimotor mapping schemes cannot achieve human level thinking because they should be much more proactive towards the future."*

Temporal thickness and counterfactual depth

The second theme follows naturally from models that generate future outcomes—that necessarily entail deep or hierarchical structure. These models induce a separation of timescales in the ensuing recognition dynamics (Tani et al., 2004), which speaks to the temporal thickness or counterfactual depth of representations that drive epistemic behaviour (Seth, 2014). In virtue of the fact that all hierarchal inference involves belief propagation (i.e. variational message passing), it seems obvious to me that the use of a recurrent neural network is necessary—because the message passing required in belief propagation cannot, by definition, be reduced to:

> *"A system composed of patchy assemblies of different computational schemes."*

Big data or big ideas?

A generative model that entertains different hypotheses about unfolding dynamics also speaks to the "*poverty of stimulus*" problem. I wholeheartedly agree with Jun that current trends towards big data and deep learning are heading in the wrong direction. To simulate epistemic foraging in sentient robots, we need to understand how they make inferences to the best explanation through a process of abduction and active inference. In other words, how can the implicit hypotheses and models entertained by a robot make use of sparse—if carefully sampled—data. In neuroscience, this is akin to trying to understand the fundamental nature of insight and aha moments. If we can formalise and reproduce this in robots I suspect that the poverty of stimulus problem will be rapidly dissolved.

Attias, H., 2003. Planning by Probabilistic Inference. In: Proc. of the 9th Int. Workshop on Artificial Intelligence and Statistics. Vol., ed.^eds.
Botvinick, M., Toussaint, M., 2012. Planning as inference. Trends Cogn Sci. 16, 485-8.
Friston, K., et al., 2015. Active inference and epistemic value. Cogn Neurosci. 1-28.
Hohwy, J., 2016. The Self-Evidencing Brain. Noûs. 50, 259-285.
Mirza, M.B., et al., 2016. Scene Construction, Visual Foraging, and Active Inference. Front Comput Neurosci. 10, 56.
Oudeyer, P.-Y., Kaplan, F., 2007. What is intrinsic motivation? a typology of computational approaches. Frontiers in Neurorobotics. 1, 6.
Schmidhuber, J., 2006. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. Connection Science. 18, 173-187.
Seth, A., 2014. The cybernetic brain: from interoceptive inference to sensorimotor contingencies. In: MINDS project. Vol., ed.^eds. MINDS, Metzinger, T; Windt, JM.
Tani, J., Ito, M., Sugita, Y., 2004. Self-organization of distributedly represented multiple behavior schemata in a mirror system: reviews of robot experiments using RNNPB. Neural Networks. 17, 1273-1289.

Response to Commentaries on "Exploring Robotic Minds by Predictive Coding Principle"

Jun Tani

Cognitive Neurorobotics Research Unit, Okinawa Institute of Science and Technology, Okinawa, Japan

tani1216jp@gmail.com

Thanks to all for the inspiring dialog. My brief response will focus on two themes. One concerns the role of meta-priors for bridging the gap between deterministic and probabilistic processing within the predictive coding framework. The other concerns development in terms of co-emergent phenomena.

Friston focuses on solving the "poverty of stimulus" and related problems, noting that active (hierarchical Bayesian) inference may entail predictive coding, but not vice versa. Clark notes that precision-weighting reflects the self-estimated reliability of specific prediction error signals for a given task in a given context. These comments are exactly right. Prior predictive coding RNNs based on deterministic dynamic systems afforded neither active Bayesian inference nor precision estimation in prediction. We have been bridging deterministic and probabilistic predictive coding using variational Bayes RNNs (Murata et al., 2015; Ahmadi and Tani, 2017). Ahmadi and Tani (2017) proposes the variational Bayes predictive coding MTRNN (VBP-MTRNN) characterized by maximizing the lower bound (negative free energy) represented by a weighted sum of the regularization term (which becomes larger when the posterior distribution of the latent variable becomes closer to its prior (given as a normal distribution))—and the likelihood term (which becomes larger by minimizing the reconstruction error). Summarily, this weighting plays the role of a meta-prior determining the quality of learned structures, affecting the learning of fluctuated temporal patterns. Heavy weighting of the regularization term causes the development of stochastic dynamics imitating probabilistic processes observed in target patterns and also makes active inference less effective because error propagates only weakly. On the other hand, simulations show that heavy weighting of the likelihood term causes the development of deterministic chaos for imitating the randomness observed in target sequences, resulting in rote learning according to the strong top-down prior. It was found that generalization in learning can be maximized between these two extremes. Crucially, in this work we see that as predictive coding models have developed from 1st order prediction, 2nd order (precision prediction), and to 3rd order including the meta-prior discussed here. It is noted that, whatever higher-order the system seeks, the settings of priors or meta-priors determine behavior so long as the Bayesian framework is used.

The simulations above may afford insight into the mechanisms underlying autism spectrum disorders (ASD). Van de Cruys et al (2014) have suggested that ASD might be caused by overly strong top-down prior potentiation to minimize prediction error, which can enhance capacities for rote learning while losing the capacity to generalize what is learned, a pathology typical of ASD. The proposed model naturally reflects such pathology with the likelihood weighted above a threshold. Furthermore, this model may afford insight into mechanisms underlying spontaneous or free action. The meta-prior arbitrates between deterministic chaos and externally sampled noise in the generation of action. Arbitration by such a meta-prior at each level in the hierarchy may thus be involved in balancing homeostatic control in the lower level with goal-directed control in the higher level (Pezzulo et al., 2015).

Next, let's consider how robotic experiments using predictive coding or free energy minimization help to understand infant development. There is a mix of optimism and pessimism on this issue. Most commentators consider curiosity-driven exploration in terms of minimizing prediction error with maximal slope (e.g., Butz and Kutter, 2017; Marshall et al., 2004; Oudeyer and Smith, 2016) to be one way to stage development from easy to difficult. If this mechanism is implemented in the aforementioned VBP-MTRNN, the balancing between more probabilistic exploration and more deterministic exploitation might be arbitrated by the meta-prior value of the weighting. Then, the question again arises - How to modulate this meta-prior in the course of development?

As recent neuroscience studies suggest, critical periods in development arise due to interactions between innate structures and epigenetic experiences, and cannot be explained with just "two words"—embodiment and hierarchical inference. For example, Takesian and Hensch (2013) suggest that the onset of the critical period in visual cortex development is determined by the maturation of specific GABA circuits balancing excitatory and inhibitory neural activity, whereas molecular "brakes" (often extracellular) close this window and limit further "rewiring". A third term is necessary, innate structure.

How might we understand the developmental process systematically, in terms of a triplet interaction between innate structure, embodiment, and hierarchical cognitive inference? One proposal is to conduct synthetic experiments scaled up to the evolution of

genomes over generations of neuronal learning as implemented in hundreds of interacting robots, in which evolving genomes provide contextual parametric constraints on neural structures and their functions during development. Meta-priors contextually regulating the developmental process—autonomously balancing between exploration and exploitation, shifting from homeostatic control in earlier stages to goal-directed control in later stages, or starting and closing of critical period for each modality—may emerge through this triplet interaction, and in this way we may investigate—what Friston calls—"the fundamental nature of insight and aha moments".

Ahmadi, A. and Tani, J (2017) Bridging the gap between probabilistic and deterministic models: A simulation study on a variational Bayes predictive coding recurrent neural network model, arxiv: 1706.10240.
Butz, M.V., Kutter, E.F., (2017). How the mind comes into being: An introduction to cognitive science from a functional and computational perspective. Oxford University Press.
Van de Cruys, S., Evers, K., Van der Hallen, R., Van Eylen, L., Boets, B., de-Wit, L., & Wagemans, J. (2014). Precise minds in uncertain worlds: predictive coding in autism. Psychological review, 121(4), 649-675
Marshall, J., Blank, D., & Meeden, L. (2004). An emergent framework for self-motivation in developmental robotics, ICDL2004, 104-111.
Murata, S., Yamashita, Y., Arie, H., Ogata, T., Sugano, S., & Tani, J. (2015). Learning to perceive the world as probabilistic or deterministic via interaction with others: a neuro-robotics experiment. IEEE transactions on neural networks and learning systems. published Online DOI: 10.1109/TNNLS.2015.2492140.
Oudeyer, P-Y. and Smith. L (2016). How Evolution may work through Curiosity-driven Developmental Process Topics in Cognitive Science, 1-11.
Pezzulo, G., Rigoli, F., & Friston, K. (2015). Active Inference, homeostatic regulation and adaptive behavioural control. Progress in Neurobiology, 134, 17–35.
Takesian, A. E., & Hensch, T. K. (2013). Balancing plasticity/stability across brain development. Prog Brain Res, 207(3), 3-34.

New Dialogue Initiation

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.

IEEE TCDS Table of Contents

Volume 9, Issue 1, March 2017

Guest Editorial Cognitive Agents and Robots for Human-Centered Systems

Alessandro Di Nuovo, Giovanni Acampora, Matthew Schlesinger

A current challenge for artificial cognitive systems is to acquire the capability to freely interact with humans in such a way that the burden of learning and adaptation lies with the machine and not with the human. This special issue presents a variety of recent developments in artificial cognitive agents, embedded in human-centric environments, with a particular focus on robotic applications. It aims to capture potential multidisciplinary research directions in order to further expand artificial system capabilities in an open-ended process while socially interacting in real environments for the humans' benefit.

A Bio-Inspired Cognitive Agent for Autonomous Urban Vehicles Routing Optimization

Giuseppe Vitello, Alfonso Alongi, Vincenzo Conti, Salvatore Vitabile

Autonomous urban vehicle prototypes are expected to be efficient even in not explicitly planned circumstances and dynamic environments. The development of autonomous vehicles for urban driving needs real-time information from vehicles and road network to optimize traffic flows. In traffic agent-based models, each vehicle is an agent, while the road network is the environment. Cognitive agents are able to reason on the perceived data, to evaluate the information obtained by reasoning, and to learn and respond, preserving their self-sufficiency, independency, self-determination, and self-reliance. In this paper, a bio-inspired cognitive agent for autonomous urban vehicles routing optimization is proposed. The use of selected bio-inspired analyzing techniques, which are commonly employed to investigate the topological and functional features of a metabolic network, allows the agent to analyze the structural aspects of a road network, find its extreme pathways and outline the balanced flow combinations. This approach optimizes traffic flows over network, minimizes road congestions, and maximizes the number of autonomous vehicles reaching their destination target. Agent behavior has been tested using data coming from Palermo urban road network, Italy, while the adopted bio-inspired analysis techniques have been compared with the A* literature algorithm. Experimental results demonstrate that the approach permits to find a better global routing optimization solution. To the best of our knowledge, it is the first time that metabolic mechanisms involved in a cell survival process have been used to design a congestion solution.

A Multilevel Body Motion-Based Human Activity Analysis Methodology

Kamrad Khoshhal Roudposhti, Jorge Dias, Paulo Peixoto, Vangelis Metsis, Urbano Nunes

Human body motion analysis is an initial procedure for understanding and perceiving human activities. A multilevel approach is proposed here for automatic human activity and social role identification. Different topics contribute to the development of the proposed approach, such as feature extraction, body motion description, and probabilistic modeling, all combined in a multilevel framework. The approach uses 3-D data extracted from a motion capture device. A Bayesian network technique is used to implement the framework. A mid-level body motion descriptor, using the Laban movement analysis system, is the core of the proposed framework. The mid-level descriptor links low-level features to higher levels of human activities, by providing a set of proper human motion-based features. This paper proposes a general framework which is applied in automatic estimation of human activities and behaviors, by exploring dependencies among different levels of feature spaces of the body movement.

Learning Robot Control Using a Hierarchical SOM-Based Encoding

Georgios Pierris, Torbjørn S. Dahl

Hierarchical representations and modeling of sensorimotor observations is a fundamental approach for the development of scalable robot control strategies. Previously, we introduced the novel hierarchical self-organizing map-based encoding algorithm (HSOME) that is based on a computational model of infant cognition. Each layer is a temporally augmented self-organizing map and every node updates a decaying activation value. The bottom level encodes sensorimotor instances while their temporal associations are hierarchically built on the layers above. In the past, HSOME has shown to support hierarchical encoding of sequential sensor-actuator observations both in abstract domains and real humanoid robots. Two novel features are presented here starting with the novel skill acquisition in the complex domain of learning a double tap tactile gesture between two humanoid robots. During reproduction, the robot can either perform a double tap or prioritize to receive a higher reward by performing a single tap instead. Second, HSOME has been extended to recall past observations and reproduce rhythmic patterns in the absence of input relevant to the joints by priming initially the reproduction of specific skills with an input. We also demonstrate in simulation how a complex behavior emerges from the automatic reuse of distinct oscillatory swimming demonstrations of a robotic salamander.

Learning From Explanations Using Sentiment and Advice in RL

Samantha Krening, Brent Harrison, Karen M. Feigh, Charles Lee Isbell, Mark Riedl, Andrea Thomaz

In order for robots to learn from people with no machine learning expertise, robots should learn from natural human instruction. Most machine learning techniques that incorporate explanations require people to use a limited vocabulary and provide state information, even if it is not intuitive. This paper discusses a software agent that learned to play the Mario Bros. game using explanations. Our goals to improve learning from explanations were twofold: (1) to filter explanations into advice and warnings and (2) to learn policies from sentences without state information. We used sentiment analysis to filter explanations into advice of what to do and warnings of what to avoid. We developed object-focused advice to represent what actions the agent should take when dealing with objects. A reinforcement learning agent used object-focused advice to learn policies that maximized its reward. After mitigating false negatives, using sentiment as a filter was approximately 85% accurate. object-focused advice performed better than when no advice was given, the agent learned where to apply the advice, and the agent could recover from adversarial advice. We also found the method of interaction should be designed to ease the cognitive load of the human teacher or the advice may be of poor quality.

Supervisory Control of Multiple Robots Through Group Communication

Alessandra Rossi, Mariacarla Staffa, Silvia Rossi

Single-human supervision of collaborative semi-autonomous multirobot teams is recently getting the attention of the robotic community. In this context, the adoption of a growing number of robots does not necessarily produce a gain in performance, due to the increased workload of the human supervisor. However, enabling human operators to communicate with groups of robots can reduce the operators' effort in guiding the team. Here, group communicating is intended not only to assign a task to a group but also as a way to identify the group members. This is particularly relevant in proximate interactions or in the necessity of freeing operator's hands. In this paper, starting from an analysis of real human utterances in selecting groups of robots, we extracted the features that are useful to define a basic vocabulary and analyzed the single robot needed awareness about its own characteristics and those of the robots in the neighborhood. Such analysis is used to develop a semi-autonomous multirobot simulated environment, where a human operator can guide groups of robots. The simulated environment is used to measure the humans' interaction effort and the task effectiveness while increasing the number of robots involved in a joint task, in the two cases where the commands are issued toward single or grouped robots.

Flexible Task Execution and Attentional Regulations in Human-Robot Interaction

Riccardo Caccavale, Alberto Finzi

A robotic system that interacts with humans is expected to flexibly execute structured cooperative tasks while reacting to unexpected events and behaviors. In this paper, we face these issues presenting a framework that integrates cognitive control, executive attention, and hierarchical plan execution. In the proposed approach, the execution of structured tasks is guided by top-down

(task-oriented) and bottom-up (stimuli-driven) attentional processes that affect behavior selection and activation, while resolving conflicts and decisional impasses. Specifically, attention is here deployed to stimulate the activations of multiple hierarchical behaviors orienting them toward the execution of finalized and interactive activities. On the other hand, this framework allows a human to indirectly and smoothly influence the robotic task execution exploiting attention manipulation. We provide an overview of the overall system architecture discussing the framework at work in different case studies. In particular, we show that multiple concurrent tasks can be effectively orchestrated and interleaved in a flexible manner; moreover, in a human-robot interaction setting, we test and assess the effectiveness of attention manipulation for interactive plan guidance.

Aesthetics Evaluation for Robotic Chinese Calligraphy

Zhe Ma, Jianbo Su

A great many efforts have been devoted to reproducing a reference image for Chinese calligraphy with a robot. However, there exist modeling errors in mapping from the presentation of a reference image to brush trajectories. In this paper, a closed-loop calligraphy system is established to decrease effects brought by modeling errors and optimize aesthetics. A strategy of robotic Chinese calligraphy is newly proposed based on aesthetics evaluation, which is characterized by three aesthetics indexes. Each stroke is planned online to achieve the best aesthetics effect of the whole calligraphy work. A constraint optimization problem is formulated to optimize the aesthetics effect, resulting in the trajectory of each stroke for robotic control. Moreover, the concept of the complete feature set is presented to describe ideal features, which simultaneously characterize a calligraphy task as well as calligraphy aesthetics. The rationality of the complete feature set is analyzed in both qualitative and quantitative way. The advantage of the proposed strategy is verified in a simulation of robotic Chinese calligraphy on a database containing 400 calligraphy references. To validate effectiveness of the proposed theory, Chinese calligraphy is performed on an NAO robot with and without the proposed aesthetics evaluation.

Volume 9, Issue 2, June 2017

Guest Editorial Sensorimotor Contingencies for Cognitive Robotics

Guillem Alenya, Ricardo Tellez, Kevin O'Regan, Cecilio Angulo

The sensorimotor approach to cognition states, that the key to bring semantics to the world of a robot, requires making the robot learn the relation between the actions that the robot performs and the change it experiences in its sensed data because of those actions. Those relations are called sensorimotor contingencies (SMCs). This special issue presents a variety of recent developments in SMCs with a particular focus on cognitive robotics applications.

Yielding Self-Perception in Robots Through Sensorimotor Contingencies

Pablo Lanillos, Emmanuel Dean-Leon, Gordon Cheng

We address self-perception in robots as the key for world understanding and causality interpretation. We present a self-perception mechanism that enables a humanoid robot to understand certain sensory changes caused by naive actions during interaction with objects. Visual, proprioceptive and tactile cues are combined via artificial attention and probabilistic reasoning to permit the robot to discern between inbody and outbody sources in the scene. With that support and exploiting intermodal sensory contingencies, the robot can infer simple concepts such as discovering potential "usable" objects. Theoretically and through experimentation with a real humanoid robot, we show how self-perception is a backdrop ability for high order cognitive skills. Moreover, we present a novel model for self-detection, which does not need to track the body parts. Furthermore, results show that the proposed approach successfully discovers objects in the reaching space, improving scene understanding by discriminating real objects from visual artifacts.

Online Multimodal Ensemble Learning Using Self-Learned Sensorimotor Representations

Martina Zambelli, Yiannis Demiris

Internal models play a key role in cognitive agents by providing on the one hand predictions of sensory consequences of motor commands (forward models), and on the other hand inverse mappings (inverse models) to realize tasks involving control loops, such as imitation tasks. The ability to predict and generate new actions in continuously evolving environments intrinsically requiring the use of different sensory modalities is particularly relevant for autonomous robots, which must also be able to adapt their models online. We present a learning architecture based on self-learned multimodal sensorimotor representations. To attain accurate forward models, we propose an online heterogeneous ensemble learning method that allows us to improve the prediction accuracy by leveraging differences of multiple diverse predictors. We further propose a method to learn inverse models on-the-fly to equip a robot with multimodal learning skills to perform imitation tasks using multiple sensory modalities. We have evaluated the proposed methods on an iCub humanoid robot. Since no assumptions are made on the robot kinematic/dynamic structure, the method can be applied to different robotic platforms.

Perception of Localized Features During Robotic Sensorimotor Development

Alexandros Giagkos, Daniel Lewkowicz, Patricia Shaw, Suresh Kumar, Mark Lee, Qiang Shen

The understanding of concepts related to objects are developed over a long period of time in infancy. This paper investigates how physical constraints and changes in visual perception impact on both sensorimotor development for gaze control, as well as the perception of features of interesting regions in the scene. Through a progressive series of developmental stages, simulating ten months of infant development, this paper examines feature perception toward recognition of localized regions in the environment. Results of two experiments, conducted using the iCub humanoid robot, indicate that by following the proposed approach a cognitive agent is capable of scaffolding sensorimotor experiences to allow gradual exploration of the surroundings and local region recognition, in terms of low-level feature similarities. In addition, this paper reports the emergence of vision-related phenomena that match human behaviors found in the developmental psychology literature.

Building a Sensorimotor Representation of a Naive Agent's Tactile Space

Valentin Marcel, Sylvain Argentieri, Bruno Gas

A new approach for robotics perception, rooted in the sensorimotor paradigm, is proposed in this paper. Making systems able to autonomously adapt themselves to changes in their own body or in their environment is still a challenging question for many different scientific communities. Multiple works propose either sophisticated adaptive model-based or learning-based techniques as a solution. Recent contributions have shown that it is possible for an agent to discover the structure of its interaction with the environment or its own body via the so-called sensorimotor flow. The presented work is based on this idea, and a method for building an internal representation of sensorimotor interaction is proposed, which does not require any a priori knowledge or model. A careful mathematical formalization is outlined, together with simulations, demonstrating the effectiveness of the approach. Several cases are considered allowing for a general discussion. Moreover, plausibility of the internal sensorimotor representation is highlighted by showing that it is possible to consider motion planning directly from it.

A Multimodal Model of Object Deformation Under Robotic Pushing

Veronica E. Arriola-Rios, Jeremy L. Wyatt

In this paper, we present a multimodal framework for offline learning of generative models of object deformation under robotic pushing. The model is multimodal in that it is based on integrating force and visual information. The framework consists of several submodels that are independently calibrated from the same data. These component models can be sequenced to provide many-step prediction and classification. When presented with a test example-a robot finger pushing a deformable object made of an unidentified, but previously learned, material-the predictions of modules for different materials are compared so as to classify the unknown material. Our approach, which consists of offline learning and combination of multiple models, goes beyond previous techniques by enabling: 1) predictions over many steps; 2) learning of plastic and elastic deformation from real data; 3) prediction of forces experienced by the robot; 4) classification of materials from both force and visual data; and 5) prediction of object behavior after contact by the robot terminates. While previous work on deformable object behavior in robotics has offered one or two of these features none has offered a way to achieve them all, and none has offered classification from a generative model. We do so through separately learned models

which can be combined in different ways for different purposes.

Analysis of Cognitive Dissonance and Overload through Ability-Demand Gap Models

Gahangir Hossain, Mohammed Yeasin

Maintaining alternative decisions in working memory (WM) can lead to accumulating high cognitive load. Some aspects of cognitive load improve attentiveness, but adding a cognitively inconsistent (conflict) situation results in a failure in cognitive task performance. This research introduces the notion of the human ability-demand gap (discrepancy between human cognitive ability and task performance) and its association with task-evoked cognitive overload and cognitive dissonance (inconsistency). By using the ability-demand gap as a 3-D response model, cognitive dissonance and overload was proposed to understand the confluence among working memory capacity, users' cognitive load, and task performance. The maximum gap was computed using Kolmogorov-Smirnov (K-S) statistics. The empirical studies show that the maximum ability-demand gap can be considered as the threshold between cognitive dissonance and overload. It was also observed that there was a cyclical and nonlinear relationship between working memory capacity, cognitive dissonance/lock-up, and cognitive overload.

Constructing a Language From Scratch: Combining Bottom–Up and Top–Down Learning Processes in a Computational Model of Language Acquisition

Judith Gaspers, Philipp Cimiano, Katharina Rohlfing, Britta Wrede

We present a computational model that allows us to study the interplay of different processes involved in first language acquisition. We build on the assumption that language acquisition is usage-driven and assume that there are different processes in language acquisition operating at different levels. Bottom-up processing allows a learner to identify regularities in the linguistic input received, while top-down processing exploits prior experience and previous knowledge to guide choices made during bottom-up processing. To shed light on the interplay between top-down and bottom-up processing in language acquisition, we present a computational model of language acquisition that is based on bootstrapping mechanisms and is usage-based in that it relies on discovered regularities to segment speech into word-like units. Based on this initial segmentation, our model induces a construction grammar that in turn acts as a top-down prior that guides the segmentation of new sentences into words. We spell out in detail these processes and their interplay, showing that top-down processing increases both understanding performance and segmentation accuracy. Our model thus contributes to a better understanding of the interplay between bottom-up and top-down processes in first language acquisition and thus to a better understanding of the mechanisms and architecture involved in language acquisition.

Behavior-Based SSVEP Hierarchical Architecture for Telepresence Control of Humanoid Robot to Achieve Full-Body Movement

Jing Zhao, Wei Li, Xiaoqian Mao, Hong Hu, Linwei Niu, Genshe Chen

The challenge to telepresence control a humanoid robot through a steady-state visual evoked potential (SSVEP) based model is to rapidly and accurately control full-body movement of the robot because a subject has to synchronously recognize the complex natural environments based on live video feedback and activate the proper mental states by targeting the visual stimuli. To mitigate this problem, this paper presents a behavior-based hierarchical architecture, which coordinates a large number of robot behaviors using only the most effective five stimuli. We defined and implemented fourteen robot behaviors for motion control and object manipulation, which were encoded through the visual stimuli of SSVEPs, and classified them into four behavioral sets. We proposed switch mechanisms in the hierarchical architecture to coordinate these behaviors and control the full-body movement of a NAO humanoid robot. To improve operation performance, we investigated the individual sensitivities of visual stimuli and allocated the stimuli targets according to frequency-responsive properties of individual subjects. We compared different types of walking strategies. The experimental results showed that the behavior-based SSVEP hierarchical architecture enabled the humanoid robot to complete an operation task, including navigating to an object and picking the object up with a fast operation time and a low chance of collision in an environment cluttered with obstacles.

EDITOR **Pierre-Yves Oudeyer,** Inria and Ensta ParisTech, France, pierre-yves.oudeyer@inria.fr
EDITORIAL ASSISTANT **Fabien Benureau,** Inria and CNRS, France
ISSN 1550-1914