Trends in Cognitive Sciences

CellPress
REVIEWS

Forum

Priors in Animal and Artificial Intelligence: Where Does Learning Begin?

Elisabetta Versace,[1,*]
Antone Martinho-Truswell,[2,*]
Alex Kacelnik,[2] and
Giorgio Vallortigara[3]

A major goal for the next generation of artificial intelligence (AI) is to build machines that are able to reason and cope with novel tasks, environments, and situations in a manner that approaches the abilities of animals. Evidence from precocial species suggests that driving learning through suitable priors can help to successfully face this challenge.

Minds from Scratch

Empiricists and nativists have clashed for centuries in understanding the architecture of the mind: the former as a *tabula rasa*, and the latter as a system designed prior to experience [1]. Today, the related burning question is how to build successful minds from scratch. The next generation of artificial minds is expected to far surpass the recent success of machine learning in specific tasks and formally structured domains, from defeating human Go masters to transcribing speech and recognising objects of particular categories. This improvement will come through accomplishments in strong AI that include reasoning, generalisation to new tasks, and advanced language processing. 'Common sense', in the sense of flexibly reasoning appropriate responses, remains an unachieved goal for current AI. The question, summarised in the debate between the nativist Gary Marcus and the pioneer of machine learning, Yann LeCun [2–4], is the following: shall we search for a unitary general learning principle able to flexibly adapt to all conditions, including novel ones, or structure artificial minds with driving assumptions, or priors, that orient learning and improve acquisition speed by imposing limiting biases? Examining the well-studied example of a newly hatched precocial bird neatly shows how research on animal cognition can nourish design principles for AI.

Predisposed Cognition

Soon after birth, a chick orients itself towards the mother hen, recognises its social partners even when partly occluded, emits pleasure and distress calls in different contexts, and performs simple but effective arithmetic calculations without explicit training [5]. Dealing with animals (human and otherwise) for a large extent of our time, we realise how varied and flexible their behaviour is. However, precocial species, namely those with a mature motor and sensory system at birth, tell us something more. Newly hatched domestic chicks, for instance, are predisposed to orient towards objects that exhibit features associated with animate objects, such as biological motion, changes in speed, and face-like configurations (Figure 1A). It has been hypothesised that this set of unlearned priors helps chicks to orient towards the mother hen and their siblings. Mere exposure triggers a rapid and robust learning process called filial imprinting (Figure 1B,C). Imprinted birds are able to recognise complex, multimodal objects such as conspecifics, and can imprint on other animals, or indeed any arbitrary stimuli (within certain constraints) to which they have been exposed – most famously Konrad Lorenz, who studied imprinting in ducks and geese. After imprinting, hatchlings not only recognise their mother from different and previously unexperienced points of view (including translational and rotational invariance; e.g., [6]) but they also recognise their siblings as they change appearance during development. We recently discovered that the imprinting mechanism allows a remarkable degree of generalisation to novel objects that have in common with the imprinting object only abstract features (such as the presence of AA, 'same'; or AB, 'different', patterns) and that this is true in different species [7,8]. While these abilities had previously been shown in extensively trained animals, these results make clear that such generalisations are a spontaneous competence that is already available at the onset of life.

How can young birds orient towards the 'right' stimuli in the absence of any previous experience, and master complex generalisation tasks without supervision and reinforcement? In contrast to machine-learning systems, chicks do not require explicit reinforcement, supervised learning, or thousands/millions of examples to feed learning. They are equipped with dedicated orienting and learning mechanisms that work as adaptive priors and architectural structures. These priors imply some assumptions about the external world that guide learning, but can, and must, allow errors, as was the case of goslings imprinted on Konrad Lorenz. Research has shown that early preferences of chicks are not strictly species-specific but apply equally to hen face-like or polecat face-like features [9], or to the biological-motion appearance of either a hen or a cat [10]. This is due to the fact that the orienting mechanisms cannot be too specific for the individual features of the mother hen, which are to some extent unpredictable from the genetic repertoire. A level of non-specificity is functional in avoiding excessive false negatives in the form of failed recognition caused by variability between adults within a species, and by changes in the appearance of even a single individual. Optimal learning mechanisms must trade being sufficiently open to allow a wide

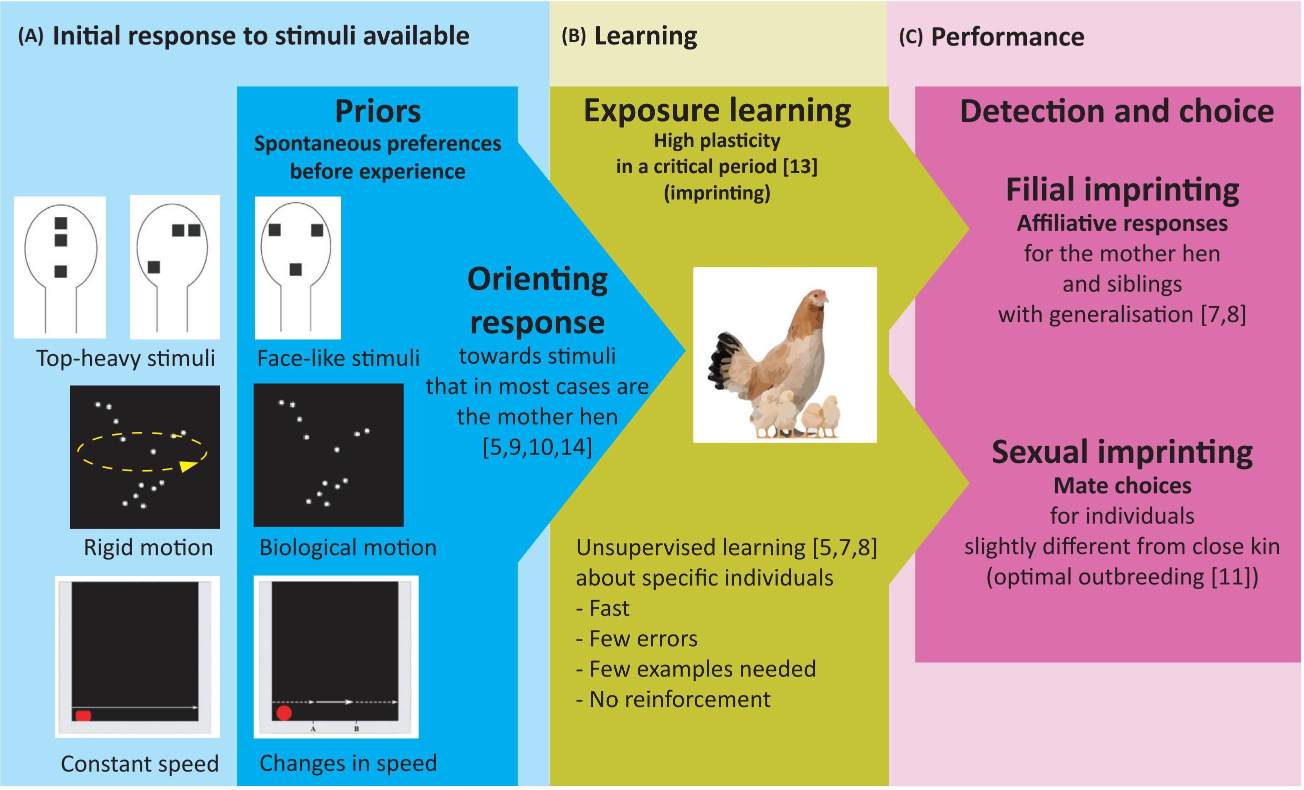

Trends in Cognitive Sciences

Figure 1. Priors Orient Fast and Efficient Learning in Precocial Hatchlings. (A) Without experience, some stimuli are more attractive to hatchlings: they preferentially orient towards face-like stimuli, biological motion, and objects of variable speed [5,9,10]. In the wild, these features are associated with the mother and siblings [14]. (B) This initial orienting response directs fast, unreinforced learning (imprinting) towards 'appropriate' stimuli. (C) What is learned during the imprinting phase, with some degree of generalisation (e.g., [7,8]), determines subsequent choices. In the short term, hatchlings affiliate with the mother hen (filial imprinting), which is crucial for survival; in the long-term, mating choices will be directed towards individuals that are different from close kin (avoiding inbreeding) but belong to the same species (see optimal outbreeding in [11]). See also [4,6,12].

range of stimuli to be stored as imprinted memories, against being sufficiently specific to avoid imprinting on inappropriate objects that in natural environments coexist with the chick's mother and siblings. In nature, this trade-off between error rates is tuned only by natural selection, but it can be emulated in artificial systems where other paths are available to solve the problem. Functional assumptions paired with plastic learning, hence, enable the hatchling to be fast and effective in learning, such that a limited number of examples is sufficient to make ecologically valid generalisations in most biologically relevant cases. Moreover, preference for a slight deviation from the memory of the imprinting object allows the chick to sample better the properties of the desirable target, exactly as for adult birds where deviation from the object of sexual imprinting (a similar process that occurs later in life) avoids inbreeding with immediate kin while promoting mating with a member of the same species (optimal outbreeding [11]), showing effective modulation of learning via *a priori* preferences (Figure 1C).

Plasticity and Priors

Learning from biological examples implies that high plasticity coupled with prior assumptions is not sufficient for strong AI: a temporal dimension, with predesigned critical or sensitive periods, is a fundamental part of the solution. When looking at biological systems, in fact, there is evidence that both early predispositions and high plasticity are transient phenomena that terminate either with some maturational processes or when the necessary information (e.g., the identity of the mother hen in the case of chicks) has been acquired. Critical and sensitive periods have been observed in avian and mammalian species in several domains and functions, including orientation to face-like stimuli [12], ocular dominance [13], and language acquisition [1]. The peak of plasticity in the wiring of the visual system observed in specific timewindows points to the costs associated with plasticity, which cannot be indefinitely sustained. This is one reason why, after a certain age, learning new languages and solving amblyopia is so difficult, and why early experience is in many cases important for subsequent stages of life. Recent evidence in rodents suggests that the spontaneous plasticity of the nervous system is actively reduced by molecular 'brakes' that promote circuit stabilisation in mature brain function [13]. Declining plasticity is also ecologically adaptive: the realities of the perinatal period mean that the most relevant potential substrate for imprinting is the mother bird. If plasticity endured beyond the critical period, numerous erroneous substrates – any moving animal – would likewise trigger imprinting, thus reducing its adaptive valence. The constraint in time also serves to constrain candidate substrates appropriately. A similar effect may be observed in bird song, wherein young altricial songbirds learn their mating songs while still confined to the nest, which prevents the incorporation of irrelevant sounds experienced after fledging.

Considering the balance between priors, plasticity, and the observed brakes to plasticity, we argue that evidence from animal research suggests that (i) AI systems could benefit from being equipped with a rich but constrained set of priors and specialised learning mechanisms similar to those seen in precocial animal species, rather than being endowed only with general purpose, unifying mechanisms, (ii) plasticity without priors and critical periods of expression for these priors might be associated with costs that prevent effective learning and stable cognitive functions. The ability to shift between priors and thus direct plasticity may speed our way to strong AI, while pursuing less-structured AI may help to identify new and potentially unexpected useful priors.

Acknowledgments

G.V. was funded by a European Research Council (ERC) advanced grant (ERC-2011-ADG_20110406: project 461 295517, PREMESOR), Fondazione Caritro Grant Biomarker DSA [40102839] and PRIN 2015 (Neural bases of animacy detection, and their relevance to the typical and atypical development of the brain). A.M. and A.K. were funded by a Leverhulme Foundation research grant (Relational Concept in Ducklings, RPG-2016-192).

[1]Queen Mary University of London, School of Biological and Chemical Sciences, Department of Biological and Experimental Psychology, Mile End Road, London E1 4NS, UK

[2]University of Oxford, Department of Zoology, New Radcliffe House, Oxford OX1 3PS, UK

[3]University of Trento, Center for Mind/Brain Sciences, Piazza Manifattura 1, Rovereto, Italy

*Correspondence:
e.versace@qmul.ac.uk (E. Versace) and
antone.martinho@zoo.ox.ac.uk (A. Martinho-Truswell).

https://doi.org/10.1016/j.tics.2018.07.005

References

1. Pinker, S. (2004) *The Blank Slate: The Modern Denial of Human Nature,* Viking
2. LeCun, Y. and Marcus, G. (2017) Does AI need more innate machinery? Debate at the NYU Center for Mind, Brain, and Consciousness, October 5, 2017. https://wp.nyu.edu/consciousness/innate-ai/
3. Marcus, G. (2018) Innateness, alphazero, and artificial intelligence. *arXiv*
4. Lake, B.M. (2017) Building machines that learn and think like people. *Behav. Brain Sci.* 40, e253
5. Vallortigara, G. (2012) Core knowledge of object, number, and geometry: a comparative and neural approach. *Cogn. Neuropsychol.* 29, 213–236
6. Wood, J.N. (2015) Characterizing the information content of a newly hatched chick's first visual object representation. *Dev. Sci.* 18, 194–205
7. Martinho, A. and Kacelnik, A. (2016) Ducklings imprint on the relational concept of 'same or different'. *Science* 353, 286–288
8. Versace, E. *et al.* (2017) Spontaneous generalization of abstract multimodal patterns in young domestic chicks. *Anim. Cogn.* 20, 521–529
9. Johnson, M.H. and Horn, G. (1988) Development of filial preferences in dark-reared chicks. *Anim. Behav.* 36, 675–683
10. Vallortigara, G. *et al.* (2005) Visually inexperienced chicks exhibit spontaneous preference for biological motion patterns. *PLoS Biol.* 3, e208
11. Bateson, P. (1978) Sexual imprinting and optimal outbreeding. *Nature* 273, 659–660
12. Johnson, M.H. *et al.* (2015) The two-process theory of face processing: modifications based on two decades of data from infants and adults. *Neurosci. Biobehav. Rev.* 50, 169–179
13. Hensch, T.K. and Quinlan, E.M. (2018) Critical periods in amblyopia. *Vis. Neurosci.* 35, E014
14. Di Giorgio, E. *et al.* (2017) Filial responses as predisposed and learned preferences: early attachment in chicks and babies. *Behav. Brain Res.* 325, 90–104

Forum

The Little Engine That Can: Infants' Persistence Matters

Kelsey Lucca[1,]* and
Jessica A. Sommerville[1]

Persistence is central to outcomes across a range of domains: the harder you try, the further you get. Yet relatively little is known about the developmental origins of persistence. Here, we highlight key reasons for a surge of interest in persistence in infancy and early childhood.

Persistence, the ability and motivation to engage in sustained effort to overcome challenges and achieve goals (Figure 1), is a key predictor of educational attainment (i.e., graduation rates, grade point average), positive qualities (e.g., resiliency), and life outcomes (e.g., job maintenance, marital success) starting in middle childhood [1]. Yet, we know little about the developmental roots of persistence in infancy and early childhood. Here, we call for a new empirical interest in persistence that capitalizes on and expands upon recent discoveries in infants' knowledge, learning, and behavior. Specifically, we argue that (i) individual differences in persistence emerge during infancy and influence later development, (ii) persistence is a valuable measure of what infants understand and care about across domains, and (iii) persistence offers a window into metacognitive and decision-making processes. Below, we marshal the evidence to support each of these claims and point to important future directions on this topic.

Persistence in Infancy Has Implications for Long-Term Outcomes

Much of what is known about early persistence comes from research on a