Computational Explorations of the Interaction of Evolution and Learning

New Models in Evolutionary Psychology

Robert A.P. Reuter
rreuter@ulb.ac.be
Séminaire de Recherche en Sciences Cognitives
Faculté des Sciences Psychologiques et de l'Education
Université libre de Bruxelles

Frequently, human psychology opposes evolution and learning on the basis of their supposed differences in flexibility. Humans are thought to be more flexibly intelligent than other animals, because they lost most of their instincts. But, as William James proposed, we think that it might turn out that just the contrary is true. Human may have more instincts, more powerful instincts to learn for instance. Indeed, the view that instincts are rigid, correspond to fixed, innate, pattern of behaviour that do not enable the organisms to adapt to its ontologically relevant environment, seems to be an ancient believe of many scientist in the social sciences. This is however, nowadays, not the view of many researchers in the field of behavioural biology and ethology. Instincts do take into account information from the environment, that is basically what they are designed for, they provide an organism with behavioural responses to certain environmental stimuli. They are surely not about acting as if no environment existed.

The Powers of Learning

On the other hand, learning is thought to be very powerful, especially when it comes to react to unpredictable (phylogenetically speaking) characteristics of the environment. We must acknowledge that learning has the power to adjust the behavioural responses of an organisms to the contingencies in its particular environment, that no reliable evolved action pattern could provide. After all, how could an new-born know (in the strong sense of having the exact mental representation) its mother. It seems impossible to have the "genetic blueprint" do something as magical as this! Of course, the new-born has to learn to recognise its mother, but this is not the point. What is important is that the new-born "knows" that there is something like a mother, that she is likely to be the individual that most "cares" about your needs, etc. (we could say that it acts as-if it had such an explicit knowledge, although the very nature of this knowledge is not the relevant question here). Indeed a certain preparation to learn seems to be essential in order to focus the motivational, attentional resources of the new-born, at least, to certain aspects of the environment, physical or social, that are "expected" to be there. The "emotional value" of such environmental characteristics might provide some information about the stimuli to attend to, to look after actively. This seems to be true with human faces and speech in new-borns. They "actively" orient their attention towards stimuli that look like faces, they "like" to be talked to. All these, and other, behaviours show that what counts as environmental stimuli relevant for the learning to take place, are not neutral. They are indeed part of the biological features of the species. What we consider as pleasant is far from being the same as what a fly finds "pleasant" (no conscious feelings need to be involved here, pleasantness may just be a anthropomorphic description of a behavioural preference). Indeed the same objects may attract a fly, although you may not even notice it. Combinatorial explosion is a problem that every individual of a species faces when it is engaged in a process of learning. What are the stimuli that should enter the information processing and which should not? This problem is especially salient when it comes to acting in response to external events, that are potentially harmful or even lethal! You should be able to decide rather quickly whether you run away or not when facing a lion, for instance. There are not many ways in which you could learn the correct response by mere reinforcement learning. You would be dead if the first answer appears to be the wrong one! Of course, you could learn it from your mother, or other members of your group. But you would also have to "find" the link between stimuli and response. One way to conceive a solution to this problem is given by the paradigms of neural network computational explorations. Indeed, when it comes to finding the needle-in-the-hay-stack, neural networks are able to learn what the relevant features (events and co-occurring of events) of an environment are, provided that they be given information about some aspects of the environment. It seems needless to say that no experimenter could claim that the input she gives her network are exhaustive of the environmental characteristics present at a certain moment. Filtering is an essential first-step operation in computation.

Neural Networks and Innateness

But neural networks can, provided they implement a powerful learning algorithm, learn reliably and rather fast what are the relevant hidden features (e.g., correlations) of an environment (or a certain domain of it). All the knowledge resides in the connections of the network, which adapt their weight in order to produce a desirable output. But what counts as a desirable response to some input, is often in the eye of the beholder, and not in the networks "knowledge base". Neural networks represent in this sense powerful models of learning and cognitive capacities. What is more, although they are often used in the domain of learning, they have the potential to provide us with a "natural" way to think about evolutionary processes involved in the design of the brain. Indeed, the weights of the connections between neurones could be initialised (given their initial values) by some mechanisms under genetic control. Connectionist models are generally mute about this possibility, but nothing prevents us from thinking about evolved design a neural network in terms of "fixing" the initial values of its connections (which may still be subject to later modification by learning mechanisms). When it comes to difficult problem solving, linked to the combinatorial explosion of potentially "good" patterns of connectivity, it seems that any network that is closer to the solution is better off, or at least it should reach an optimal solution faster. And speed is what counts in ontogenetic contexts when viewed from the evolutionary perspective. Any genotype, coding for the initial connectivity of a network, that is co-responsible for a good starting point should increase its frequency in the gene pool, compared to those that are less reliable or less close to the solution. Learning, in the sense of changing the connectivity pattern towards an optimal solution, enables networks that do not represent the optimal pattern right from the start to "find" the solution nevertheless. Learning thus makes evolution of a highly improbable connectivity pattern possible, where a fixed connectivity pattern (purely determined by genetic coding) would easily fail to represent the solution. One wrong connection would make it fail, since no changing can occur. Learning softens the fitness landscape as we approach the optimal solution. Next-to-good solutions can change into good solutions, which can converge on the optimal solution. This is however a reasonably fast process when the initial pattern is not to far away from the solution to be found. There is a trade-off between efficiency and flexibility. If the network is too flexible, having no correctly fixed connections at all, it fails to find the solution fast enough to be competitive with the other networks that have a better start-off. On the other hand, being inflexible on some connections closes many pathways towards the solution.

Enter the Baldwin Effect

This is where the so-called Baldwin Effect enters the play. It describes a process of interaction between learning and evolution. Evolution often could just not find the solution to a problem (one in a million of a million of possible states of connectivity), because it would be too improbable. What is more, any organism that approaches the solution but does not reach it (assuming no ontogenetic learning takes place), would have NO advantage over other organisms that approach it less. What good is, after all half-a-wing? Thus learning can make it possible for individuals to have differential fitness that is related to the distance between their initial state and the optimal state. Even if you fail to have the solution right from the start, you can find it faster than someone else whose initial state is further away from it. Hence your fitness is higher than that of your competitor.

From Learning to Inheritated Traits

On the other hand, the Baldwin Effect tells us that, in some conditions, learning (ontogenetic process) tends to be replaced by genetic determination (phylogenetic scale). Consider the following reasoning. If a genome provides an organism with a start-off closer to the solution (by having its neural connections better tuned in the beginning) than another organism, then it tends to produce differentially more offspring. Thus, over evolutionary time, it should tend to increase its frequency in the gene pool. Learning a certain behaviour in response to the environment tends to be replaced by genetically coded behaviour responses over evolutionary relevant intervals, given a more or less stable environment (or certain aspects of it).

Modelling Evolutionary Processes

How could such considerations be tested experimentally? We must remember that evolutionary time intervals are quite large and can hardly be manipulated, if they can at all, in the lab! Artificial Life (ALIFE) simulations have the potential to help us understand the interactions between evolution and learning, since they are designed to represent the mechanisms of life, like evolution by natural selection. It is possible to make virtual creatures evolve certain

features, and to learn certain behaviours by means of certain learning algorithms that yield ontological modifications. Neural Networks seem to be especially powerful tools to study the interaction between evolution and learning, since the two complementary processes can involve exactly the same elements of the virtual organism, namely the weights of the neural connections. Whether they are initially determined by the artificial genome or result from changes due to the exposure to the external (and eventually internal) stimuli makes no qualitative differences. None of the two causes of a connection's weight can be detected by simply looking at its value. The mere value does not reveal anything about its origin nor about the history of the weight changes - whether they take place during a life-time or during evolution. So applying genetic algorithms to the initial connectivity pattern of artificial neural networks (ANN) provides us with a "natural" - although still simplified - way to conceive the Baldwin Effect. It also enables us to see how controlled computational models can be designed in order to study such a phenomenon that otherwise cannot be examined experimentally in nature.
In the following, we present two quite simple, but rather ingenious simulations of the Baldwin Effect using ANN and Genetic Algorithms (GA), namely Hinton & Nowlan (1987) and Ackley & Littman (1992) (quoted in Mitchell & Forrest, 1997).

Computational Exploration - Part I

Hinton et al.'s model was designed in order to be very simple. They, indeed, instead of using a complex hebbian or error-driven learning algorithm, used random guessing. On each learning trial, a network guesses a 1 or a 0 at random for each of its learnable connections. In the GA population, each network is represented by a string of 20 characters which can be 0, 1, or ? - the ? corresponding to a learnable connection, 0 to absent, and 1 to present. Note that 1 and 0, for a particular connection can correspond to a correct connection, as well as an incorrect one. There is only one pattern of connectivity that was, arbitrarily, postulated as the one in 220 possible patterns corresponding to the solution to be found by the ANN. This implies a "needle-in-the-hay-stack" search problem because there is only one correct state in a immensely large space of possibilities. Learning, as seen before, changes the fitness landscape, from a single spike to a smoother "zone of increased fitness" (Michtell & Forrest, 1995). Fitness can indeed be computed quite easily, corresponding to an inverse function of the number of trials needed to find the correct solution. Having all the connections set correctly means maximum fitness, never finding the correct setting means lowest fitness. There is thus a trade-off between efficiency and flexibility, as seen above.

Hinton et al.'s simulation showed that learning enables individual networks to receive partial credit instead of engaging in an all-or-none situation. They concluded from their results that learning helps ANN adapt to predictable but difficult aspects of an environment, in the sense that only one setting of connections represents the only solution to some problem. Learning systematically increases population fitness by giving partial credits to solutions close to the optimum peek. Furthermore, learning leads to adaptations (optimal or close to optimal solutions) becoming genetically fixed. Indeed, reproduction rate of initial patterns of connectivity (determined by an AG) varied according to the explicit fitness function, mentioned above. Hinton et al. found that the initial patterns tend to become more and more fixed genetically as evolution took place.

Computational Exploration - Part II

Let us now turn to another computational exploration by Ackley & Littman where fitness is no longer such an explicitly and a priori determined function of the pattern of connectivity. Fitness of real organisms tend to be determined by various actions engaged in and interactions with other organisms. Fitness can thus better be regarded as some emergent function of behaviour, co-determined by genetic design and ontogenetic learning processes - not to forget chance: being at the right place at the right time. Ackley et al. implemented so-called adaptive agents in a virtual environment with different objects placed in it: trees to hide in, food to eat, and predators to avoid of course. The agents consist of two neural networks, one intended to evaluate the actual state of the agent, in terms of its overall energy and the actions that lead to this state, the other one intended to control the actions of the agent. Briefly said, each agent comes with innate goals about what is good to do and not to do, and is able to change its actions by reinforcement learning based on the mismatch between desired state and actual state. This is achieved by fixing the connection pattern of the evaluative network based on the specifications of a sting generated by an AG, and by initially specifying the connections of the action network which are subject to modification by reinforcement learning. Reinforcement learning can be seen as some kind of intermediate between supervised and narrow-sense unsupervised learning: there is no precise information about the amount of difference between the actual state and the desired one, there is only information about the success or the failure of a certain action pattern. In the virtual environment designed by Ackley et al. the struggle for survival and reproduction encompasses different factors. Agents have to act in order to prevent death, by eating and avoiding predators, by reproducing - by means of simple cloning or by cross-over with a nearby agent, reproduction depending on the amount of energy they call their own. Ackley et al. found that learning lead to considerable improvement of the current state of an agent. Furthermore, by letting compete different types of agents, different in terms of ability to learn and intervention of evolution on their initial connections in the action network, they found differential success in the following order: Agent with learning and evolution surpassed agents with learning only, surpassing themselves agents merely designed by evolution. What is more, they observed an evolution from "goal instincts" towards "behaviour instincts". This point is worth some discussion. Indeed, whereas in the beginning agents tend to present some initial variability in the connectivity pattern of their action networks, after "some time" they tend to show less variability in their action networks but greater variability in their evaluation networks. In the beginning of evolution it seems to be important to know what the desirable outcomes are and to change one behaviour accordingly, and after some evolutionary changes can occur, genetic specifications of adaptive actions become sufficient, maybe even preferable. Indeed behaviour improving your fitness is sufficient for differential reproductive success. This however can reliably take place only in agents that are descents of agents who "evolved" these action network specifications "helped" (or "guided") by the specifications provided by their evaluation network.

Conclusions

We would like to end this essay by concluding that learning and evolution are by no means opposed processes. They are better seen as complementary processes. More evolved design enables not less flexible learning, but more (Pinker 1997), especially in survival conditions where time is an essential variable of success! Although a tabula rasa organism could learn anything quite reliably - for instance if it has a very strong learning mechanisms or if it has a lot of time - but knowing a lot right from the start - this knowledge may just simply be encoded in the initial connection weights and represent a preferential hypothesis about the environment, and nothing more - enables an organisms to be better off. Finding a solution a bite quicker than all the others gives you some better chance to increase the frequency of your genes, even if your genes are only partially and not totally responsible for the solution you found!

We have also seen that ALIFE modelling has the powerful potential to enable us to achieve controlled computational experiments of such processes like evolution, competition, learning in non-static environments, and so on. ALIFE models could play the same role in biology and evolutionary psychology as ANN did in classical experimental psychology - namely to re-energise theoretical and empirical work by means of providing experimental conditions to refute and verify certain model of evolutionary processes involved, for instance, in the emergence of behaviour.
Not to forget that GA could be applied to ANN in order to solve more practical engineering problems that are hard to resolve by classical artificial intelligence methods, because they involve a combinatorial explosion of the potentially successful (or unsuccessful!) pathways in a vast solution space. Contributions to the field of cognitive sciences could be made by providing a better way to "resolve" the old problem of "innate versus learned", in terms of complementary and not opposing processes. After all, the characteristics of the environment our hominid ancestors lived in clearly favoured the evolution of some of our specific flexible learning mechanisms that no other species has! And this environment of adaptive fitness can impossibly not have had some impact, still measurable today, on the design of our modern mind. Ninety-nine percent of their history, our hominid ancestors lived as hunter-gatherers, in ecological conditions clearly different from those today. It seems reasonable to assume, that this evolutionary history can be detected in the functional - if not structural - design of our mind-brain. After all, evolution is quite a slow process.

References:

- Mitchell, M., & Forrest, S. (1997). Genetic Algorithms and Artificial Life IN Langton, C., G. (editor), *Artificial Life - an Overview*, A Bradford Book, MIT Press.
- Pinker, S. (1997). *How the Mind Works*, NY: W.W. Norton & Company.

UNPUBLISHED: For your information only. Contact the author: rreuter@ulb.ac.be

---