How Evolution May Work Through Curiosity-Driven Developmental Process

Pierre-Yves Oudeyer, Linda Smith

HAL Id: hal-01404334
https://hal.inria.fr/hal-01404334

Submitted on 28 Nov 2016

How Evolution may work through Curiosity-driven Developmental Process

Pierre-Yves Oudeyer
Inria and Ensta ParisTech, France
http://www.pyoudeyer.com

Linda B. Smith
Indiana University, US
http://www.iub.edu/~cogdev

Abstract:
Infants' own activities create and actively select their learning experiences. Here we review recent models of embodied information seeking and curiosity-driven learning and show that these mechanisms have deep implications for development and evolution. We discuss how these mechanisms yield self-organized epigenesis with emergent ordered behavioral and cognitive developmental stages. We describe a robotic experiment that explored the hypothesis that progress in learning, in and for itself, generates intrinsic rewards: the robot learners probabilistically selected experiences according to their potential for reducing uncertainty. In these experiments, curiosity-driven learning led the robot learner to successively discover object affordances and vocal interaction with its peers. We explain how a learning curriculum adapted to the current constraints of the learning system automatically formed, constraining learning and shaping the developmental trajectory. The observed trajectories in the robot experiment share many properties with those in infant development, including a mixture of regularities and diversities in the developmental patterns. Finally, we argue that such emergent developmental structures can guide and constrain evolution, in particular with regards to the origins of language.

Learning experiences do not passively "happen" to infants. Rather, infants' own activities create and select these experiences. Piaget (1952) described a pattern of infant activity that is highly illustrative of this point. He placed a rattle in a four-month-old infant's hands. As the infant moved the rattle, it would both come into sight and also make a noise, arousing and agitating the infant and causing more body motions, and thus causing the rattle to move into and out of sight and to make more noise. The infant had no prior knowledge of the rattle but discovered through activity the task and goal of rattle shaking. As the infant accidentally moved the rattle, and saw and heard the consequences, becoming captured by the activity and outcomes, the infant may be said to have gained intentional control over the shaking of the rattle and the goal of making noise (Thelen, 1994). This example —a body movement that leads to an interesting outcome and thus more activity and the re-experience and building of expectations about the outcome—may be foundational, not just to developmental process, but also to how evolution works through developmental process.

Infants' exploration of rattles are instances of action and learning motivated by curiosity and may reflect intrinsic motivations that select "interesting" sensorimotor activities (Lowenstein, 1994; Gottlieb et al., 2013). However, what is interesting is history dependent. Once all the variations in rattle shaking are easily predicted, and all the outcomes expected, then playing with a rattle is not very interesting, as evident in the disinterest of older infants in rattles. What is interesting depends on what one knows and what one does not know. Across many

different fields, theorists have suggested that "interest" is engaged by what is just beyond current knowledge, neither too well known nor too far beyond what is understandable. This theoretical idea has been offered many times in psychology, through concepts like cognitive dissonance (Kagan, 1972), optimal incongruity (Hunt, 1965), intermediate novelty (Berlyne, 1960; Kidd, Piantadosi, & Aslin, 2012) and optimal challenge (Csikszenthmihalyi, 1991). There have been several recent theoretical advances in these ideas in developmental robotics (Baldassarre and Mirolli, 2013; Gottlieb et al., 2013; Oudeyer et al., 2007), in models about the evolutionary origins of intrinsic motivation systems (Singh et al., 2010; Barto, 2013), and in recent findings in neuroscience linking intrinsic motivation with attention (Gottlieb et al., 2013), as well as in new formal models of infant visual attention (Kidd, Piantadosi, & Aslin, 2012).

In general, learning in these curiosity driven activities *progresses* to yield an *improvement* of prediction or control over a repeated activity and thus a *reduction* of uncertainty (Friston, 2012; Kidd et al, 2012; Oudeyer and Kaplan, 2007; Schmidhuber, 1991). Such intrinsically motivating activities have been called "progress niches" (Oudeyer et al., 2007): *progress in learning in and for itself* generates intrinsic rewards and an action selection system directly aims to maximize this reward. In reinforcement learning frameworks, this search for progress niches leads to spontaneous exploration (Gottlieb et al., 2013; Kidd et al, 2012; Oudeyer et al., 2007; Schmidhuber, 1991). In this view, *progress* in prediction or control is a primary driver (and accordingly, intrinsic rewards for learning progress/uncertainty reduction may be *primary* rewards). These theoretical advances lead to a definition of curiosity as an epistemic motivational mechanism that pushes an organism to explore activities for the primary sake of *gaining information* (as opposed to searching for information in service of achieving an external goal like finding food or shelter). Such a motivational mechanism of curiosity will often be only one of several motivational mechanisms operating in any living being, and at any given time curiosity may interact, complement or conflict with other motivations. From a machine learning perspective, mechanisms of information seeking are called *active learning*, where the learner probabilistically and through its own activity selects experiences according to their potential for reducing uncertainty (Cohn et al., 1996; Lopes and Oudeyer, 2010; Lopes and Montesano, 2014). Information seeking mechanisms have been used either as an "exploration bonus" mechanism in service of efficient maximization of a task-specific reward, or as primary rewards driving models of curiosity-driven learning (Gottlieb et al., 2013).

Here we outline a robotic model that illustrates these ideas about curiosity-driven learning and how an active search for learning progress can lead to a system that first explores simple activities, and then progressively shifts to more complex learning experiences, effectively self-generating and self-organizing a learning curriculum adapted to the current constraints of the learning system, and at the same time constraining learning and shaping the developmental trajectory. We then turn to the implications of these results for how evolution may work through development and curiosity to yield constrained and adaptive outcomes within a species.

A Developmental Playground

In this section, we describe a series of robot experiments that illustrate how mechanisms of curiosity-driven exploration, dynamically interacting with learning, physical and social constraints, can self-organize developmental trajectories and in particular lead a learner to successively discover two important functionalities, object affordances and vocal interaction with its peers.

In these Playground Experiments, a quadruped "learning" robot (the learner) is placed on an infant play mat with a set of nearby objects and is joined by an "adult" robot (the teacher), see Figure **1 (A)** (Oudeyer and Kaplan, 2006; Kaplan and Oudeyer, 2007b; Oudeyer et al., 2007). On the mat and near the learner are objects for discovery: an elephant (which can be bitten or "grasped" by the mouth), a hanging toy (which can be "bashed" or pushed with the leg). The teacher is pre-programmed to imitate the sound made by the learner when the learning robot looks to the teacher while vocalizing at the same time.

The learner is equipped with a repertoire of motor primitives parameterized by several continuous numbers that control movements of its legs, head and a simulated vocal tract. Each motor primitive is a dynamical system controlling various forms of actions: (a) turning the head in different directions; (b) opening and closing the mouth while crouching with varying strengths and timing; (c) rocking the leg with varying angles and speed; (d) vocalizing with varying pitches and lengths. These primitives can be combined to form a large continuous space of possible actions. Similarly, sensory primitives allow the robot to detect visual movement, salient visual properties, proprioceptive touch in the mouth, and pitch and length of perceived sounds. For the robot, these motor and sensory primitives are initially black boxes and he has no knowledge about their semantics, effects or relations.

The learning robot learns how to use and tune these primitives to produce various effects on its surrounding environment, and exploration is driven by the maximization of learning *progress*, by choosing physical experiences ("experiments") that improve the quality of predictions of the consequences of its actions.

Figure **1 (B)** outlines a computational architecture, called R-IAC (Oudeyer, Kaplan et al. 2007; Moulin-Frier et al., 2014). A prediction machine (M) learns to predict the consequences of actions taken by the robot in given sensory contexts. For example, this module might learn to predict (with a neural network) which visual movements or proprioceptive perceptions result from using a leg motor primitive with certain parameters. A meta-cognitive module estimates the evolution of errors in prediction of M in various regions of the sensorimotor space. This module estimates how much errors decrease in predicting an action, for example, in predicting the consequence of a leg movement when this action is applied towards a particular area of the environment. These estimates of error reduction are used to compute the intrinsic reward from progress in learning. This reward is an internal quantity that is proportional to the decrease of prediction errors, and the maximization of this quantity is the goal of action selection within a computational reinforcement-learning architecture (Oudeyer and Kaplan, 2007; Oudeyer et al., 2007). Importantly, the action selection system chooses most often to explore activities where the estimated reward from learning progress is high. However, this choice is probabilistic, which leaves the system open to learning in new areas and open to discovering other activities that may also yield progress in learning[1]. Since the sensorimotor flow does not come pre-segmented into activities and tasks, a system that seeks to maximize differences in learnability is also used to progressively categorize the sensorimotor space into regions. This categorization thereby models the incremental creation and refining of cognitive categories differentiating activities/tasks.

---

[1] (technically the decision on how much time to spend on high learning progress activities and other activities is achieved using Multi-Armed Bandit algorithms for the so-called exploration/exploitation dilemma (Audibert et al., 2009)).
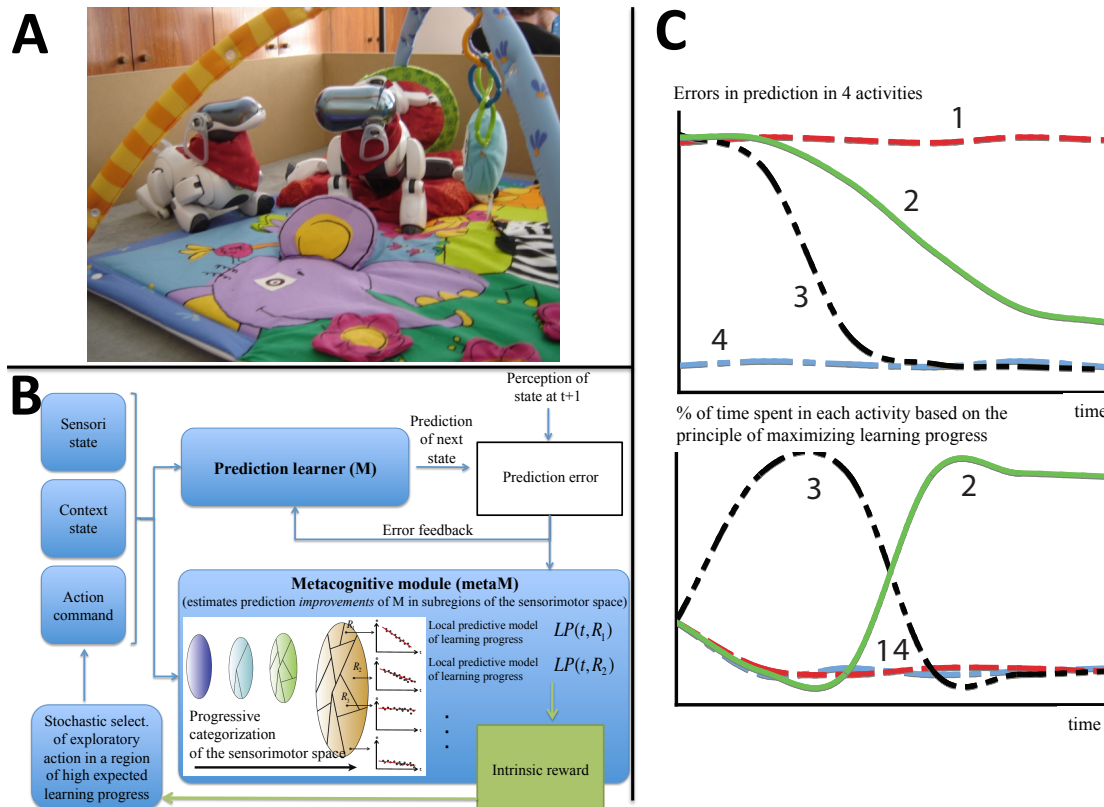
Figure 1 The Playground Experiment (Oudeyer and Kaplan, 2006; Oudeyer et al., 2007) (A) The learning context; (B) The computational architecture for curiosity-driven exploration in which the robot learner probabilistically selects experiences according to their potential for reducing uncertainty, i.e. for learning progress; (C) Illustration of a self-organized developmental sequence where the robot automatically identifies, categorizes and shifts from simple to mode complex learning experiences. Figure adapted with permission from Gottlieb et al. (2013).

To illustrate how such an exploration mechanism can automatically generate ordered learning stages, let us first imagine a learner confronted with four categories of activities, as shown on figure **1 (C)**. The practice of each of these four activities, which can be of varying difficulty, leads to different learning rates at different points in time (see the top curves, which show the evolution of prediction errors in each activity if the learner were to focus full-time and exclusively on each). If, however, the learner uses curiosity-driven exploration to decide what and when to practice by focusing on progress niches, it will avoid activities already predictable (curve 4) or too difficult to learn to predict (curve 1), in order to focus first on the activity with the fastest learning rate (curve 3) and eventually, when the latter starts to reach a "plateau" to switch to the second most promising learning situation (curve 2). Thus, such robots will show a regular developmental course –one that will be "universal" for learners with similar internal processes learning in similar environments. Embodied exploration driven by learning progress creates an organized exploratory strategy: the system systematically achieves these learning experiences in an order and does so because they yield (given the propensities of the learner and the physical world) different patterns of uncertainty reduction.

In the Playground experiment described earlier, multiple experimental runs lead to two general categories of results: self-organization and a mixture of regularities and diversities in the developmental patterns (Oudeyer and Kaplan, 2006; Oudeyer et al., 2007).

Self-organization. In all of the runs, one observes the self-organization of structured developmental trajectories, where the robot explores objects and actions in a progressively more complex stage-like manner while acquiring autonomously diverse affordances and skills that can be reused later on and that change the learning progress in more complicated tasks. The following developmental sequence was typically observed:

1. In a first phase, the learner achieves unorganized body babbling;
2. In a second phase, after learning a first rough model and meta-model, the robot stops combining motor primitives, exploring them one by one, but each primitive is explored itself in a random manner;
3. In a third phase, the learner begins to experiment with actions towards zones of its environment where the external observer knows there are objects (the robot is not provided with a representation of the concept of "object"), but in a non-affordant manner (e.g. it vocalizes at the non-responding elephant or tries to bash the teacher robot which is too far to be touched);
4. In a fourth phase, the learner now explores the affordances of different objects in the environment: typically focussing first on grasping movements with the elephant, then shifting to bashing movements with the hanging toy, and finally shifting to explorations of vocalizing towards the imitating teacher.
5. In the end, the learner has learnt sensorimotor affordances with several objects, as well as social affordances, and has mastered multiple skills. None of these specific objectives were pre-programmed. Instead, they self-organized through the dynamic interaction between curiosity-driven exploration, statistical inference, the properties of the body, and the properties of the environment.

These playground experiments do not simply simulate particular skills (such as batting at toys to make them swing or vocalizations) but simulate an ordered and systematic developmental trajectory, with a universality and stage-like structure that may be mistakenly taken to indicate an internally-driven process of maturation. However, the trajectory is created through activity and through the general principle that sensorimotor experiences that reduce uncertainty in prediction are rewarding. In this way, developmental achievements can build on themselves without specific pre-programmed dependencies but nonetheless – like evolution itself –create structure (see Smith & Breazeal, 2007; and Smith 2013, for related findings and arguments).

Regularities and diversity. Because these are self-organizing developmental processes, they generate not only strong regularities but also diversity across individual developmental trajectories. For example, in most runs one observes successively unorganized body babbling, then focused exploration of head movements, then exploration of touching an object, then grasping an object, and finally vocalizing towards a peer robot (pre-programmed to imitate). This can be explained as gradual exploration of new progress niches, and those stages and their ordering can be viewed as a form of attractor in the space of developmental trajectories. Yet, with the same mechanism and same initial parameters, individual trajectories may invert stages, or even generate qualitatively different behaviours. This is due to stochasticity, to even small variability in the physical realities and to the fact that this developmental dynamic system has several attractors with more or less extended and strong domains of attraction (characterized by amplitude of learning progress). We see this diversity as a positive outcome since individual development is not identical across different individuals but is always, for each individual, unique in its own ways. This kind of approach, then, offers a way to understand individual differences as emergent in developmental process itself and makes clear how developmental process might vary across contexts, even with an identical learning mechanism.

A further result to be highlighted is the early development of vocal interaction. With a single generic mechanism, the robot both explores and learns how to manipulate objects and how to vocalize to trigger specific responses from a more mature partner (Oudeyer and Kaplan, 2006; Kaplan and Oudeyer, 2007a). Vocal babbling and language play have been shown to be key in infant language development; however, the motivation to engage in vocal play has often been associated with hypothesized language specific motivation. The Playground Experiment makes it possible to see how the exploration and learning of communicative behavior might be at least partially explained by general curiosity-driven exploration of the body affordances, as also suggested by Oller (2000).

Other robotic models have explored how social guidance can be leveraged by an intrinsically motivated active learner and dynamically interact with curiosity to structure developmental trajectories (Thomaz and Breazeal, 2008; Nguyen and Oudeyer, 2013). Focusing on vocal development, Moulin-Frier et al. conducted experiments where a robot explored the control of a realistic model of the vocal tract in interaction with vocal peers through a drive to maximize learning progress (Moulin-Frier et al., 2014). This model relied on a physical model of the vocal tract, its motor control and the auditory system. The experiments showed how such a mechanism can explain the adaptive transition from vocal self-exploration with little influence from the speech environment, to a later stage where vocal exploration becomes influenced by vocalizations of peers. Within the initial self-exploration phase, a sequence of vocal production stages self-organizes, and shares properties with infant data: the vocal learner first discovers how to control phonation, then vocal variations of unarticulated sounds, and finally articulated proto-syllables. As the vocal learner becomes more proficient at producing complex sounds, the imitating vocalizations of the teacher provide high learning progress resulting in a shift from self-exploration to vocal imitation.

An evo-devo perspective: from curiosity to the evolution of language pre-requisites

The developmental patterns in the Playground experiments exhibit a form of behavioral and cognitive epigenesis, as proposed by Gottlieb (1991). Developmental structures in these models are neither learnt "tabula rasa" nor a pre-determined result of an innate "program": instead, they self-organize out of the dynamic interaction between constrained cognitive mechanisms (including curiosity, learning, and abstraction), the morphological properties of the body, and the physical and social environment that itself is constrained and ordered by the developmental level of the organism (Oudeyer, 2011; Oudeyer, 2010; Smith, 2013; Thelen & Smith, 1996). This self-organization includes the dynamic and automatic formation of behavioral and cognitive stages of progressively increasing complexity, sharing many properties with infant development (Piaget, 1952).

Such robotic models and experiments operationalize theories of epigenesis in behavioral development (e.g., Gottlieb, 1991; West & King, 1987; Lickliter & Honeycutt, 2009), as well as the dynamic systems conceptualization of development (Thelen and Smith, 1996). In particular, they allow us to see in detail how the interaction of heterogeneous mechanisms and constraints at several scales could form a dynamical system where developmental structures emerge. Such emergent developmental structures have deep implications for evolution. In particular, they constitute a reservoir of behavioral and cognitive innovations which can be later on recruited for functions not yet anticipated, and at both developmental and evolutionary scales: this is *exaptation* (Gould, 1991).

First, the results show that modality-independent generic mechanisms for curiosity-driven exploration of the body and its interactions can lead to the emergence of basic speech skills, vocal interaction, and vocal imitation. This suggests that in principle the infant may develop initial speech capabilities without an innate specific bias for learning speech and vocal interaction, and without teleological knowledge that such skills will be recruited later on for the onset of language.

Second, such mechanisms could be related to state-of-the-art models of the formation and evolution of shared vocalization systems in populations of individuals (de Boer, 2001; Oudeyer, 2006; Moulin-Frier et al, 2011). Such population models have shown that when vocal learners are equipped with mechanisms of spontaneous vocal self-exploration and progressively tune their vocalizations to match those of their neighbors, conventionalized systems of sounds can be formed automatically at the group level from an initial state where each individual only produces random vocalizations. In these models, mechanisms for systematic babbling were ad hoc and hand programmed. Yet, the robotic experiments show that principled and modality-independent mechanisms for curiosity-driven exploration can drive a learner to explore its vocal tract, and also drive it to imitate the vocalizations of its peers. Combining these models leads to a startling hypothesis: conventional patterns of vocalizations at the group level could self-organize as a result of the interaction of individuals intrinsically motivated to learn how their body, and their vocal tract in particular, can produce effects on objects and social peers. Similar analyses have been offered for the origins of joint visual attention (Yu & Smith, 2013; Deak, et al; 2014).

Further, arguments presented in Barto (2013) have shown that the evolutionary origins of such an intrinsic motivation to learn can be explained because it maximizes long-term evolutionary fitness under rapidly changing environmental conditions (e.g., due to human social and cultural structures, which can evolve much faster than the phylogenetic scale). In their computer simulations, and under conditions of rapidly changing environments and with fitness/reward functions based on reproduction, Singh et al. (Singh, Lewis et al. 2010) have shown that it can be more efficient to evolve a control architecture that rewards learning *per se* rather than domain-specific adaptations. In these ways, mechanisms of information seeking can thus evolve independently of language, but yet may have spontaneously bootstraped vocal structures, both at the individual and population level, later on recruited through exaptation for a language function that is not already foreseen.

These arguments suggest a potentially strong role for curiosity-driven developmental process in the evolution of language prerequisites. They also raise interesting issues with respect to the comparison of humans with other animals. As argued earlier, the motivational mechanism of curiosity interacts with other motivational mechanisms like food or mate search, and its weight in motivational arbitration may vary widely across species equipped with such curiosity. Hence, the high degree of competition for survival in many species can be expected to promote avoidance of risk, where aversive motivational systems overcome strongly the expression of curiosity-driven exploration. The multimodal systematicity and the extent to which open-ended free play and curiosity is expressed in humans, where children are comparatively highly protected for a long period, is unrivalled in the animal kingdom (Power, 1999). The dominance of curiosity in the motivational hierarchy may be key to the emergence of wide ranging domain-specific knowledge in humans as compared to other species. Yet, this remains an open question and points to the next critical challenge in understanding curiosity driven learning: how does curiosity compete and couple with other forms of motivation (possibly aversive)? The nature of this dynamic coupling is likely to have significant impact

on how we understand the developmental dynamics and may also provide insight into individual differences. Do contexts in which there are limited resources for such basic drives as food and bodily safety alter the developmental trajectory by limiting curiosity-driven learning?

Conclusion

Play offers young learners a way to harness the complexity of their learning environment through exploration. The world presents learners with massive amounts of data relevant to many different kinds of tasks and problems. However, much of the data is available only when the world is probed through the learners' own actions. Importantly, for developing learners, the kinds of play –and the kinds of probes possible – change systematically, and are at any given point in development, constrained both by the learning opportunities provided by the physical and social environment and by current physical and cognitive abilities of the learner. Human infants are motorically altricial and thus for their first few months dependent on caretakers: they can only observe or physically interact with objects and events in the environment that their caretakers set for them. The ordered constraint of experiences – ordered by developmental changes in the learners morphology and abilities and by the learners active probes of the environment – may play a critical role in constraining and channeling species-typical development and seems likely to have been exploited by evolutionary processes to ensure functional developmental outcomes that fit the environment (Gottlieb, 1991; West & King, 1987).

References

Audibert, J-Y., Munos, R. and C. Szepesvari (2009) Exploration-exploitation tradeoff using variance estimates in multi-armed bandits. Theor. Comput. Sci., 410(19):1876–1902.

Baldassare, G. and M. Mirolli (2013). Intrinsically motivated learning in natural and artificial systems, Berlin: Springer-Verlag.

Barto, A. (2013) Intrinsic motivation and reinforcement learning In G. Baldassarre and M. Mirolli, editors, Intrinsically Motivated Learning in Natural and Artificial Systems, pp. 17-47, Springer.

Berlyne, D. (1960). Conflict, Arousal and Curiosity (McGraw-Hill).

Cohn, D. A., Ghahramani, Z., & Jordan, M. I. (1996) Active learning with statistical models. Journal of artificial intelligence research, 4 : 129-145.

Deak, G., Krasno, A. ; Triesch, J. ; Lewis, J., & Sepeta, L. (2014) Watch the hands: infants can learn to follow gaze by seeing adults manipulate objects. Developmental Science, early on-line. DOI: 10.1111/desc.12122

de Boer, B. (2001) The Origins of Vowel Systems. Oxford Linguistics. Oxford University Press.

Csikszenthmihalyi, M. (1991). Flow-the Psychology of Optimal Experience (Harper Perennial).

Friston, K., Adams, R. A., Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: saccades as experiments. Frontiers in psychology, 3.

Gottlieb, J., Oudeyer, P-Y., Lopes, M., Baranes, A. (2013) Information Seeking, Curiosity and Attention: Computational and Neural Mechanisms Trends in Cognitive Science, , 17(11), pp. 585-596.

Gottlieb, G. (1991). Experiential canalization of behavioral development: Theory. Developmental Psychology, 27(1), 4-13. DOI: 10.1037/0012-1649.27.1.4

Gould, S. J. (1991). Exaptation: A crucial tool for an evolutionary psychology. Journal of social issues, 47(3), 43-65.

Hunt, J. M. (1965). Intrinsic motivation and its role in psychological development. Nebraska Symposium on Motivation, 13, 189–282.

Kagan, J. (1972). Motives and development. J. Pers. Soc. Psychol. 22, 51–66.

Kaplan, F. and Oudeyer, P-Y. (2007a) The progress-drive hypothesis: an interpretation of early imitation. In Dautenhahn, K. and Nehaniv, C., editor, Models and mechanisms of imitation and social learning: Behavioural, social and communication dimensions, pp. 361--377, Cambridge University Press.

Kaplan F. and Oudeyer P-Y. (2007b) In search of the neural circuits of intrinsic motivation, Frontiers in Neuroscience, 1(1), pp.225--236.

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.

Lickliter, R., & Honeycutt, H. (2009). Rethinking Epigenesis and Evolution in Light of Developmental Science. Handbook of Behavioral and Comparative Neuroscience: Epigenetics, Evolution, and Behavior.

Lowenstein, G. (1994). The psychology of curiosity: a review and reinterpretation, Psychological Bulletin 116(1): 75-98.

Lopes, M., & Montesano, L. (2014). Active Learning for Autonomous Intelligent Agents: Exploration, Curiosity, and Interaction. *arXiv preprint arXiv:1403.1497*.

Lopes, M., Oudeyer, P-Y. (2010) Active Learning and Intrinsically Motivated Exploration in Robots: Advances and Challenges (Guest editorial) IEEE Transactions on Autonomous Mental Development, 2(2), pp. 65-69.

Moulin-Frier, C., Schwartz, J. L., Diard, J., Bessière, P., Gilet, E., Simonin, E., ... & Millán, J. D. R. (2011). Emergence of articulatory-acoustic systems from deictic interaction games in a" vocalize to Localize" framework. Primate Communication, ad Human Language Vocalisation, gestures, imitation and deixis in humans and non-humans, 193-220.

Moulin-Frier, C., Nguyen, M. and Oudeyer, P-Y. (2014) Self-Organization of Early Vocal Development in Infants and Machines: The Role of Intrinsic Motivation, Frontiers in Cognitive Science, doi: 10.3389/fpsyg.2013.01006

Nguyen, M., Oudeyer, P-Y. (2013) Active Choice of Teachers, Learning Strategies and Goals for a Socially Guided Intrinsic Motivation Learner, Paladyn Journal of Behavioural Robotics, 3(3):136–146.

Oller, D.K. (2000) The Emergence of the Speech Capacity. Lawrence Erlbaum and Associates, Inc.

Oudeyer, P-Y. (2011) Developmental Robotics, Encyclopedia of the Sciences of Learning, N.M. Seel ed., Springer Reference Series, Springer.

Oudeyer, P-Y. (2010) On the impact of robotics in behavioral and cognitive sciences: from insect navigation to human cognitive development, IEEE Transactions on Autonomous Mental Development, 2(1), pp. 2--16.

Oudeyer P-Y, Kaplan, F. and Hafner, V. (2007) Intrinsic Motivation Systems for Autonomous Mental Development, IEEE Transactions on Evolutionary Computation, 11(2), pp. 265--286.

Oudeyer P-Y. and Kaplan F. (2007) What is intrinsic motivation? A typology of computational approaches Frontiers in Neurorobotics, 1:6, doi: 10.3389/neuro.12.006.2007

Oudeyer P-Y., Kaplan F. (2006) Discovering Communication, Connection Science, 18(2), pp. 189--206.

Oudeyer, P-Y. (2006) Self-Organization in the Evolution of Speech, Studies in the Evolution of Language, Oxford University Press.

Piaget, J. (1952) The Origins of Intelligence in Children, New York: International University Press.

Power, T.G. (1999) Play and Exploration in Children and Animals, Lawrence Erlbaum Associates.

Schmidhuber, J. (1991). Curious model-building control systems. IEEE International Joint Conference on Neural Networks.

Singh, S., R. L. Lewis, et al. (2010). "Intrinsically motivated reinforcement learning: An evolutionary perspective." Autonomous Mental Development, IEEE Transactions on 2(2): 70-82.

Smith L.B., Thelen E. (1994). A Dynamic Systems Approach to Development: Applications. MIT Press.

Smith, L. B., & Breazeal, C. (2007). The dynamic lift of developmental process. Developmental Science, 10(1), 61-68.

Smith, L. B. (2013). It's all connected: Pathways in visual object recognition and early noun learning. American Psychologist, 68(8), 618.

Thelen, E. (1994). Three-month-old infants can learn task-specific patterns of interlimb coordination. Psychological Science, 5(5), 280-285.

Thelen, E. S., & Smith, L. B. (1996). Dynamic systems approach to the development of cognition and action. MIT press.

A. L. Thomaz and C. Breazeal (2008) Experiments in Socially Guided Exploration: Lessons learned in building robots that learn with and without human teachers. Connection Science, Special Issue on Social Learning in Embodied Agents, 20(2-3), pp. 91-110.

West, M. J., & King, A. P. (1987). Settling nature and nurture into an ontogenetic niche. Developmental Psychobiology, 20(5), 549-562. PMID: 3678619

Yu, C., & Smith, L. B. (2013). Joint Attention without Gaze Following: Human Infants and Their Parents Coordinate Visual Attention to Objects through Eye-Hand Coordination. PloS one, 8(11), e79659.