IEEE CIS

CDS NEWSLETTER

The Newsletter of the Technical Committee on Cognitive and Developmental Systems

Developmental Robotics
Machine Intelligence
Neuroscience
Psychology

Volume 15, number 1
Spring 2018

Editorial

The Mysteries of Human Curiosity-Driven Learning and the Challenges of Translational Educational Sciences

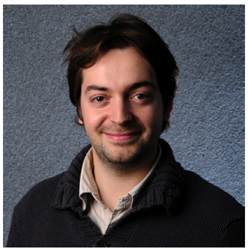

Pierre-Yves Oudeyer

Inria and Ensta ParisTech, France
Editor

pierre-yves.oudeyer@inria.fr

Curiosity-driven learning is probably one of the most fundamental mechanisms in human learning, and yet it is also probably one of the least understood. Broadly construed as spontaneous exploration and engagement with activities or material without any extrinsic goal (as opposed to searching for information useful for an extrinsic goal), many mysteries remain to be uncovered. What are the causal links between curiosity and learning? How does prior knowledge about a topic or an activity relates to curiosity about this topic? What is the role of curiosity in life-span development? Can human curiosity explain the apparently unique tendency of humans for extreme specialization? Reversely, how do different forms of curiosity (diversive or specific) evolve as children grow up and become adults? While early computational models of curiosity propose theoretical approaches to understand their cognitive mechanisms, how can we understand the affective/emotional dimensions of curiosity? And how has the linguistic concept of "curiosity" evolved in occidental culture?

These outstanding questions are discussed and developed in this issue's dialog of the newsletter, initiated by Celeste Kidd, with inspiring contributions from researchers in neuroscience, psychology, computational modelling and AI, education and philosophy: Elizabeth Bonawitz, Maya Zhe Wang, Brian Sweis, Benjamin Hayden, Susan Engel, Abigail Hsiung, Shabnam Hakimi, Alison Adcock, Moritz Daum, Arjun Shankar, Tobias Hauser, Goren Gordon and Perry Zurn.

Understanding the fundamental principles of child learning is also a key pre-requisite for designing principled educational technologies and interventions. For example, understanding how to stimulate and leverage curiosity to implement motivating educational activities that have long term learning impact is a paramount challenge. However, as argued in the new dialog initiation proposed by Georges Kachergis, other very difficult challenges need to be addressed for educational impact: how can one achieve efficiently "translational educational sciences" and get these principles used in real-world large-scale educational technologies? In this dialog entitled "Leveraging adaptive games to learn how to help children learn effectively", Georges Kachergis highlights challenges related to collaborations between cognitive scientists and game developers, how to deploy real world experiments, and how to enable scientific understanding when many variables cannot easily be controlled? Those of you interested in reacting to this dialog initiation are welcome to submit a response by December 15th, 2018. 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 Kasmarik

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

Dear colleagues, I'm writing this message from Tokyo on the eve of ICDL-Epirob, 2018, one of our community's premier annual events. I'm looking forward to the coming week of workshops and talks, including the keynote talks on Developmental Theories of AI by Prof. Oliver Brock, What can we learn from the brain from AI and robotics by Prof. Kenji Doya, Embodiment and Human Development by Prof. Peter J. Marshall and AI and robotics at Sony by Mr. Masahiro Fujita.

I'm also looking forward to the oral and poster sessions across a broad range of topics in epigenetic robotics and developmental learning, and the opportunity to meet with everyone at the conference.

Later this year I'm looking forward to the completion of the Frontiers special issue on Intrinsically Motivated Open-Ended Learning in Autonomous Robots, an ongoing project that has followed on from the International Workshop on Intrinsically Motivated Open Ended Learning in Rome late last year. This will also bring together work on a wide range of topics from the community.

I know that there are also many other ongoing projects, workshops and special issues, and I commend to you the work of the community.

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

Dialogue

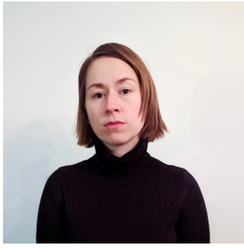

Celeste Kidd

Assistant Professor of Psychology, University of California, Berkeley

celestekidd@gmail.com

Curiosity as Driver of Extreme Specialization in Humans

The features that make us uniquely and distinctly human have been of interest to many people, from psychologists to philosophers to religious scholars, for centuries. Typical candidate traits include things like speech (Lieberman, 1991), upright posture (Clarke & Tobias, 1995), protracted childhoods (Jolly, 1972), helpless infants (Piantadosi & Kidd, 2016), sophisticated social cooperation (Melis & Semmann, 2010), and creativity (Carruthers, 2002).

There is, however, an essential human trait that has received far less recognition: the capacity for extreme specialization. Many humans spend a lifetime perfecting a single niche skill, such as a musical instrument, art medium, or style of dance. Others specialize in trades with economic roles (e.g., butchers, bakers, and candlestick makers). And while some other species exhibit certain forms of specialization—ants, for example, exhibit increased task specialization as the colony size increases (Amador-Vargas et al., 2015)—none approach the breadth and depth of specialization found in humans. In particular, specialization in species usually seems to hinge on abilities that are directly relevant to survival. Human specialization, in contrast, knows no limits or bounds and seems applicable to virtually any domain of existence. Here we will argue that this extreme specialization is enabled in large part due to key mechanisms within the human attentional system—specifically those mechanisms that bias learners towards material for which they already possess some background knowledge. More broadly, this extreme specialization is enabled by the driving pressures that underlie human curiosity.

Curiosity can be thought of as the force behind the acquisition of new knowledge (James, 1913; Pavlov, 1927; Skinner, 1938; Oudeyer & Kaplan, 2007; Gottlieb et al., 2013). It is a strong determinant of how we spend our days, and influences not just our intellectual interests, but also a myriad of recreational decisions, from who we speak to and what we discuss, to what we listen to and watch, to what we fixate on in a scene and what we learn about the world. It is a key driving force behind the grandest human innovations, yet less sophisticated, purpose-specific forms of curiosity can be observed in more primitive intelligences (e.g., C. elegans). Curiosity, or intrinsic motivation, is likely a necessary feature of intelligent systems generally. Even robotic and artificial intelligence systems must possess a mechanism to seek out and learn material that is relevant to their present and future goals (Oudeyer & Kaplan, 2007).

Human curiosity is known to relate to our existing knowledge. For example, work from the infant attention literature suggests that infants prefer novel stimuli, defined as distinct from what the infant already knows (Sokolov, 1963) or partially encoded representations over either entirely known or entirely novel ones (Dember & Earl, 1957; Kinney & Kagan, 1976; Berlyne, 1978; Kidd et al., 2013). More contemporary theories observe that curiosity is triggered when a gap is detected between what a learner currently knows, and what they could know (Loewenstein, 1994). This suggests the involvement of metacognition, since a learner must first identify that there is a gap to be filled before curiosity should be piqued. Yet little work to date has explored the relationship between metacognitive processes and curiosity. Are people who possess more metacognitive abilities pertaining to their own knowledge more curious? Can you make someone more curious by calling attention to what they do not know?

While we know that there exists some relationship between existing knowledge about a stimulus and the learner's degree of interest in that stimulus, we still do not fully understand precisely how those two factors relate to each other, nor do we understand the cognitive or neural mechanisms underlying how and why the learner's curiosity is piqued (for a review of what we don't know, see Hayden & Kidd, 2015). For example, we do not understand how neural reward systems treat information and weigh it in decision-making, though it is clear that humans and monkeys are willing to sacrifice some reward to gain even useless information (Blanchard, Hayden, & Bromberg-Martin, 2015). Is there a common currency for reward and information, and how is the value of information determined, represented, and integrated neurally?

We have limited evidence to suggest that being in a curious state could facilitate learning (Gruber et al., 2014; Stahl & Feigenson, 2015); however, we also have evidence that learners are more curious when they possess information that is partially encoded, and thus on the verge of being learned (Kang et al., 2009). Thus, we must be sensitive to the fact that some of the apparent boosts to learning attributed to curiosity in the literature may have the direction of causality wrong—being on the verge-of-learning may induce greater curiosity, rather than curiosity inducing better learning. How do we understand curiosity and the biological mechanisms underlying it in a way that reasonably accounts for these two apparently opposing causal mechanisms?

What is the purpose of this curiosity system, and why does it yield the sort of specialization that we see in humans but not other species? How does it function, and what purposes does it serve? Why are there humans that become compelled to acquire information about fictitious worlds (e.g., Harry Potter, Star Wars)? What might be the connection between curiosity, creativity, and specialization?

Amador-Vargas, S., Gronenberg, W., Wcislo, W. T., & Mueller, U. (2015). Specialization and Group Size: Brain and behavioural correlates of colony size in ants lacking morphological castes. Proceedings of the Royal Society B, 282(1801), 20142502.
Berlyne, D. E. (1978). Curiosity and learning. Motivation and Emotion, 2(2), 97-175.
Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. Neuron, 85(3), 602-614.
Carruthers, P. (2002). The cognitive functions of language. Behavioral and Brain Sciences, 25(6), 657-674.
Clarke, R. J., & Tobias, P. V. (1995). Sterkfontein Member 2 foot bones of the oldest South African hominid. Science, 269(5223), 521.
Dember, W. N., & Earl, R. W. (1957). Analysis of exploratory, manipulatory, and curiosity behaviors. Psychological Review, 64(2), 91.
Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information-seeking, Curiosity, and Attention: Computational and Neural Mechanisms. Trends in Cognitive Sciences, 17(11), 585-593.
Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. Neuron, 84(2), 486-496.
James, W. (1913). The Principles of Psychology, Volume II. New York, NY, United States: Henry Holt and Company.
Jolly, A. (1972). The Evolution of Primate Behavior. New York, NY, United States: MacMillan.
Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T. Y., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. Psychological Science, 20(8), 963-973.
Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. Neuron, 88(3), 449-460.
Kidd, C., Piantadosi, S.T., & Aslin, R.N. (2012.) The Goldilocks Effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. PLOS ONE, 7(5), e36399.
Kinney, D. K., & Kagan, J. (1976). Infant attention to auditory discrepancy. Child Development, 155-164.
Lieberman, P. (1991). Uniquely Human: The evolution of speech, thought, and selfless behavior. Cambridge, MA, United States: Harvard University Press.
Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological Bulletin, 116(1), 75.
Melis, A. P., & Semmann, D. (2010). How is human cooperation different?. Philosophical Transactions of the Royal Society of London B: Biological Sciences, 365(1553), 2663-2674.
Oudeyer, P. Y., & Kaplan, F. (2007). What is intrinsic motivation? A typology of computational approaches. Frontiers in Neurorobotics, 1, 6.
Pavlov, I.P. (1927). Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex. Oxford, England: Oxford University Press.
Piantadosi, S. T., & Kidd, C. (2016). Extraordinary intelligence and the care of infants. Proceedings of the National Academy of Sciences, 113(25), 6874-6879.
Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. Science, 348(6230), 91-94.
Skinner, B. F. (1938). The Behavior of Organisms: An Experimental Analysis. Oxford, England: Appleton–Century.
Sokolov, E. (1963). Perception and the Conditioned Reflex. Oxford, England: Pergamon Press.

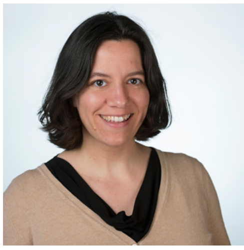

Elizabeth Bonawitz

Computational Cognitive Development Lab, Rutgers University, Newark, USA

elizabeth.bonawitz@rutgers.edu

Towards Computational Models of Curiosity in Cognitive Development

Kidd raises an interesting chicken and egg problem pertaining to curiosity and the specialization of knowledge. One the one hand, epistemic curiosity likely depends on some prior knowledge being in place; it is piqued given the realization of a gap between known and unknown information (Lowenstein, 1994). On the other hand, curiosity drives knowledge enrichment. For example, binding association and intervention is a key component of causal knowledge in early development (Bonawitz et al., 2010). Curiosity may help drive this link between observed associative relationships and the outcomes of our interventions by motivating action. One way to gain insight into these self-reinforcing roles of curiosity and developing knowledge is to examine them in early childhood.

Computational modeling further helps to make precise the role of prior knowledge and opportunities for information gain. By bridging these approaches we can come to better understand the contributions of curiosity to our uniquely human traits. There is a long tradition of modeling curiosity (e.g. see Oudeyer 2018 for a review) and also of modeling the role of prior knowledge in human learning (e.g. see Tenenbaum, Griffiths, & Kemp, 2006). Taking the theory-based probabilistic perspective as a starting point, we might conceptualize of curiosity as an artifact that falls out from "running" certain inferential processes. Indeed, we can think of the mind as carrying out simulations (Battaglia, Hamrick, & Tenenbaum, 2013), search (Ullman, Goodman, & Tenenbaum, 2012), and sampling (Bonawitz et al., 2014; Griffiths, Vul, & Sanborn, 2012) over intuitive theories. Curiosity may exist as a state during this inferential process and be greatest when information is likely to be gained, when information will likely resolve conflict or uncertainty, when the reward of knowledge is high and that cost of carrying out the information seeking action is low, and so forth.

Curiosity is also often encoded simply as a drive or utility in a learning system, but modeling can also help specify the causes of curiosity and quantify their contribution towards this drive in early development. For example, models can specify the role of prior knowledge and various utilities, providing a framework to build a utility calculus of curiosity. Recent research suggests that even very young children are already capable of carrying out this intuitive calculus, and that curiosity and prior knowledge are deeply intertwined. For example, young children are more motivated to explore when events violate the predictions of intuitive theories (Bonawitz et al., 2012; Stahl & Feigenson, 2015), suggesting that even the young mind is driven to reduce uncertainty and learn more following a conflict of beliefs and evidence. Preschoolers are also sensitive to information gain, exploring over exploiting rewards when the knowledge gained will serve later use (Bonawitz, Bass, & Lapidow, 2018). Curiosity may be piqued when it is brought to the attention of a learner that knowledge is incomplete. For example, research with preschoolers has found that pedagogical questions simultaneously point to the importance of particular features, while also encouraging further exploration and discovery (Yu, Landrum, Bonawitz, & Shafto, 2018).

Computational approaches provide an important starting point for understanding the intertwined roles of curiosity and knowledge acquisition, but current frameworks do not yet have a meaningful way to incorporate affect. Discovery can feel good and thus rewards curiosity (e.g. "Explanation as Orgasm, Gopnik, 2000), but affect also cyclically drives exploration (e.g. as in dopaminergic neuroscientific models of prediction error, learning, and further motivation, Pessiglione et al., 2006). Awe has been defined as a simultaneous affective and cognitive experience related to curiosity (Valdesolo & Graham, 2014). Being "induced" to feel awe stimulates curious exploration in preschoolers (Colantonio & Bonawitz, 2018). Despite the clear role of emotion in experiencing reward and motivating further learning, we currently do not have a way to capture manipulations of affect in our computational models of curiosity. That is, just as emotions can give us a sense of success or failure critical to learning, we can also hard code rewards and costs in our models, but this does not necessarily inform our understanding of curiosity as an affective drive. Physiological and emotional components may be the defining features of curiosity, yet they are currently ill-defined.

Computational models of learning in early childhood may help us to formalize and understand the intertwined roles of curiosity and specialization of knowledge. However, this approach also makes apparent that our models must consider emotional states. Children are both intensely emotional and curious creatures, providing yet another benefit of studying curiosity in early childhood. Perhaps by recognizing this gap in our current understanding, we will be more likely to curiously seek out new models and empirical studies, helping further refine our knowledge.

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. Proceedings of the National Academy of Sciences, 110(45), 18327-18332.
Bonawitz, E., Bass, L., & Lapidow, E. (2018) Choosing to Learn: Evidence Evaluation for Active Learning and Teaching in Early Childhood. In: Saylor M., Ganea P. (eds) Active Learning from Infancy to Childhood. Springer, Cham, 213-231.
Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. Cognitive psychology, 64(4), 215-234.
Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: Sampling in cognitive development. Trends in Cognitive Sciences, 18(10), 497-500.
Bonawitz, E. B., Ferranti, D., Saxe, R., Gopnik, A., Meltzoff, A. N., Woodward, J., & Schulz, L. E. (2010). Just do it? Investigating the gap between prediction and action in toddlers' causal inferences. Cognition, 115(1), 104-117.
Colantonio, J., & Bonawitz, E. (in press) Awesome play: Awe increases preschoolers exploration and discovery. In Kalish, C., Rau, M., Zhu, J., & Rogers, T.T. (Eds.) Proceedings of the 40th Annual Conference of the Cognitive Science Society. Madison, WI: Cognitive Science Society.
Gopnik, A. (2000). Explanation as orgasm and the drive for causal knowledge: The function, evolution, and phenomenology of the theory formation system. In F. Keil & R. Wilson (Eds.), Cognition and explanation (299-323). Cambridge, MA: MIT Press.
Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. Current Directions in Psychological Science, 21(4), 263-268.
Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological Bulletin, 116(1), 75.
Oudeyer, P. Y. (2018). Computational Theories of Curiosity-Driven Learning. arXiv preprint arXiv:1802.10546.
Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. Nature, 442(7106), 1042.
Stahl, A. E., & Feigenson, L. (2015). Observing the unexpected enhances infants' learning and exploration. Science, 348(6230), 91-94
Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. Trends in cognitive sciences, 10(7), 309-318.
Ullman, T. D., Goodman, N. D., & Tenenbaum, J. B. (2012). Theory learning as stochastic search in the language of thought. Cognitive Development, 27(4), 455-480.
Valdesolo, P., & Graham, J. (2014). Awe, uncertainty, and agency detection. Psychological science, 25(1), 170-178.
Yu, Y., Landrum, A., Bonawitz, E., & Shafto, P. (in press) Questioning supports effective transmission of knowledge and increased exploratory learning in pre-kindergarten children. Developmental Science.

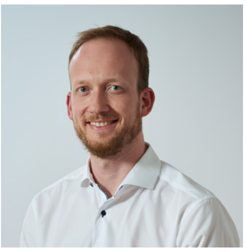

Tobias U. Hauser

Max Planck UCL Centre for Computational Psychiatry and Ageing Research,
University College London
United Kingdom

t.hauser@ucl.ac.uk

Is Human Curiosity Neurobiologically Unique?

Is there something uniquely human about curiosity, and does this allow us to focus all our energy to become experts in one particular field? In her Dialogue Initiation, Celeste Kidd suggested that curiosity is a key driver for human specialisation that makes us humans unique. In what follows, I will approach this hypothesis from a computational neurobiological standpoint and raise the question whether there is any neurobiological uniqueness to human curiosity.

What is curiosity? Although we all have an implicit concept of what curiosity constitutes or who a curious person is, a formal definition is not trivial. Rather than investigating curiosity per se, it is mostly broken down to behaviours that we interpret as being curious. In particular, since the early days of curiosity research, scientists studied what is termed 'exploratory' decisions (Berlyne, 1960). Exploratory decisions can be seen as the counterpart to exploitative decisions, forming a well-known exploration-exploitation trade-off in artificial intelligence (Sutton and Barto, 1998). While exploitation describes decisions that optimise an objective outcome such as reward (i.e. selecting the option that you think provides you with the most reward), exploration characterises actions that are suboptimal in the light of the objective outcome, but may convey additional information that can be useful in the future (i.e. choosing an object where you know very little about how good it is). Exploration is thereby critical for both animals and machines, especially in situations with much uncertainty, such as environments with changing incentives.

So how can we solve this exploration-exploitation trade-off? In the remainder of this response, I will focus on situations with explicit decisions and where objective outcomes are directly quantifiable (e.g. in the form of reward), rather than 'spontaneous' exploration in environments with unspecified outcome structures, because much less is known about the neural mechanisms of the latter. A detailed discussion of the computational bases of the latter situations can be found elsewhere (e.g., Loewenstein, 1994; Gottlieb et al., 2013; Friston et al., 2015; Oudeyer, 2018).

Solutions to dilemmas with specified objective outcomes can be broadly categorised in three different groups. The first and computationally simplest solution is to choose a random action completely ignoring its expected value. This can be interpreted as entirely ignoring all knowledge about the world every so often and is implemented in algorithms such as ε-greedy (Sutton and Barto, 1998). A second, somewhat more sophisticated solution is to again introduce noise into a system, but let this be guided by our knowledge. This is implemented in softmax decision functions (Luce, 1959), or more sophisticatedly in Thompson sampling (Thompson, 1933). Such algorithms let an agent choose actions in proportion to their expected value, which means that we have the ability to sample from all possible choice options, but very unattractive ones will hardly ever be chosen. Lastly, directed exploration takes the uncertainty about each action into account and chooses the option where the agent can learn most, i.e. where it has the biggest information gain (Loewenstein, 1994; Wilson et al., 2014). Prominent examples for directed exploration are upper confidence bound algorithms (UCB; Kaelbling, 1994; Auer, 2002).

Although these algorithms have mainly been used in artificial intelligence, evidence indicates that humans and non-human animals use similar strategies. Importantly, humans seem to apply a mixture of strategies (Wilson et al., 2014; Gershman, 2018), but their relative contribution changes over development (Somerville et al., 2017). It is thus likely that different brain systems are in control of different exploration strategies, and maybe a dominance of a directed exploration system resembling an UCB algorithm is what distinguishes human from non-human curiosity.

Directed exploration is most likely governed by a network including frontopolar cortex. Transient inhibition of this brain area in humans indeed impairs directed exploration, but leaves random forms of exploration intact (Zajkowski et al., 2017). Given that this brain area is one of the last to mature in humans (Ziegler et al., under review), this might also explain why directed exploration only really surfaces during late adolescence (Somerville et al., 2017). A critical, but unanswered question is whether and to what extent humans and non-human animals differ in how they employ this directed exploration strategy. It is possible that only through a dominance of directed exploration humans can use an efficient way to extend their knowledge. This seems particularly relevant in environments where much is already know, but where a directed exploration allows to expand knowledge in a targeted way by aiming for the subtle gaps, which in turn will propel specialisation and expertise.

In sum, it is interesting to conjecture whether there is a neurobiological basis for unique patterns of exploration in humans and whether a stronger reliance of a directed exploration strategy allows human to become highly specialised in one discipline. Such a dominance in itself could also be a result of a developmental process, where intrinsicially guided learning may be assigned a high reward in itself (epistemic value; Friston et al., 2015), and where a metacognitive exploration module optimises the exploration strategies throughout development (cf. Gottlieb et al., 2013; Oudeyer, 2018).

Lastly, it is important to address the question of why within our species we differ so strongly in curiosity and specialisation. Despite having a reasonable junk of humankind reaching some level of specialisation, probably a majority of our species shows only humble specialisation, but forms the economic backbone for those who venture into extreme specialisation and who are thus unable to divert much of their energy to ensure basic needs, such as gathering food. Taking a societal rather than an individual evolutionary perspective could provide answers to this question, because a distribution of exploration traits in a society may indeed provide an evolutionary advantage (Williams and Taylor, 2006).

Auer P (2002) Using Confidence Bounds for Exploitation-Exploration Trade-offs. J Mach Learn Res 3:397–422.
Berlyne DE (1960) Conflict, arousal, and curiosity. New York, NY, US: McGraw-Hill Book Company.
Friston K, Rigoli F, Ognibene D, Mathys C, Fitzgerald T, Pezzulo G (2015) Active inference and epistemic value. Cogn Neurosci:1–28.
Gershman SJ (2018) Deconstructing the human algorithms for exploration. Cognition 173:34–42.
Gottlieb J, Oudeyer P-Y, Lopes M, Baranes A (2013) Information-seeking, curiosity, and attention: computational and neural mechanisms. Trends Cogn Sci 17:585–593.
Kaelbling LP (1994) Associative Reinforcement Learning: Functions in k-DNF. Mach Learn 15:279–298.
Loewenstein G (1994) The Psychology of Curiosity: A Review and Reinterpretation. Psychol Bull 116:75–98.
Luce RD (1959) Individual choice behavior: a theoretical analysis. Wiley.
Oudeyer P-Y (2018) Computational Theories of Curiosity-Driven Learning. ArXiv180210546 Cs Available at: http://arxiv.org/abs/1802.10546 [Accessed May 1, 2018].
Somerville LH, Sasse SF, Garrad MC, Drysdale AT, Abi Akar N, Insel C, Wilson RC (2017) Charting the expansion of strategic exploratory behavior during adolescence. J Exp Psychol Gen 146:155–164.
Sutton RS, Barto AG (1998) Reinforcement learning: An introduction. MIT Press.
Thompson WR (1933) On the likelihood that one unknown probability exceeds another in the view of the evidence of two samples. Biometrika 25:285–294.
Williams J, Taylor E (2006) The evolution of hyperactivity, impulsivity and cognitive diversity. J R Soc Interface 3:399–413.
Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD (2014) Humans use directed and random exploration to solve the explore-exploit dilemma. J Exp Psychol Gen 143:2074–2081.
Zajkowski WK, Kossut M, Wilson RC (2017) A causal role for right frontopolar cortex in directed, but not random, exploration. eLife 6.
Ziegler G, Hauser TU, Moutoussis M, Bullmore ET, Goodyer IM, Fonagy P, Jones PB, Consortium the N, Lindenberger U, Dolan RJ (under review) Compulsivity and impulsivity are linked to distinct aberrant developmental trajectories of fronto-striatal myelination.

Maya Zhe Wang

Dept. of Brain and Cognitive Sciences and Center for Visual Sciences, University of Rochester, Rochester, NY, USA

mwang@bcs.rochester.edu

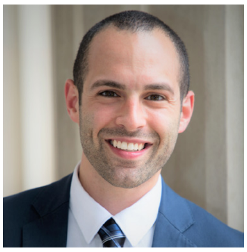

Brian M. Sweis

Dept. of Neuroscience, University of Minnesota, Minneapolis, MN, USA

sweis001@umn.edu

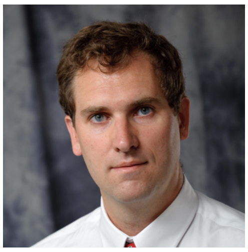

Benjamin Y. Hayden

Dept. of Neuroscience, University of Minnesota, Minneapolis, MN, USA

benhayden@gmail.com

A Testable Definition of Curiosity

Curiosity is a fascinating and perplexing cognitive phenomenon that likely results from evolutionary selective pressure (Kidd and Hayden, 2016; Gottlieb et al., 2013; Golman et al., 2016). We believe that it is not unique to humans and, thus, that animal models can be an important tool for understanding its evolutionary and neuronal origins. Critical to advancing the study of curiosity is the ability to validly measure it with laboratory tasks (Kang et al., 2009; Gruber et al., 2014). These tasks need to be work in multiple species (Kidd and Hayden, 2016).

One limitation to developing such tasks is the lack of a standard cross-species definition of the term curiosity. For example, curiosity can refer to either a desire for information that has no instrumental benefit, or for strategically beneficial information. A growing consensus, however, holds that the term curiosity should refer specifically to demand for information with no clear strategic benefit (Golman et al., 2016; Loewenstein, 1994). Humans can introspect and provide reports about their own curiosity-driven motivations. Other animals cannot. As such, animal models of curiosity demand special operationally defined criteria. We propose three here. First, operationally defined curiosity requires a willingness to sacrifice reward in order to obtain additional information (e.g. Kang et al., 2009; Blanchard et al., 2015a). Second, the amount the subject is willing to pay must scale with the amount of available additional information (e.g. Wang et al., 2018). Third, that provided information must provide no obvious instrumental or strategic benefit, even subjectively (e.g. Bromberg-Martin et al., 2009).

This subjective element may be tricky. It is easy to misinterpret animal behavior by assuming that the animal understands the task in the same way the experimenter does (e.g. Hayden, 2016; Sweis et al., 2018). Consider, for example, the seemingly simple case of risk-seeking. Decision-makers may prefer a risk options in part because choosing it reduces uncertainty about its value. A carefully controlled experiment may involve options that have stochastic outcomes – meaning its outcomes are uninformative and cannot satisfy any curiosity– but the animal may falsely believe the outcomes are patterned (Blanchard et al., 2014).

By practically defining a testable definition of curiosity, it will be possible to design better tasks that work across species to explore the behavioral and neurophysiological determinants of curiosity. One important way to do this is to embed differentially informative options in naturalistic foraging tasks, in which animals are asked to make decisions similar to those they have evolved to make (e.g. Blanchard et al., 2015b). Through such frameworks, we can better examine the underlying neural mechanisms involved in curiosity and its uniqueness in humans.

Blanchard, T. C., Hayden, B. Y. & Bromberg-Martin, E. S. Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. Neuron 85, 602–614 (2015).
Blanchard, T. C., Wilke, A. & Hayden, B. Y. Hot hand bias in rhesus monkeys. Journal of Experimental Psychology: Animal Learning and Cognition 40(3) 280-286 (2014).
Blanchard, T. C., Strait, C. E., and Hayden, B. Y. (2015) Ramping ensemble activity in dorsal anterior cingulate cortex neurons during persistent commitment to a decision. Journal of Neurophysiology 114(4) p 2439-2449.
Bromberg-Martin, E. S. & Hikosaka, O. Midbrain Dopamine Neurons Signal Preference for Advance Information about Upcoming Rewards. Neuron 63, 119–126 (2009).
Gottlieb, J., Oudeyer, P.-Y., Lopes, M. & Baranes, A. Information-seeking, curiosity, and attention: computational and neural mechanisms. Trends in Cognitive Sciences 17, 585–593 (2013).
Golman, R. & Loewenstein, G. Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information. (2016).
Gruber, M. J., Gelman, B. D. & Ranganath, C. States of Curiosity Modulate Hippocampus-Dependent Learning via the Dopaminergic Circuit. Neuron 84, 486–496 (2014).
Hayden, B. Y. Time discounting and time preference in animals: a critical review. Psychon Bull Rev 23, 39–53 (2016).
Kang, M. J. et al. The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. Psychological Science 20, 963–973 (2009).
Kidd, C. & Hayden, B. Y. The Psychology and Neuroscience of Curiosity. Neuron 88, 449–460 (2015).
Loewenstein, G. The psychology of curiosity: A review and reinterpretation. Psychological bulletin 116, 75 (1994).
Sweis, B.M., Thomas, M.J., Redish, A.D. Mice learn to avoid regret. PLOS Biology (in press, 2018a).
Wang, M. Z. & Hayden, B. Monkeys are Curious about Counterfactual Outcomes. (2018). bioRxiv. doi:10.1101/291708

Why We Say "Why" When We Do

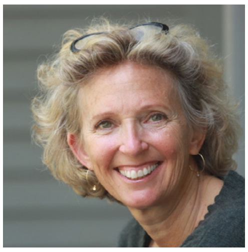

Susan Engel

Psychology Dept.
Williams College,
Williamstown, MA, USA

sengel@williams.edu

Not long after birth, human infants begin expressing curiosity. Anything that surprises them causes them to study further and explore. Because their schemas for understanding and predicting experience are, relatively speaking, quite spare, they are surprised by a lot of what they encounter, a lot of the time. A different liquid in their breakfast cup, a new route to day care, a stranger at the door, and a thousand episodes and objects in between, all spark uncertainty, which often leads to inquiry. Moreover, by the time they are two, their curiosity takes on an exceptional characteristic. Young children not only seek more information about the unfamiliar, they also seek to explain unexpected phenomena, exhibiting something uniquely human: epistemic curiosity. They want to know not only what, when, who and where, but how and why. Observation, manipulation and questions all help them discover what they want to know (Gopnik, Meltzoff & Kuhl 2001, Magid, Sheskin & Schulz 2015, Bonawitz, Shafto, Gweon, Goodman, Spelke, & Schulz 2011).

By early childhood however, this omnivorous curiosity tends to wane. Less of daily life arouses children's uncertainty and demands explanation. As it becomes less ubiquitous, curiosity becomes more specific (Engel 2015). 18-month-old children explore their environment nearly constantly, while four-year-olds express curiosity only some of the time under some circumstances. They are more likely to seek information in depth about one domain, or group of objects, and not others. For instance, most four-year-olds are more interested in seeking explanations for people's behavior (their own and others) and biological phenomena than they are in the non-living natural world (Kelemen, Callanan, Casler & Pérez-Granados, 2005). But their interests become even more specialized than that. Imagine three children who, on a given day, come in contact with a variety of bugs, hear several stories about family life, and are given the chance to examine some small machines. All three will probably know a little about each of these domains. But given the opportunity to gain more information, one child will notice a detail about one of the bugs and want to inspect further, another will zero in on the interpersonal conflict described in the stories, and a third may begin taking apart one of the machines to figure out why it makes an unusual sound. Individual differences in the focus, as well as intensity, of inquiry, become the rule. Why, beginning at such a young age, do we say why to some events, and not others? Moreover, why do these specific interests often seem to lead children deeper and deeper into a topic, leading one child to be "surprised" about a detail another will never notice?

It would be easy to think, based on laboratory research and brief observations in natural settings, that children's interest in particular objects or events is only sparked by the appearance of ambiguous or surprising stimuli, and is always fleeting. Even when we acknowledge that children have specific interests, we tend to conceptualize their expressions of curiosity and acts of inquiry as a series of brief moments- each instance of novelty invites exploration, which in turn results in a resolution of uncertainty, then the child moves on to other experiences and activities. But it would be a mistake to confuse the necessary time constraints of laboratory research with our conceptual picture of the psychological processes we seek to understand. Though most data on curiosity are gathered during short periods of time, curiosity itself may unfold over a more continuous and extended stretch of time. There is evidence to support such a view.

Beginning at about age 2 ½ (around the time children begin using language to satisfy their curiosity about things they cannot see or physically manipulate), many appear to pursue a cluster of questions about a particular conceptual topic (Engel 2018). One child asks questions about death, another about infinity, and another about the mind. These sustained lines of questioning unfold over a period of weeks or months. Sometimes questions or speculations are sparked by a particular event (the death of a pet) but other times they just keep popping up (a child who frequently asks questions about examples of infinity). These lines of inquiry show that at least some of the time, children gather information to answer an internal, and often abstract, question. The problem they want to solve is one they have mentally constructed. It takes time to construct those problems, and then more time to seek answers.

We have yet to understand what is entailed when a child begins to show sustained interest in a given phenomenon, moving from a fairly simple collection of discrete pieces of information to the articulation of a complex and/or fairly abstract problem (for example, "What happens to people when they die?", "What is encompassed by the universe?", "What are thoughts made of?"). We know little about why some children appear to pursue such topics and others do not. To better understand the development of curiosity, and its cousin, specialization, we need to find ways to study the emergence of intellectual pursuits that unfold over time. Even four-year- olds mull things over.

Bonawitz, E, Shafto, P. Gweon, H. Goodman, N. Spelke, E. and Schulz, L. (2011) "The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery." Cognition 120 (3), 322-330.
Engel, S. (2015) The Hungry Mind: The origins of curiosity in childhood, Cambridge MA: HUP.
Engel, S. (2018) Do Children Have Ideas? Invited Talk: Harvard Graduate School of Education, Cambridge MA
Gopnik, A., Meltzoff, A. N., & Kuhl, P. K. (2001). The scientist in the crib: What early learning tells us about the mind. New York, NY: Perennial.
Kelemen, D., Callanan, M. A., Casler, K., & Pérez-Granados, D. (2005). Why things happen: teleological explanation in parent-child conversations. Developmental Psychology, 41(1), 251-264.
Magid, R., Sheskin, M., & Schulz, L. (2015). Imagination and the generation of new ideas. Cognitive Development, 34, 99-110.

Unpacking the Different Flavors of Curiosity

Abigail Hsiung

Center for Cognitive Neuroscience, Duke University, Durham, NC, USA

abigail.hsiung@duke.edu

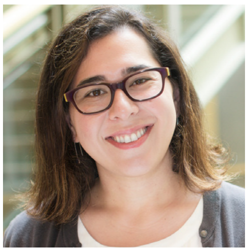

Shabnam Hakimi

Center for Cognitive Neuroscience, Duke University, Durham, NC, USA

shabnam.hakimi@duke.edu

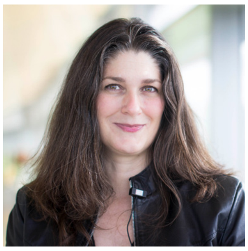

R. Alison Adcock

Center for Cognitive Neuroscience, Duke University, Durham, NC, USA

alison.adcock@duke.edu

Disciplines as diverse as psychology, education, and neuroscience strive to understand curiosity and harness its power for human thriving. Despite interest and decades of research, however, it is noteworthy that even a consensus on the definition of curiosity remains elusive. Critically for scientists, however, the absence of a definition or taxonomy of curiosity means we lack the power to predict when curiosity occurs, and how it might alter behavior. Because of this, there is a challenge for understanding the functions of curiosity within and across species.

We propose that to fully characterize the dimensions of curiosity, we must understand how external elicitors of curiosity interact with an agent's internal state and goals, to produce not only varied information-seeking behavior but also distinct learning episodes. Building on Berlyne's initial characterizations of curiosity (Berlyne, 1954), we argue that curiosity is comprised of multiple dimensions; the interaction amongst these dimensions influences how curiosity is elicited, its associated subjective experience, and the learning behaviors and plasticity it produces. An understanding of how curiosity is induced, what creates a drive to seek information that is not directly instrumental, provides a useful starting point.

Previous work has suggested that curiosity arises from a perceived gap in one's knowledge, referred to as the 'information gap' (Loewenstein, 1994). However, this theory makes no prediction about how the size of the information gap impacts induction or the subsequent search for information. Information gaps can vary widely in size, ranging from the considerable gap associated with novelty, where stimuli have not previously been encountered and the space to acquire information is quite large, to gaps so small that a single cue can relieve tension induced by the question at hand (Blanchard, Hayden, & Bromberg-Martin, 2015). These varying sizes of gaps in knowledge imply varied behaviors for an agent navigating the potential search space.

In addition to the quantitative amount of missing information, another variable to be considered is how the state of missing information is perceived by an agent—that is, how it maps onto the agent's internal representation of the world and its affective or neuromodulatory state. According to this view, the awareness of an information gap might give rise to multiple—and possibly conflicting—motivational drives and behavioral goals. For example, we have previously demonstrated that the same external information gap gives rise to distinct brain and behavioral responses according to the valence of motivational states (Murty, LaBar, & Adcock 2016; Chiew et al., 2018).

This principle can be refined to suggest specific predictions related to motivational states, as follows: if an agent's goal in acquiring new information is to reduce environmental or decision uncertainty for an action imperative, curiosity might be precisely targeted, functioning to limit actions to only those serving to obtain goal-relevant information (Blanchard et al., 2015). Conversely, if the goal is to interrogate a novel space, resulting actions may be more exploratory, or diffuse and varied (Baranes et al., 2014). We have argued elsewhere for distinguishing these motivational states (see Murty & Adcock 2017; Dickerson & Adcock in press). Would both of these example behaviors be expected to drive behavioral specialization? Or only the first?

Understanding curiosity from a function-first perspective requires that we characterize the complex taxonomies, at multiple levels of analysis, that shape both the induction of curiosity and behavioral pursuit of information. By characterizing different stimuli, search processes, and learning episodes associated with curiosity, we may start to understand its functions in our adaptive exploration, knowledge acquisition, and skill development. Only after we delineate these multiple different flavors of curiosity can we begin to assess whether and how curiosity may drive extreme specialization, in humans and other animals.

Baranes, A. F., Oudeyer, P.-Y., & Gottlieb, J. (2014). The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. Frontiers in Neuroscience, 8, 317.
Berlyne, D. E. (1954). A Theory of Human Curiosity. British Journal of Psychology, 45(3).
Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. Neuron, 85(3), 602–614.
Chiew, K. S., Hashemi, J., Gans, L. K., Lerebours, L., Clement, N. J., Vu, M.-A. T., et al. (2018). Motivational valence alters memory formation without altering exploration of a real-life spatial environment. PLoS ONE 13(3): e0193506.
Dickerson, K. C. & Adcock, R. A. (in press). Motivation and Memory. To appear in The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience, Fourth Edition.
Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological Bulletin.
Murty, V. P., & Adcock, R. A. (2017). Distinct Medial Temporal Lobe Network States as Neural Contexts for Motivated Memory Formation. In The Hippocampus from Cells to Systems (pp. 467–501). Springer, Cham.
Murty, V. P., LaBar, K. S., Adcock, R. A. (2016). Distinct medial temporal networks encode surprise during motivation by reward versus punishment. Neurobiology of learning and memory. 2016;134:55-64.

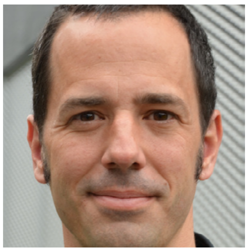

Moritz M. Daum

Dept. of Psychology, Neuroscience Center Zurich, University of Zurich and ETH Zurich, Germany

daum@psychologie.uzh.ch

The Social Side of Curiosity

The point I would like to make with this response is that if you communicate contingently with your child, you provide a "social glue" that increases the child's self-confidence and fosters curiosity. This idea has its roots in anecdotal reports that parents dress up like their children to show their support. In the following, focusing on early childhood, I will first talk about the importance of the quality of social interactions, and of imitation in particular, for children's development. I will then briefly describe how this might be linked to the development of children's curiosity.

Imitation is an important source for children's learning from and communicating with the social environment. Early in life, children start imitating other's behaviour. Soon after, they selectively imitate and take into account contextual information. Children are more likely to imitate a model who has been reliable in the past (Zmyj, Buttelmann, Carpenter, & Daum, 2010) or who is more familiar to them (Buttelmann, Zmyj, Daum, & Carpenter, 2013).

Children use imitation to communicate when verbal competences are still limited (Nadel, Guérini, Péze, & Rivet, 1999), a behaviour that is still observable in adults in the form of mimicry (Chartrand & Bargh, 1999). This and other forms of contingent interactions result in an increase in the sense of affiliation between two interaction partners. Children who perceive a person as acting temporally contingent to the child more often use social referencing with her when exploring novel events (Striano, Henning, & Vaish, 2006). Being contingently mimicked by another person results in an increase of prosocial behaviour (Carpenter, Uebel, & Tomasello, 2013). Above that, the quality of nonverbal and verbal interactions with children, for example through fluent and connected communication has a substantial impact on language development (Hirsh-Pasek et al., 2015). To summarise, imitating and being imitated helps increasing the social bond between two interaction partners (Lakin, Chartrand, & Arkin, 2008).

How can this be translated to the development of curiosity? Being curious results in an increased likelihood of experiencing novel situations. Whether or not such a situation is perceived as a learning opportunity that should be approached is often not clear to the child.

To decide whether or not to approach, children make use of information provided by their social interaction partners, usually their primary caregivers. Children are more prone to learning when they can use their primary caregiver as a secure base (Waters & Cummings, 2000), they are more likely to approach an unfamiliar situations when positively supported by caregivers' emotions (e.g., Klinnert, Emde, Butterfield, & Campos, 1986), they develop better self-regulation and effortful control (Drake, Belsky, & Fearon, 2014). Thus, (contingent) reactions of parents to their children's behaviour are important for the development of trust in their parents as reliable sources of information and a safe base for the curious exploration of the world.

Based on this line of argumentation, successful extreme specialisation, a form of sustained curiosity, is more likely to develop within a high quality social network where a curious individual receives the social support to approach novel learning opportunities as often as possible.

Buttelmann, D., Zmyj, N., Daum, M. M., & Carpenter, M. (2013). Selective imitation of in-group over out-group members in 14-month-old infants. Child Development, 84(2), 422–428. https://doi.org/10.1111/j.1467-8624.2012.01860.x

Carpenter, M., Uebel, J., & Tomasello, M. (2013). Being mimicked increases prosocial behavior in 18-month-old infants. Child Development, 84(5), 1511–8. https://doi.org/10.1111/cdev.12083

Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: the perception-behavior link and social interaction. Journal of Personality and Social Psychology, 76(6), 893–910.

Drake, K., Belsky, J., & Fearon, R. M. P. (2014). From early attachment to engagement with learning in school: The role of self-regulation and persistence. Developmental Psychology, 50(5), 1350–1361. https://doi.org/10.1037/a0032779

Hirsh-Pasek, K., Adamson, L. B., Bakeman, R., Owen, M. T., Golinkoff, R. M., Pace, A., ... Suma, K. (2015). The contribution of early communication quality to low-income children's language success. Psychological Science, 26(7), 1071–1083. https://doi.org/10.1177/0956797615581493

Klinnert, M. D., Emde, R. N., Butterfield, P., & Campos, J. J. (1986). Social referencing: The infant's use of emotional signals from a friendly adult with mother present. Developmental Psychology, 22(4), 427–432. https://doi.org/10.1037/0012-1649.22.4.427

Nadel, J., Guérini, C., Péze, A., & Rivet, C. (1999). The evolving nature of imitation as a format for communication. In J. Nadel & G. Butterworth (Eds.), Imitation in infancy (pp. 209–234). Cambridge, UK: Cambridge University Press.

Striano, T., Henning, A., & Vaish, A. (2006). Selective looking by 12-month-olds to a temporally contingent partner. Interaction Studies, 7(2), 233–250. https://doi.org/10.1075/is.7.2.07str

Waters, E., & Cummings, E. M. (2000). A secure base from which to explore close relationships. Child Development, 71(1), 164–172. https://doi.org/10.1111/1467-8624.00130

Zmyj, N., Buttelmann, D., Carpenter, M., & Daum, M. M. (2010). The reliability of a model influences 14-month-olds' imitation. Journal of Experimental Child Psychology, 106(4), 208–220.

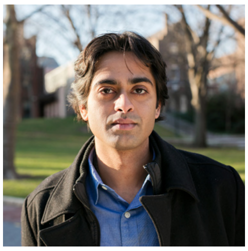

Arjun Shankar

Dept. of Anthropology,
Hamilton College,
Clinton, NY, USA

ashankar@hamilton.edu

Capitalism, Curiosity, and Specialization

Curiosity, as many have pointed out, is never singular: it comes in many types/forms and can direct itself in a number of directions. It can be diversive or specific, perceptual or epistemic (see, for example, the RSA Project, The Power of Curiosity, 2012). Indeed, in the idealized university of the early 20th Century, university learning, graduate training and the like was driven by the belief in an idle curiosity (Handler, 2018): a form of curiosity that was wandering, exploratory, knew no specific intellectual boundary or had a specific goal or path in mind. It was thought that this form of curiosity created broad-minded, critical thinkers who had the capacity to solve multifaceted problems given their ability to connect across many different fields of knowledge. We know then that, historically at least, humans have not been naturally or necessarily more predisposed to a curiosity that leads to specialization. And yet, as the current theme of the IEEE CIS newsletter makes clear, specialized knowledge and how to derive it has taken a central role in defining and explaining what curiosity is, how it functions, and even what it ought to be. What I want to argue here is that this form of curiosity as specialization has been cultivated within the sociocultural processes tied to capitalism in the 20th and 21st Century.

Foucault famously wrote that disciplining was a mechanism of control that "characterize, classify, specialize; they distribute along a scale, around a norm, hierarchize individuals in relation to one another and, if necessary, disqualify and invalidate." In the late-20th and 21st-Century this form of hierarchy is tied to one's value as a productive member of society – productive here being a gloss for economic valuations (Graeber, 2011). Knowledge production, as those within the university know quite well, is tied directly to the precepts of funding. In such an economic system, curiosity is forced to follow fundable paths and the knowledge university scholars can produce, the kinds of questions they ask, and the problems they seek to solve are defined within this system (Chatterjee and Maira, 2014). Indeed, the disciplines themselves facilitate this process, functioning as vocational areas that delimit knowledge production within specialized silos that ostensibly provide a "more efficient" means by which to derive insights, but which also prevent the slow undisciplined forms of exploration that can lead to the unexpected and the unique. Another primary outgrowth of capitalism's role in university knowledge production is the need to have outputs, goals, and specific purposes in mind when one begins an activity of inquiry. In this understanding of curiosity, it is a means to an end, a way to ask questions to derive insights which, in turn, can produce a knowledge that can be used by those outside of the university, and most specifically those in various corporate sectors. Indeed, one example of this process is the oft-discussed topic of "innovation", a process of knowledge production whose goal is to bring new, unique, creative products to market, but which always begins with the question "why not" – i.e. why can't this product be brought to market – instead of "why" and which has subsumed much of scholarly curiosity into its gambit (Nowotny, 2010). Indeed, in this model the niche market and the niche academic specialization fit together quite nicely.

We know that curiosity, whether it is towards specialized knowledge or path-independent exploration is cultivatable. The question, of course, is how we cultivate multiple types of curiosity, especially in the wake of new inter- and trans-disciplinary forms of knowledge production that take seriously the need for a transgressive, radical curiosity that moves beyond the (capitalist) specializations of the past.

Chatterjee, Piya and Sunaina Maira. 2014. The Imperial University: Academic Repression and Scholarly Dissent. Minneapolis: Minnesota Press.
Foucault, Michel. 1975. Discipline and Punish: The Birth of the Prison. New York: Pantheon Books.
Graeber, David. 2001. Towards an Anthropological Theory of Value: The False Coin of Our Own Dreams. New York: Palgrave.
Handler, Richard. 2018. "Undergraduate Research in Veblen's Vision: Idle Curiosity, Bureaucratic Accountancy and Pecuniary Emulation in Contemporary Higher Education," in The Experience of Neoliberal Education, ed. Bonnie Urciuoli. New York: Bergahn.
Nowotny, Helga. 2010. Insatiable Curiosity: Innovation in a Fragile Future. Cambridge: MIT Press.
RSA Project. 2012. Power of Curiosity: How Linking Inquisitiveness to Innovation Could Help to Address Our Energy Challenges.

Goren Gordon

Curiosity Lab,
Department of
Industrial Engineering,
Tel-Aviv University,
Israel

gorengor@post.tau.ac.il

A Reinforcement Learning Perspective on Curiosity-Based Extreme Specialization

To reach extreme specialization, an individual must select to continue with the same activity. Hence, during the decision-making process, the selected activity should have a higher value than other activities. This value function, according to the reinforcement learning paradigm (Sutton and Barto, 1998), is learned. How does curiosity fit into this framework? Within the same reinforcement learning paradigm, curiosity essentially means that new information, which can be measured in various ways (Little and Sommer, 2013), generates rewards. Hence, in the traditional Q-value formalism the states and actions that generated the new information will have higher values. However, and this is the crux of the current contribution, the scope and context of these states and actions dramatically influence future decision making, which is based on these Q-values (Gordon and Ahissar, 2012; Kulkarni et al. 2016). How the individual perceives the current context, expressed computationally by the internal definition of the state-action spaces, can determine the course of extreme specialization.

Let us first consider the state-space. Considering one extreme, the state can represent a very broad context the individual is in, when the new information arrives. The most extreme scenario relates to a non-existent specification of the current state, resulting in information for its own sake, without any context, being rewarding. This may lead to seeking out new information, regardless of its content, context or relation to other extrinsic variables. This situation may foster *diverse curiosity* (Litman and Spielberger, 2003), the opposite of extreme specialization, in which individuals seek and explore whatever comes their way, and not focusing on any specific topic. On the other end of the specification scale, the state may represent a very specific and unique situation, e.g. a specific question about a specific topic answered by a specific teacher. While new information still results in rewards, because of the uniqueness of the state, it will probably not be encountered again, hence not directly influence future decision making. The middle-ground, which is more common, refers to the situation wherein states represent a specific context or area of interest. Thus, while exploring a specific scenario, the individual receives new information, i.e. rewards, and associates the scenario with pleasure. While the information supplied the reward, the state/scenario/context will be sought after in the future. The association between a specific self-perceived context and the reward which was generated by the scenario's novelty can result in extreme specialization. The individual will seek the same scenario again, and if new information is supplied, this association will grow resulting in further learning and expertise in that specific field.

Next, let us consider the action space (Gordon et al. 2014). Again, considering the broad extreme, a generic activity which generates new information can have higher value. Examples of these are question asking, tinkering and general exploration, without the accompanying context. If a student asks questions and is given answers that satiate her curiosity, then the act of asking questions, regardless of the question itself, will be rewarded and that student will continue to ask questions. On the other hand, the unique extreme refers to a specific unique action, e.g. asking that specific question or opening that specific box. While that action will generate information and thus rewards, performing it again will not cause new learning and hence will not be selected in future events. The middle-ground of action-space specification refers to a specific *type of activity* that generates new information, new learning and rewards. It can include the act of scientific experimentation, traveling to new places and reading books on specific topics. However, it may also include creating art: if, while painting a new piece, the artist feels the elation of learning something new about either herself or her art, that learning generates rewards that may be associated with the act of painting, and not necessarily the piece itself. Thus, the artist will select painting in future events, causing extreme specialization in that art-form.

To conclude, the link between curiosity and extreme specialization may lie in the *self-perceived context* of the information/reward-generating activity of the individual, and not necessarily in the information itself.

R. S. Sutton and A. G. Barto, Reinforcement learning : an introduction. Cambridge, Mass.: MIT Press, 1998.
D. Y. Little and F. T. Sommer, "Learning and exploration in action-perception loops," Front Neural Circuits, vol. 7, p. 37, 2013.
G. Gordon and E. Ahissar, "Hierarchical curiosity loops and active sensing," Neural Netw., vol. 32, pp. 119–129, 2012.
T. D. Kulkarni, K. R. Narasimhan, A. Saeedi, and J. B. Tenenbaum, "Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation," ArXiv160406057 Cs Stat, Apr. 2016.
J. A. Litman and C. D. Spielberger, "Measuring Epistemic Curiosity and Its Diversive and Specific Components," J. Pers. Assess., vol. 80, no. 1, pp. 75–86, Feb. 2003.
G. Gordon, E. Fonio, and E. Ahissar, "Learning and control of exploration primitives," J. Comput. Neurosci., vol. 37, no. 2, pp. 259–280, 2014.

Perry Zurn

Department of Philosophy, American University, DC, USA

pzurn@american.edu

Is Curiosity Uniquely Human?

In the centuries-old debate over what distinguishes the human, curiosity has played a significant, if controversial, role. In her provocation, "Curiosity as Driver of Extreme Specialization in Humans," Celeste Kidd suggests that while curiosity itself is not uniquely human (but likely belongs to all intelligent systems), curiosity preconditions the uniquely human capacity for extreme specialization. Since curiosity is a drive to fill gaps in knowledge—circumscribed by partially encoded information—it necessarily results in increasingly sophisticated knowledge networks. While compelling, Kidd's account leaves me wondering what is at stake in assessing curiosity as uniquely human or as contributing to the uniquely human? And what indeed is curiosity's relationship to extreme specialization?

Early in the development of Western philosophical thought, curiosity was perceived as distinctly animal. For Saint Augustine, the serpent is the symbol of curiosity: it creeps on its belly, consumed by an interest in earthly things (Saint Augustine, 388). As an animal, appetitive passion, curiosity keeps humans from becoming the divinely superior creatures they were destined to be. In the modern era, curiosity made an about-face and became the distinguishing mark of the human. Thomas Hobbes, for instance, argues it is precisely through curiosity that humans "surmount the nature of beasts." (Hobbes, 1650) Curiosity here becomes rational and teleological; it is the desire to know causes and effects, all for the sake of future planning and development. More recently, Jacques Derrida argues that curiosity can only be thought of as uniquely human when it is reduced to its linguistic expression. If, however, one takes curiosity to be simply an "exploratory comportment," it is traceable not only in animal but vegetal life (Derrida, 1991). After all, plants and roots probe.

Are some forms of human curiosity still disciplined or dismissed for being too appetitive? Correlatively, how has an attachment to human exceptionalism—and to exceptionally human curiosity—contributed to the Anthropocene epoch? In what way would our environmental policies and ecological thinking change if we recognized curiosity not only among other creatures, but in the very earth itself?

And what of human curiosity's relationship to extreme specialization? If we understand knowledge to be a network, the edge between each node is infinitely divisible; that is, knowledge gaps go all the way down. Or, as Socrates might have said, the more I know, the more I don't know. Nevertheless, different modes of curiosity capitalize on the gaps—or on gaping architecture—differently. I argue that there are at least three models of curiosity thematized in Western philosophical history: the busybody, the hunter, and the dancer (Zurn, in press). Each figure is driven by a different question: What's new? What is? What if? And each model carries a different kinesthetic signature, a different set of movements across knowledge networks (Zurn and Bassett, 2018).The busybody collects discrete bits of information and therefore constructs a loose knowledge network. The hunter focuses on a particular trajectory of inquiry, producing targeted connections and tight knowledge networks. The dancer takes leaps of creative imagination, courting the discontinuous concepts that result in exceptionally disparate and dynamic knowledge networks. These modes of curiosity may function independently or in tandem to produce a variegated architecture of interest and information.

The curiosity that supports extreme specialization—that is, the sort of curiosity that systematically identifies and fills in gaps—is hunter-like. But this is not the only sort. The curiosity that collects Harry Potter trivia—and, for that matter, the curiosity that creates the Harry Potter world—are more than this. Their relationship to existing knowledge and knowledge gaps is structurally distinct.

Curiosity functions within an ecology of knowledge. It is not enough that we recognize curiosity in C. elegans and robotic or artificial intelligences. Nor is it adequate to assign a singular function to curiosity in human intellectual life. There are multiple systems of curiosity across species networks and evident in human knowledge network growth. The challenge is to honor the beauty of this complexity.

Saint Augustine, On Genesis Against the Manichees (388), trans. Roland J. Teske (Washington DC: Catholic University of America Press, 2010).
Thomas Hobbes, Human Nature. The English Works of Thomas Hobbes, Vl. 4 (1650), ed. Sir William Molesworth (Germany: Scientia Verlag Aalen, 1966), 1-76.
Jacques Derrida, Répondre du secret, December 11, 1991), University of California, Irvine, Special Collections, MS-C01, Box: Folder: 21:4-9; manuscript pg. 12 (my translation). See also Perry Zurn, "The Curiosity at Work in Deconstruction," Journal of French and Francophone Philosophy 26.1 (2018): 65-87.
Perry Zurn, "Busybody, Hunter, Dancer: Three Historico-Philosophical Models of Curiosity," Curious about Curiosity: Toward New Philosophical Explorations of the Desire to Know, ed. Marianna Papastefanou (Cambridge Scholars Press, forthcoming).
Perry Zurn and Danielle Bassett, "On Curiosity: A Fundamental Aspect of Personality, a Practice of Network Growth," Personality Neuroscience 1 (2018): 1-10. In press.

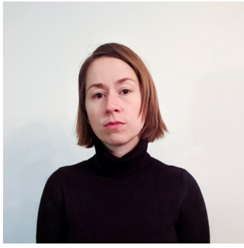

Celeste Kidd

Assistant Professor of Psychology,
University of California, Berkeley

celestekidd@gmail.com

Summary of responses to dialog initiation "Curiosity as driver of extreme specialization in humans"

In this issue of the IEEE CIS Newsletter on Cognitive and Developmental Systems, thirteen researchers from highly varied backgrounds—neuroscience, psychology, education, philosophy, and cognitive science—respond to my position statement about how curiosity could enable extreme knowledge specialization in humans. Interestingly, almost every statement pointed to a nascent formal understanding of the cognitive and neural systems that govern curiosity as a major impediment to research progress in this area. The researchers' proposals for next steps, however, varied as widely as their backgrounds. Some discuss whether humans are neurobiologically unique (Tobias Hauser, for example), while another takes a reinforcement learning approach to tackling the same question (Goren Gordon).

Many responses pointed out relevant components of learning that science has yet to integrate into existing broader theories of curiosity. The responses discuss these understudied components, and lay out the empirical evidence for how we know that they impact curiosity and, ultimately, the specialization of knowledge in humans. These understudied components include emotional states (as discussed by Elizabeth Bonawitz), motivation (Abigail Hsiung, Shabnam Hakimi, and R. Alison Adcock), and social connection (Moritz M. Daum).

The responses also indicate a lack of consensus on what should count as curiosity, a debate as old as scientific study of the topic (see Perry Zurn's response for a compelling philosophical perspective on this question). Namely, researchers disagree as to whether a strict line should be drawn between an organism seeking information for utilitarian purposes (e.g., to solve a particular task, now or in the future) as opposed to for the sake of the information itself. The term curiosity is broad enough that it can be used to describe a wide range of behaviors—from the motivating force behind exploration during play, to the desire for answers to trivia questions, to the strategic deployment of gaze in free-viewing, as a few examples. It may even be applicable to describing the probing behavior of plants (Perry Zurn), and a driving market force behind capitalism (Arjun Shankar).

Several responses point out that sharply delineating between curiosity and information-seeking is difficult for a number of different reasons (see the response by Maya Zhe Wang, Brian M. Sweis, and Benjamin Y. Hayden). First, it is not always possible to know the beliefs and motivations of organisms who clearly possess curiosity. For example, researchers may design tasks intended to present options for exploration that offer no specific utility, but it is difficult if not impossible to guarantee that participants (especially when they are children and monkeys) will share the same understanding of the tasks as the researchers. More generally, even if participants in these tasks do know that selecting certain options won't help increase overt rewards within the task set-up, they cannot know that the information they obtain from exploring will never again be useful in any circumstance.

Another common response theme is that we still need to better understand the relationship between existing knowledge and curiosity, and the role of perceived knowledge utility on the part of the learner (see Goren Gordon's response). Our knowledge in this areas is clearly hindered by a lack of longitudinal data across development (see, for discussion, the response by Susan Engel). In general, infants possess less knowledge of the world and subsequently are surprised by many things they encounter on a daily basis, which could be the explanation for why they appear to exhibit more intense curiosity than adults. By contrast, surprise is less ubiquitous in adulthood, and curiosity has become more specific. These dynamics are almost certainly relevant to understanding human specialization. However, we lack a precise understanding of these dynamics, and this is an area ripe for future research.

New Dialogue Initiation

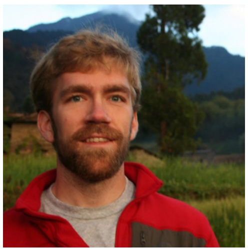

George Kachergis

Dept. of Artificial Intelligence
Donders Institute for Brain, Cognition, and Behaviour
Radboud University

g.kachergis@donders.ru.nl

Leveraging Adaptive Games to Learn How to Help Children Learn Effectively

Despite more than a century of rigorous psychological research on learning and memory which has discovered robust methods for improving learning, relatively few large shifts in education have come about due to insights gained about how people learn (Roediger, 2013). With an eye towards changing this, Dunlosky et al. (2013) reviewed ten of the most well-studied learning techniques that are ripe for translation and found that the most effective and generally useful techniques for a variety of materials and learning contexts are: using spaced (i.e. distributed) practice as opposed to massed repetitions, retrieval practice (aka testing), and interleaved practice—mixing topics of study over short periods. While these principles might be difficult to apply in a standard classroom, with only a single teacher to facilitate interleaved study and retrieval practice, software has the potential to adapt to individual learners, recording responses and adjusting the content based on learning progress. Indeed, research on computer-based cognitive tutors has promoted the translation of memory and learning research into educational practice in at least some classrooms, and such systems have proven effective (for a review, see Koedinger & Corbett, 2005).

However, with the growing ubiquity of tablets and smartphones, software developers are already offering tens of thousands of educational games and applications on the app stores. While parents often flock to the well-designed apps that may promise to teach basic literacy or numeracy skills and that are most engaging and entertaining for their children, the bulk of these apps do not draw on principles from the science of learning (see Hirsh-Pasek et al. 2015 for a review). The main question of this dialogue initiation is thus, "How can we work with app developers--or develop our own apps--to help translate our research into educational practice, while also using data collected from these apps to learn more about development?" The challenges can be organized into a few categories: 1) app development, 2) organizational, and 3) scientific.

Software Challenges

Most researchers do not have the resources to hire an app developer, and the reverse is also generally the case. As researchers, can we reach out to app developers to advise on learning principles they may want to incorporate? Would app developers then agree to help us in collecting data to enable research and further improvements? A few edTech startups have been built with this reciprocal arrangement in mind, with both researchers and app developers on board (e.g., Cognitive Toybox and egoTeach, my own attempt), and can find support from industry and other funding sources that are not possible without an industry partner. I also encourage researchers to look at popular education apps, evaluate their scientific basis (for examples, see Hirsh-Pasek et al. 2015) and to reach out to app developers and discuss collaboration.

Organizational Challenges

When a researcher is faced with deploying an app for testing, there are diverse approaches that result in distinct types of datasets. Deploying in the competitive environment of the app stores will often result in slow (or no) adoption and perhaps not very extended use, unless the app is quite engaging. Even a large naturalistic dataset can present its own difficulties: most users contribute little data, so usage- and age-matched comparisons can be a challenge and long learning trajectories are hard to come by. Aside from recruiting and running individual participants for short-term studies, are there ways to engage schools for larger-scale, longer-term testing? What types of apps and content would teachers most welcome? How can we work together as teams of researchers along with educators to create more universal learning aids that also allow us to gather data?

Scientific Challenges

Aside from further testing strategies such as distributed practice and retrieval practice or investigating new techniques, there is still much to be learned about which methods can be synergistically combined, and which might interfere with each other. For example, some recent work has investigated the relationship between curiosity, motivation, and learning in educational apps (Oudeyer, Gottlieb, & Lopes, 2016). Which learning techniques can be implemented and studied jointly in an app? Another issue is how to measure generalization of the skills outside the app--in the classroom, or in other apps. Once there are multiple researchers field testing, say, numeracy apps, how can they coordinate recruitment to ensure that they have different participants? With a population of diverse participants on the app stores, with varied backgrounds and interests, it is also interesting to consider if individual differences in cognitive abilities might influence the effectiveness of some strategies, as well as participants' behavior and motivation.

Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J., and Willingham, D. T. (2013). Improving Students' Learning With Effective Learning Techniques: Promising Directions From Cognitive and Educational Psychology. Psychological Science in the Public Interest, 14(1), pp. 4–58.

Hirsh-Pasek, K., Zosh, J. M., Golinkoff, R. M., Gray, J. H., Robb, M. B., and Kaufman, J. (2015). Putting education in "educational" apps: Lessons from the science of learning. Psychological Science in the Public Interest, 16(1), pp. 3–34.

Koedinger, K., & Corbett, A. (2005). Cognitive Tutors. In R. Sawyer (Ed.), The Cambridge Handbook of the Learning Sciences (Cambridge Handbooks in Psychology, pp. 61-78). Cambridge: Cambridge University Press. doi:10.1017/CBO9780511816833.006

Oudeyer, P.-Y., Gottlieb, J., and Lopes, M. (2016). Intrinsic motivation, curiosity, and learning: Theory and applications in educational technologies. Progress in Brain Research, 229, pp. 257–284.

Roediger, H. L. (2013). Applying cognitive psychology to education: Translational educational science. Psychological Science in the Public Interest, 14(1), pp. 1–3.

IEEE TCDS Table of Contents

Volume 10, Issue 1, March 2018

Guest Editorial Special Issue on Affordances

Lorenzo Jamone, Emre Ugur, José Santos-Victor

The concept of affordances appeared in psychology during the late 60s, as an alternative perspective on the visual perception of the environment. More precisely, the term affordances was introduced by J. J. Gibson in 1966 [items 1)-3) of the Appendix]. Gibson defines the affordances of an object as "what it offers the animal, what it provides or furnishes, either for good or ill." What the objects afford to the agent are action possibilities that are directly perceived through vision, in a precategorical and subconscious way, without the need to construct a fully detailed model of the world or to perform semantic reasoning or explicit object recognition. Central to Gibson's theory is the notion that the sensorimotor capabilities of the agent dramatically influence perception: the concept of affordances is "something that refers to both the environment and the animal in a way that no existing term does. It implies the complementarity of the animal and the environment." Notably, the human ability to perceive affordances emerges gradually during development, and it is the outcome of exploratory and observational learning; as this ability appears, predicting the effects of the actions becomes possible as well, eventually leading to problem solving skills.

Affordances in Psychology, Neuroscience, and Robotics: A Survey

Lorenzo Jamone, Emre Ugur, Angelo Cangelosi, Luciano Fadiga, Alexandre Bernardino, Justus Piater, José Santos-Victor

The concept of affordances appeared in psychology during the late 60s as an alternative perspective on the visual perception of the environment. It was revolutionary in the intuition that the way living beings perceive the world is deeply influenced by the actions they are able to perform. Then, across the last 40 years, it has influenced many applied fields, e.g., design, human-computer interaction, computer vision, and robotics. In this paper, we offer a multidisciplinary perspective on the notion of affordances. We first discuss the main definitions and formalizations of the affordance theory, then we report the most significant evidence in psychology and neuroscience that support it, and finally we review the most relevant applications of this concept in robotics.

A Logic-Based Computational Framework for Inferring Cognitive Affordances

Vasanth Sarathy, Matthias Scheutz

The concept of "affordance" refers to the relationship between human perceivers and aspects of their environment. Being able to infer affordances is central to commonsense reasoning, tool use and creative problem solving in artificial agents. Existing approaches to inferring affordances have focused on functional aspects, relying on either static ontologies or statistical formalisms to extract relationships between physical features of objects, actions, and the corresponding effects of their interaction. These approaches do not provide flexibility with which to reason about affordances in the open world, where affordances are influenced by changing context, social norms, historical precedence, and uncertainty. We develop a computational framework comprising a probabilistic rules-based logical representation coupled with a computational architecture (cognitive affordances logically expressed) to reason about affordances in a more general manner than described in the existing literature. Our computational architecture allows robotic agents to make deductive and abductive inferences about functional and social affordances, collectively and dynamically, thereby allowing the agent to adapt to changing conditions. We demonstrate our approach with experiments, and show that an agent can successfully reason through situations that involve a tight interplay between various social and functional norms.

Toward Lifelong Affordance Learning Using a Distributed Markov Model

Arren J. Glover, Gordon F. Wyeth

Robots are able to learn how to interact with objects by developing computational models of affordance. This paper presents an approach in which learning and operation occur concurrently, toward achieving lifelong affordance learning. In a such a regime a robot must be able to learn about new objects, but without a general rule for what an "object" is, the robot must learn about everything in the environment to determine their affordances. In this paper, sensorimotor coordination is modeled using a distributed semi-Markov decision process; it is created online during robot operation, and performs continual action selection to reach a goal state. In an initial experiment we show that this model captures an object's affordances, which are exploited to perform several different tasks using a mobile robot equipped with a gripper and infrared "tactile" sensor. In a secondary experiment, we show that the robot can learn that the marker is the only visual feature that can be gripped and that walls and floor do not have the affordance of being "grip-able." The distributed mechanism is necessary for the modeling of multiple sensory stimuli simultaneously, and the selection of the object with the necessary affordances for the task is emergent from the robot's actions, while other parts of the environment that are perceived, such as walls and floors, are ignored.

Bootstrapping Relational Affordances of Object Pairs Using Transfer

Severin Fichtl, Dirk Kraft, Norbert Krüger, Frank Guerin

Robots acting in everyday environments need a good knowledge of how a manipulation action can affect pairs of objects in a relationship, such as "inside" or "behind" or "on top." These relationships afford certain means-end actions such as pulling a container to retrieve the contents, or pulling a tool to retrieve a desired object. We investigate how these relational affordances could be learned by a robot from its own action experience. A major challenge in this approach is to reduce the number of training samples needed to achieve accuracy, and hence we investigate an approach which can leverage past knowledge to accelerate current learning (which we call bootstrapping). We learn random forest-based affordance predictors from visual inputs and demonstrate two approaches to knowledge transfer for bootstrapping. In the first approach [direct bootstrapping (DB)], the state-space for a new affordance predictor is augmented with the output of previously learned affordances. In the second approach [category-based bootstrapping (CB)], we form categories that capture underlying commonalities of a pair of existing affordances and augment the state-space with this category classifier's output. In addition, we introduce a novel heuristic, which suggests how a large set of potential affordance categories can be pruned to leave only those categories which are most promising for bootstrapping future affordances. Our results show that both bootstrapping approaches outperform learning without bootstrapping. We also show that there is no significant difference in performance between DB and CB.

A Modular Dynamic Sensorimotor Model for Affordances Learning, Sequences Planning, and Tool-Use

Raphael Braud, Alexandre Pitti, Philippe Gaussier

This paper proposes a computational model for learning robot control and sequence planning based on the ideomotor principle. This model encodes covariation laws between sensors and motors in a modular fashion and exploits these primitive skills to build complex action sequences, potentially involving tool-use. Implemented for a robotic arm, the model starts with raw unlabeled sensor and motor vectors and autonomously assigns functions to neutral objects in the environment. Our experimental evaluation highlights the emergent properties of such a modular system and we discuss their consequences from ideomotor and sensorimotor-theoretic perspectives.

Artificial Life Environment Modeled by Dynamic Fuzzy Cognitive Maps

Lúcia Valéria R. Arruda, Márcio Mendonça, Flávio Neves, Ivan Rossato Chrun, Elpiniki I. Papageorgiou

This paper presents an artificial life environment based on dynamic fuzzy cognitive maps (DFCMs) and inspired by multiagent systems, machine learning, and concepts from classical fuzzy cognitive map theory. The proposed architecture includes features such as a reinforcement learning algorithm to dynamically fine-tune the weights of the DFCM, a finite states machine, governing the behavior of the creatures by adding/removing concepts into the DFCM, and others. These features are used to add adaptability to the artificial creatures (agents) in a simulated hunter-prey environment with synthetic data. Some experiments carried out in a simulated virtual environment have shown promising results for further research in the subject of this paper.

Bootstrapping Q-Learning for Robotics From Neuro-Evolution Results

Matthieu Zimmer, Stephane Doncieux

Reinforcement learning (RL) problems are hard to solve in a robotics context as classical algorithms rely on discrete representations of actions and states, but in robotics both are continuous. A discrete set of actions and states can be defined, but it requires an expertise that may not be available, in particular in open environments. It is proposed to define a process to make a robot build its own representation for an RL algorithm. The principle is to first use a direct policy search in the sensori-motor space, i.e., with no predefined discrete sets of states nor actions, and then extract from the corresponding learning traces discrete actions and identify the relevant dimensions of the state to estimate the value function. Once this is done, the robot can apply RL: 1) to be more robust to new domains and, if required and 2) to learn faster than a direct policy search. This approach allows to take the best of both worlds: first learning in a continuous space to avoid the need of a specific representation, but at a price of a long learning process and a poor generalization, and then learning with an adapted representation to be faster and more robust.

Volume 10, Issue 2, June 2018

Guest Editorial Special Issue on Neuromorphic Computing and Cognitive Systems

Huajin Tang, Tiejun Huang, Jeffrey L. Krichmar, Garrick Orchard, Arindam Basu

In recent years, neuromorphic computing has become an important emerging research area. Neuromorphic computing takes advantage of computer architectures and sensors whose design and functionality are inspired by the brain. There has been rapid progress in computational theory, spiking neurons, learning algorithms, signal processing, circuit design and implementation, which have shown appealing computational advantages over conventional solutions [1]. The low size, weight, and power of these hardware architectures shows great potential for embedded cognitive systems [2].

Adaptive Robot Path Planning Using a Spiking Neuron Algorithm With Axonal Delays

Tiffany Hwu, Alexander Y. Wang, Nicolas Oros, Jeffrey L. Krichmar

A path planning algorithm for outdoor robots, which is based on neuronal spike timing, is introduced. The algorithm is inspired by recent experimental evidence for experience-dependent plasticity of axonal conductance. Based on this evidence, we developed a novel learning rule that altered axonal delays corresponding to cost traversals and demonstrated its effectiveness on real-world environmental maps. We implemented the spiking neuron path planning algorithm on an autonomous robot that can adjust its routes depending on the context of the environment. The robot demonstrates the ability to plan different trajectories that exploit smooth roads when energy conservation is advantageous, or plan the shortest path across a grass field when reducing distance traveled is beneficial. Because the algorithm is suitable for spike-based neuromorphic hardware, it has the potential of realizing orders of magnitude gains in power efficiency and computational gains through parallelization.

Neuro-Activity-Based Dynamic Path Planner for 3-D Rough Terrain

Azhar Aulia Saputra, Yuichiro Toda, János Botzheim, Naoyuki Kubota

This paper presents a natural mechanism of the human brain for generating a dynamic path planning in 3-D rough terrain. The proposed paper not only emphasizes the inner state process of the neuron but also the development process of the neurons in the brain. There are two algorithm processes in this proposed model, the forward transmission activity for constructing the neuron connections to find the possible way and the synaptic pruning activity with backward neuron transmission for finding the best pathway from current position to target position and reducing inefficient neuron with its synaptic connections. In order to respond and avoid the unpredictable obstacle, dynamic path planning is also considered in this proposed model. An integrated system for applying the proposed model in the real cases is also presented. In order to prove the effectiveness of the proposed model, we applied it in the pathway of a four-legged robot on rough terrain in both computer simulation and real cases. Unpredictable collision is also performed in those experiments. The model can find the best pathway and facilitate the safe movement of the robot. When the robot found an unpredictable collision, the path planner dynamically changed the pathway. The proposed path planning model is capable to be applied in further advance implementation.

EMPD: An Efficient Membrane Potential Driven Supervised Learning Algorithm for Spiking Neurons

Malu Zhang, Hong Qu, Ammar Belatreche, Xiurui Xie

The brain-inspired spiking neurons, considered as the third generation of artificial neurons, are more biologically plausible and computationally powerful than traditional artificial neurons. One of the fundamental research in spiking neurons is to transform streams of incoming spikes into precisely timed spikes. Due to the inherent complexity of processing spike sequences, the formulation of efficient supervised learning algorithm is difficult and remains an important problem in the research area. This paper presents an efficient membrane potential driven (EMPD) supervised learning method capable of training neurons to generate desired sequences of spikes. The learning rule of EMPD is composed of two processes: 1) at desired output times, the gradient descent method is implemented to minimize the error function defined as the difference between the membrane potential and the firing threshold and 2) at undesired output time, synaptic weights are adjusted to make the membrane potential below the threshold. For efficiency, at undesired output times, EMPD calculates the membrane potential and makes a comparison with firing threshold only at some special time points when the neuron is most likely to cross the firing threshold. Experimental results show that the proposed EMPD approach has higher learning efficiency and accuracy over the existing learning algorithms.

Robotic Homunculus: Learning of Artificial Skin Representation in a Humanoid Robot Motivated by Primary Somatosensory Cortex

Matej Hoffmann, Zdeněk Straka, Igor Farkaš, Michal Vavrečka, Giorgio Metta

Using the iCub humanoid robot with an artificial pressure-sensitive skin, we investigate how representations of the whole skin surface resembling those found in primate primary somatosensory cortex can be formed from local tactile stimulations traversing the body of the physical robot. We employ the well-known self-organizing map algorithm and introduce its modification that makes it possible to restrict the maximum receptive field (MRF) size of neuron groups at the output layer. This is motivated by findings from biology where basic somatotopy of the cortical sheet seems to be prescribed genetically and connections are localized to particular regions. We explore different settings of the MRF and the effect of activity-independent (input-output connections constraints implemented by MRF) and activity-dependent (learning from skin stimulations) mechanisms on the formation of the tactile map. The framework conveniently allows one to specify prior knowledge regarding the skin topology and thus to effectively seed a particular representation that training shapes further. Furthermore, we show that the MRF modification facilitates learning in situations when concurrent stimulation at nonadjacent places occurs ("multitouch"). The procedure was sufficiently robust and not intensive on the data collection and can be applied to any robots where representation of their "skin" is desirable.

A Novel Parsimonious Cause-Effect Reasoning Algorithm for Robot Imitation and Plan Recognition

Garrett Katz, Di-Wei Huang, Theresa Hauge, Rodolphe Gentili, James Reggia

Manually programming robots is difficult, impeding more widespread use of robotic systems. In response, efforts are being made to develop robots that use imitation learning. With such systems a robot learns by watching humans perform tasks. However, most imitation learning systems replicate a demonstrator's actions rather than obtaining a deeper understanding of why those actions occurred. Here we introduce an imitation learning framework based on causal reasoning that infers a demonstrator's intentions. As with imitation learning in people, our approach constructs an explanation for a demonstrator's actions, and generates a plan based on this explanation to carry out the same goals rather than trying to faithfully reproduce the demonstrator's precise motor actions. This enables generalization to new situations. We present novel causal inference algorithms for imitation learning and establish their soundness, completeness and complexity characteristics. Our approach is validated using a physical robot, which successfully learns and generalizes skills involving bimanual manipulation. Human performance on similar skills is reported. Computer experiments using the Monroe Plan Corpus further validate our approach. These results suggest that causal reasoning is an effective unifying principle for imitation learning. Our system provides a platform for exploring neural implementations of this principle in future work.

Predicting Spike Trains from PMd to M1 Using Discrete Time Rescaling Targeted GLM

Dong Xing, Cunle Qian, Hongbao Li, Shaomin Zhang, Qiaosheng Zhang, Yaoyao Hao, Xiaoxiang Zheng, Zhaohui Wu, Yiwen Wang, Gang Pan

The computational model for spike train prediction with inputs from other related cerebral cortices is important in revealing the underlying connection among different cortical areas. To evaluate goodness-of-fit of the model, the time rescaling Kolmogorov-Smirnov (KS) statistic is usually applied, of which the calculation is separated from optimization procedure of the model. If the KS statistic could be embedded into objective function of the optimization procedure, precision of the firing probability series generated by the model would be increased directly. This paper presents a linear-nonlinear-Poisson cascade framework for prediction of spike trains, whose objective function is changed from maximizing log-likelihood of the spike trains to minimizing the penalization of discrete time rescaling KS statistic to eliminate the separation between optimization and evaluation of the model. We apply our model on the task of predicting firing probability of neurons from primary motor cortex with spike trains from dorsal premotor cortex as input, which are two cerebral cortices associated with movements planning and executing. The experimental results show that by introducing the goodness-of-fit metric into the objective function, results of the model will gain a significant improvement, which outperforms the state of the art.

Visual Pattern Recognition Using Enhanced Visual Features and PSD-Based Learning Rule

Xiaoliang Xu, Xin Jin, Rui Yan, Qiming Fang, Wensi Lu

This paper proposes a feedforward visual pattern recognition model based on a spiking neural network (SNN). The proposed model mainly includes four functional layers: 1) feature extraction; 2) encoding; 3) learning; and 4) readout. A modified HMAX model is first presented to extract features from external stimuli. In order to reduce the computational cost, we simplify the S1 layer using a single Gabor filter window. To simulate biological vision's sensitivity to the vertical direction, we strengthen the feature of filtered orientation in 90° by adding a sharpened replica of the filtered image in 90° before max pooling in C1 layer. Then the phase encoding approach is used to convert the extracted visual features into spike patterns. These spike patterns will be learned by precise-spike-driven-based learning rules in an SNN. Finally, experimental results on benchmark datasets including MNIST, Caltech 101, and optical characters demonstrate the efficiency and robustness in noisy environments of the proposed model.

Multimodal Functional and Structural Brain Connectivity Analysis in Autism: A Preliminary Integrated Approach With EEG, fMRI, and DTI

Bogdan Alexandru Cociu, Saptarshi Das, Lucia Billeci, Wasifa Jamal, Koushik Maharatna, Sara Calderoni, Antonio Narzisi, Filippo Muratori

This paper proposes a novel approach of integrating different neuroimaging techniques to characterize an autistic brain. Different techniques like electroencephalography (EEG), functional magnetic resonance imaging (fMRI), and diffusion tensor imaging (DTI) have traditionally been used to find biomarkers for autism, but there have been very few attempts for a combined or multimodal approach of EEG, fMRI, and DTI to understand the neurobiological basis of autism spectrum disorder (ASD). Here, we explore how the structural brain network correlate with the functional brain network, such that the information encompassed by these two could be uncovered only by using the latter. In this paper, source localization from EEG using independent component analysis and dipole fitting has been applied first, followed by selecting those dipoles that are closest to the active regions identified with fMRI. This allows translating the high temporal resolution of EEG to estimate time varying connectivity at the spatial source level. Our analysis shows that the estimated functional connectivity between two active regions can be correlated with the physical properties of the structure obtained from DTI analysis. This constitutes a first step toward opening the possibility of using pervasive EEG to monitor the long-term impact of ASD treatment without the need for frequent expensive fMRI or DTI investigations.

Observing and Modeling Developing Knowledge and Uncertainty During Cross-Situational Word Learning

George Kachergis, Chen Yu

Being able to learn word meanings across multiple scenes consisting of multiple words and referents (i.e., cross-situationally) is thought to be important for language acquisition. The ability has been studied in infants, children, and adults, and yet there is much debate about the basic storage and retrieval mechanisms that operate during cross-situational word learning. It has been difficult to uncover the learning mechanics in part because the standard experimental paradigm, which presents a few words and objects on each of a series of training trials, measures learning only at the end of training, after several occurrences of each word-object pair. Diverse models are able to match the final level of performance of the standard paradigm, while the rich history and context of the learning trajectories remain obscured. This paper examines accuracy and uncertainty over time in a version of the cross-situational learning task in which words are tested throughout training, as well as in a final test. With similar levels of performance to the standard task, we examine how well the online response trajectories match recent hypothesis- and association-based computational models of word learning.

Prediction Error in the PMd As a Criterion for Biological Motion Discrimination: A Computational Account

Yuji Kawai, Yukie Nagai, Minoru Asada

Neuroscientific studies suggest that the dorsal premotor area is activated by biological motions, and is also related to the prediction errors of observed and self-induced motions. We hypothesize that biological and nonbiological motions can be discriminated by such prediction errors. We therefore propose a model to verify this hypothesis. A neural network model is constructed that learns to predict the velocity of the self's next body movement from that of the present one and produces a smooth movement. Consequently, a property of the input sequence is represented. The trained network evaluates observed motions based on the prediction errors. If these errors are small, the movements share a representation with the self-motor property, and therefore, are regarded as biological ones. To verify our hypothesis, we examined how the network represents the biological motions. The results show that predictive learning, supported by a recurrent structure, helps to obtain the representation that discriminates between biological and nonbiological motions. Moreover, this recurrent neural network can discriminate the ankle and wrist trajectories of a walking human as biological motion, regardless of the subject's sex, or emotional state.

Learning 4-D Spatial Representations Through Perceptual Experience With Hypercubes

Takanobu Miwa, Yukihito Sakai, Shuji Hashimoto

Imagine a day when humans can form mental representations of higher-dimensional space and objects. These higher-dimensional spatial representations may enable us to gain unique insights into scientific and cultural advancements. To augment human spatial cognition from three to four dimensions, we have developed an interactive 4-D visualization system for acquiring an understanding of 4-D space and objects. In this paper, we examine whether humans are capable of formulating 4-D spatial representations through perceptual experience in 4-D space with 4-D objects. Participants learn about 4-D space and hypercubes through an interactive system, and are then examined on a series of 4-D spatial ability tests. They demonstrate the ability to perform perspective taking, navigation, and mental spatial transformation tasks in 4-D space. The results provide empirical evidence that humans are capable of learning 4-D spatial representations. Moreover, the results support the interpretation that humans form a cognitive coordinate system, consisting of an origin and four directional axes, to understand 4-D space and objects.

Fuzzy Feature Extraction for Multichannel EEG Classification

Pei-Yuan Zhou, Keith C. C. Chan

EEG signals (EEGs) are usually collected by placing multiple electrodes at various positions along the scalp as multichannel data. Given that many channels are collected for each single-trial, the multichannel EEG classification problem can be treated as multivariate time series classification problem. For multichannel EEG data to be more accurately classified, we propose an algorithm, called the fuzzy multichannel EEG classifier (FMCEC). This algorithm can take into consideration the interaction among different signals collected at different time instants and locations on the skull when constructing a classifier. The FMCEC first preprocess raw EEG data by eliminating noise by discretization of the data. It then performs fuzzification of the resulting discretized data to capture imprecision and vagueness in the data. Given the fuzzified data, FMCEC then discovers intrachannel patterns within each channel and then interchannel patterns between different channels of EEGs. The discovered patterns, which are represented as fuzzy temporal patterns, are then used to characterize and differentiate between different classes of multichannel EEG data. To evaluate the effectiveness of FMCEC, we tested it with several sets of real EEG datasets. The results show that the algorithm can be a promising tool for the classification of multichannel EEG data.

Orthogonal Principal Coefficients Embedding for Unsupervised Subspace Learning

Xinxing Xu, Shijie Xiao, Zhang Yi, Xi Peng, Yong Liu

As a recently proposed method for subspace learning, principal coefficients embedding (PCE) method can automatically determine the dimension of the feature space and robustly handle various corruptions in real-world applications. However, the projection matrix learned by PCE is not orthogonal, so the original data may be reconstructed improperly. To address this issue, we proposed a new method termed orthogonal PCE (OPCE). OPCE cannot only automatically determine the dimension of the feature space, but also additionally considers the orthogonal property of the projection matrix for better discriminating ability. Moreover, OPCE can be solved in closed-form, thus making it computational efficient. Extensive experimental results from multiple benchmark data sets demonstrate the effectiveness and computational efficiency of the proposed method.

A Basal Ganglia Network Centric Reinforcement Learning Model and Its Application in Unmanned Aerial Vehicle

Yi Zeng, Guixiang Wang, Bo Xu

Reinforcement learning brings flexibility and generality for machine learning, while most of them are mathematical optimization driven approaches, and lack of cognitive and neural evidence. In order to provide a more cognitive and neural mechanisms driven foundation and validate its applicability in complex task, we develop a basal ganglia (BG) network centric reinforcement learning model. Compared to existing work on modeling BG, this paper is unique from the following perspectives: 1) the orbitofrontal cortex (OFC) is taken into consideration. OFC is critical in decision making because of its responsibility for reward representation and is critical in controlling the learning process, while most of the BG centric models do not include OFC; 2) to compensate the inaccurate memory of numeric values, precise encoding is proposed to enable working memory system remember important values during the learning process. The method combines vector convolution and the idea of storage by digit bit and is efficient for accurate value storage; and 3) for information coding, the Hodgkin-Huxley model is used to obtain a more biological plausible description of action potential with plenty of ionic activities. To validate the effectiveness of the proposed model, we apply the model to the unmanned aerial vehicle (UAV) autonomous learning process in a 3-D environment. Experimental results show that our model is able to give the UAV the ability of free exploration in the environment and has comparable learning speed as the Q learning algorithm, while the major advances for our model is that it is with solid cognitive and neural basis.

Biologically Inspired Self-Organizing Map Applied to Task Assignment and Path Planning of an AUV System

Daqi Zhu, Xiang Cao, Bing Sun, Chaomin Luo

An integrated biologically inspired self-organizing map (SOM) algorithm is proposed for task assignment and path planning of an autonomous underwater vehicle (AUV) system in 3-D underwater environments with obstacle avoidance. The algorithm embeds the biologically inspired neural network (BINN) into the SOM neural networks. The task assignment and path planning aim to arrange a team of AUVs to visit all appointed target locations, while assuring obstacle avoidance without speed jump. The SOM neuron network is developed to assign a team of AUVs to achieve multiple target locations in underwater environments. Then, in order to avoid obstacles and speed jump for each AUV that visits the corresponding target location, the BINN is utilized to update weights of the winner of SOM, and achieve AUVs path planning and effective navigation. The effectiveness of the proposed hybrid model is validated by simulation studies.

Autonomous Discovery of Motor Constraints in an Intrinsically Motivated Vocal Learner

Juan Manuel Acevedo-Valle, Cecilio Angulo, Clément Moulin-Frier

This paper introduces new results on the modeling of early vocal development using artificial intelligent cognitive architectures and a simulated vocal tract. The problem is addressed using intrinsically motivated learning algorithms for autonomous sensorimotor exploration, a kind of algorithm belonging to the active learning architectures family. The artificial agent is able to autonomously select goals to explore its own sensorimotor system in regions, where its competence to execute intended goals is improved. We propose to include a somatosensory system to provide a proprioceptive feedback signal to reinforce learning through the autonomous discovery of motor constraints. Constraints are represented by a somatosensory model which is unknown beforehand to the learner. Both the sensorimotor and somatosensory system are modeled using Gaussian mixture models. We argue that using an architecture which includes a somatosensory model would reduce redundancy in the sensorimotor model and drive the learning process more efficiently than algorithms taking into account only auditory feedback. The role of this proposed system is to predict whether an undesired collision within the vocal tract under a certain motor configuration is likely to occur. Thus, compromised motor configurations are rejected, guaranteeing that the agent is less prone to violate its own constraints.

Bio-Inspired Model Learning Visual Goals and Attention Skills Through Contingencies and Intrinsic Motivations

Valerio Sperati, Gianluca Baldassarre

Animal learning is driven not only by biological needs but also by intrinsic motivations (IMs) serving the acquisition of knowledge. Computational modeling involving IMs is indicating that learning of motor skills requires that autonomous agents self-generate tasks/goals and use them to acquire skills solving/leading to them. We propose a neural architecture driven by IMs that is able to self-generate goals on the basis of the environmental changes caused by the agent's actions. The main novelties of the model are that it is focused on the acquisition of attention (looking) skills and that its architecture and functioning are broadly inspired by the functioning of relevant primate brain areas (superior colliculus, basal ganglia, and frontal cortex). These areas, involved in IM-based behavior learning, play important functions for reflexive and voluntary attention. The model is tested within a simple simulated pan-tilt camera robot engaged in learning to switch on different lights by looking at them, and is able to self-generate visual goals and learn attention skills under IM guidance. The model represents a novel hypothesis on how primates and robots might autonomously learn attention skills and has a potential to account for developmental psychology experiments and the underlying brain mechanisms.

Seamless Integration and Coordination of Cognitive Skills in Humanoid Robots: A Deep Learning Approach

Jungsik Hwang, Jun Tani

This paper investigates how adequate coordination among the different cognitive processes of a humanoid robot can be developed through end-to-end learning of direct perception of visuomotor stream. We propose a deep dynamic neural network model built on a dynamic vision network, a motor generation network, and a higher-level network. The proposed model was designed to process and to integrate direct perception of dynamic visuomotor patterns in a hierarchical model characterized by different spatial and temporal constraints imposed on each level. We conducted synthetic robotic experiments in which a robot learned to read human's intention through observing the gestures and then to generate the corresponding goal-directed actions. Results verify that the proposed model is able to learn the tutored skills and to generalize them to novel situations. The model showed synergic coordination of perception, action, and decision making, and it integrated and coordinated a set of cognitive skills including visual perception, intention reading, attention switching, working memory, action preparation, and execution in a seamless manner. Analysis reveals that coherent internal representations emerged at each level of the hierarchy. Higher-level representation reflecting actional intention developed by means of continuous integration of the lower-level visuo-proprioceptive stream.

Learning Temporal Intervals in Neural Dynamics

Boris Duran, Yulia Sandamirskaya

Storing and reproducing temporal intervals is an important component of perception, action generation, and learning. How temporal intervals can be represented in neuronal networks is thus an important research question both in study of biological organisms and artificial neuromorphic systems. Here, we introduce a neural-dynamic computing architecture for learning temporal durations of actions. The architecture uses a dynamic neural fields (DNFs) representation of the elapsed time and a memory trace dynamics to store the experienced action duration. Interconnected dynamical nodes signal beginning of an action, its successful accomplishment, or failure, and activate formation of the memory trace that corresponds to the action's duration. The accumulated memory trace influences the competition between the dynamical nodes in such a way that the failure node gains a competitive advantage earlier if the stored duration is shorter. The model uses neurally based DNF dynamics and is a process model of how temporal durations may be stored in neural systems, both biological and artificial ones. The focus of this paper is on the mechanism to store and use duration in artificial neuronal systems. The model is validated in closed-loop experiments with a simulated robot.

Quantifying Cognitive Workload in Simulated Flight Using Passive, Dry EEG Measurements

Justin A. Blanco, Michael K. Johnson, Kyle J. Jaquess, Hyuk Oh, Li-Chuan Lo, Rodolphe J. Gentili, Bradley D. Hatfield

A reliable method for quantifying cognitive workload in pilots could find uses in flight training and scheduling, cockpit design, and improving flight safety. Many proposed methods for monitoring cognitive workload in this population rely on measuring physiological responses to externally delivered probe stimuli and/or use traditional gel-based electroencephalography (EEG) sensors. Here we develop passive, probe-independent algorithms for classifying three levels of flight task complexity based on 4-channel, gel-free EEG during simulated flight. Using a library of 168 input features drawn from different data science application domains, we evaluated 13 different classifiers, using a nested tenfold cross-validation procedure to estimate generalization performance. The best subsets of features yielded a median classification accuracy of 90.17% across subjects, with perfect accuracy in one subject and greater than 75% in 16 of 21 subjects. Though EEG line length and linear discriminant analysis were generally among the most effective features and classifiers, respectively, we find that to maximize prediction accuracy, feature set-classifier combinations should be individualized. No single channel proved more valuable than another in predicting flight task complexity, but combining EEG features across channels maintained or improved performance in 81% of subjects.

Enhanced Robotic Hand/Eye Coordination Inspired From Human-Like Behavioral Patterns

Fei Chao, Zuyuan Zhu, Chih-Min Lin, Huosheng Hu, Longzhi Yang, Changjing Shang, Changle Zhou

Robotic hand-eye coordination is recognized as an important skill to deal with complex real environments. Conventional robotic hand-eye coordination methods merely transfer stimulus signals from robotic visual space to hand actuator space. This paper introduces a reverse method. Build another channel that transfers stimulus signals from robotic hand space to visual space. Based on the reverse channel, a human-like behavior pattern: "Stop-to-Fixate," is imparted to the robot, thereby giving the robot an enhanced reaching ability. A visual processing system inspired by the human retina structure is used to compress visual information so as to reduce the robot's learning complexity. In addition, two constructive neural networks establish the two sensory delivery channels. The experimental results demonstrate that the robotic system gradually obtains a reaching ability. In particular, when the robotic hand touches an unseen object, the reverse channel successfully drives the visual system to notice the unseen object.

Covariate Conscious Approach for Gait Recognition Based Upon Zernike Moment Invariants

Himanshu Aggarwal, Dinesh Kumar Vishwakarma

Gait recognition, i.e., identification of an individual from his/her walking pattern is an emerging field. While existing gait recognition techniques perform satisfactorily in normal walking conditions, their performance tend to suffer drastically with variations in clothing and carrying conditions. In this paper, we propose a novel covariate cognizant framework to deal with the presence of such covariates. We describe gait motion by forming a single 2-D spatio-temporal template from video sequence, called average energy silhouette image (AESI). Zernike moment invariants are then computed to screen the parts of AESI infected with covariates. Following this, features are extracted from spatial distribution of oriented gradients and novel mean of directional pixels methods. The obtained features are fused together to form the final well-endowed feature set. Experimental evaluation of the proposed framework on three publicly available datasets, i.e., CASIA Dataset B, OU-ISIR Treadmill Dataset B, and USF Human-ID challenge dataset with recently published gait recognition approaches, prove its superior performance.

EEG-Based Emotion Recognition Using Hierarchical Network With Subnetwork Nodes

Yimin Yang, Q. M. Jonathan Wu, Wei-Long Zheng, Bao-Liang Lu

Emotions play a crucial role in decision-making, brain activity, human cognition, and social intercourse. This paper proposes a hierarchical network structure with subnetwork nodes to discriminate three human emotions: 1) positive; 2) neutral; and 3) negative. Each subnetwork node embedded in the network that are formed by hundreds of hidden nodes, could be functional as an independent hidden layer for feature representation. The top layer of the hierarchical network, like the mammal cortex in the brain, combine such features generated from subnetwork nodes, but simultaneously, recast these features into a mapping space so that the network can be performed to produce more reliable cognition. The proposed method is compared with other state-of-the-art methods. The experimental results from two different EEG datasets show that a promising result is obtained when using the proposed method with both single and multiple modality.

A Novel Biologically Inspired Visual Cognition Model: Automatic Extraction of Semantics, Formation of Integrated Concepts, and Reselection Features for Ambiguity

Peijie Yin ; Hong Qiao ; Wei Wu, Lu Qi, Yinlin Li, Shanlin Zhong, Bo Zhang

Techniques that integrate neuroscience and information science benefit both fields. Many related models have been proposed in computer vision; however, in general, the robustness and recognition precision are still key problems in object recognition models. In this paper, inspired by the process by which humans recognize objects and its biological mechanisms, a new integrated and dynamic framework is proposed that mimics the semantic extraction, concept formation and feature reselection found in human visual processing. The main contributions of the proposed model are as follows: 1) semantic feature extraction: local semantic features are learned from episodic features extracted from raw images using a deep neural network; 2) integrated concept formation: concepts are formed using the local semantic information and structural information is learned through a network; and 3) feature reselection: when ambiguity is detected during the recognition process, distinctive features based on the differences between the ambiguous candidates are reselected for recognition. Experimental results on four datasets show that-compared with other methods-the new proposed model is more robust and achieves higher precision for visual recognition, especially when the input samples are semantically ambiguous. Meanwhile, the introduced biological mechanisms further strengthen the interaction between neuroscience and information science.

Zero-Shot Image Classification Based on Deep Feature Extraction

Xuesong Wang, Chen Chen, Yuhu Cheng, Z. Jane Wang

The attribute-based zero-shot learning methods generally use low-level features of images to train attribute classifiers, and the corresponding classification accuracy heavily depends on specific low-level features. Because deep networks can automatically extract features from original unlabeled images and the extracted features can better represent the nature of original images, we proposed a zero-shot image classification method based on deep feature extraction. In the image preprocessing step, in order to reduce the computational complexity and the correlations between pixels, image patches extraction and zero-phase component analysis whitening are performed. The compressed feature representations of unlabeled image patches are learned through a stacked sparse autoencoder and a feature mapping matrix can be obtained. Further, we use the feature mapping matrix as a convolution kernel to convolve with image patches. Since the convolution operation results in the feature vector with huge dimensionality, the convolution features will be pooled to reduce the number of network parameters and to reduce the spatial resolution of the network to prevent over-fitting. Finally, the exacted image features are used to train the conventional indirect attribute prediction model to predict image attributes and classify images under the zero-shot setting. Experimental results on the shoes, outdoor scene recognition, and a-Yahoo datasets show that, compared with several popular zero-shot learning methods, the proposed method can yield more accurate attribute prediction and better zero-shot image classification.

A Hormone-Driven Epigenetic Mechanism for Adaptation in Autonomous Robots

John Lones, Matthew Lewis, Lola Cañamero

Different epigenetic mechanisms provide biological organisms with the ability to adjust their physiology and/or morphology and adapt to a wide range of challenges posed by their environments. In particular, one type of epigenetic process, in which hormone concentrations are linked to the regulation of hormone receptors, has been shown to have implications for behavioral development. In this paper, taking inspiration from these biological processes, we investigate whether an epigenetic model based on the concept of hormonal regulation of receptors can provide a similarly robust and general adaptive mechanism for autonomous robots. We have implemented our model using a Koala robot, and tested it in a series of experiments in six different environments with varying challenges to negotiate. Our results, including the emergence of varied behaviors that permit the robot to exploit its current environment, demonstrate the potential of our epigenetic model as a general mechanism for adaptation in autonomous robots.

Heteroscedastic Regression and Active Learning for Modeling Affordances in Humanoids

Francesca Stramandinoli, Vadim Tikhanoff, Ugo Pattacini, Francesco Nori

We investigate the learning of object and tool affordances in the iCub robot. We adopt the setup proposed in a previous experiment, using a Bayesian network (BN) in place of a least square support vector machine (LSSVM). The collected data, consisting of continuous and discrete variables, are used for learning the structure of the BN. Hence, the model is leveraged to: 1) identify a regression function for the prediction of the effects of actions on objects, calculated as the mean of the observed values; and 2) provide information on the reliability of the predicted values through the estimation of the variance for subsets of local observations. The information on the input-dependent variance is used to guide the learning algorithm in order to improve the performance of the robot, and hence to reduce the variance from the predicted values. The replacement of the LSSVM with the BN model provides a general probabilistic framework for dependencies among random variables; we perform conditional probability queries that enable the robot to choose the actions to perform on objects and select the most appropriate tool to obtain desired effects. The capability to make inference enables the robot to gather a better understanding of the world.

Artificial Cognitive Systems That Can Answer Human Creativity Tests: An Approach and Two Case Studies

Ana-Maria Olteţeanu, Zoe Falomir, Christian Freksa

Creative cognitive systems are rarely assessed with the same tools as human creativity. In this paper, an approach is proposed for building cognitive systems which can solve human creativity tests. The importance of using cognitively viable processes, cognitive knowledge acquisition and organization, and cognitively comparable evaluation when implementing creative problem-solving systems is emphasized. Two case studies of artificial cognitive systems evaluated with human creativity tests are reviewed. A general approach is put forward. The applicability of this general approach to other creativity tests and artificial cognitive systems, together with ways of performing cognitive knowledge acquisition for these systems are then explored.

A Practical SSVEP-Based Algorithm for Perceptual Dominance Estimation in Binocular Rivalry

Kazuo Tanaka, Motoyasu Tanaka, Toshiya Kajiwara, Hua O. Wang

This paper presents a practical algorithm for perceptual dominance estimation in binocular rivalry. The algorithm using steady state visual evoked potentials (SSVEP) effectively realizes a real-time estimation of perceptual dominance in binocular rivalry from electroencephalogram (EEG) signals. For accuracy of estimation, the algorithm utilizes the EEG spectrogram obtained via the short-time Fourier transformation with a short time analysis window. More importantly, the algorithm focuses on the acceleration (second-order time differential) of the spectrogram difference between SSVEP of left-eye stimulus and that of right-eye stimulus, where the left-eye stimulus and the right-eye stimulus are separately displayed for left-eye and right-eye, respectively, with different-flashing-frequency visual stimuli of dissimilar images. Experimental results demonstrate the utility of the algorithm, particularly, the importance of introducing the acceleration of the spectrogram difference. With the simplicity of the algorithm, it is in fact suitable to accurately measure binocular rivalry without verbal communications in real-time environments.

Volume 10, Issue 3, September 2018

A Comparison Study of Saliency Models for Fixation Prediction on Infants and Adults

Ali Mahdi, Mei Su, Matthew Schlesinger, Jun Qin

Various saliency models have been developed over the years. The performance of saliency models is typically evaluated based on databases of experimentally recorded adult eye fixations. Although studies on infant gaze patterns have attracted much attention recently, saliency-based models have not been widely applied for prediction of infant gaze patterns. In this paper, we conduct a comprehensive comparison study of eight state-of-the-art saliency models on predictions of experimentally captured fixations from infants and adults. Seven evaluation metrics are used to evaluate and compare the performance of saliency models. The results demonstrate a consistent performance of saliency models predicting adult fixations over infant fixations in terms of overlap, center fitting, intersection, information loss of approximation, and spatial distance between the distributions of saliency map and fixation map. In saliency and baselines models performance ranking, the results show that graph-based visual saliency model and Itti model are among the top three contenders, infants and adults have bias toward the centers of images, and all models and the center baseline model outperformed the chance baseline model.

Toward Electroencephalographic Profiling of Player Motivation: A Survey

Xuejie Liu, Kathryn Merrick, Hussein Abbass

Understanding and profiling player motivation complements and extends research on gameflow, player profiling, and game artificial intelligence, which helps us design entertaining games. However, automated identification of a player's motive profile remains an open challenge. An emerging technology that shows promise as a novel technique for identifying cognitive phenomena is electroencephalography (EEG). This paper begins with a survey of literature applying EEG to measure cognitive characteristics relevant to player motivation types. Then we present conceptual models that link motivation theory to mental states that can be identified using EEG including emotion, risk-taking, and social attitudes. We conclude this paper by examining the research challenges associated with using EEG to validate these models.

Decision Making in Multiagent Systems: A Survey

Yara Rizk, Mariette Awad, Edward W. Tunstel

Intelligent transport systems, efficient electric grids, and sensor networks for data collection and analysis are some examples of the multiagent systems (MAS) that cooperate to achieve common goals. Decision making is an integral part of intelligent agents and MAS that will allow such systems to accomplish increasingly complex tasks. In this survey, we investigate state-of-the-art work within the past five years on cooperative MAS decision making models, including Markov decision processes, game theory, swarm intelligence, and graph theoretic models. We survey algorithms that result in optimal and suboptimal policies such as reinforcement learning, dynamic programming, evolutionary computing, and neural networks. We also discuss the application of these models to robotics, wireless sensor networks, cognitive radio networks, intelligent transport systems, and smart electric grids. In addition, we define key terms in the area and discuss remaining challenges that include incorporating big data advancements to decision making, developing autonomous, scalable and computationally efficient algorithms, tackling more complex tasks, and developing standardized evaluation metrics. While recent surveys have been published on this topic, we present a broader discussion of related models and applications.
Note to Practitioners:
Future smart cities will rely on cooperative MAS that make decisions about what actions to perform that will lead to the completion of their tasks. Decision making models and algorithms have been developed and reported in the literature to generate such sequences of actions. These models are based on a wide variety of principles including human decision making and social animal behavior. In this paper, we survey existing decision making models and algorithms that generate optimal and suboptimal sequences of actions. We also discuss some of the remaining challenges faced by the research community before more effective MAS deployment can be achieved in this age of Internet of Things, robotics, and mobile devices. These challenges include developing more scalable and efficient algorithms, utilizing the abundant sensory data available, tackling more complex tasks, and developing evaluation standards for decision making.

Robots That Say “No” Affective Symbol Grounding and the Case of Intent Interpretations

Frank Förster, Joe Saunders, Chrystopher L. Nehaniv

Modern theories on early child language acquisition tend to focus on referential words, mostly nouns, labeling concrete objects, or physical properties. In this experimental proof-of-concept study, we show how nonreferential negation words, typically belonging to a child's first ten words, may be acquired. A child-like humanoid robot is deployed in speech-wise unconstrained interaction with naïve human participants. In agreement with psycholinguistic observations, we corroborate the hypothesis that affect plays a pivotal role in the socially distributed acquisition process where the adept conversation partner provides linguistic interpretations of the affective displays of the less adept speaker. Negation words are prosodically salient within intent interpretations that are triggered by the learner's display of affect. From there they can be picked up and used by the budding language learner which may involve the grounding of these words in the very affective states that triggered them in the first place. The pragmatic analysis of the robot's linguistic performance indicates that the correct timing of negative utterances is essential for the listener to infer the meaning of otherwise ambiguous negative utterances. In order to assess the robot's performance thoroughly comparative data from psycholinguistic studies of parent–child dyads is needed highlighting the need for further interdisciplinary work.

Learning for Goal-Directed Actions Using RNNPB: Developmental Change of “What to Imitate”

Jun-Cheol Park, Dae-Shik Kim, Yukie Nagai

"What to imitate" is one of the most important and difficult issues in robot imitation learning. A possible solution from an engineering approach involves focusing on the salient properties of actions. We investigate the developmental change of what to imitate in robot action learning in this paper. Our robot is equipped with a recurrent neural network with parametric bias (RNNPB), and learned to imitate multiple goal-directed actions in two different environments (i.e., simulation and real humanoid robot). Our close analysis of the error measures and the internal representation of the RNNPB revealed that actions' most salient properties (i.e., reaching the desired end of motor trajectories) were learned first, while the less salient properties (i.e., matching the shape of motor trajectories) were learned later. Interestingly, this result was analogous to the developmental process of human infant's action imitation. We discuss the importance of our results in terms of understanding the underlying mechanisms of human development.

"To Approach Humans?": A Unified Framework for Approaching Pose Prediction and Socially Aware Robot Navigation

Xuan-Tung Truong, Trung-Dung Ngo

We propose a unified framework for approaching pose prediction, and socially aware robot navigation, which enables a mobile service robot to safely and socially approach a dynamic human or human group in a social environment. The proposed framework is composed of four major functional blocks: 1) human detection and human features extraction to estimate the human states, and the social interaction information from the socio-spatio-temporal characteristics of a human and a group of humans; 2) a dynamic social zone (DSZ) consisting of an extended personal space and a social interaction space is modeled by the human states and social interaction information to represent space around the human and human group; 3) the approaching pose of the robot to a human or a human group is predicted using the DSZ and the environmental surroundings; and 4) the DSZ and the estimated approaching pose are incorporated into a motion planning system, comprising a local path planner and dynamic window approach technique, to generate the motion control commands for the mobile robot. We evaluate the developed framework through both simulation and real-world experiments under the newly proposed human safety and comfort indices, including the social individual index, social group index, and social direction index. The results show that the unified framework is fully capable of driving a mobile robot to approach both stationary and moving humans and human groups in a socially acceptable manner while guaranteeing human safety and comfort.

Synergetic Learning Control Paradigm for Redundant Robot to Enhance Error-Energy Index

Mitsuhiro Hayashibe, Shingo Shimoda

In order to perform energetically efficient motion as in human control, so-called optimization-based approach is commonly used in both robotics and neuroscience. Such an optimization approach can provide optimal solution when the prior dynamics information of the manipulator and the environment is explicitly given. However, the environment, where the robot faces with in a real world rarely has such a situation. The dynamics conditions change by the contact situation or the hand load for the manipulation task. Simple computational paradigm to realize both adaptability and learning is essential to bridge the gap between learning and control process in redundancy. We verify a novel synergetic learning control paradigm in reaching task of redundant manipulator. The performance in handling different dynamics conditions is evaluated in dual criteria of error-energy (E-E) coupling without prior knowledge of the given environmental dynamics and with model-optimization-free approach. This paper aims at investigating the ability of phenomenological optimization with the proposed human-inspired learning control paradigm for environmental dynamics recognition and adaptation, which is different from conventional model optimization approach. E-E index is introduced to evaluate not only the tracking performance, but also the error reduction rate per the energy consumption.

Evolutionary Wall-Following Hexapod Robot Using Advanced Multiobjective Continuous Ant Colony Optimized Fuzzy Controller

Chia-Feng Juang, Yue-Hua Jhan, Yan-Ming Chen, Chi-Ming Hsu

This paper proposes an evolutionary wall-following hexapod robot, where a new multiobjective evolutionary fuzzy control approach is proposed to control both walking orientation and speed of a hexapod robot for a wall-following task. According to the measurements of four distance sensors, a fuzzy controller (FC) controls the walking speed of the robot by changing the common swing angles of its six legs. At the same time, the FC controls the orientation of the robot by applying additional changes to the swing angles of the three legs in each side. In addition to the basic requirement of walking along the wall in an unknown environment, the control objectives are that the robot should maintain a proper robot-to-wall distance and walk at a high speed. This paper formulates the control problem as a constrained multiobjective FC optimization problem. A data-driven advanced multi-objective front-guided continuous ant colony optimization (AMO-FCACO) is proposed to address the problem and find a Pareto set of optimal solutions of different FCs. The performance of the AMO-FCACO-based fuzzy wall-following control approach is verified through simulations and comparisons with various multiobjective optimization algorithms. Experiments on controlling a real robot in an unknown environment using two software-designed FCs are performed to view the control performance in practice.

What Can I Do With This Tool? Self-Supervised Learning of Tool Affordances From Their 3-D Geometry

Tanis Mar, Vadim Tikhanoff, Lorenzo Natale

The ability to use tools can significantly increase the range of activities that an agent is capable of. Humans start using external objects since an early age to accomplish their goals, learning from interaction and observation the relationship between the objects used, their own actions, and the resulting effects, i.e., the tool affordances. Robots capable of autonomously learning affordances in a similar self-supervised way would be far more versatile and simpler to design than purpose-specific ones. This paper proposes and evaluates an approach to allow robots to learn tool affordances from interaction, and generalize them among similar tools based on their 3-D geometry. A set of actions is performed by the iCub robot with a large number of tools grasped in different poses, and the effects observed. Tool affordances are learned as a regression between tool-pose features and action-effect vector projections on respective self-organizing maps, which enables the system to avoid categorization and keep gradual representations of both elements. Moreover, we propose a set of robot-centric 3-D tool descriptors, and study their suitability for interaction scenarios, comparing also their performance against features derived from deep convolutional neural networks. Results show that the presented methods allow the robot to predict the effect of its tool use actions accurately, even for previously unseen tool and poses, and thereby to select the best action for a particular goal given a tool-pose.

A Human–Robot-Environment Interactive Reasoning Mechanism for Object Sorting Robot

Yunhan Lin, Huasong Min, Haotian Zhou, Feilong Pei

In this paper, we design an object sorting robot system, which is based on robot operating system distributed processing framework. This system can communicate with human beings; can percept the 3-D environment by Kinect sensor; has the ability of reasoning; can transfer the natural language intention to machine instruction to control the movement of manipulator. In particular, in order to improve the intelligence and usability of our robot, we propose a human–robot-environment interactive reasoning mechanism. In our method, a “dialogue and 3-D scene interaction module” is added into the traditional case-based reasoning–belief-desire-intention mechanism. Our proposed mechanism not only realizes the traditional function of map matching but also achieves the function of desire analysis and guidance. When the user’s desire is incomplete and/or mismatched with the actual scene, our robot will take the initiative to guide users through dialogue, and the user’s input information will be used to replenish the user’s desire. Experimental results prove the advantages of our mechanism.

Reconstruction of Visual Image From Functional Magnetic Resonance Imaging Using Spiking Neuron Model

Yongqiang Ma, Hao Wu, Mengjiao Zhu, Pengju Ren, Nanning Zheng, Badong Chen

Spikes are the basic elements of information dissemination in the organic brain. A spike train, a series of discrete action potentials, incorporates the conception of time. Nowadays, the majority of existing artificial neural networks use numerical information to realize machine learning and cognition tasks. The advantage of numerical computation is its precision and remarkable performance on well-structured problems. However, cognition tasks have semi-structured or ill-structured information, even worse, some cannot reach a conclusion. According to biological experiments, the spiking neuron model is potentially more suitable for dealing with undetermined and unstructured problems. Researchers are trying to decode functional magnetic resonance imaging (fMRI) data, and this is a typical task to extract meaningful information from relatively undetermined and unstructured data. In this paper, we selected the Tempotron neuron model to analyze and reconstruct visual images from fMRI data when subjects received several kinds of visual stimuli; the preliminary results of pattern reconstruction were unpolished, but somewhat effective. Thus, a new structure for the spiking neural network was built to achieve better results. Several important aspects of the proposed model were discussed in this paper: the encoding of external stimulating signals, the extraction of effective features, the relative position of spike trains on the timeline, the back propagation of error, and the rationality of parameter selection. These aspects are crucial in the spiking neuron model’s implementation, and are worth further investigation in future studies.

A Model of Multisensory Integration and Its Influence on Hippocampal Spatial Cell

Responses

Karthik Soman, Vignesh Muralidharan, V. Srinivasa Chakravarthy

Head direction (HD) cells, grid cells, and place cells, often dubbed spatial cells, are neural correlates of spatial navigation. We propose a computational model to study the influence of multisensory modalities, especially vision, and proprioception on responses of these cells. A virtual animal was made to navigate within a square box along a synthetic trajectory. Visual information was obtained via a cue card placed at a specific location in the environment, while proprioceptive information was derived from curvature-modulated limb oscillations associated with the gait of the virtual animal. A self-organizing layer was used to encode HD information from both sensory streams. The sensory integration (SI) of HD from both modalities was performed using a continuous attractor network with local connectivity, followed by oscillatory path integration and lateral anti-Hebbian network, where spatial cell responses were observed. The model captured experimental findings which investigated the role of visual manipulation (cue card removal and cue card rotation) on these spatial cells. The model showed a more stable formation of spatial representations via the visual pathway compared to the proprioceptive pathway, emphasizing the role of visual input as an anchor for HD, grid, and place responses. The model suggests the need for SI at the HD level for formation of such stable representations of space essential for effective navigation.

Information-Driven Multirobot Behavior Adaptation to Emotional Intention in Human–Robot Interaction

Luefeng Chen, Min Wu, Mengtian Zhou, Jinhua She, Fangyan Dong, Kaoru Hirota

To adapt robots' behavior to human emotional intention, an information-driven multirobot behavior adaptation mechanism is proposed in human–robot interaction (HRI). In the mechanism, optimal policy of behavior is selected by information-driven fuzzy friend-Q-learning (IDFFQ), and facial expression with identification information are used to understand human emotional intention. It aims to make robots be capable of understanding and adapting their behaviors to human emotional intention, in such a way that HRI runs smoothly. Simulation experiments are performed according to a scenario of drinking at a bar. Results show that the proposed IDFFQ reduces 51 learning steps compared to the fuzzy production rule-based friend-Q-learning (FPRFQ), and computational time is about 1/4 of the time consumed in FPRFQ. In Addition, the accuracy of emotion recognition and emotional intention understanding are 80.36% and 85.71%, respectively. The preliminary application experiments are carried out to the developing emotional social robot system, and the basic experimental results are shown in the scenario of drinking at a bar with three emotional robots and 12 volunteers.

Let Us Not Play It by Ear: Auditory Gating and Audiovisual Perception During Rapid Goal-Directed Action

Gerome A. Manson, Damian Manzone, John de Grosbois, Rachel Goodman, Joanne Wong, Connor Reid, Arindam Bhattacharjee, Valentin Crainic, Luc Tremblay

Susceptibility to audiovisual illusions is altered during rapid upper-limb reaching movements. Although, this change in perception appears to be linked to the movement's real-time characteristics, the mechanism for this alteration remains unknown. In this paper, we examined whether this modulation of multisensory perception could be explained by a decrease in auditory perceptual sensitivity. In two protocols, participants were instructed to perform rapid reaches to a visual target and report the number of auditory events (beeps) they perceived. One or 2 brief beeps were presented with either 0, 1, or 2 brief visual flashes at 0, 100, or 200 ms (protocol 1), or 0, 200, or 400 ms (protocol 2) relative to movement onset. The results revealed a significant, stable, decrease in auditory perception when the stimuli were presented during the movement, as compared to after the movement, and no-movement trials. Overall, the findings of this paper suggest that auditory perceptual sensitivity is monotonically reduced as one engages in visually guided goal-directed actions.

Artificial Intelligent System for Automatic Depression Level Analysis Through Visual and Vocal Expressions

Asim Jan, Hongying Meng, Yona Falinie Binti A. Gaus, Fan Zhang

A human being's cognitive system can be simulated by artificial intelligent systems. Machines and robots equipped with cognitive capability can automatically recognize a humans mental state through their gestures and facial expressions. In this paper, an artificial intelligent system is proposed to monitor depression. It can predict the scales of Beck depression inventory II (BDI-II) from vocal and visual expressions. First, different visual features are extracted from facial expression images. Deep learning method is utilized to extract key visual features from the facial expression frames. Second, spectral low-level descriptors and mel-frequency cepstral coefficients features are extracted from short audio segments to capture the vocal expressions. Third, feature dynamic history histogram (FDHH) is proposed to capture the temporal movement on the feature space. Finally, these FDHH and audio features are fused using regression techniques for the prediction of the BDI-II scales. The proposed method has been tested on the public Audio/Visual Emotion Challenges 2014 dataset as it is tuned to be more focused on the study of depression. The results outperform all the other existing methods on the same dataset.

Semiasynchronous BCI Using Wearable Two-Channel EEG

Yubing Jiang, Nguyen Trung Hau, Wan-Young Chung

Motor imagery (MI)-based electroencephalogram can be modulated voluntarily by users through feedback provided by a brain–computer interface (BCI). In MI-based BCIs, feedback design plays an important role in MI training, especially in BCIs that use a limited number of electrodes. This paper presents a semiasynchronous MI-based BCI system that uses discrete and continuous feedback together, in order to improve practicability and training efficiency. Results show that the proposed method can reduce command generation time to 3.77 s with the accuracy of 77%, using our proposed continuous feedback with a two-channel MI-based BCI system. Our proposed system is noninvasive and wearable, featuring improvement in the practical utility, and operational convenience of a BCI system.

CA-NEAT: Evolved Compositional Pattern Producing Networks for Cellular Automata Morphogenesis and Replication

Stefano Nichele, Mathias Berild Ose, Sebastian Risi, Gunnar Tufte

Cellular Automata (CA) are a remarkable example of morphogenetic system, where cells grow and self-organize through local interactions. CA have been used as abstractions of biological development and artificial life. Such systems have been able to show properties that are often desirable but difficult to achieve in engineered systems, e.g., morphogenesis and replication of regular patterns without any form of centralized coordination. However, cellular systems are hard to program (i.e., evolve) and control, especially when the number of cell states and neighborhood increase. In this paper, we propose a new principle of morphogenesis based on compositional pattern producing networks (CPPNs), an abstraction of development that has been able to produce complex structural motifs without local interactions. CPPNs are used as CA genotypes and evolved with a neuroevolution of augmenting topologies (NEATs) algorithm. This allows complexification of genomes throughout evolution with phenotypes emerging from self-organization through development based on local interactions. In this paper, the problems of 2-D pattern morphogenesis and replication are investigated. Results show that CA-NEAT is an appropriate means of approaching cellular systems engineering, especially for future applications where natural levels of complexity are targeted. We argue that CA-NEAT could provide a valuable mapping for morphogenetic systems, beyond CA systems, where development through local interactions is desired.

Active Learning on Service Providing Model: Adjustment of Robot Behaviors Through Human Feedback

Shih-Huan Tseng, Feng-Chih Liu, Li-Chen Fu

As robots are put into humans' daily life, the assigned tasks to robots are varied, and the different needs of people interacting with robots are immense. As a result, when facing different users, it is important for robots to personalize the interactions and provide user-desired services. This paper, therefore, proposes a learning strategy on the service-providing model. Through human feedback, the strategy enables the robot to learn the users' needs, as well as preferences, and adjust its behaviors. Here, we assume that users' needs and preferences may vary with time; hence the goal of this paper is to let the adjustment of robot behaviors be able to adapt to those variations. In turn, the service-providing model of the robot could adjust online as well. That is, it can select a new action from those favorable actions that have already been selected or an action that is not an unfavorable action but has annoyed humans recently. To implement our system, the service robot under discussion is applied to the home environment. For performance evaluation, we have performed extensive experiments that satisfactorily demonstrate that our robot can provide services to different users and adapt to their preference change.

Learning to Achieve Different Levels of Adaptability for Human–Robot Collaboration Utilizing a Neuro-Dynamical System

Shingo Murata, Yuxi Li, Hiroaki Arie, Tetsuya Ogata, Shigeki Sugano

Intelligent robots are expected to collaboratively work with humans in dynamically changing daily life environments. To realize successful human–robot collaboration, robots need to deal with latent spatiotemporal complexity in the workspace and the task. To overcome this crucial issue, three levels of adaptability—motion modification, action selection, and role switching—should be considered. This paper demonstrates that a single hierarchically organized neuro-dynamical system called a multiple timescale recurrent neural network can achieve these levels of adaptability by utilizing hierarchical and bidirectional information processing. The system is implemented in a humanoid robot and the robot is required to learn to perform collaborative tasks in which some parts must be performed by a human partner and others by the robot. Experimental results show that the robot can perform collaborative tasks under dynamically changing environments, including both learned and unlearned situations, thanks to different levels of adaptability acquired in the system.

Data-Structures for Multisensory Information Processing in an Embodied Machine-Mind

Romi Banerjee, Sankar K. Pal

The real-world is a medley of multisensory information and so are our experiences, memories, and responses. As embodied beings, we respond to the information endogenously and in ways derived from self-defining factors. Thus inspired, we attempt formalization of data-structures to facilitate generation of system-bespoke comprehension-granules of the real-world. The conceptualized structures encapsulate multisensory inputs (sourced from the real-world or memories), intrinsic and deliberate emotions, messages (bearing intermittent process-results, queries, multimodal data, etc.) across system modules and memory units, and sensorimotor responses to the inputs. The structural-schematics are anthropomorphic. These variable-length constructs are theoretically platform-independent, support genericity across data-modality and information-inclusion, and provide for representation of novel sensory-data. An epigenome-styled header node for the afferent data-units provides for the activation of intuitive "fight-flight" behavior. The documentation includes a flow-graph, depicting the translation of information across the data-structures and the different ways of thinking while interpreting a real-world scene or a mind-generated event. Applicability of the structures has been analyzed in the context of comprehension in an embodied mind-machine framework and other similar architectures. Studies herein target contribution to the design of generally intelligent man–machine symbiotic systems.

An Extended Reinforcement Learning Framework to Model Cognitive Development With Enactive Pattern Representation

Zhenping Xie, Yaochu Jin

In order to make machines more intelligent, it is inevitable to understand human-like cognitive development, in which adaptive, autonomous and progressive evolution of cognitive decision-making in interacting with the environment plays a key role. Inspired by enactive artificial intelligence and evolutionary sampling learning, a new cognitive development learning model termed evolutionary enactive learning (EEL) is proposed in this paper. The proposed model is constructed by extending the reinforcement learning framework and introducing the utility-selection theory to guide the coevolution of pattern representation and decision-making policies. Theoretical analysis on the model's validity of EEL is given. To further demonstrate the effectiveness of the proposed method, two simulated cognitive decision-making tasks are designed, in which pattern representation and decision-making must be jointly developed to achieve good cognitive performance. Our experimental results clearly demonstrate that the resulting learning process is rational and effective. Finally, we indicate that the proposed EEL could be readily further extended by introducing existing machine learning techniques to solve more practical applications.

Cognitive Navigation by Neuro-Inspired Localization, Mapping, and Episodic Memory

Huajin Tang, Rui Yan, Kay Chen Tan

One of the important topics in the study of robotic cognition is to enable robot to perceive, plan, and react to situations in a real-world environment. We present a novel angle on this subject, by integrating active navigation with sequence learning. We propose a neuro-inspired cognitive navigation model which integrates the cognitive mapping ability of entorhinal cortex (EC) and episodic memory ability of hippocampus to enable the robot to perform more versatile cognitive tasks. The EC layer is modeled by a 3-D continuous attractor network structure to build the map of the environment. The hippocampus is modeled by a recurrent spiking neural network to store and retrieve task-related information. The information between cognitive map and memory network are exchanged through respective encoding and decoding schemes. The cognitive system is applied on a mobile robot platform and the robot exploration, localization, and navigation are investigated. The robotic experiments demonstrate the effectiveness of the proposed system.

Modeling Development of Multimodal Emotion Perception Guided by Tactile Dominance and Perceptual Improvement

Takato Horii, Yukie Nagai, Minoru Asada

Humans recognize others' emotional states such as delight, anger, sorrow, and pleasure through their multimodal expressions. However, it is unclear how this capability of emotion perception is acquired during infancy. This paper presents a neural network model that reproduces the developmental process of emotion perception through an infant–caregiver interaction. This network comprises hierarchically structured restricted Boltzmann machines (RBMs) that receive multimodal expressions from a caregiver (visual, audio, and tactile signals in our current experiment) and learn to estimate her/his emotional states. We hypothesize that emotional categories of multimodal stimuli are acquired in a higher layer in the network owing to two important functions: 1) tactile dominance and 2) perceptual improvement. The former refers to that tactile sensors can detect emotional valence of stimuli such as positive, negative, and zero valence more directly than can other sensors due to characteristics of the nerve systems of the skin. This function was implemented as semisupervised learning in the model. The latter refers to developmental changes in the perceptual acuity, which was replicated by refining the variance parameters of the low-layered RBMs. Experimental results demonstrated that tactile dominance and perceptual improvement have the role of facilitating the differentiation of emotional states of multimodal expressions; however, the influences only appear when both functions are included in the model together. Considering our results from the psychological perspective may help to elucidate the neural and social mechanisms of the development of emotion perception.

Using the Partial Directed Coherence to Assess Functional Connectivity in Electroencephalography Data for Brain–Computer Interfaces

Jorge Antonio Gaxiola-Tirado, Rocio Salazar-Varas, David Gutiérrez

In this paper, we propose a statistical selection procedure by which various mental tasks can be characterized by specific brain functional connectivity. Different connectivity patterns are identified by the partial directed coherence (PDC) which is a frequency-domain metric that provides information about directionality in the interaction between signals recorded at different sensors. The basis of our selection is a statistical analysis of the directed connectivities revealed by their repeated appearance and larger PDC magnitudes in sets of electroencephalography (EEG) sensors treated as networks. Hence, our proposed method identifies significant differences between directed connectivities on EEG-sensor networks that are specific to the mental tasks involved. A combinatory analysis of different possible networks allows us to find those that characterize and discriminate the tasks and, as proof-of-concept, we analyze the connectivities of movement imageries (MIs) used in the operation of a brain–computer interface. The directed interconnections revealed by our proposed method are in agreement with brain functional connectivities already reported for MIs, and good classification rates are achieved when such interconnections are used as features in a Mahalanobis-distance-based classifier.

Mapping Language to Vision in a Real-World Robotic Scenario

Karla Štepánová, Frederico Belmonte Klein, Angelo Cangelosi, Michal Vavrečka

Language has evolved over centuries and was gradually enriched and improved. The question, how people find assignment between meanings and referents, remains unanswered. There are many of computational models based on the statistical co-occurrence of meaning-reference pairs. Unfortunately, these mapping strategies show poor performance in an environment with a higher number of objects or noise. Therefore, we propose a more robust noise-resistant algorithm. We tested the performance of this novel algorithm with simulated and physical iCub robots. We developed a testing scenario consisting of objects with varying visual properties presented to the robot accompanied by utterances describing the given object. The results suggest that the proposed mapping procedure is robust, resistant against noise and shows better performance than one-step mapping for all levels of noise in the linguistic input, as well as slower performance degradation with increasing noise. Furthermore, the proposed procedure increases the clustering accuracy of both modalities.

How Does a Robot Develop Its Reaching Ability Like Human Infants Do?

Dingsheng Luo, Fan Hu, Tao Zhang, Yian Deng, Xihong Wu

In this paper, we address the issue robots developing a reaching ability, as human infants do. Specifically, a novel infant-inspired framework is proposed based on recent findings that the emergence of reaching is the product of a deeply embodied process. The methodology behind our framework is to develop proprioception, fixation, and internal models (including both the forward and inverse models) in three sequential phases, with the whole developing process being driven by self-produced motor babbling. In the first phase, an autoencoder-based proprioception model is proposed. In the second phase, a simplified strategy for imitating the function of fixation is developed. In the third phase, a new forward model and two new inverse models are further proposed. We evaluated our framework and associated models with the PKU-HR6.0II physical robot and two simulated robots. Experiments confirm that the PKU-HR6.0II could successfully develop its reaching ability in a manner similar to that of human infants through our proposed framework. Better performance and adaptability are also demonstrated in comparison to existing benchmarks. It is also shown that the proposed framework has the potential to achieve other manipulation abilities such as grasping and placing.

A Human-Vehicle Collaborative Simulated Driving System Based on Hybrid Brain–Computer Interfaces and Computer Vision

Wenyu Li, Feng Duan, Shili Sheng, Chengguang Xu, Rensong Liu, Zhiwen Zhang, Xue Jiang

Automatic driving vehicles have been developed to provide more convenient and comfortable driving experiences. However, these vehicles failed in satisfying the variance of human intentions. Recently, the strategy of collaborating brain–computer interface (BCI) controlling and automatic driving receives attention. Since the BCI system remained some limitation in real-time controlling, a fusion method has been proposed to explore and verify the feasibility of human-vehicle collaborative driving in this paper. A hybrid BCI was developed to interpret human intentions. In addition, a computer vision-based automatic driving component was developed to maintain the vehicle on the road. A system for fusing these two kinds of vehicle driving decisions was first proposed in this paper. This system can simultaneously obtain the visual data and the hybrid electroencephalograph (EEG) signals. The hybrid EEG signals consist of steady-state visual evoked potentials and motor imagery. The obtained multisource information can be fused to make the final decision to drive a simulated vehicle. The proposed system was evaluated with different destinations. The experimental results verify the feasibility of fusing both human intention and computer vision. The task success rate reached 91.1% and the information transfer rate was 85.80 bit/min.

Building a Spiking Neural Network Model of the Basal Ganglia on SpiNNaker

Basabdatta Sen-Bhattacharya, Sebastian James, Oliver Rhodes, Indar Sugiarto, Andrew Rowley, Alan B. Stokes, Kevin Gurney, Steve B. Furber

We present a biologically inspired and scalable model of the basal ganglia (BG) simulated on the spiking neural network architecture (SpiNNaker) machine, a biologically inspired low-power hardware platform allowing parallel, asynchronous computing. Our BG model consists of six cell populations, where the neuro-computational unit is a conductance-based Izhikevich spiking neuron; the number of neurons in each population is proportional to that reported in anatomical literature. This model is treated as a single-channel of action-selection in the BG, and is scaled-up to three channels with lateral cross-channel connections. When tested with two competing inputs, this three-channel model demonstrates action-selection behavior. The SpiNNaker-based model is mapped exactly on to SpineML running on a conventional computer; both model responses show functional and qualitative similarity, thus validating the usability of SpiNNaker for simulating biologically plausible networks. Furthermore, the SpiNNaker-based model simulates in real time for time-steps $\geq {1}$ms; power dissipated during model execution is $\approx {1.8}$W.

Editor Pierre-Yves Oudeyer, Inria and Ensta ParisTech, France, pierre-yves.oudeyer@inria.fr
Editorial assistant Fabien C. Y. Benureau, CNRU, Okinawa Institute of Science and Technology (OIST), Japan
ISSN 1550-1914