Postdoctoral Research Fellowship in Theoretical Biology/Computer Science

HOW EXPLORATORY AND SELECTIVE DEVELOPMENTAL MECHANISMS GENERATE FACILITATED VARIATION, AND IMPOSE DIRECTION ON EVOLUTION

Kevin Laland[1] & Richard Watson[2]

[1] School of Biology, University of St Andrews, U.K.
[2] Department of Computer Science, University of Southampton, U.K.

BACKGROUND

Living organisms must produce functional responses to highly diverse, complex and constantly changing inputs (e.g. immunological responses to rapidly evolving bacteria or viruses, or learned adjustments to changing situations). Often organisms respond to such challenges through 'exploratory mechanisms' (Gerhart & Kirschner, 1997; West-Eberhard 2003; Kirschner & Gerhart, 2005), which are complex developmental systems that operate by generating variation (i.e. 'exploring' possibilities), largely at random, testing variants' functionality, and selecting the best solutions for regeneration, in an iterative developmental process. The process resembles adaptation by natural selection (a.k.a. '*somatic selection*'), except that it allows for *ontogenetic information gain* within a lifetime rather than conventional *genetic information gain* across multiple generations.

Diverse biological processes function in this way. For instance, the adaptive immune system generates antibodies and T cells with initially random variation, then internal selection multiplies and refines those that bind successfully to antigens, with a memory of effective molecules retained (Klenerman, 2017). The vascular and tracheal systems, animal learning, much collective animal behavior (e.g. central-place foraging), and micro-tubular systems all operate on similar principles – exploiting exploratory and selective mechanisms to generate novel functional responses in development (Campbell 1960; Gerhart & Kirschner, 1997). Other biological processes, such as the remodeling of bone and soft tissue (muscles, tendons), are known to be responsive to functional demands (Hall, 2015), and these processes have also been characterized as reliant on somatic selection (West-Eberhard, 2003). The anatomical organization of brains exhibits similar adaptability. There are estimated to be over 100 trillion neural connections (synapses) in the human brain, several orders of magnitude more than could be specified by the *c.* 21,000 genes in our genome (Edelman, 1987; Kirschner & Gerhart, 2005). While the gross structure of the vertebrate nervous system is thought to be set up by the demarcation of pathways of nerve growth by genes (Kirschner & Gerhart 2005), experiments show that brains depend in part on exploratory mechanisms to establish their anatomical organization (Edelman 1987; Gerhart & Kirschner, 1997; Kirschner & Gerhart 2005). During development, the nervous system generates excess neurons, excess neuronal connections, and excessively distributed neuronal connections, through random exploration. It then prunes these, retaining solely those required. Much of the patterning of the brain depends on exploratory mechanisms' use of functional interactions to sort out connectivity. The final anatomy of vertebrate brains thus depends heavily on experience.

Exploratory mechanisms are adaptive because rapid exploration of a large space of possibilities combined with feedback (e.g. reward/punishment) allows for information gain from the current environment. Crucially this occurs at timescales faster than genetic evolution (i.e. within individual's lifetimes). Challenges arise from the internal environment too. Organisms must cope effectively with very large numbers of individual-specific 'internal failures' in somatic genome, epigenome and microbiome that are too numerous and/or unique to be anticipated by genetically coded plasticity, and the self-organization of random variation is critical to this form of 'adaptive improvisation' (Soen et al., 2015). As a result, across a very broad range of conditions, including unanticipated circumstances, organisms are often capable of producing highly functional responses. Experiments and theory both show that randomness in exploration is especially useful in confronting a variable or novel environment (Deneubourg et al 1983; Strickland et al 1995; Soen et al, 2015; Richerson 2019). Exploratory mechanisms have a major advantage in flexibility (Gerhart & Kirschner 1997), being self-correcting and adaptable to changes in other parts of the organism – e.g. resizing cortical areas to match sensory fields (Gerhart & Kirschner, 1997).

Exploratory mechanisms can be costly systems because they are wasteful – to find effective solutions they must generate a very large number of variants, only a fraction of which will be retained (Gerhart & Kirschner, 1997). Hence, natural selection should favor biases in the operation of exploratory mechanisms (e.g. allowing the immune system rapidly and cheaply to target a specific reliably present antigen). In principle, exploratory mechanisms can adjust to new challenges during ontogeny, and later these phenotypes can be stabilized by natural selection (e.g. generating probabilistic biases in exploration through shaping precursor cell numbers, nerve growth factor concentrations, or induction factor concentrations in particular regions). Yet there must be limits to the extent of such biases if the benefits of exploration are to be retained. That is, there is an inherent trade-off between being biased to produce particular responses efficiently and the cost of being able to explore novel responses when needed. Currently, the nature of those trade-offs is not understood.

The role of auxiliary processes. There is also another important means by which systems reliant on exploratory mechanisms can evolve. Exploratory mechanisms are often associated with auxiliary processes with which they constantly interact, and coevolve. For instance, adaptive immunity operates in concert with innate immunity, whilst learning interacts with perceptual, motivational and other cognitive systems. These auxiliary processes often evolve by conventional natural selection to generate genetic information gain on much slower timescales.

For illustration, to make effective memories the adaptive immune response needs to be induced (or 'primed'), which happens during the early stages of infection. This requires activation of the innate immune system, without which the adaptive immune system can be overwhelmed, as studies using genetically deficient mice have shown (Klenerman, 2017). The same holds for learning, where adaptive specializations have long been recognized (Hinde & Stevenson-Hinde, 1973). For instance, our experimental work has uncovered an adaptive specialization in the learning of stickleback fishes, with 9-spined but not 3-spined sticklebacks able to exploit public information concerning the richness of food patches (Coolen et al 2003). Further experiments reveal a shift in reliance on public information use in 9-spines in reproductive condition, and imply that this adaptive specialization likely evolved through endocrinal changes, without comprising the general learning ability of the animals (Webster et al 2011; Webster et al, 2019). Adaptive specialization in learning can also occur through up-regulating animals' perceptual systems (Mineka & Cook, 1988; Olsson & Phelps, 2007). Refinement of auxiliary systems makes sense: individuals experience multiple immunological and learned challenges during their lifetimes, and cannot afford to comprise their future flexibility by over-adjusting to a single specific challenge (a response analogous to overfitting in learning theory; Kouvaris et al 2017). Our theoretical work paints a similar picture. We have found that both evolved inductive biases (Kouvaris et al, 2017; Kouvaris 2018) and evolved learning strategies (Kendal et al, 2009, 2018; Toyokawa et al, 2019) can make some behaviors easier to learn, or more readily acquired than others, but without compromising the general learning capability. Formal equivalences between how individual animals learn and how populations evolve (Watson & Szathmary 2016), shed light on how organisms can evolve sensitivities to environmental regularities that allow them to generalize to produce adaptive variation in novel circumstances; reducing variability in some dimensions but increasing variability that enhances evolvability (Kouvaris et al 2017).

Modeling developmental plasticity. The manner in which organisms respond to variable environmental conditions is usually conceptualized within evolutionary biology by norms of reaction, which specify the phenotypes of a single genotype expressed over a range of environments (Griffiths et al 2000). From this standpoint, the reaction norm is a property of the genotype, and maps the phenotype and hence evolutionary fitness of each genotype across environmental conditions. This allows the evolution of plastic traits to be modeled using standard population genetic and quantitative genetic tools, which in turn allows predictions to be made as to how plastic traits will evolve (Schlicting & Pigliucci, 1998; Via & Lande 1985; Sultan & Spencer 2002; Lande 2009).

While any form of phenotypic plasticity (including exploratory processes) can be characterized as a reaction norm, the evolution of plasticity depends strongly on the mechanistic basis of that 'reaction'. Formal models for the evolution of plasticity based on reaction norm concepts are almost always modelled as simple genetically-determined relationships between a range of environments and a range of responses (e.g. a linear function). Whilst that is a sensible starting point for modelling some forms of plasticity, it omits all exploratory behaviour. Importantly, that modelling approach assumes that, whatever adaptive phenotype an organism expresses, past evolution has already had the opportunity to select genotypes for their ability to produce this phenotype in this circumstance. In short, on this assumption, although the phenotype is plastic, it cannot be new

(aside from limited forms of interpolation or extrapolation). In contrast, when considering a plastic reaction produced by an exploratory process, the range of environmental possibilities and phenotypic responses may be very large; far too large for past selection to have sampled even a fraction of the possibilities. In this case, an exploratory process may find phenotypes that are genuinely novel, never before expressed in evolutionary history. Accordingly, this offers the possibility that significantly novel adaptive phenotypes are produced by exploratory processes that have not already been sampled by past selection. This demands a shift in explanatory thinking: how can these phenotypes be adaptive if they have not been previously exposed to selection?

This raises the second, and arguably more important, difference from formal reaction norm models; exploratory processes are iterative processes of improvement utilizing within-lifetime feedback from the environment. The phenotype resulting from an exploratory process can be a 'within-lifetime adaptation', tuned to the current environment 'on the fly', and not merely a genetically predetermined response to (or simple function of) an environmental cue. Selection on genetic variants may then favour a *strategy* for producing adaptive phenotypes in a much more general way than tuning the slope and intercept of a reaction norm. This strategy is constituted by the mechanistic details of the generation, and within-lifetime selection, of phenotypic variation. This feedback is entirely missing in evolutionary genetic models that specify reaction norms based on non-exploratory mechanisms (e.g. Lande 2009). For an exploratory process, the targets of evolution are thus the parameters of the 're/generate' process and the 'test' process, not a simple mapping between input and output.

Exploratory processes thus afford the possibility that phenotypes discovered by within-lifetime adaptive plasticity may precede genetic exploration of phenotype space (Baldwin 1896, Hinton and Nowlan 1987, West-Eberhard 2003) – an active 'adaptability driver' (Bateson 2006) rather than a passive repository of genetic information gain. The interaction between these two adaptive timescales may be complex. For illustration, selection on immune responses could either lead to evolution of the innate immune system (for instance through specific targeting of reliably present molecules derived from microbes, such as Toll-like receptors), or to the selection of alleles of adaptive immunity genes (for instance, that reliably produce useful T cell or B cell variants) (Du Pasquier & Litman 2011). Some theory exploring adaptive dynamics across two timescales exists (e.g. Feldman & Cavalli-Sforza, 1976; Boyd & Richerson, 1985; Kouvaris et al, 2017; Watson et al 2014, 2016). The interaction of lifetime learning and genetic evolution is an example that is relatively well-developed, and one of the best known is Hinton and Nowlan's (1987) model. This work confirms that lifetime learning can lead genetic evolution to find solutions much faster than a model without lifetime learning. However, in Hinton and Nowlan's model, genetic evolution has the effect of prohibiting further exploration. This limitation is inevitable in their model because each trait in the phenotype vector is either entirely genetically controlled or entirely plastic. Although there are many such traits, and some fraction of traits may be genetically controlled whilst others remain plastic, genetic evolution and plastic exploration operate in the same space (a vector of traits) and necessarily compete for control of each phenotypic trait in a direct one-for-one manner. This precludes the possibility of genetic accommodation that does not result in a loss of plasticity (genetic assimilation). It is therefore not possible in this model for genetic evolution to bias lifetime learning without removing the exploratory process.

At the other extreme, if genetic evolution and lifetime exploration work in entirely separate spaces then there is no consistent selective pressure for genetic evolution to accommodate lifetime learning since incremental genetic changes cannot incrementally reduce the cost of lifetime learning (Mayley 1996). How can the exploration operate in a phenotypic space that genetic evolution can subsequently accommodate without the two processes directly competing for control of the phenotype? The biology tells us there is a middle ground. One possibility is that genetic evolution controls the correlations between traits, rather than the traits themselves. A second is that exchanges between exploratory and auxiliary processes also allow for separate but interacting spaces. A third is that the costs and benefits of exploration trade off to maintain a frequency dependent balance with partial genetic accommodation. In the programme of work below, we describe a new modelling approach to address these ideas.

The absence of a dedicated body of theory focused on the two-way interaction of evolution and exploratory mechanisms makes it more difficult for evolutionary biologists to envision how organisms can possess agency or impose direction on evolution, leading to much unresolved debate within the field (Laland et al 2014; Wray et al 2014; Laland et al 2015). Partly because they possess exploratory mechanisms, living organisms are not just passively pushed around by genes or external forces, but rather are self-regulating, purposive entities, capable of seeking out information, and thereby generating novel adaptive phenotypes during development. When

plasticity is characterized as genotypic norms of reaction it is difficult to envision plasticity playing an active evolutionary role, as from such a standpoint, plastic traits evolve in a manner little different from non-plastic traits (Laland et al 2015; Uller et al., 2019). When plasticity recognizes truly exploratory processes, acting as evolutionary processes on ontogenetic timescales, it becomes possible to envision how adaptive phenotypic solutions can arise during development, how such solutions can be truly novel (rather than immediate products of earlier selection) and how developmental plasticity can, not just precede in time but also, impose direction on genetic evolution (West-Eberhard 2003; Uller et al., 2019). To model this in a productive way that captures the biology, it is necessary for new models where exploration is not inevitably and inflexibly opposed by genetic assimilation.

PROGRAM OF WORK

We propose a novel interdisciplinary collaboration between a biologist (Laland, Project Leader) and computer scientist (Watson, Project Co-Leader), each with outstanding track records in this field. This ground-breaking study will pioneer computational models of adaptation in biological systems that exhibit exploratory behavior in interaction with genetic evolution. The novelty of our approach is to exploit the deep functional isomorphism between multiple adaptive timescales in learning systems, well-understood in (machine) learning theory (Watson & Szathmáry 2016). This has already been highly successful in modelling the interaction of variation and selection processes occurring on different timescales, namely, in the evolution of evolvability (Kouvaris et al, 2017; Kouvaris 2018; Watson et al 2014, 2016; Watson & Szathmáry 2016; see also Parter et al 2008). Here we build on this earlier work to study the interaction of variation and selection processes both within and between generations – the evolution of 'explorability'. This enables (*1*) the formal evolutionary analysis of exploratory mechanisms, (*2*) the simultaneous modelling of adaptive processes across two timescales, developmental and evolutionary, without the latter removing the former, and (*3*) analyses of interactions between exploratory mechanisms and auxiliary processes. Each extension plausibly enables forms of adaptation impossible with simpler formulations.

WORK PACKAGE 1. EXPLORATORY MECHANISMS AND EVOLVABILITY

We will devise computational models of exploratory mechanisms with the novel feature that genetic changes control *correlations* between phenotypic trait values that remain variable through lifetime exploration ('evolving exploratory correlations model', M3). We will implement directly comparable models of exploration where both genetic and lifetime variation operate in the same space (in the style of Hinton and Nowlan) ('direct exploration', M2), an 'instructional' (or 'genetic reaction norm', M1) model and, for completeness, a non-plastic model (M0).

Watson's earlier work on the relationship between evolution and correlation learning (Watson et al, 2014, 2016; Kouvaris et al, 2017), investigated the genetic evolution of correlations that affect the variability of genetic evolution (i.e. the evolution of evolvability). This work studied the interaction of evolution on two timescales – the relatively fast evolution of direct effects on traits and the relatively slow evolution of correlations affecting the co-variation of those traits – and showed that it is possible for the evolution of correlations to facilitate evolvability by biasing subsequent genetic exploration but not removing variability. By working in correlations, the slow adaptive process does not directly compete with the fast adaptive process, but instead can bias the combinations of traits that the fast process can produce without preventing variability in any trait. This allows the interaction of the two processes to solve adaptive problems that cannot be solved by either process alone. Here, M3 converts the timescales of these prior models to investigate the genetic evolution of correlations that affect within-lifetime exploration (i.e. the evolution of 'explorability' rather than evolution of evolvability).

In the exploratory models (M2 and M3), phenotypic variants will be generated through an iterative process of (biased) random phenotypic variation and developmental selection in response to environmental inputs. In M2, genetic evolution speeds up learning by removing variability on individual traits (in the style of Hinton & Nowlan) whereas in M3, genetic evolution controls the co-variation of exploratory variation. The latter will utilize a simple mechanistic model of trait interactions (Watson et al 2014). The 'reaction norm model' (M1) will generate phenotypes in response to environmental inputs just once (non-iteratively) according to genetic instructions, whilst M0 has genetically specified, non-plastic traits, insensitive to environment.

In addition to these basic scenarios, we will explore genes that specify structural features of the exploratory/selective process (rate of generation of variants, intensity of selective pruning, investment in memory/auxiliary processes). Gene influences will be modeled on known genetic influences on learning and/or immunity (e.g. large numbers of alleles, high capability for recombination, random mixing of elements from

different genes). Organisms will pay a fitness cost that is a positive function of the developmental time necessary to generate a suitable variant, or set of variants and/or the total number of variants generated, which will allow us to evaluate how the costs of exploration may affect evolutionary outcomes. Information gained by the developmental system potentially changes the functionality, and hence selection acting on, the genetic system, allowing information flow between them. Advanced work will extend the spectrum of models from univariate (M2), bivariate (associative) (M3), to higher-order (tensor) information systems, as well as complex restrictions (i.e. deep developmental networks). This allows genetic evolution to bias exploration in increasingly subtle ways representative of different biological cases whilst retaining generality afforded by machine learning theory. We will test the following hypothesis:

H1: ***Exploratory mechanisms evolve in a qualitatively different manner from other (i.e. instructional) forms of plasticity, with predictable directional differences in (a) evolvability, (b) degree of plasticity, (c) adaptability, (d) global adaptation, (e) amount of genetic variation, and (f) memory (as specified below).***

H1a. Evolvability. Exploratory mechanisms are expected to confer on developmental systems a greater capacity to evolve, without themselves evolving much, by eliciting genetic change in the auxiliary process (genetic accommodation). In keeping with findings from previous work on the evolution of evolvability, this is expected to be greater for M3 than M2 (Watson et al., 2014, 2016). In contrast, instructional forms of plasticity should primarily evolve through genetic change in plasticity (genetic assimilation). In the absence of auxiliary processes, rates of genetic change should be lower in exploratory compared to instructional plasticity. Conversely, when coevolving with exploratory mechanisms, rates of genetic change in auxiliary processes should be greater than amongst genes underlying both exploratory and instructional plasticity. Exploratory mechanism M3 should primarily evolve through adaptive biases in auxiliary systems, and maintain a high level of variability in phenotypes, and fitness differences, facilitating further phenotypic adaptation, rather than limiting it.

H1b. Plasticity. Instructional forms of plasticity will gain and lose plasticity to a greater extent than exploratory mechanisms, which should maintain high levels of plasticity whilst adapting to changing environments. This effect is expected to be greater, and sustained longer, for M3 than M2.

H1c. Adaptability. Correlation-based exploratory mechanisms, M3, will maintain a higher level of adaptation across a broader range of conditions than non-correlational exploration, M2, or instructional plasticity M1.

H1d. Global adaptation. Systems reliant on exploratory mechanism M3 will attain higher fitness phenotypes (e.g. globally optimal) across a broader range of conditions (C.f. M2, M1) (by analogy with Kounios et al 2016).

H1e. Genetic variation. Where genetic variation is necessary to generate the high levels of phenotypic variation required for exploratory mechanisms' functionality, developmental systems reliant on exploratory mechanisms will maintain greater genetic variation than systems reliant on instructional forms of plasticity.

H1f. Properties of memories. Exploratory systems repeatedly exposed to environmental contingences or threats should retain (associative) memories of effective phenotypes [Watson et al 2014]. The presence and duration of memory should co-vary with rate of recurrence. However, in the presence of auxiliary processes, memories will be offloaded without precluding further exploration (see below).

Outputs: Two papers, targeted for top journals (e.g. *Nature*, *Science*, *Nature Ecology & Evolution*, *Evolution*).

WORK PACKAGE 2. THE ROLE OF AUXILIARY PROCESSES

We will extend the models to incorporate auxiliary processes, comparing the evolutionary dynamics of auxiliary processes operating alone and in conjunction with exploratory mechanisms. Auxiliary processes will be modeled as genetically specified and non-plastic. We will test the following hypotheses:

H2: ***(a) Exploratory mechanisms operating in conjunction with auxiliary processes will evolve less rapidly than exploratory mechanisms operating alone, and (b) auxiliary processes operating in conjunction with exploratory mechanisms will evolve more rapidly than auxiliary processes operating alone.***

H3: ***Information gained from exploratory (developmental) mechanisms will often be 'out-sourced' to auxiliary (genetic) systems, allowing developmental events to impose direction on biological evolution in a 'one-shot learning' style rather than, or in addition to, an associative learning style.***

We predict that genetic information specifying the auxiliary process will primarily comprise information initially

gained by the ontogenetic (exploratory) process, and that auxiliary processes operating in conjunction with exploratory mechanisms will evolve more rapidly (i.e. plasticity led evolution) than auxiliary processes alone.

H4: ***Where 'out-sourcing' of information to the auxiliary process is available in conjunction with biasing the exploratory mechanism (M3) or incrementally fixing it (M2), out-sourcing will be favored by natural selection.***

This implies that constraints on learning or biases in the generation of variation by the adaptive immune system should be rare, relative to biases in perceptual/motivational systems, or evolution of innate immune responses.

H5: ***Exploratory mechanisms operating in conjunction with auxiliary processes will allow for coordinated ontogenetic and genetic adaptation and thereby exhibit enhanced (a) evolvability, (b) adaptability, and (c) global adaptation, compared to exploratory mechanisms alone and instructional plasticity.***

Outputs: This work package will generate three further papers, again targeted for top journals.

WORK PACKAGE 3. DESIGNING EMPIRICAL TESTS

Design features of the theoretical work are modeled on two comparatively well-studied exploratory mechanisms, namely animal learning and adaptive immunity, with which the applicants and their collaborators have some familiarity. Animal learning is a topic that both PIs have studied for many years, including extensive experimental work in the Laland laboratory, whilst a parallel (separately funded) project has been established with an immunology lab that will allow some provisional empirical testing of the models' applicability to acquired immunity. The response variables subjected to analyses have been selected for their potential to afford empirical tests of the models' findings. This will ensure that the model abstractions remain closely aligned with specific biological systems. WP3 will develop designs for potential experimental, statistical or comparative phylogenetic work that evaluates these findings once the analyses are complete. The projects' empirically testable predictions and recommendations for methodological approaches will be written up as a synthetic paper, and discussed at the workshop, as well as with other members of the wider consortium. A primary objective will be to develop a full-blown empirical test of the theoretical findings, suitable for submission as a future grant application, by the end of the current grant.

Additional outputs: We will also write a synthetic paper and organize a workshop on exploratory process.

SIGNIFICANCE OF PROPOSED WORK.

1. Interdisciplinary exchange. The project has the potential to have a major impact on the fields of evolutionary biology, immunology, learning theory, neuroscience and psychology. Thus far each exploratory mechanism in different biological domains has primarily been studied independently, and with limited interaction with evolution, but their common underlying principles raises the possibility of general rules or patterns that apply across very different biological domains. In addition, in the longer term there are likely to be nontrivial practical uses for a model of exploratory mechanism dynamics (e.g. in predicting how the adaptive immune system will respond to microbial evolution and what would break it, or through facilitating new approaches and significant improvements to machine learning).

2. Computational tools. The highly novel computational modelling will pioneer new tools for diverse researchers that study exploratory systems, forging connections between hitherto distinct fields, and leading to a deeper understanding of these complex forms of plasticity. The new modeling approach will be of value to other researchers who wish to study how purposive systems evolve.

3. Explanatory tools. The project will not only lead to a more detailed understanding of how complex organisms evolve but also help to demystify current debates within evolutionary biology, including the notion of phenotype-led evolution, the concept of organismal agency, and the evolutionary significance of active phenotypes. The comparison between models of exploratory and instructional plasticity will highlight limitations of the genetically specified reaction norm perspective common in the field. The programme of work offers the prospect of a major step forward in understanding how complex animals, including humans, evolve and adapt to novel threats.

CAPACITY FOR SUCCESS

This project is a novel interdisciplinary collaboration between an evolutionary biologist (Laland) and a computer scientist (Watson), each with outstanding track records in this field. Laland has studied the interplay between

learning and evolution for >30 years, producing 12 books and *c.*300 scientific articles, and regularly publishing in leading journals (e.g. Science, Nature, PNAS). He is an elected Fellow of the Royal Society of Edinburgh, has been the recipient of over £16m in grant income, including an ERC Advanced grant, and is widely cited (h-index=102, lifetime cites c. 42,000). Watson has investigated and modelled evolutionary theory using neural networks and computational biology for >20 years. His award-winning work has established that evolution by natural selection is capable of more 'intelligent' problem solving than previously realized – findings that led to a cover article of New Scientist and many other features (e.g. BBC Earth, Guardian podcast). This work has consistently expanded understanding of evolutionary adaptation. Watson has over 100 publications (h-index=35, lifetime cites >5,000).

REFERENCES

Boyd R & Richerson PJ 1985. Culture and the evolutionary process. Chicago University Press.

Campbell, DT 1960. Blind variation and selective retentions in creative thought as in other knowledge processes. Psych. Rev. 67: 380.

Coolen, I... & Laland, KN 2003. Species difference in adaptive use of public information in sticklebacks. *Proceedings of the Royal Society London B* 270: 2413–2419.

Curio E 1988. Cultural transmission of enemy recognition by birds. In Zentall TR & Galef BG jr (eds) Social Learning. Psychological and Biological Perspectives. Hillsdale, NJ: Erlbaum.

Deneubourg JL et al 1983. Probabilistic behavior in ants: a strategy of errors? J. Theor. Biol. 105: 259-71.

Du Pasquier L & Litman GW (eds.) 2011. Origin and Evolution of the Vertebrate Immune System. Springer.

Edelman, GM 1987. Neural Darwinism. The theory of neuronal group selection. Basic Books.

Feldman MW & Cavalli-Sforza LL 1976. Cultural and biological evolutionary processes, selection for a trait under complex transmission. Theor. Pop. Biol. 9: 238-59.

Gerhart J & Kirschner M 1997. Cells, Embryos, and Evolution. Blackwell.

Griffiths AJF et al 2000. An introduction to genetic analysis. 7th Edition. New York, Freeman.

Hinde RA & Stevenson-Hinde J 1973. Constraints on learning. Academic Press.

Hinton GE & Nowlan SJ 1987 How learning can guide evolution. *Complex Syst.* 1: 495–502.

Kendal, J.R., Giraldeau, L.A. & Laland, K.N. 2009. The evolution of social learning rules: payoff-biased and frequency-dependent biased transmission. *Journal of Theoretical Biology* 260(2): 210-219.

Kendal R … Laland KN et al 2018. Social learning strategies: Bridge building between fields. Trends Cog. Sci. 22(7): 651-665.

Kirschner MW & Gerhart JC 2005. The Plausibility of Life. Resolving Darwin's Dilemma. New Haven: Yale UP.

Klenerman P 2017. The Immune System. A Very Short Introduction. Oxford UP.

Kounios, L., Clune, J., Kouvaris, K., Wagner, G. P., Pavlicev, M., Weinreich, D. M., & Watson, R. A. (2016). Resolving the paradox of evolvability with learning theory: How evolution learns to improve evolvability on rugged fitness landscapes. *arXiv preprint arXiv:1612.05955*.

Kouvaris, K… & Watson, RA (2017). How evolution learns to generalise: Using the principles of learning theory to understand the evolution of developmental organisation. *PLoS Computational Biology*, *13*(4), e1005358.

Kouvaris, K. (2018). *How evolution learns to evolve: principles of induction in the evolution of adaptive potential* (Doctoral dissertation, University of Southampton).

Laland KN et al. 2014. Does evolutionary theory need a rethink? *Nature*. 514: 161-4.

Laland KN et al. 2015 The extended evolutionary synthesis: its structure, assumptions and predictions. *Proc R Soc B* 282: 20151019.

Laland KN et al 2019. Understanding niche construction as an evolutionary process. In Uller T & Laland KN (eds) Evolutionary Causation. Biological and Philosophical Reflections. MIT Press.

Lande R 2009. Adaptation to an extraordinary environment by evolution of phenotypic plasticity and genetic assimilation. J. Ev. Biol. 22: 1435-46.

Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, *4*(3), 213-234.

Mineka S & Cook M. 1988. Social learning and the acquisition of snake fear in monkeys. In Zentall TR & Galef BG jr (eds) Social Learning. Psychological and Biological Perspectives. Hillsdale, NJ: Erlbaum.

Moczek A et al. 2011. The role of developmental plasticity in evolutionary innovation. Proc. R. Soc. B 278: 2705-13.

Olsson A & Phelps E 2007. Nat. Neurosci. 10: 1095-1102.

Oppenheim RW 1989. The neurotrophic theory and naturally occurring motoneuron death. Trends Neurosci. 12(7): 252-5.

Parter M, Kashtan N & Alon U 2008. Facilitated variation. How evolution learns from past environments to generalize to new environments. PLoS Comp Biol 4(11): e1000206

Pulliam H & Dunford C. 1980. Programmed to Learn. Columbia UP.

Richerson PJ 2019. An integrated Bayesian theory of phenotypic flexibility. Behavioural Processes. 161: 54-64.

Schlicting CD & Pigliucci M 1998. Phenotypic Evolution. A Reaction Norm Perspective. Sinauer: Sunderland, Mass.

Skinner BF 1938. The behavior of organisms. Appleton Century Crofts.

Soen, Y., Knafo, M., & Elgart, M. (2015). A principle of organization which facilitates broad Lamarckian-like adaptations by improvisation. *Biology direct*, *10*(1), 68

Staddon, JER 2016. Adaptive Behavior and Learning. 2nd Edition. Cambridge University Press.

Stephenson G 1967. Cultural acquisition of a specific learned response among rhesus monkeys. In Starck D et al (eds) Progress in Primatology. Stuttgart: Gustav Fisher Verlag.

Strickland et al. 1995. Complex trails and simple algorithms in ant foraging. Proc. R. Soc. B 260(1357): 53-8.

Sultan, S. E. & Spencer, H. G. 2002 Metapopulation structure favors plasticity over local adaptation. *Am. Nat.* 160: 271– 283.

Toyokawa W, Whalen A & Laland KN. 2019. Social learning strategies regulate the wisdom and madness of interactive crowds. *Nature Human Behaviour* 1, https://www.nature.com/articles/s41562-018-0518-x.

Uller T et al. 2019. Developmental plasticity and evolutionary explanations. Evol. & Devel. doi.org/10.1111/ede.12314.

Via S. & Lande, R. 1985 Genotype–environment interaction and the evolution of phenotypic plasticity. *Evolution* 39: 505–522.

Walsh D 2015. Organism, Agency & Evolution. Cambridge: Cambridge UP.

Watson R A & Szathmáry E. 2016. How can evolution learn? *Trends in Ecology & Evolution*, *31*(2), 147-157.

Watson RA & Wagner GP Pavlicev M, Weinreich DM, & Mills R. 2014. The evolution of phenotypic correlations and "developmental memory". *Evolution*, *68*(4), 1124-1138.

Watson RA et al 2016 Evolutionary connectionism: algorithmic principles underlying the evolution of biological organization in evo-devo, evo-eco and evolutionary transitions. *Evol. Biol.* 43: 553-81.

Webster, MM & Laland KN 2011. Reproductive state affects reliance on public information in sticklebacks. *Proceedings of the Royal Society Series B* 278: 619-27. doi:10.1098/rspb.2010.1562.

Webster M; Chouinard-Thuly L, Herczeg G, … & Laland KN 2019. A Four-Questions Perspective on Public Information Use in Sticklebacks (Gasterosteidae). *Royal Society Open Science.* 6: 181735.

West-Eberhard MJ 2003. Developmental Plasticity and Evolution. Oxford.

Wray G et al 2014. Does evolutionary theory need a rethink. No, all is well. Nature 514: 161-4.