

Editorial

Explaining and Communicating About Developmental Systems: How To Change Representations



Pierre-Yves Oudeyer

Inria and Ensta
ParisTech, France

Editor
Chair of the Technical
Technical Committee
on Cognitive and
Developmental Systems
pierre-yves.oudeyer@inria.fr

Human infants have an extraordinary capability to discover new ways of representing and manipulating knowledge, and a key dimension of development is the progressive transformation of sensorimotor and cognitive representations. Such adaptation of representations provides high flexibility and creativity to humans, and for this reason researchers have also tried to study how analogous mechanisms could be implemented in developing robots. However, how these mechanisms are working in humans, and what are the possible approaches to implement them in machines, are still largely open questions.

This issue's dialog, initiated by Stéphane Doncieux, explores various challenges of how and why representational redescription capabilities (reusing the terms of Annette Karmiloff-Smith) could be happening in robots, with contributions from Jessica Kosie and Dare Baldwin, Georges Konidaris, Freck Stulp and Timothy Hospedales, Paul Verschure and Giovanni Pezzulo, Frank Guerin, Paul Abelha and Bipin Indurkha. In particular, they discuss how new representations can be formed out of the dynamic interaction between learning algorithms, cognitive architecture and their physical and social embodiment. The dialog highlights the need for a multiplicity of processes, happening at different time scales and levels of abstraction, ranging from fast opaque low-level sensorimotor learning to slow, more transparent rule-based learning.

This dialog illustrates, as previous dialogs in this newsletter, a major finding of modern development sciences: the biology and behavior of cognitive minds cannot anymore be conceptualized through the reductionist nature/nurture lens, but rather should be

studied as the result of the history of dynamical interaction between biological, cognitive and social structures in particular contexts. The shape of legs or skills in running or writing cannot be reduced to certain genes or certain neurons, but their origins depend on the full context which drives their activation and connection with other processes.

This is challenging many representations of human biology and cognition that science has been building in the 20th century, and in particular the representations organized around "innate" and "acquired" that are now popular in the general public. In the new dialog proposed in this newsletter, John Spencer, Mark Blumberg and David Shenk observe that in spite of numerous scientific discoveries supporting the view of development as a complex multi-factored process, the discussions of development in several scientific fields and in the general public are still strongly organized around the nature/nurture distinction. Thus, they ask a crucial question: is this because there is not yet sufficient scientific evidence, or is this because the simplicity of the nature/nurture framework is much easier to communicate (or just better communicated by its supporters)? Those of you interested in reacting to this dialog initiation are welcome to submit a response by April 15th, 2015. The length of each response must be between 600 and 800 words including references (contact pierre-yves.oudeyer@inria.fr).

CDS TC Community News

As announced in the previous issue, the name of the newsletter, as well as the name of the technical committee and of the companion IEEE CIS journal, have changed. They are now

respectively the [IEEE CIS Newsletter](#), [TC and Transactions](#) on “Cognitive and Developmental Systems” (CDS). This evolution in the name reflects the goal of contextualizing the study and modeling of developmental systems, which remain at the core, with related interdisciplinary issues on cognitive systems, evolutionary-developmental processes, and the processes that give rise to structure at different scales ranging from body and brain growth to the formation of social structures in groups of interacting individuals.

In this context, I would like to address a warm thank to Angelo Cangelosi who has achieved, as editor-in-chief of the Transactions on AMD/CDS, an outstanding job in providing a new dynamics to the journal, working on the consolidation of the community and soliciting several special issues who have already been published or are forthcoming. In 2016, the new editor of the IEEE CIS TCDS journal will be Yaochu Jin, who will continue the work to gather a growing interdisciplinary community

around the study of the mechanisms of development and cognition in natural and artificial systems.

One particularly stimulating yearly event where many gather to share their advances in these areas is the [IEEE ICDL-Epirob conference](#), which will happen this year in Cergy-Pontoises/Paris, on 19th-22th September. The general chairs are Philippe Gaussier and Minoru Asada, and the program chairs are Verena Hafner and Alexandre Pitti. The conference will feature invited talks by Karl Friston, Julie Grezes and Tamim Asfour. The deadline for paper submission is 1st April. I would like to bring special attention to the Babybot Challenge organized for the second time at the conference, and implementing a competition of computational models addressing the findings of selected infant studies. The deadline is the 15th of June, and more information can be found at: <http://www.icdl-epirob.org>

Table of Contents

Editorial

Pierre-Yves Oudeyer Representational Redescription: the Next Challenge?	4
---	---

Dialogue

Stéphane Doncieux Representational Redescription: the Next Challenge?	4
Jessica E. Kosie and Dare Baldwin Flexibility Is Key in Representational Redescription	5
George Konidaris What Are Representations For?	7
Freek Stulp and Timothy Hospedales Dual-Process Representational Redescriptions	8
Paul Verschure and Giovanni Pezzulo Embodied Knowledge Acquisition: Examples From the Distributed Adaptive Control Architecture	9
Frank Guerin, Paulo Abelha and Bipin Indurkha Analogy May Be the Key	10
Stéphane Doncieux Summary and Reply	11

New Dialogue Initiation

John P. Spencer, Mark Blumberg, David Shenk Moving beyond nature-nurture: a problem of science or communication?	13
--	----

IEEE TAMD Table of Contents

Volume 7, Issue 3, September 2015	15
Volume 7, Issue 4, December 2015	18

Dialogue

Representational Redescription: The Next Challenge?



Stéphane Doncieux

ISIR,
Université Pierre et
Marie Curie,
Paris, France

doncieux@isir.upmc.fr

Drawing inspiration from developmental psychology, it has been suggested to build cognitive architectures that allow robots to progressively acquire abstract representations (Guerin et al., 2013). Humans don't have a single optimal representation of the problems they solve. They can redescribe the information they have acquired in different formats (Karmiloff-Smith, 1995). It allows them to explore different representations and use multiple problem solving strategies, from lowlevel systematic search to abstract reasoning (Evans, 2003).

Representational redescription is the ability to change the way information is stored and manipulated, to make further treatments easier and more efficient. A representation is the description of some data in a given format. The lowest level possible for formats is the raw format of sensors and effectors. Some examples of high level representation can be drawn from artificial intelligence and machine learning communities: markov decision processes formalism, first order logic or neural networks. Changing the representation allows usage of different problem solving strategies. Adapted representations make computations easier by relying on a small set of relevant primitives instead of a big set of unstructured data.

Use a single representation or change representations over time?

Humans may use representational redescription because of physiological constraints. The genome contains twenty thousand genes to describe the whole body, including the brain with its hundred billions of neurons. Such a small number of genes may not be enough for a genetic transmission of sophisticated representations. Does it necessarily mean that robots should follow the same path? Human representational redescription may also be an advantage rewarded by evolutionary pressure because of the adaptation ability it has resulted in. Would it help robots to face open environments? This would undoubtedly be an interesting feature. In the following, we will consider the questions that it raises.

Where to start?

Sensorimotor data first need to be observed before they can be redescribed in a format that allows an agent to better understand what happened and eventually to reproduce it. Babies have grasping or sucking reflexes

that allow them to start interacting with surrounding objects before they can perform more complex actions. Guerin et al. suggest using a similar set of innate sensorimotor schemas to bootstrap the process (Guerin et al., 2013). How to choose this set of primitive schemas and where to stop? If we, as roboticians, do know how to implement an efficient grasping behavior, why should we start with an inefficient grasping reflex? A sophisticated grasping behavior may allow the robot to rapidly and efficiently interact with objects, thus generating a lot of useful data to learn about them. Where should we then put the frontier between the schemas that are provided to the robot and the ones that should be discovered? Providing efficient behaviors is clearly a convenient way to bootstrap the process. Are there other alternatives?

Evolution shaped development, but could it be also involved in the representational redescription process?

Evolution has shaped, over millions of years, living beings and their development process. But beyond this first evo-devo relation, evolutionary mechanisms may also be at play during development and learning. The principles of variation and selection have contributed to the success of evolutionary computation because of their simplicity, robustness and versatility. They have been used in a robotics context for more than twenty years (Doncieux et al., 2015), and were notably able to generate non trivial behaviors with neural networks. They are also believed to be the primary mechanisms in development, both for learning motor schemas and for selecting problem solving strategies (Guerin et al., 2013). They could then have a significant role to play in the representational redescription process, in particular thanks to their ability to generate controllers relying on the most simple representations, i.e. sensorimotor data. Furthermore, this hypothesis may be biologically plausible, as evolutionary principles can be implemented along with neural mechanisms (Fernando et al., 2012). Evolution could then be involved in brain functions and thus in development and learning. Should representation formats be given a priori or should it emerge from the developmental process?

Representational redescription requires the availability of the representation formats in which the redescription is expected to occur. A first possibility would be to provide the

agent with different representation formats, like first order logic or markov decision processes formalism, for instance. Dedicated machine learning algorithms could extract them from a lower level representation, e.g. the sensorimotor flow. An alternative would be to use a versatile connectionist formalism and rely on deep learning algorithms to re-describe lower layers representations to more abstract ones. The first alternative is a somewhat top-down approach in which learning and decision algorithms are available from the very beginning. The developmental process "just" needs to represent sensorimotor data in the corresponding format for the system to exhibit high level cognitive abilities. The second is a bottom-up approach in which higher level representations emerge progressively and where the corresponding problem solving strategies will also need to emerge.

Does provided knowledge limit developmental abilities?

Providing knowledge allows one to take shortcuts in the developmental process: no need

Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. (2015). Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*, 2.

Evans, J. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, 7(10):454-459.

Fernando, C., Szathmari, E., and Husbands, P. (2012). Selectionist and evolutionary approaches to brain function: a critical appraisal. *Frontiers in Computational*

to discover what is provided and the corresponding developmental time is then saved. Providing sensorimotor schemas or representation formats constrains what the agent can do, what it will observe and what it will extract from these observations. If the agent is expected to face an open environment, isn't it a limit to its adaptive abilities? Are there conflicts between shortening developmental time and having an open-ended developmental process?

How to make a robot endowed with representational redescription transparent?

Giving a robot the ability to change its representations and problem solving strategies may make it difficult to understand for a human. A non-expert may have trouble predicting what the system will actually do and what it understands from its environment. Making such robots transparent may then be critical for them to be used in practice, in particular if they are to enter our everyday environment. How could it be achieved?

Neuroscience, 6(April):1-28.

Guerin, F., Kruger, N., and Kraft, D. (2013). A Survey of the Ontogeny of Tool Use : from Sensorimotor Experience to Planning. *IEEE Transactions on Autonomous Mental Development*, 5(1):18-45.

Karmiloff-Smith, A. (1995). *Beyond Modularity: A Developmental Perspective on Cognitive Science*. The MIT Press.

Flexibility Is Key in Representational Redescription



Jessica E. Kosie



Dare Baldwin

Adaptive flexibility is a hallmark of human functioning. In fact, we can view the field of robotics as yet another illustration of humans' ability to adaptively construct novel solutions to problems the world presents. Should a major goal of robotics be to design robotic systems that display the kind and degree of adaptive flexibility characteristic of humans? Clearly there are a myriad possible ways that robots can (and already do!) assist humans in the absence of such flexibility. However, many theorists share the strong intuition that robots' ability to assist humans in a sizable variety of contexts will be maximized if they can interface with humans in a natural manner (Metta et al., 2010). To the extent that one endorses this goal, it seems that achieving parity in robot-human adaptive flexibility will be important. Key to this endeavor will be gaining a full understanding of the human side of this flexibility equation.

Human adaptive flexibility seems to arise, at least in part, out of a propensity for representational redescription. As just one example, the way humans apprehend events as experience

unfolds illustrates the mind's deeply redescriptive bent. In event processing, streaming sensory information is transformed into discrete events that are categorized, sequenced, integrated across, and infused with inferences about motives and causes and other unobservables. A given sensory stream can be redescribed in any number of ways, with the optimal redescription being a matter of fit to the perceiver's processing goals. Often it is adequate and efficient to redescribe events only in very general terms, as when one idly notices another tying a shoelace without processing the motion stream in any detail. In contrast, one would engage in a much more fine-grained redescriptive analysis of the same motion stream if one's processing goal were to try to learn how to tie shoe laces. As this example suggests, adult humans are fluent in shifting the level or scale at which they redescribe unfolding events, as well as selecting an optimal perspective from which to redescribe events. As well, humans readily update their event representations via integration of new information as it comes online. Stéphane Doncieux asks whether we should

Jessica E. Kosie

Department of Psychology
University of Oregon
Oregon, USA

jkosie@uoregon.edu

build robots that use a single representation or that change representations over time. The observations we have just offered seem to suggest clear value in developing robotic systems capable of flexible event redescription skills that generate representations open to change.

Dare Baldwin

Department of Psychology
University of Oregon
Oregon, USA

baldwin@uoregon.edu

Redescription seems to be the mind's solution to what Christiansen and Chater (in press) call the "now-or-never bottleneck." Faced with a barrage of incessantly streaming sensation, compressing and recoding that flow as it unfolds enables the mind to harvest more of the available information to enhance learning and guide future action. This comes at a cost, however: information is lost in the compression process. Imagine the advantage to human functioning if robot partners could simultaneously both engage "humanistically" by means of human-like flexible, redescriptive mechanisms (thereby facilitating interaction) and record veridical (and even an expanded range of) sensory information as it indeed occurred (thereby enabling access to the prior sensory data that to a human alone would be lost).

Doncieux also asks: where should we start in building robotic systems? That is, which representational primitives are most appropriate to gift the system with? We offer several ideas here. First, if one goal is to achieve systems that both a) self-organize via interaction with the world, and b) achieve representations similar to those that humans acquire, then gifting robotic systems with a starting state like that of infant humans seems a viable strategy. Working from a human model in robotic design in essence capitalizes on the effort that evolutionary mechanisms have already put in across millennia to achieve a self-organizing system that flexibly and responsively represents a changing environment. Although there remains much to be learned about the precise nature of such human starting states, recent research compellingly demonstrates that human newborn "primitives" aren't as primitive as one might think. In particular, a seminal body of work by Von Hofsten and colleagues (e.g., von Hofsten, 1982, 2004, 2007) clarifies that so-called "reflexes" are far from

the fragile, fixed, involuntary behavioral patterns that they were long assumed to be. For example, reaching actions in newborns appear to be both intentional and to display flexible adaptation to the position and velocity of the object targeted by the reach. Similarly, eye movements in newborns display systematic patterns of selectivity that progressively adapt across time as visual learning occurs (e.g., Haith, 1980; Johnson, Dziurawiec, & Ellis, 1991). This body of evidence suggests that representational flexibility and openness to change is a hallmark of the human self-organizing mind from the ground up.

Lastly, systems that self-organize responsively to the environment face a fundamental tension: On the one hand, they must act, both in order to survive and to elicit feedback from the environment to learn. On the other hand, the more their actions are driven endogenously (initiated and guided by their own internal, representational mechanisms, that is) the less sensitive they are to exogenous factors (the sensory information they are encountering as they move through the world). This lack of sensitivity to the environment potentially undercuts certain kinds of learning. This is one way in which knowledge can limit developmental ability, a point that Doncieux asks that we consider. Recent evidence spawns the hypothesis (e.g., Toyozumi, et al, 2013) that critical periods (periods during which learning via cortical plasticity promotes high levels of ultimate sensory functioning) may be a biological solution to the tension between knowledge and learning. In particular, critical periods may offer a limited time window in which the mind is maximally responsive to external input, with internally-generated neural activation inhibited during this period. Developmental robotics holds potential to provide a unique domain within which to investigate this hypothesis, and to examine whether there are alternative possible resolutions to the knowledge/learning tradeoff. Investigation of this kind holds promise for ultimately augmenting human learning capacity, both directly and via assistance from robotic systems with enhanced learning potential.

Christiansen, M.H., & Chater, N. (in press). The now-or-never bottleneck: A fundamental constraint on language. *Behavioral and Brain Sciences*.

Haith, M.M. (1980). Rules that babies look by: The organization of newborn visual activity. Erlbaum.

Johnson, M.H., Dziurawiec, S., Ellis, H., & Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition*, 40, 1-19.

Metta, G., Natale, L., Nori, F., Sandini, G., Vernon, D., Fadiga, L.,...Montesano, L. (2010). The iCub humanoid robot: An open-systems platform for research in cognitive

development. *Neural Networks*, 23, 1125-1134.

Toyozumi, T., Miyamoto, H., Yazaki-Sugiyama, Y., Atapour, N., Hensch, T.K., & Miller, K.D. (2013). A theory of the transition to critical period plasticity: Inhibition selectively suppresses spontaneous activity. *Neuron*, 80, 51-63.

von Hofsten, C. (1982). Eye-hand coordination in the newborn. *Developmental Psychology*, 18, 450-461.

von Hofsten, C. (2004). An action perspective on motor development. *Trends in Cognitive Sciences*, 8, 266-272.

von Hofsten, C. (2007). Action in development. *Developmental Science*, 10, 54-60.

What Are Representations For?



George Konidaris

Intelligent Robot Lab
Departments of
Computer Science and
Electrical & Computer
Engineering
Duke University
Durham, NC, USA
gdk@cs.duke.edu

Stéphane Doncieux poses the question of whether or not an agent should construct multiple different representations of the world. The answer to this question is surely yes. It has been said that problem solving in computer science is the art of relentless abstraction, and that given the right representation, all hard problems become easy. These things are true. But in AI we are concerned about behavior, not representations. Consequently, I think we need to ask a slightly different question.

AI researchers tend to think in terms of models, which are really a set of representational assumptions. A subfield agrees on a model—Markov decision processes, or STRIPS descriptions, for example—establishes it as a convention, and then focuses on algorithms for solving problems posed in that model. This can be tremendously positive: if the model is sufficiently general to capture all problems of interest, and sufficiently specific to encode the structure necessary for making progress, then it provides a common ground upon which real progress can be made. However, it can also impede: many subfields of AI have been using particular models for so long that they have come to believe that their models are *real*. Those of us who work with robots know that things are not quite so simple. Models are not real; only robots are real: only sensors and actuators. Models, and therefore representations, are fictional—they are hallucinations that agents are free to invent in order to improve their behavior. Because of this freedom the question of *whether or how many* (re-)representations an agent should construct is necessarily under-specified. The question is rather, *what are representations for?* Since we are ultimately interested in intelligent behavior, more precisely: *which representations should an agent have to support a specific behavior?*

My recent work (Konidaris and al., 2014; 2015; Konidaris, 2015) (in collaboration with Leslie Kaelbling and Tomas Lozano-Perez) has applied this approach to constructing high-level representations for planning. The idea is broadly to learn representations that allow an agent to move from a reinforcement learning scenario—where it has high-dimensional sensors and actuators and must learn by trial and error—to a high-level planning scenario

where it has an abstract representation with which it can construct long-range plans. The major question is then what symbols an agent should have—what abstract vocabulary should it construct in order to plan?

This question has been asked before, with varying levels of success. The key difference in our work is that we first asked the question: *what are these symbols for?* In this case, they are for determining whether plans composed of sequences of the robot's action are feasible, or not. We therefore formalized the agent's *plan space*—the set of all plans it may wish to reason about—and found a symbolic vocabulary that is provably sufficient for determining whether any plan in a robot's plan space is feasible, or not (Konidaris and al., 2014). We were also able to show that an agent can autonomously learn those symbols by simply executing its actions and observing the results (Konidaris and al., 2015). My current work extends this approach to construct symbolic hierarchies (Konidaris, 2015), where successive layers of high-level actions lead to successively more abstract re-descriptions of the environment—exactly the sort of re-representation Doncieux describes.

A key result of our work is that not only do the appropriate symbols depend on the agent's actions, but so does the *type* of representation: the appropriate abstract representation for an agent equipped with actions that reach subgoals is a graph, whereas actions that reach abstract subgoals (which change only some low-level variables, and leave others unmodified) demand a STRIPS-style factored representation. Here we have an abstract representation that is grounded in an agent's real sensorimotor interaction with the world, has precise properties, and is provably suitable for a specific behavioral competency. Moreover, we have shown that adding higher levels necessarily entails a re-description of the world. We were able to do this only because we started with a precise formulation of what that representation was for. My answer to Doncieux's question of *how many representations, and which?* is therefore to instead ask the question, *how should my robot behave, and which representations support that behavior?* The second question will lead to an answer to the first.

G.D. Konidaris and L.P. Kaelbling and T. Lozano-Perez. Constructing Symbolic Representations for High-Level Planning. In Proceedings of the Twenty-Eighth Conference on Artificial Intelligence, pages 1932-1940, 2014.
G.D. Konidaris and L.P. Kaelbling and T. Lozano-Perez. Symbol Acquisition for Probabilistic High-Level Planning,

in Proceedings of the Twenty Fourth International Joint Conference on Artificial Intelligence, pages 3619-3627, 2015.

G.D. Konidaris. Constructing Abstraction Hierarchies Using a Skill-Symbol Loop. ArXiv:1509.07582.

Dual-Process Representational Redescriptions



Freerk Stulp

Robotics and Computer Vision
Computer Science and System Engineering Dept.,
ENSTA - ParisTech,
FLOWERS Team,
Palaiseau, France
freerk.stulp@ensta-paristech.fr



Timothy Hospedales

RIM Group
Electronic Engineering and Computer Science
Queen Mary University,
London, UK
t.hospedales@qmul.ac.uk

An important distinction, which Stéphane Doncieux mentions in the first paragraph, is whether to consider representational redescription (RR) at the low sensori-motor level, or at a higher abstract reasoning level. Dual-process theories pose that there are indeed two distinct processes for these two levels. From a variety of nomenclatures (Stanovich and West, 2000; Table 3) we choose to denote these systems “Implicit Cognition” and “Explicit Cognition”. Implicit Cognition is rapid, parallel, automatic and associative, and includes innate behaviors. Explicit Cognition is slow, abstract, and rule-based, and its hypothetical reasoning mechanisms have only been acquired by humans. This duality is also at the heart of the “Central Paradox of Cognition”: is cognition of a connectionist or a symbolic nature?

At the lowest level, representational redescrptions can be hard-coded into the system, either through evolution (e.g. edge detectors in the retina), or by engineers (e.g. dedicated VLSI implementations for edge detection). Although this conforms to Stéphane’s definition of RR, we prefer to exclude such innate redescrptions, because they are not modifiable.

At an intermediate level, but still in Implicit Cognition, we see RR as corresponding to a search within the parameters of a given model. Representation search can occur through local refinement (e.g. backpropagation), or more globally through evolutionary processes (Fernando, 2012). Corresponding methods in artificial intelligence include NEAT (“NeuroEvolution of Augmenting Topologies”, Stanley and Miikkulainen, 2002). NEAT adapts the topology of a neural network, which allows it to find very compact representations. For instance, the simplicity and elegance of the network that NEAT discovers for controlling a double pole is striking: it has only one hidden node! NEAT has redescrbed the topology of the network such that the connection weights required to solve the task are quite easy to find.

The capability of deep neural networks to represent increasingly complex and abstract features and concepts at increasingly higher layers is also a form of RR. In such networks, representations are stored in a distributed, implicit way. From a dual-process perspective, we will denote RR at this level as “Implicit

Representational Redescription”. The implicitness of the representations used at this level means it is difficult to reason about such representations at a higher level of abstraction, which makes reuse more difficult. For instance, did a low-level search in feature space enable Isaac Newton’s to redescrbe the reason why apples fall, and that this is the same reason why the moon orbits the earth?

We consider such symbolic redescrptions based on abstract reasoning to be “Explicit Representational Redescrptions”. This form of redescrption is more closely related to common notions of insight and epiphanies. A classical example of this approach is the cognitive architecture Soar, which is able to find novel symbolic representations by “chunking” simpler representations. Such abstract mechanisms could be used to determine which low-level representations are to be dynamically used at different times and in different contexts. How to achieve such symbolic redescrptions is less clear, but our intuition is that evolutionary processes are less likely to play a role in Explicit than in Implicit RR (Third question - Evolution in RR?). In response to the last question (How to make RR transparent?), we believe Implicit RR is, almost by definition, opaque, whereas Explicit RR yields abstract representations that are amenable to human inspection.

Perhaps the most profound open question in this context is: how can opaque implicit representations be redescrbed to acquire transparent explicit representations? Possible strategies may be perhaps seen in recent deep and recurrent networks such as Memory Networks (Weston et al, 2015) and Neural Turing Machines (Graves et al, 2014), which tightly integrate distributed and explicit compact data stores. Such networks may also potentially allow different representations to be dynamically used at different times and in different contexts.

In summary, and in response to the first question (single or changing representations?), we believe that (continually) changing representations are essential to achieving life-long autonomy. In this context, the questions that Stéphane raises have quite different answers, depending on whether we consider Implicit RR, Explicit RR, or the even deeper question of redescrbing across these two levels.

Stanovich, K.E. and West, R.F. (2000). Individual differences in reasoning: implications for the rationality debate? *Behavioral and Brain Sciences*, 23(5): 645-65

Fernando, C., Szathmari, E., and Husbands, P. (2012). Selectionist and evolutionary approaches to brain function: a critical appraisal. *Frontiers in Computational Neuroscience*, 6(April):1-28.

Kenneth O. Stanley, K.O., and Miikkulainen, R. (2002). Evolving Neural Networks Through Augmenting Topologies. *Evolutionary Computation*, 10(2): 99-127

J. Weston, S. Chopra, and A. Bordes. (2015). Memory networks. In *Int. Conf. on Learning Representations (ICLR)*.

A. Graves, G. Wayne, and I. Danihelka. (2014). Neural Turing machines. *arXiv preprint: 1410.5401*.

Embodied Knowledge Acquisition: Examples From the Distributed Adaptive Control Architecture



Paul Verschure

University Pompeu Fabra
and Catalan Institute
of Advanced Studies
(ICREA), Barcelona, Spain
paul.verschure@upf.edu



Giovanni Pezzulo

Institute of Cognitive
Sciences and
Technologies,
National Research Council,
Rome, Italy
giovanni.pezzulo@istc.cnr.it

A longstanding objective of (developmental) robotics is realizing cognitive architectures that acquire representations autonomously – and at various levels of abstraction. One source of insights to understand knowledge acquisition, discussed in Stéphane Doncieux's dialogue, is developmental psychology. Here we discuss another, biologically grounded methodology that incorporates neurobiological knowledge into the design of robots to emulate the way living organisms acquire knowledge by interacting with the external environment.

The Distributed Adaptive Control (DAC) theory of mind and brain (Verschure et al., 2014, 2003; Verschure, 2012) exemplifies this methodology. It proposes that the brain is organized as a three-layered control structure with tight coupling within and between these layers: the Soma (SL, comprising the body with its sensors, organs and actuators) controlled by the Reactive (RL), Adaptive (AL) and Contextual (CL) layers.

The Reactive Layer (RL) comprises a repertoire of sensorimotor, drive-reduction mechanisms, which mediate the selection and execution of basic behaviors (e.g., approach food; avoid obstacles or dangers), thus providing a capability for allostatic control, on top of which the other two (AL and CL) layers can operate and learn.

The Adaptive Layer (AL) acquires a state space of the agent-environment interaction and shapes more complex action patterns. Importantly, the learning dynamics of AL are defined by the sensorimotor contingencies implicitly produced by the RL while the robot interacts with the environment. Following the paradigm of classical conditioning, sensory states (e.g. distal cues) and motor patterns (e.g., approach behavior) become linked through the valence states signaled by the RL, and form the robot's state space. In the brain, such sensori-motor chunking might be realized by the hippocampus, by integrating information from the lateral and the medial entorhinal cortex that represent sensory and action information, respectively (Rennó-Costa et al., 2010).

Finally, the Contextual Layer (CL) uses two (sequential short and long-term) memory systems to learn sequences of integrated sensorimotor representations that are

(continuously) generated by the AL, and which achieve goals. Every learned sequence is "labeled" with respect to the goal it achieves and its valence, and can be successively selected to mediate (or plan) instrumental action. The CL thus acquires an internal model (possibly corresponding to prefrontal function, Duff et al., 2010) that is constrained by the robot's sensorimotor interactions, and which expands the time horizon in which the robot can operate and "reason". An additional, self model component of the CL monitors task performance and develops (re)descriptions of task dynamics, generating meta-representational knowledge that forms autobiographical memory (Verschure, 2012).

As this (simplified) description highlights, DAC bootstraps representation and cognition from sensorimotor contingencies, in the sense that knowledge acquisition and representation uses the same behavioural patterns that the robot enacts during (instrumental) action. This acquisition process strongly constrains the kind of representations DAC develops. For example, the CL's internal model is control-oriented (i.e., it ultimately mediates instrumental control, not knowledge per se) and valence-oriented (i.e., links to drive-reduction RL systems). The overall DAC representational framework is organized around the goals of the agent and modulated by perception, memory and value. It thus includes those "grounded abstractions" that a situated agents needs to forage and survive, not "Linnaean taxonomies" (sometimes associated to visual hierarchies or prefrontal cortex); and not even hidden representations of the kind developed by deep nets that cannot act or move but have access to practically unbounded data (Hinton, 2007). Finally, knowledge is in a sense "re-described" at every layer - from simple sensorimotor contingencies to complex action plans and conscious experience - but the DAC methodology offers a clear rationale for this re-description: serving the increasing demands of action control, see also (Cisek, 1999; Pezzulo et al., 2011; Pezzulo and Castelfranchi, 2009).

DAC's action-oriented approach can be potentially used to study knowledge acquisition during development. Understanding whether it is appropriate to model all forms of knowledge remains an open objective for future research.

Cisek, P., 1999. Beyond the computer metaphor: Behavior as interaction. *J. Conscious. Stud.* 6, 125–142.

Duff, A., Fibla, M.S., Verschure, P.F.M.J., 2010. A biologically based model for the integration of sensory-motor contingencies in rules and plans: A prefrontal cortex based extension of the Distributed Adaptive Control architecture. *Brain Res. Bull.*

Hinton, G.E., 2007. Learning multiple layers of representation. *Trends Cogn Sci* 11, 428–434. doi:10.1016/j.tics.2007.09.004

Pezzulo, G., Castelfranchi, C., 2009. Thinking as the Control of Imagination: a Conceptual Framework for Goal-Directed Systems. *Psychol. Res.* 73, 559–577.

Pezzulo, G., LW, B., Cangelosi, A., Fischer, M.H., McRae, K., Spivey, M., 2011. The Mechanics of Embodiment: A Dialogue on Embodiment and Computational Modeling.

Front. Cogn. 2, 1–21.

Rennó-Costa, C., Lisman, J.E., Verschure, P.F.M.J., 2010. The Mechanism of Rate Remapping in the Dentate Gyrus. *Neuron* 68, 1051–1058. doi:10.1016/j.neuron.2010.11.024

Verschure, P.F.M.J., 2012. Distributed Adaptive Control: A theory of the Mind, Brain, Body Nexus. *Biol. Inspired Cogn. Archit.* 1, 55–72. doi:10.1016/j.bica.2012.04.005

Verschure, P.F.M.J., Voegtlin, T., Douglas, R.J., 2003. Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature* 425, 620–624. doi:10.1038/nature02024

Verschure, P., 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.



Frank Guerin

Computing Science
Department,
University of Aberdeen
Scotland, UK
f.guerin@abdn.ac.uk



Paulo Abelha

Computing Science
Department,
University of Aberdeen
Scotland, UK
p.abelha@abdn.ac.uk



Bipin Indurkha

Jagiellonian University
Kraków, Poland
bipin.indurkha@uj.edu.pl

Analogy May Be the Key

The ability to change representations is a crucial challenge to cognitive systems, because Representational Redescription (RR) is something very natural for humans. All normal humans, during development, pass through stages where different representations are in use (see Karmiloff-Smith book cited in dialog initiation). RR is alien to typical Artificial Intelligence (AI) approaches, which are based on one “correct” or “best” representation. Therefore, RR may be in the category of problems that are easy for humans but hard for machines, like natural language processing, or vision. These “hard for machines” problems required decades of effort to reach the current state-of-the-art. Given the amount of effort needed to make serious advances with practical applications in these other areas, one might project perhaps a thirty-year timeline for RR to deliver a useful system.

RR is about taking the representation a child has at a particular stage of development and in a particular domain (say representation A), and transforming it into a more sophisticated representation (B) acquired by the same child at a later stage. A key problem is that we seem to have almost no idea what would constitute a plausible A or B: we do not have a good idea of the representation in use by a human child at any age. Human children display abilities to transfer, and in commonsense knowledge (both linked), that greatly exceed any existing AI system. This is why the “commonsense knowledge problem” is still unsolved. Decades of research have not yet figured out how to implement child-like representations. RR is a higher-order problem on top of this. RR requires as a basis child-like representations where knowledge is implicit and linked to contexts, and yet flexible enough to be transferred to similar situations; but existing AI representations are not like this. It might be necessary to solve this representation problem before RR.

Minsky said “... commonsense is knowing maybe 30 or 60 million things about the world and having them represented so that

when something happens, you can make analogies with others” (Dreifus, 1998). Children (or humans in general) must have representations that facilitate analogies (which we see as synonymous with transfer). We believe that analogy may be key to the representation problem and also the RR problem, because the process of analogy does involve the creation of novel representations. This has been called the “vertical” view of analogy (Morrison and Dietrich, 1995), but most work on analogy in AI does not address this aspect. Mitchell and Hofstadter’s Copycat is one nice example of analogy work which creates representations. Its internal bits of representations in formation have been described as bearing “a close resemblance to the shifting enzyme population of a cell” (Hofstadter et al., 1995), so there is a role for evolutionary mechanisms in representation formation. So far Copycat was only applied to toy domains; we are interested in seeing this type of technique applied in a real domain such as robotics or language understanding. Given how ubiquitous analogy is in human cognition it is remarkable how neglected it is in these AI areas. This reinforces the view that the road ahead is long.

Can Deep Learning give us the appropriate representations which facilitate analogy? The problem with existing deep nets is clear for object recognition from different viewpoints. In order to do this easily the network would need to learn the correct parts of the object and relationships among them. This could potentially be learned, but requires further research. Existing deep nets tend to have the object knowledge and viewpoints buried in complex nonlinear relationships. We can extrapolate beyond visual objects to learning concepts from text corpora to see similar difficulties. With concepts such as Hofstadter’s example of an “airline hub” (Hofstadter et al., 2003), it is clear one needs some sort of symbolic handle on mid-level concepts and relationships to handle human-like analogy.

Hofstadter, Douglas, and Melanie Mitchell. (1995) "The copycat project: A model of mental fluidity and analogy-making." Fluid concepts and creative analogies, Basic Books

Hofstadter, Douglas, and Emmanuel Sander (2013). Surfaces and Essences. Basic Books.

Dreifus, C. (1998): 'A Conversation: With Dr. Marvin Minsky;

Why Isn't Artificial Intelligence More Like the Real Thing?, The New York Times (August 1998).

Morrison, C. T., and Dietrich, E. (1995). Structure-mapping vs. high-level perception: The mistaken fight over the explanation of analogy. In Cognitive Science Conf., 678-682. Psychology Press.

Summary and Reply: "Representational Redescription: the Next Challenge?"



Stéphane Doncieux

ISIR,
Université Pierre et
Marie Curie,
Paris, France
doncieux@isir.upmc.fr

Many different approaches have been developed to control robot behavior and provide them with learning, planning or reasoning abilities. Each comes with dedicated algorithms and specific representations. The question I have raised is to know whether these different representations, or new ones, could compete as a new and unified representation in developmental robotics, or if we should consider them as complimentary and focus on how to bootstrap them or switch from one to another. Representational redescription is this ability to change the way knowledge is stored and manipulated. I have also highlighted some of the questions that such an approach raises.

Kosie and Baldwin argue about the advantages of robots able to redescribe their representations. As it is an important feature of human beings, a robot with similar abilities would easily engage in "humanistic" interactions. They would also offer the possibility to record veridical sensory information allowing later access to objective data that are lost in the human redescription process.

As highlighted by Konidaris, each sub-field of artificial intelligence has focused on one particular formalism and tried to get the best of it. Many different problem solving algorithms and the corresponding representation formalisms are then available now. Are these representations adapted to a representational redescription cognitive system or does the redescription ability create new requirements that would make them inadequate? Guerin et al. defend this last point of view. They consider the ability to make analogies as a key to solve the redescription problem. Flexibility and ability to consider the context are, according to them, required, but current formalisms do not include such features. They predict then that the road ahead is still long before a cognitive system with representational redescription can be developed.

Kosie and Baldwin comment on the question of where such systems should start. They summarize in particular recent works showing that what was thought to be primitives in human development actually reveals to be much more flexible than what was previously thought. They also highlight that biological systems demonstrate critical

periods in their development during which internal information is clearly neglected to maximize the impact of external inputs. This suggests that biological systems don't rely on fixed primitives to start exploring and bootstrap development, but on more adaptive mechanisms.

Stulp and Hospedales introduce the notions of implicit and explicit cognition, that respectively correspond to fast, parallel and associative learning and slow, abstract and rule-based learning. Implicit cognition typically correspond to connexionist approaches, may it be neuroevolution or deep learning. Verschure and Pezzulo present DAC (Distributed Adaptive Control), a connectionist architecture and thus an example of what an implicit cognition system may look like. It has the specificity to let the robot decide and choose its behavior. The notion of explicit cognition proposed by Stulp and Hospedales relies on symbolic representations. The discovery of such symbols remains an open question. Finally, Stulp and Hospedales oppose the opacity of implicit cognition -- also mentioned in Guerin et al. comment on deep learning -- to the transparency of explicit cognition. Implicit cognition can be bootstrapped from low sensori-motor level. It seems difficult if not impossible to directly bootstrap explicit cognition. One of the most critical open question for representational redescription for them may then be to redescribe implicit representations in explicit representations.

Konidaris, by putting forward the question of what a representation is for, suggests a pragmatic method to answer my questions: if we know what the behavior of the robot should be, we should be able to identify the representation or set of representations that fit the needs. The potential of this methodology is clear for explicit cognitive systems as the system is transparent. It seems less clear for an implicit cognitive system for which internal representations will be particularly difficult to predict or understand. We have for instance solved with neuroevolution a ball collecting task in which a robot has to directly control its wheels to make the robot collect balls and put them into a basket (Mouret and Doncieux, 2011). Some of the generated neural networks had no hidden neurons at all,

what was unexpected for this task (Ollion, 2013). It suggests that our intuitions may be misleading for implicit cognitive systems and that Konidaris method may not be adapted in this case.

To summarize, there is a consensus among the people who participated to this dialog concerning the potential and importance of studying representational redescription in developmental robotics. Two critical issues have been identified: redescrbing implicit representations to explicit representations and building a cognitive system able to perform analogies.

Going toward explicit representations is actually very promising. Such systems are more transparent, what is particularly important if such robots have to interact with humans. I also expect them to be faster and more robust for solving new problems with respect to learning based methods, as they would allow to find a solution without the need to explore. This will be true only if the representation used by these planning or reasoning algorithms is adapted to the situation and contains only the relevant symbols. I don't expect that such representations can be developed without an efficient implicit cognitive system that will make the robot accumulate the required experience out of which abstract concepts can be extracted.

I agree that the flexibility of human analogy is still out of reach for artificial intelligence and machine learning. If many methods have been proposed to transfer knowledge while learning (Pan and Yang, 2010), notably in the field of reinforcement learning (Taylor and Stone, 2009), finding out what is to be transferred or not and in what circumstances, i.e. finding analogies between available knowledge and

a current situation, remains an open issue. I also agree with the importance of analogy to reach human-like representational redescription, but my first guess is that it is not the current bottleneck. As said by Guerin et al. this process requires "child-like representations where knowledge is implicit and linked to context". The cognitive process needs then to be bootstrapped first to make it build implicit representations. My second guess is that rich implicit representations could make analogy much easier. Neural networks can be given self-organizing properties that allow them to regroup perceptions and thus find similarities (Johnsson and Balkenius, 2011). This could be a basis for implementing analogy.

My position is then that rich implicit representations is a prerequisite to build explicit representations and to support analogies, but this process raises a new critical issue : how to make the robot acquire the experience that is required for further representational redescription? In the DREAM project (European FET proactive project No 640891, <http://www.robotsthatdream.eu/>), we focus on the bootstrap of an implicit cognitive system. The goal is to make the robot acquire experience with algorithms that do not need problem and environment specific representations. Our approach relies on evolutionary algorithms for their versatility and robustness. Inspired by the impact of sleep on the human representational redescription process (Wagner et al., 2004), we are following a process in which active phases in the real world ("daytime") are separated by phases of analysis or exploration in simulation ("nighttime") during which the restructuring occurs. The new representations are to be used during the next "daytime" phase to solve problems or go on exploring. We will then try to make robots sleep and "dream" to make them more adaptive.

[1] **Magnus Johnsson and Christian Balkenius.** Sense of touch in robots with self-organizing maps. *Robotics, IEEE Transactions on*, 27(3):498–507, 2011.

[2] **J-B Mouret and S Doncieux.** Encouraging Behavioral Diversity in Evolutionary Robotics: An Empirical Study. *Evolutionary computation*, 20(1):91–133, aug 2012.

[3] **C. Ollion.** Emergence of Internal Representations in Evolutionary Robotics: influence of multiple selective pressures. These de doctorat, EDITE, ISIR/UPMC - Paris-Descartes University, 2013.

[4] **Sinno Jialin Pan and Qiang Yang.** A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, October 2010.

[5] **ME Taylor and Peter Stone.** Transfer learning for reinforcement learning domains: A survey. *The Journal of Machine Learning Research*, 10:1633–1685, 2009.

[6] **Ulrich Wagner, Steffen Gais, Hilde Haider, Rolf Verleger, and Jan Born.** Sleep inspires insight. *Nature*, 427(6972):352–5, jan 2004.

New Dialogue Initiation

Moving beyond nature-nurture: a problem of science or communication?



John P. Spencer

School of Psychology,
University of East Anglia
United Kingdom

j.spencer@uea.ac.uk



Mark Blumberg

Dept. of Psychological &
Brain Sciences,
University of Iowa,
Iowa, USA

mark-blumberg@uiowa.edu



David Shenk

DeLTA Center,
University of Iowa,
Iowa, USA

homepage@davidshenk.com

In his classic essay, “Seven Wonders,” the physician and essayist Lewis Thomas wrote that childhood was one of life’s great mysteries. For Thomas, childhood led to a sense of wonder, not because it is a magical time, but because it might have been avoided. Why, he pondered, didn’t evolution allow us to skip childhood altogether, “to jump catlike from our juvenile to our adult [and] productive stage of life?” It is indeed extraordinary how long it takes for humans to develop into capable adults.

How does this individual development work? The question of how a tiny clump of cells slowly becomes a person with a particular physique, intellect, personality, and emotional reservoir has long challenged scientists, let alone the general public. For centuries, this question has been phrased in terms of nature vs. nurture—trying to determine what portion of development is dictated by inborn, innate forces such as our genes versus what portion is shaped by experience. In recent decades, evidence has updated our understanding of genes and their relationship to the individual; further research on fetal development, neuroplasticity, the functional organization of the brain, the nature of intelligence, and studies of expertise have all come together to suggest that the old debates about nature and nurture should be thrown out, in favor of something new—a unified “developmental systems” perspective.

The new understanding starts with a new conception of the gene. Out of Gregor Mendel’s 19th-century pea-plant experiments came a century-long popular and scientific belief that genes were effectively blueprints with elaborate predesigned instructions for all traits—eye color, thumb size, mathematical aptitude, musical sensitivity, and so on. But with increasing knowledge about the actual mechanics of development, the orthodox Mendelian view has been thoroughly upgraded into a more sophisticated understanding of how traits actually emerge. Genes are not like robot actors who always say the same lines in the exact same way. Instead, they interact with their surroundings from moment to moment in complex and interesting ways.

The developmental systems view also builds on recent advances in behavioral neuroscience. Researchers historically viewed the brain as a modular system, hardwired for specialized abilities. But recent data have revealed tremendous plasticity, particularly

early in development. At one extreme, infant plasticity can enable complex cognitive functioning even in infants with atypically developing brains or after substantial brain damage. Simply put, then, none of us is hardwired, programmed, or preordained. Each of us develops.

Although a wealth of scientific data are consistent with the developmental systems perspective, the nature-nurture debate continues to be the predominant framework for talking about development. This is the case within scientific disciplines where words like ‘innate’ and ‘inborn’ are still commonly used (see, e.g., Root, Denny, Hen & Axel, 2014. *Nature*, 515, 269-273); it is also the case beyond academia where journalists, practitioners, policy makers, teachers, and parents continue to think about a person’s traits as a direct result of genes (“Infidelity lurks in your genes”, *New York Times*, May 22, 2015).

In our view, the scientific evidence for a new view of development is overwhelming. Why, then, does the centuries-old framing persist? Is this a question of a lack of convincing scientific evidence? That is, given more time and accumulating knowledge, more and more people will come to accept and espouse a developmental systems perspective? Or is this fundamentally a question of communication? That is, is it the case that the nature-nurture framework is easier to describe, easier to sell? Perhaps the communication advantage of the nativist perspective underlies its impressive resilience.

This dialog raises timely questions for us. In an effort to improve communication of the developmental systems perspective, we have undertaken an ambitious project to be published in the spring of 2016 as part of Wiley’s WIREs series. Our goal was to create an on-line collection that presents the developmental systems perspective to a broad audience in an accessible and scientifically rigorous way. The collection offers an overview of the developmental systems perspective, spanning molecular and cultural levels, from nanoseconds to millennia, addressing both development and evolution. The themes explored should be of interest to students as well as parents, teachers, and policy makers who wish to understand and foster the development of individual children. Wiley has generously agreed to make the collection free for the first year and, wherever possible, the contributors to the collection

have strived to present their complex material using straightforward language. This collection represents one perspective on

how to move beyond the false nature-nurture dichotomy. We look forward to hearing other perspectives—what do you think?

IEEE TAMD Table of Contents

Volume 7, Issue 3, September 2015

Editorial Announcing the Title Change of the IEEE Transactions on Autonomous Mental Development in 2016

A. Salah

Guest Editorial Multimodal Modeling and Analysis Informed by Brain Imaging—Part I

J. Han, T. Liu, C. C. Guo, J. Weng

Investigating Critical Frequency Bands and Channels for EEG-Based Emotion Recognition with Deep Neural Network

W.-L. Zheng, B.-L. Lu

To investigate critical frequency bands and channels, this paper introduces deep belief networks (DBNs) to constructing EEG-based emotion recognition models for three emotions: positive, neutral and negative. We develop an EEG dataset acquired from 15 subjects. Each subject performs the experiments twice at the interval of a few days. DBNs are trained with differential entropy features extracted from multichannel EEG data. We examine the weights of the trained DBNs and investigate the critical frequency bands and channels. Four different profiles of 4, 6, 9, and 12 channels are selected. The recognition accuracies of these four profiles are relatively stable with the best accuracy of 86.65%, which is even better than that of the original 62 channels. The critical frequency bands and channels determined by using the weights of trained DBNs are consistent with the existing observations. In addition, our experiment results show that neural signatures associated with different emotions do exist and they share commonality across sessions and individuals. We compare the performance of deep models with shallow models. The average accuracies of DBN, SVM, LR, and KNN are 86.08%, 83.99%, 82.70%, and 72.60%, respectively.

What Strikes the Strings of Your Heart?—Multi-Label Dimensionality Reduction for Music Emotion Analysis via Brain Imaging

Yang Liu, Yan Liu, C. Wang, X. Wang, P. Zhou, G. Yu, G. ; K.C.C. Chan,

After 20 years of extensive study in psychology, some musical factors have been identified that can evoke certain kinds of emotions. However, the underlying mechanism of the relationship between music and emotion remains unanswered. This paper intends to find the genuine correlates of music emotion by exploring a systematic and quantitative framework. The task is formulated as a dimensionality reduction problem, which seeks the complete and compact feature set with intrinsic correlates for the given objectives. Since a song generally elicits more than one emotion, we explore dimensionality reduction techniques for multi-label classification. One challenging problem is that the hard label cannot represent the extent of the emotion and it is also difficult to ask the subjects to quantize their feelings. This work tries utilizing the electroencephalography (EEG) signal to solve this challenge. A learning scheme called EEG-based emotion smoothing (E2S) and a bilinear multi-emotion similarity preserving embedding (BME-SPE) algorithm are proposed. We validate the effectiveness of the proposed framework on standard dataset CAL-500. Several influential correlates have been identified and the classification via those correlates has achieved good performance. We build a Chinese music dataset according to the identified correlates and find that the music from different cultures may share similar emotions.

Emotion Recognition with the Help of Privileged Information

S. Wang, Y. Zhu, L. Yue, Q. Ji

In this article, we propose a novel approach to recognize emotions with the help of privileged information, which is only available during training, but not available during testing. Such additional information can be exploited during training to construct a better classifier. Specifically, we recognize audience's emotion from EEG signals with the help of the stimulus videos, and tag videos' emotions with the aid of electroencephalogram (EEG) signals. First, frequency features are extracted from EEG signals and audio/visual features are extracted from video stimulus.

Second, features are selected by statistical tests. Third, a new EEG feature space and a new video feature space are constructed simultaneously using canonical correlation analysis (CCA). Finally, two support vector machines (SVM) are trained on the new EEG and video feature spaces respectively. During emotion recognition from EEG, only EEG signals are available, and the SVM classifier obtained on EEG feature space is used; while for video emotion tagging, only video clips are available, and the SVM classifier constructed on video feature space is adopted. Experiments of EEG-based emotion recognition and emotion video tagging are conducted on three benchmark databases, demonstrating that video content, as the context, can improve the emotion recognition from EEG signals and EEG signals available during training can enhance emotion video tagging.

Decoding Semantics Categorization during Natural Viewing of Video Streams

X. Hu, L. Guo, J. Han, T. Liu

Exploring the functional mechanism of the human brain during semantics categorization and subsequently leverage current semantics-oriented multimedia analysis by functional brain imaging have been receiving great attention in recent years. In the field, most of existing studies utilized strictly controlled laboratory paradigms as experimental settings in brain imaging data acquisition. They also face the critical problem of modeling functional brain response from acquired brain imaging data. In this paper, we present a brain decoding study based on sparse multinomial logistic regression (SMLR) algorithm to explore the brain regions and functional interactions during semantics categorization. The setups of our study are two folds. First, we use naturalistic video streams as stimuli in functional magnetic resonance imaging (fMRI) to simulate the complex environment for semantics perception that the human brain has to process in real life. Second, we model brain responses to semantics categorization as functional interactions among large-scale brain networks. Our experimental results show that semantics categorization can be accurately predicted by both intrasubject and intersubject brain decoding models. The brain responses identified by the decoding model reveal that a wide range of brain regions and functional interactions are recruited during semantics categorization. Especially, the working memory system exhibits significant contributions. Other substantially involved brain systems include emotion, attention, vision and language systems.

An Iterative Approach for EEG-Based Rapid Face Search: A Refined Retrieval by Brain Computer Interfaces

Y. Wang, L. Jiang, Y. Wang, B. Cai, Y. Wang, W. Chen, S. Zhang, X. Zheng

Recent face recognition techniques have achieved remarkable successes in fast face retrieval on huge image datasets. But the performance is still limited when large illumination, pose, and facial expression variations are presented. In contrast, the human brain has powerful cognitive capability to recognize faces and demonstrates robustness across viewpoints, lighting conditions, even in the presence of partial occlusion. This paper proposes a closed-loop face retrieval system that combines the state-of-the-art face recognition method with the powerful cognitive function of the human brain illustrated in electroencephalography signals. The system starts with a random face image and outputs the ranking of all of the images in the database according to their similarity to the target individual. At each iteration, the single trial event related potentials (ERP) detector scores the user's interest in rapid serial visual presentation paradigm, where the presented images are selected from the computer face recognition module. When the system converges, the ERP detector further refines the lower ranking to achieve better performance. In total, 10 subjects participated in the experiment, exploring a database containing 1,854 images of 46 celebrities. Our approach outperforms existing methods with better average precision, indicating human cognitive ability complements computer face recognition and contributes to better face retrieval.

Age Effect in Human Brain Responses to Emotion Arousing Images: The EEG 3D-Vector Field Tomography Modeling Approach

C.D. Papadaniil, V.E. Kosmidou, A.C. Tsolaki, L.J. Hadjileontiadis, M. Tsolaki, I. Y. Kompatsiaris

Understanding of the brain responses to emotional stimulation remains a great challenge. Studies on the aging effect in neural activation report controversial results. In this paper, pictures of two classes of facial affect, i.e., anger and fear, were presented to young and elderly participants. High-density 256-channel EEG data were recorded and an innovative methodology was used to map the activated brain state at the N170 event-related potential component. The methodology, namely 3D Vector Field Tomography, reconstructs the electrostatic field within the head volume and requires no prior modeling of the individual's brain. Results showed that the elderly exhibited greater N170

amplitudes, while age-based differences were also observed in the topographic distribution of the EEG recordings at the N170 component. The brain activation analysis was performed by adopting a set of regions of interest. Results on the maximum activation area appeared to be emotion-specific; the anger emotional conditions induced the maximum activation in the inferior frontal gyrus, while fear activated more the superior temporal gyrus. The approach used here shows the potential of the proposed computational model to reveal the age effect on the brain activation upon emotion arousing images, which could be further transferred to the design of assistive clinical applications.

Perceptual Experience Analysis for Tone-mapped HDR Videos Based on EEG and Peripheral Physiological Signals

S.-E. Moon, J.-S. Lee

High dynamic range (HDR) imaging has been attracting much attention as a technology that can provide immersive experience. Its ultimate goal is to provide better quality of experience (QoE) via enhanced contrast. In this paper, we analyze perceptual experience of tone-mapped HDR videos both explicitly by conducting a subjective questionnaire assessment and implicitly by using EEG and peripheral physiological signals. From the results of the subjective assessment, it is revealed that tone-mapped HDR videos are more interesting and more natural, and give better quality than low dynamic range (LDR) videos. Physiological signals were recorded during watching tone-mapped HDR and LDR videos, and classification systems are constructed to explore perceptual difference captured by the physiological signals. Significant difference in the physiological signals is observed between tone-mapped HDR and LDR videos in the classification under both a subject-dependent and a subject-independent scenarios. Also, significant difference in the signals between high versus low perceived contrast and overall quality is detected via classification under the subject-dependent scenario. Moreover, it is shown that features extracted from the gamma frequency band are effective for classification.

Predicting Purchase Decisions Based on Spatio-Temporal Functional MRI Features Using Machine Learning

Y. Wang, V. Chattaraman, K. Hyejeong, G. Deshpande

Machine learning algorithms allow us to directly predict brain states based on functional magnetic resonance imaging (fMRI) data. In this study, we demonstrate the application of this framework to neuromarketing by predicting purchase decisions from spatio-temporal fMRI data. A sample of 24 subjects were shown product images and asked to make decisions of whether to buy them or not while undergoing fMRI scanning. Eight brain regions which were significantly activated during decision-making were identified using a general linear model. Time series were extracted from these regions and input into a recursive cluster elimination based support vector machine (RCE-SVM) for predicting purchase decisions. This method iteratively eliminates features which are unimportant until only the most discriminative features giving maximum accuracy are obtained. We were able to predict purchase decisions with 71% accuracy, which is higher than previously reported. In addition, we found that the most discriminative features were in signals from medial and superior frontal cortices. Therefore, this approach provides a reliable framework for using fMRI data to predict purchase-related decision-making as well as infer its neural correlates.

A Robust Gradient-Based Algorithm to Correct Bias Fields of Brain MR Images

Q. Ling, Z. Li, Q. Huang, X. Li

We developed a novel algorithm to estimate bias fields from brain magnetic resonance (MR) images using a gradient-based method. The bias field is modeled as a multiplicative and slowly varying surface. We fit the bias field by a low-order polynomial. The polynomial's parameters are directly obtained by minimizing the sum of square errors between the gradients of MR images (both in the x-direction and y-direction) and the partial derivatives of the desired polynomial in the log domain. Compared to the existing retrospective algorithms, our algorithm combines the estimation of the gradient of the bias field and the reintegration of the obtained gradient polynomial together so that it is more robust against noise and can achieve better performance, which are demonstrated through experiments with both real and simulated brain MR images.

Volume 7, Issue 4, December 2015**Guest Editorial Multimodal Modeling and Analysis Informed by Brain Imaging—Part II**

J. Han, T. Liu, C. C. Guo, J. Weng

Types, Locations, and Scales from Cluttered Natural Video and Actions

X. Song, W. Zhang, J. Weng

We model the autonomous development of brain-inspired circuits through two modalities—video stream and action stream that are synchronized in time. We assume that such multimodal streams are available to a baby through inborn reflexes, self-supervision, and caretaker's supervision, when the baby interacts with the real world. By autonomous development, we mean that not only that the internal (inside the "skull") self-organization is fully autonomous, but the developmental program (DP) that regulates the computation of the network is also task nonspecific. In this work, the task-nonspecificity is reflected by the fact that the actions associated with an attended object in a cluttered, natural, and dynamic scene is taught after the DP is finished and the "life" has begun. The actions correspond to neuronal firing patterns representing object type, object location and object scale, but learning is directly from unsegmented cluttered scenes. Along the line of where-what networks (WWN), this is the first one that explicitly models multiple "brain" areas—each for a different range of object scales. Among experiments, large natural video experiments were conducted. To show the power of automatic attention in unknown cluttered backgrounds, the last experimental group demonstrated disjoint tests in the presence of large within-class variations (object 3-D-rotations in very different unknown backgrounds), but small between-class variations (small object patches in large similar and different unknown backgrounds), in contrast with global classification tests such as ImageNet and Atari Games.

Randomized Structural Sparsity-Based Support Identification with Applications to Locating Activated or Discriminative Brain Areas: A Multicenter Reproducibility Study

Y. Wang, S. Zhang, J. Zheng, Heng Chen, Huafu Chen

In this paper, we focus on how to locate the relevant or discriminative brain regions related with external stimulus or certain mental disease, which is also called support identification, based on the neuroimaging data. The main difficulty lies in the extremely high dimensional voxel space and relatively few training samples, easily resulting in an unstable brain region discovery (or called feature selection in context of pattern recognition). When the training samples are from different centers and have between-center variations, it will be even harder to obtain a reliable and consistent result. Corresponding, we revisit our recently proposed algorithm based on stability selection and structural sparsity. It is applied to the multicenter MRI data analysis for the first time. A consistent and stable result is achieved across different centers despite the between-center data variation while many other state-of-the-art methods such as two sample t-test fail. Moreover, we have empirically showed that the performance of this algorithm is robust and insensitive to several of its key parameters. In addition, the support identification results on both functional MRI and structural MRI are interpretable and can be the potential biomarkers.

Beyond Subjective Self-Rating: EEG Signal Classification of Cognitive Workload

P. Zarjam, J. Epps, N. H. Lovell

Cognitive workload is an important indicator of mental activity that has implications for human-computer interaction, biomedical and task analysis applications. Previously, subjective rating (self-assessment) has often been a preferred measure, due to its ease of use and relative sensitivity to the cognitive load variations. However, it can only be feasibly measured in a post-hoc manner with the user's cooperation, and is not available as an online, continuous measurement during the progress of the cognitive task. In this paper, we used a cognitive task inducing seven different levels of workload to investigate workload discrimination using electroencephalography (EEG) signals. The entropy, energy, and standard deviation of the wavelet coefficients extracted from the segmented EEGs were found to change very consistently in accordance with the induced load, yielding strong significance in statistical tests of ranking accuracy. High accuracy for subject-independent multichannel classification among seven load levels was achieved, across the twelve subjects studied. We compare these results with alternative measures such as performance, subjective ratings, and reaction time (response time) of the subjects and compare their reliability with the EEG-based method introduced. We also investigate test/re-test reliability of the

recorded EEG signals to evaluate their stability over time. These findings bring the use of passive brain-computer interfaces (BCI) for continuous memory load measurement closer to reality, and suggest EEG as the preferred measure of working memory load.

Local Multimodal Serial Analysis for Fusing EEG-fMRI: A New Method to Study Familial Cortical Myoclonic Tremor and Epilepsy

L. Dong, P. Wang, Y. Bin, J. Deng, Y. Li, L. Chen, C. Luo, D. Yao

Integrating information of neuroimaging multimodalities, such as electroencephalography (EEG) and functional magnetic resonance imaging (fMRI), has become popularly for investigating various types of epilepsy. However, there are also some problems for the analysis of simultaneous EEG-fMRI data in epilepsy: one is the variation of HRFs, and another is low signal-to-noise ratio (SNR) in the data. Here, we propose a new multimodal unsupervised method, termed local multimodal serial analysis (LMSA), which may compensate for these deficiencies in multimodal integration. A simulation study with comparison to the traditional EEG-informed fMRI analysis which directly implemented the general linear model (GLM) was conducted to confirm the superior performance of LMSA. Then, applied to the simultaneous EEG-fMRI data of familial cortical myoclonic tremor and epilepsy (FCMTE), some meaningful information of BOLD changes related to the EEG discharges, such as the cerebellum and frontal lobe (especially in the inferior frontal gyrus), were found using LMSA. These results demonstrate that LMSA is a promising technique for exploring various data to provide integrated information that will further our understanding of brain dysfunction.

Discriminating Bipolar Disorder From Major Depression Based on SVM-FoBa: Efficient Feature Selection With Multimodal Brain Imaging Data

N.-F. Jie, M.-H. Zhu, X.-Y. Ma, E.A. Osuch, M. Wammes, J. Theberge, H.-D. Li, Y. Zhang, T.-Z. Jiang, J. Sui, V.D. Calhoun

Discriminating between bipolar disorder (BD) and major depressive disorder (MDD) is a major clinical challenge due to the absence of known biomarkers; hence a better understanding of their pathophysiology and brain alterations is urgently needed. Given the complexity, feature selection is especially important in neuroimaging applications, however, feature dimension and model understanding present serious challenges. In this study, a novel feature selection approach based on linear support vector machine with a forward-backward search strategy (SVM-FoBa) was developed and applied to structural and resting-state functional magnetic resonance imaging data collected from 21 BD, 25 MDD and 23 healthy controls. Discriminative features were drawn from both data modalities, with which the classification of BD and MDD achieved an accuracy of 92.1% (1000 bootstrap resamples). Weight analysis of the selected features further revealed that the inferior frontal gyrus may characterize a central role in BD-MDD differentiation, in addition to the default mode network and the cerebellum. A modality-wise comparison also suggested that functional information outweighs anatomical by a large margin when classifying the two clinical disorders. This work validated the advantages of multimodal joint analysis and the effectiveness of SVM-FoBa, which has potential for use in identifying possible biomarkers for several mental disorders.

Design of a Multimodal EEG-based Hybrid BCI System with Visual Servo Module

F. Duan, D. Lin, W. Li, Z. Zhang

Current EEG-based brain-computer interface technologies mainly focus on how to independently use SSVEP, motor imagery, P300, or other signals to recognize human intention and generate several control commands. SSVEP and P300 require external stimulus, while motor imagery does not require it. However, the generated control commands of these methods are limited and cannot control a robot to provide satisfactory service to the user. Taking advantage of both SSVEP and motor imagery, this paper aims to design a hybrid BCI system that can provide multimodal BCI control commands to the robot. In this hybrid BCI system, three SSVEP signals are used to control the robot to move forward, turn left, and turn right; one motor imagery signal is used to control the robot to execute the grasp motion. In order to enhance the performance of the hybrid BCI system, a visual servo module is also developed to control the robot to execute the grasp task. The effect of the entire system is verified in a simulation platform and a real humanoid robot, respectively. The experimental results show that all of the subjects were able to successfully use this hybrid BCI system with relative ease.

EEG-Based Perceived Tactile Location Prediction

D. Wang, Y. Liu, D. Hu, G. Blohm

Previous studies have attempted to investigate the peripheral neural mechanisms implicated in tactile perception, but the neurophysiological data in humans involved in tactile spatial location perception to help the brain orient the body and interact with its surroundings are not well understood. In this paper, we use single-trial electroencephalogram (EEG) measurements to explore the perception of tactile stimuli located on participants' right forearm, which were approximately equally spaced centered on the body midline, 2 leftward and 2 rightward of midline. An EEG-based signal analysis approach to predict the location of the tactile stimuli is proposed. Offline classification suggests that tactile location can be detected from EEG signals in single trial (four-class classifier for location discriminate can achieve up to 96.76%) with a short response time (600 milliseconds after stimulus presentation). From a human-machine-interaction (HMI) point of view, this could be used to design a real-time reactive control machine for patients, e.g., suffering from hypoesthesia.

An Adaptive Motion-Onset VEP-Based Brain-Computer Interface

R. Zhang, P. Xu, R. Chen, T. Ma, X. Lv, F. Li, P. Li, T. Liu, D. Yao

Motion-onset visual evoked potential (mVEP) has been recently proposed for EEG-based brain-computer interface (BCI) system. It is a scalp potential of visual motion response, and typically composed of three components: P1, N2, and P2. Usually several repetitions are needed to increase the signal-to-noise ratio (SNR) of mVEP, but more repetitions will cost more time thus lower the efficiency. Considering the fluctuation of subject's state across time, the adaptive repetitions based on the subject's real-time signal quality is important for increasing the communication efficiency of mVEP-based BCI. In this paper, the amplitudes of the three components of mVEP are proposed to build a dynamic stopping criteria according to the practical information transfer rate (PITR) from the training data. During online test, the repeated stimulus stopped once the predefined threshold was exceeded by the real-time signals and then another circle of stimulus newly began. Evaluation tests showed that the proposed dynamic stopping strategy could significantly improve the communication efficiency of mVEP-based BCI that the average PITR increases from 14.5 bit/min of the traditional fixed repetition method to 20.8 bit/min. The improvement has great value in real-life BCI applications because the communication efficiency is very important.