Natural and artificial meta-configured altricial information-processing systems

JACKIE CHAPPELL[1]*, AARON SLOMAN[2]

[1] *School of Biosciences, University of Birmingham, B15 2TT, UK*
[2] *School of Computer Science, University of Birmingham, B15 2TT, UK*

Abstract:

The full variety of powerful information-processing mechanisms 'discovered' by evolution has not yet been re-discovered by scientists and engineers. By attending closely to the diversity of biological phenomena, we may gain new insights into (a) how evolution happens, (b) what sorts of mechanisms, forms of representation, types of learning and development and types of architectures have evolved, (c) how to explain ill-understood aspects of human and animal intelligence, and (d) new useful mechanisms for artificial systems. We analyse tradeoffs common to both biological evolution and engineering design, and propose a kind of architecture that grows itself, using, among other things, genetically determined meta-competences that deploy powerful symbolic mechanisms to achieve various kinds of discontinuous learning, often through play and exploration, including development of an 'exosomatic' ontology, referring to things in the environment — in contrast with learning systems that discover only sensorimotor contingencies or adaptive mechanisms that make only minor modifications within a fixed architecture.

Key words: behavioural epigenetics, biologically inspired robot architectures, development of behaviour, exosomatic ontology, evolution of behaviour, nature/nurture tradeoffs, precocial-altricial spectrum, preconfigured/meta-configured competences sensorimotor contingencies.

---

* email: j.m.chappell@bham.ac.uk

CONTENTS

1 THE DIVERSITY OF BIOLOGICAL INFORMATION PROCESSING

It is widely acknowledged that we have much to learn by emulating biological information processing systems and this has inspired theories and models in Artificial Intelligence and Cognitive science. Most biologically inspired research on 'unconventional' computation is concerned with the kinds of low level *general purpose* computational machinery found in biology, including chemical computations, neural computations, and evolutionary computation, along with use of stigmergy, flocking and swarming mechanisms to coordinate behaviours of 'dumb' individuals. There is an excellent overview by Susan Stepney in [45]. We are attempting to identify important forms of biologically-inspired computation not included there. Our main aim, expanding on [38], is to draw attention to important forms of development and learning that occur either during evolution or in individual development, which appear not to have been taken into account by scientists and engineers interested in biologically inspired forms of computation, though several of the relevant biological facts were noted and discussed long ago by Cushing [8], Dennett [12], and others.

Understanding the origins of the naturally-occurring variety of forms of information processing, and the tradeoffs between the alternatives, may be important for various kinds of engineering applications, including robot designs. Moreover, it may turn out that further analysis of the detailed features of these designs will point to requirements for neural, chemical and evolutionary computations that cannot be met by existing models, thereby drawing attention to previously unnoticed biological mechanisms. At this stage that is merely a possibility, for which we have no strong evidence apart from the fact that artificial systems based on current biologically inspired mechanisms do not come close to matching the competences of many familiar animals, including not only humans, but many mammals, birds, and even insects. Examples are given below. However there are important differences between the types of competences of different sorts of animals, related to differences in the ways those competences arise. We shall show how to explain some of that diversity in terms of different interactions between evolution and development, which we now describe briefly.

1.1 Sources of behaviour in organisms

Some behaviours are genetic in origin and some are learnt. As proposed in [38], we expand that simple dichotomy, by identifying different sorts of relations between evolution and learning. On that basis we can identify some important tradeoffs in biological evolution, which are also likely to be relevant to artificial systems.

Figure 1 presents diagrammatically four main ways in which the genome can produce behaviour in an organism. In the first two cases the behaviour is *preconfigured* insofar as the genetic mechanisms, possibly interacting with the environment during development, create structures that produce the behaviour, for example, most of the behaviours of invertebrates and some of the behaviours even of humans, such as sucking and breathing

Multiple routes from genome to behaviours
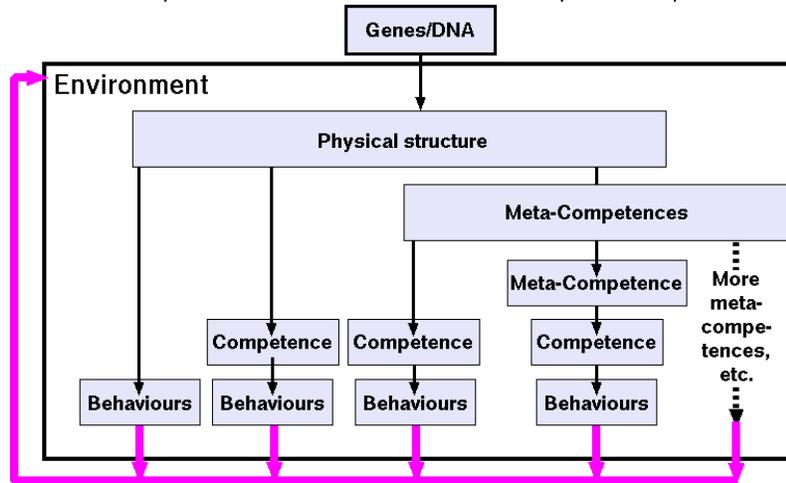### (Environment affects all embedded processes)

FIGURE 1

Schematic illustration of alternative routes from genome to behaviour. From left to right: (a) physical structures produced by the genome generate behaviour (e.g. innate reflexes); (b) the genome generates competences *which produce behaviours influenced by goals/needs and the environment, including most insect behaviours, behaviours of precocial species; (c) the genome produces meta-competences (learning mechanisms), which generate competences partly under the influence of the environment, and those competences produce behaviours partly under the influence of the environment; (d), like (c), except that two layers of staggered, or cascaded, meta-competence production precede the production of competences which eventually produce behaviours (learning to learn). The fifth column, (e), indicates that in principle arbitrary layers of meta-competences can be built, which then produce competences, producing behaviours. (a) and (b) correspond to* precon-figured *behaviours and competences. (c) to (e) correspond to* meta-configured *competences. These schematic specifications have different instantiations in different parts of a multi-functional architecture, e.g. in reactive and deliberative components, as defined in [33] and [41]. The environment, including parents and other conspecifics, prey, predators, competitors, etc. can influence all of the processes.*

in infants. Behaviours produced before development is complete can influence how the genes produce further structures that produce later behaviours, for example, calibration of control mechanisms while they develop. Some of the genetically determined structures produce *competences* that are to some extent general, insofar as they produce behaviours partly under the control of the environment, e.g. obstacle avoidance or attacking behaviours.

In the third and fourth cases the genes produce *meta-competences*, and these interact with the environment (including possibly parents that provide succour, shelter, protection from predators, and training), to produce behaviours that we label *meta-configured*. They are not entirely genetically determined, but realise genetically determined potential

to produce different sorts of competences, under the influence of the environment. Moreover, if the development of meta-competences is staggered, early behaviours produced by meta-configured competences can influence the development of later meta-competences, as shown in Figure 1. This can produce even greater diversity in the competences and behaviours produced by the same genome. This mechanism allows the environment to have a major influence on the information processing architectures in adults, subject of course to the constraints of the original genetically determined meta-competences and competences.

In this framework, we can contrast training processes that gradually shape existing competences, with processes that create new competences or new meta-competences, while possibly leaving old competences available. (This is related to Piaget's distinction between *accommodation* and *assimilation* which we have no space to discuss.) Further subdivisions remain to be investigated.

Our hypothesis can be contrasted with theories proposing *one* general-purpose mechanism (e.g. a reinforcement learning mechanism) responsible for all learning during interactions with the environment. In contrast we suggest that there can also be different genetically determined learning mechanisms tailored to different aspects of the environment and suited to acquiring different competences, using implicit innate high level knowledge about what sorts of things can be learnt in the environment and how they can be learnt. A specialised meta-competence in humans for learning language may turn out be an example, but for now we leave open the possibility [32] that language-learning uses more general meta-competences.

These distinctions will now be illustrated with biological examples and related to the distinction between *precocial* and *altricial* species.

1.2 Biological examples: precocial and altricial species

The motivation for the above distinctions comes from the diversity found in biological examples, including these two extremes: (a) some very sophisticated behavioural competences seem to be largely genetically determined, since they are available at or shortly after birth (e.g. new-born deer running with the herd and cliff-reared birds flying at the first attempt); (b) in contrast, other competences that are superficially similar appear only much later, after rich interaction with the environment (e.g. hunting mammals pursuing prey, and skilled hunting in hawks and other raptors). In some instances of the second type there are significant differences in adult behaviours that result from differences in the environments, notably in humans but also in some other species.

Type (a) competences fit our definition of 'preconfigured'. Type (b) competences are 'meta-configured' insofar as they result from the operation of various sorts of genetically determined meta-competences tailored to the environment and the agent's body. An example is learning of new competences through creative play and exploration possibly combined with hypothesis generation, testing and debugging. In some cases the learner acquires new meta-competences, e.g. learning to acquire new competences by asking for

help or advice, or by practising.

These two forms of learning and development map very broadly at the species level onto the developmental categories 'precocial' (species with highly competent, independent young), and 'altricial' (species with relatively helpless and incompetent young). However, most species have a mixture of pre- and meta-configured competences. Paradoxically, even though precocial species are much more competent initially, some altricial species achieve far greater behavioural sophistication in adult life. We suggest in section 3.3 that this is explained by the use of staggered development of meta-competences.

Preconfigured competences account for many behavioural competences found in precocial species at or soon after birth or hatching as well as competences that develop later but are shared between all members of the species irrespective of their environments. Meta-configured competences are characteristic of altricial species that start off helpless and underdeveloped but acquire cognitively sophisticated competences related to specific features of their environment, such as hunting, nest-building, problem-solving or linguistic competences. In some animals, notably humans, that process can produce competences involving types of objects never encountered by their ancestors (e.g. computer games). However, different species can have different combinations of both types of competence, where the precise mixture will be a result of tradeoffs found in their evolutionary history.

1.3 The remainder of this paper

These ideas will be illustrated in connection with four main themes: (1) the difficulty of identifying the varieties of types of learning and competence that need to be explained, (2) learning processes that use rapid, discontinuous forms of learning based on successive, increasingly complex, information structures, rather than gradual adjustment alone, (3) development that includes the growth of an information-processing architecture, instead of merely modifying existing contents of a fixed architecture, and (4) ontology development that refers to objective ('exosomatic') features of the environment that can exist independently of the organism, unlike ontologies restricted to patterns and contingencies discovered in sensorimotor signals ('somatic ontologies' referring only to what occurs within the animal's body). It should be clear that we are not referring to the ontology used by scientists and engineers studying or modelling these animals or machines. Rather, we use the label 'ontology' here to refer to what the animal or machine is capable of representing, distinguishing, reacting to, or reasoning about.

We hope to show that deep, evolutionarily old and largely unexplored mechanisms lie behind many characteristically human capabilities such as language, and also some abilities that we share with other intelligent species including other animals that create and manipulate spatial objects, such as nests. If we can develop AI systems with similar developmental mechanisms, we may be in a much stronger position both to explain and model aspects of human learning and development and also to achieve artificial systems

that approach 'human-level' competences.

2 THE PROBLEM OF IDENTIFYING WHAT NEEDS TO BE EXPLAINED

Much attention within AI has focused on a research goal labelled as 'human-level AI' e.g. by McCarthy in [26]. The specifications and requirements for such a system are currently very poorly characterised, because the competences of humans and other animals are surprisingly difficult to identify: as explained in [40], researchers often suffer from 'ontological blindness'. E.g. people who think the function of vision is mainly recognition and tracking will design visual systems that cannot see what Gibson in [14] referred to as 'affordances'. There are differences between recognising household objects like cups and seeing their affordances for manipulative actions. For example understanding different possible ways of grasping an object, in different places, does not require it to be recognised as a cup, even if it is a cup, as argued in [43]. Moreover, machines that can recognise cups don't necessarily provide any information useful for grasping them.

Likewise, there are researchers who write as if all learning is about finding patterns and correlations at various levels of abstraction in multi-modal combinations of sensor and motor signals (all occurring within the body) i.e. discovering and using 'somatic' sensorimotor contingencies. For example [23] claims

> "As organisms interact with their environment, their sensory inputs are transformed into motor outputs and their motor outputs determine what is sensed next. The continuous and dynamic coupling between sensory, neural, and motor variables defines sensorimotor networks that describe the informational embedding of organisms within their ecological niches at multiple time scales. The comparison of the relative influence such variables exert on each other helps extract (functional and structural) patterns of interaction between the networks' elements that may support biological information processing."

We do not dispute that this is a part of what goes on in organisms. However, many dynamical systems theorists ignore the question of what *other* forms of representation and learning might be needed. Similar questions arise about the excellent survey [4], which provides deep analyses for dynamic interaction with the environment but never considers what might be required for *thinking* about the environment, e.g. wondering what will happen tomorrow, making a plan to get to a conference, trying to explain non-delivery of a parcel, or even solving an equation in one's head. The "sensorimotor" dynamical systems approach fails to recognise the need for some animals and robots to acquire and use an objective, 'exosomatic', amodal, ontology, e.g. referring to 3-D structures, processes and causal relationships in the environment, independently of whether they are being perceived or acted on. The need for such an ontology was noted as long ago as 1780 by Immanuel Kant in [20].

This point can be illustrated as follows. A child that never learns what is common to processes in which an object is grasped with its mouth, with its left hand, with its right hand, with both hands, or grasped using pliers or tweezers, or grasped by someone else, will not be able to transfer information learnt about one sort of grasping to another sort, and when planning future actions involving grasping will always have to specify the precise mode of grasping – an enormous hindrance to planning. Understanding what is common to the different sorts of grasping involves using an 'objective', exosomatic ontology that refers not to patterns of changing sensory and motor signals and their relationships, but to the motions of 3-D surfaces of various kinds and their changing relationships. Some of the differences between requirements for learning somatic and exosomatic ontologies are discussed in [36] and other documents cited there.

A different sort of ontological blindness afflicts researchers who design mechanisms providing narrowly focused functionality, e.g. object or speech recognition, while ignoring requirements for that functionality to exist in an architecture combining many different kinds of functionality. Their systems cannot cope with unexpected developments that require processes to be modulated, suspended, or aborted, or which require information from several concurrently active processes to be combined in order to solve a problem that none of them can solve alone, for instance in linguistically guided visual search.

Such failures to identify what needs to be explained or modelled limit the usefulness of the common *forward-chaining* approach to AI research, namely incrementally improving current designs, e.g. using fixed benchmark tests. We have elsewhere [34] suggested using an approach to research planning, in which detailed requirements are derived by working *backwards* from a collection of remote scenarios and constructing a partially ordered network of scenarios and associated requirements. Progressively less demanding partially ordered prerequisite scenarios can then be specified, until we reach a subset that can be attempted practically using current knowledge and technology. This method of deriving milestones and benchmark tests makes it more likely that systems produced will contribute to the long-term research objectives, than systems that merely perform better on benchmarks that are not based on a backward-chaining analysis.

Some scenarios should include learning and development. In particular we should identify preconfigured competences that are required because there is no opportunity to learn them before they are first needed, and meta-competences capable of generating a range of meta-configured competences because the environment can vary in unpredictable ways.

3 DIVERSITY IN EVOLUTION OF ANIMAL COMPETENCES

This section elaborates on the differences and tradeoffs between systems designed only with preconfigured competences and those that also include meta-configured

competences, as defined in section 1.1.

Preconfigured competences: If a competence is very important to the subsequent survival and reproduction of the individual (i.e. its evolutionary fitness), and it is required early in the animal's life, and if the evolutionary history has provided appropriate opportunities for natural selection to occur, then it could pay for the competence to be relatively 'hard-wired' (in the genome), i.e. pre-configured. This may be why the vast majority of animal species have almost all their behaviours preconfigured and few or none meta-configured, for example, typical insect behaviours. All animal species have some preconfigured competences; for example, all new-born mammals instinctively seek out a nipple and suck when they find it because an early feed of colostrum milk is so crucial for their survival and development [22]. Surprisingly complex and rich behaviours are possible using this 'preconfigured' mode of development. For example, caribou (*Rangifer tarandus*) calves are able to co-ordinate their motor activity sufficiently to stand up, run for long distances over very uneven ground, and evade predators only a few hours after birth [31]. Web construction in spiders is another well known example. (The existence of such highly competent precocial species with preconfigured capabilities indicates a flaw in the widely believed theory that 'symbol-grounding' based on individual learning is required for animals or machines to use information structures with semantic content.)

Preconfigured competences need not be fully defined by the genome, if it is safe to leave some of the specification to factors influencing the developmental process or to the environment after birth. For example, consider imprinting in domestic chicks. Newly hatched chicks will follow the first, biggish, moving thing they see, which in normal circumstances is their mother. Where the mother is absent, imprinting stays active for a bit longer, but eventually finishes with the chicks imprinting on each other, so that they at least stay in a group. So under normal circumstances, chicks almost invariably imprint on their mother, even though the genetic specification probably says nothing specific about the appearance of the hen. For more on this see [2].

Thus, some precocial competences are probably only quite vaguely defined genetically, but use mechanisms which allow a combination of strong developmental channelling along with 'sensible' default settings that almost always work, at least in a certain variety of environments.

Meta-configured competences: Evolution of a meta-configured competence can occur when selection pressure on very early expert rendering of some competence is reduced (e.g. because of parental caring behaviours), or where the kinds of mechanisms which support preconfigured development do not produce sufficiently complex or flexible competences, or where the target of the competence changes a great deal between generations, or where it would take very much longer to evolve a competence than to evolve a mechanism for acquiring it.

For example, many mammalian predators learn to hunt prey. Prey behaviour is highly erratic (and itself under selection pressure to allow escape from the predator) so young

predators must learn the detailed skills of stalking, chasing, capturing killing, and eating (e.g. ripping open, dismembering) prey. It is often assumed that such changes in a prey species will lead to genetic changes in the predator species over many generations, as in so-called 'evolutionary arms-races' [10]. However it seems that at least some predators have evolved the ability to adapt required behaviours within the lifetime of each individual, instead of depending on much slower evolutionary processes. Of course, even in those cases, some aspects of the learning mechanism are relatively constrained. For example, young predators instinctively visually track and chase fast-moving objects, as anyone who has a domestic cat will know. But the eventual competence is the outcome of several stages of skill development during the animal's interaction with the environment. Young predators are usually initially very poor at hunting. Female cheetahs bring maimed but live prey back for their offspring to practice their capture skills [18], in some hawk species the parents train the young to hunt [8]; and of course human children learn for several years using toys, make-believe, games, etc., before they are ready to take on adult responsibilities.

Animals that need to re-identify locations (e.g. nest sites) may use preconfigured competences based on pheromone trails. A flying animal, cannot do this, for instance the solitary wasps that Tinbergen showed use land-marks to locate the nests in which they have laid eggs, [47]. It would be impossible for evolution to predict the precise environment in which each wasp will lay its eggs, so learning landmarks must be a meta-configured competence.

Although some of the skills developed after birth may appear to result from gradual re-shaping and speeding up of physical competences already present from an early age, some meta-configured competences are developed through discontinuous learning, such as learning a new grammatical form, a new semantic category or a new problem-solving technique. Insects learning landmarks are another example of discontinuous learning, as are marsh tits that remember where they have buried nuts, and which nuts have not yet been eaten.

Both scientists trying to model and explain what occurs in nature and engineers designing artificial intelligent systems are faced with extremely difficult problems, whether considering sophisticated preconfigured behavioural competences or complex meta-configured competences that need to be tailored to problems and opportunities in the environment that are not specified in advance. This does not imply that preconfigured meta-competences are totally general learning mechanisms: they are often tailored for learning specific things, e.g. learning landmarks, or imprinting. This ability to learn landmarks is specific to a type of 3-D environment. Many song-bird species have a meta-competence that enables them to learn songs, but as Marler points out in [24], that is also not a totally general acoustic recording and playback mechanisms, since different species learn different sorts of songs when played the same tape recordings: each prefers conspecific songs. This implies that even for a single species many different meta-competences may be required, for learning different sorts of things. As indicted in

Figure 1, some of them may have to be learnt from meta-competences developed earlier. Precisely what sorts of species specific preconfigured learning mechanisms there are is a topic requiring much further research, which may, in turn, guide construction explanatory models of such learning.

3.1 Explaining preconfigured competences

The facts we have listed about complex preconfigured (precocial) behavioural competences that are determined innately leave much unexplained. The fact that some competence is useful does not explain (a) how it evolved, (b) how it is represented in the genome, (c) how the information in the genome gets transferred into animal brains, or (d) how the brain mechanisms produce the behaviour, which may involve coordinating several sensors and effectors for an extended period of time.

As shown in Figure 1 the links between DNA and behaviours are more remote than the links with physical structure, and they are increasingly remote as more intervening meta-competences are involved. Consider only the first two cases, the preconfigured behaviours and competences. What chance does an evolutionary change-generator based on manipulating chemical structures have of producing *useful* high level genomic changes related to *complex behaviours* in a reasonable time by using only evolutionary hill-climbing search mechanisms? Of course, if successive changes in the genome do produce advantageous variations in behaviour then normal Darwinian selection can make use of the changes. But under what conditions could random molecular changes produce useful changes in behaviour?

A partial answer may be that evolution 'discovered' a way to separate modules in the genome that determine a general type of competence from parts that function as 'parameters' specifying which specific type is produced. In that case small random changes in the parameter section could produce slight modifications in the behaviour while retaining the overall function of the behaviour, giving natural selection the opportunity to search by hill-climbing, even though random changes in the 'function' part are likely to be catastrophic. This is a program design strategy whose power was also discovered fairly early in the history of computer programming.

The evolution of behaviour modifiers (analogous to 'control knobs') that are separable from the structures whose behaviours they modify is illustrated by the fact that hormonal changes can have large effects on behaviour and the fact that some parasites can change mammalian behaviour in subtle ways for their own purposes: for example, Berdoy and colleagues [3] showed that toxoplasma gondii can reduce cat-avoiding behaviour in infected rats, since the parasite needs to infect a felid in order to reproduce. In these cases it seems that the separation of function was achieved by evolution of neural structures whose behaviours are sensitive to chemical influences.

If there are also parts of the genome specifying how larger behaviour-generating modules can be combined, then that may be structured in such a way that a small chemical change in certain regions of the genome have a good chance of producing a new

11

meaningful combination using different parts. This may depend on the genome encoding a 'syntax' for complex structures and organised so that many of the random chemical changes do not disrupt well-formedness. This could enormously reduce the search for combinations that enhance viability.

In artificial evolution, this sort of mechanism has already been explored using 'genetic programming' (GP) techniques, which assume a more structured genome than standard genetic algorithms (GAs). More information about the similarities and differences between GP and GA techniques can be obtained via this web site [1]. This seems to be an important topic for research in epigenetics.

3.2 Explaining meta-configured (altricial) competences

In sections 1.1, 1.2 and 3, we noted in connection with meta-configured competences that some species have innately determined generic learning capabilities which, through various kinds of interaction with the environment, produce dramatic kinds of learning and development in individual animals such as we see in human infants, and infants of some other species, learning to use and play with toys and devices that did not exist in their evolutionary history. Similar dramatic discontinuities and flexibility are illustrated by language learning and absorption of a culture: any child is capable of learning any one (or in some cases several) of the thousands of very different languages used around the globe, and also picking up during a life-time a significant subset of the concepts, knowledge, skills and artistic forms that have taken centuries for societies to develop. Some learners need only perceive a mathematical proof, or a demonstration of a mechanical or gymnastic technique, or an interesting melody *once* in order to understand it and be able to apply it. However the same learner would not have been able to make that leap at an earlier stage of development: so we need to understand what mechanisms change so as to make new forms of one-shot learning possible.

The mechanisms that make that possible are not yet understood. People who assume that imitation can explain learning do not always notice that what an individual can or cannot learn by imitation usually depends on very specific cognitive structures and perceptual capabilities that must already have components that are required for that new competence. You can only learn by imitation what you could have learnt without it, though possibly only after a long search for a solution to a problem – and what is learnt by imitation in one generation may originally have been discovered only after a long search in an earlier generation, by an individual with the same cognitive mechanisms as the one who learns by imitation. In some cases the young learner ends up contributing further major developments to that cultural process, by substantially modifying what was learnt by imitation.

Meta-configured competences pose all the unsolved problems (a) to (d) mentioned in section 3.1, and in addition raise the following new questions about how the mechanisms that provide individual development work:

(e) what are the genetically determined architectures, forms of representation,

algorithms and other information processing mechanisms in a newborn individual,

(f) how do those interact with the environment so as both to learn so many specific facts and competences and also so as to enhance the architecture, apparently providing new forms of representation and algorithms to enable new kinds of things to be learnt at different stages of development.

As far as we know there are no learning/development mechanisms known in AI that have been shown to be capable of mimicking the development of biological meta-configured skills (although there are several AI learning mechanisms that are able to perform fragments of the tasks, some of which are mentioned in section 3.4); and despite many advances in brain science (e.g. [4], [23]) there are no known neural mechanisms that support all the detailed processes of perception of structure, perception of processes and actions, ontology extension, problem solving, plan formation, purposeful manipulation both of 3-D objects and of abstract ideas, doing mental arithmetic and the intellectual motivations, such as curiosity and puzzlement, apparently involved in these processes, though there has been much work on brain mechanisms explaining subsets of these processes, especially those shared between many species. Of course, how brains do all these things is not a new question: Haler, in [16], seems to suggest that it could be answered through more advanced brain imaging techniques. However, those may show where activities occur in the brain, but not necessarily what is happening (e.g. what information is being processed and in what way) and how it works. For that we shall need a theory of what processes are involved in various tasks. At present we can only refine requirements for such a theory.

3.3 Staggered/cascaded brain development

In section 3.1 we suggested that the evolution of preconfigured (precocial) competences might have benefited from separation between parts of the genome determining generic behaviours and parts acting as 'parameters' controlling instantiation of those types. If so, it is possible that one of the changes that allowed the evolution of meta-configured (altricial) competences was a mechanism to postpone the combination of type and parameter to later post-natal stages in individual development and further mechanisms to allow the 'parameters' to be adjusted rapidly by the environment, e.g. during play and possibly also social interactions. And likewise mechanisms for combining old competences into new more complex ones might have had their operation delayed so that instead of all such combinations occurring in a mostly genetically controlled way prior to birth or hatching, some of those construction processes were delayed till later stages in development, so that new combinations could build on meta-configured competences that have already been found useful. The individual would then do the learning that previously had to be done by the species (or genome), This would enormously speed up accommodation to changes in the environment, but would depend on cultural transmission to reduce the need for rediscovery in each generation.

It is an old idea in AI, e.g. in [46] [30], that combinatorial complexity of learning and

problem-solving processes can be reduced by postponing acquisition of the most complex concepts or skills until layers of simpler ones have been learnt as 'macro operators' based on an initial set of simple competences. Compared with the set of possible combinations of a pile of bricks made from sand the combinations of all the grains of sand in those bricks is astronomically large. Likewise, solving a problem that requires an arrangement of bricks held together with mortar, e.g. to make a bridge, is much easier than finding the solution among arrangements of the grains of sand held together with mortar. So cognitive development may be accelerated by repeatedly building new collections of 'macro operators' built on previous ones. Each new collection of macro operators defines a space of combinations which is relatively small compared with the search space of combinations of the initial set of operators. There is a price paid for that gain, namely that some combinations of the smallest operators are excluded: a smoothly curved object cannot be made from rectangular bricks.

This is related to the fact that adults and older children are sometimes more limited than younger children in imagining functions for objects [11]. This appears to be because, having built up a high level ontology, adults often restrict searching within spaces defined by that ontology (such as the designed function of an object), at the cost of missing something useful, whereas a child may explore a wider space generated by a lower level ontology. Of course the child may fail to notice more complex possibilities that the adult understands.

However, if the macro operators are selected by an individual because they have been found to be widely applicable in the environment for satisfying goals of that individual (including possibly goals suggested by a teacher), then the loss of generality may be more than compensated for by the speed with which new problems can be solved, and by tailoring of the system to address the kinds of problems that occur most frequently in a particular environment. The benefits of such construction of new collections of reusable building blocks for physical structures (namely, body parts) were clearly 'discovered' by evolution. We are proposing that this was expanded to building blocks for behavioural competences.

3.4 Implications for brain development

Many years ago Selfridge pointed out that a learning system can enormously reduce its search spaces by repeatedly making use of modules at a certain level of granularity to develop new modules at a higher level of granularity. His ideas were recently summarised in [30]. An early presentation of his ideas inspired an implementation in Pop-11, still available as part of the Poplog system [42]. It shows how search spaces for a learning system can be shrunk dramatically if results of earlier prior explorations are stored as new units for as new reusable units.

The idea that learning has to be staggered to control search spaces has been reinvented several times by AI researchers and others, for example in Winston's work on learning structural descriptions [52], Sussman's work on automatic plan construction

[46], Fahlman's 'Cascade correlation' architecture for neural nets that grow new layers as they learn more complex things [13], its successor 'Knowledge based cascade correlation' (KBCC), case-based and explanation-based learning (CBL/EBL), and the use of hierarchies in the HMOSAIC neural model of movement generation [17].

We suggest that evolution discovered, long before human researchers did, that this principle, previously used in the evolution of successive species (if the hypothesis in section 3.1 is correct), can also be applied to individual learning and problem solving, in the formation of concepts and cognitive competences.

The full benefit of this strategy will require some brain development to be delayed till after birth, and cascaded, allowing several layers of competences to be acquired successively using previously developed lower level competences which become more or less 'frozen' so that they can be relied on as units in larger structures. This is consistent with the suggestion that the rightmost column in Figure 1 can have several layers of meta-competences, each developed after other layers. If some of the meta-competences are implemented by growing new brain-structures then later structures will be grown after earlier ones.

For each layer the learning is controlled partly by the environment as a result of processes of play and exploration perhaps accompanied by generation, testing and debugging of hypotheses about what sorts of things are in the environment and what their causal properties are. (Of course we are not saying infants *know* they are doing this. If such learning of early competences takes weeks or months, that will mean that brain development must be correspondingly delayed. This may be one of the important reasons for humans (and, perhaps some other species) being born cognitively as well as physically underdeveloped.

This powerful form of meta-configured environmentally driven adaptation may have been re-implemented several times for different parts of the brain, e.g. producing layered growth of competences for different sensory modalities and also for cross-modal competences, along with layered growth of action sub-systems. A special case of this found in humans seems to be the basis of language learning.

3.5 Meta-level modularity

These ideas are related to discussions about modularity which have loomed large in evolutionary psychology in recent years, e.g. in the work of Cosmides and Tooby [7]. However in such discussions it is often assumed that only very specific modules are produced by evolution, e.g. modules for selecting certain kinds of nutritious food on the basis of taste and smell, and modules for detecting cheats in a cooperative system. If these modules are largely genetically determined then the competences they provide are of the kind we label preconfigured. We are, in contrast, postulating innate mechanisms that create such specific modules as a result of exploring and interacting with the environment, so that some of the meta-configured modules produced by innate bootstrapping mechanisms may differ substantially from one generation to the next, if

the environment changes. However we do not assume this uses totally general learning mechanisms: the learning mechanisms evolved to work in limited sorts of environments.

It is important that we are not claiming that all meta-configured rapid learning is found only in altricial species with a wide variety of competences learnt by play and exploration. It may also be found in specialised behaviours of precocial species (like the insects that use landmarks, mentioned above) or altricial species which do not undertake a lot of play or exploration. What we are claiming is that in some species, and especially in humans, a number of general versions of mechanisms supporting acquisition of a wide variety of layered meta-configured competences evolved and account for the cognitive sophistication and flexibility of those species.

4 BIOLOGICAL/EVOLUTIONARY TRADEOFFS

Apart from the questions discussed in 3.1 and 3.2 about *how* the preconfigured and meta-configured competences are acquired by individuals, there are also questions about *why* the different patterns of learning and development and different mixtures of those patterns evolved in different species. We can partially answer these questions by analysing some of the tradeoffs that can influence evolutionary processes.

Where learnt capabilities involve collaboration with conspecifics, rapid cultural changes can cause additional pressures favouring mechanisms capable of rapidly acquiring complex non-innate knowledge, including novel ontologies — as shown by very young human children picking up concepts their parents never had to learn at that age, such as mouse-controlled computer games. Such learning mechanisms, in turn can speed up cultural change: a form of positive feedback. A special case is language learning, where phonology, syntax and vocabulary learnt by a child born in one country may be very different from what the parents learnt as children in other countries.

One might liken the first extreme, precocial development, to a canal system: evolution constructs a system of channels, locks and sluices, through which the water tends to flow, and individual learning may amount to altering the speed of flow or shutting some channels and opening others according to need. In contrast, the second extreme, altricial development, is more like what we might call a 'meta-configured' canal system where initially water flows over a topographical surface. Its course is highly individual, and while the contours of the land may constrain where the water tends to flow, small perturbations in the surface and possibly other factors such as where and when the rains fall and the winds blow can redirect the flow in new directions, which in turn can modify the environment so as to produce new channels, to which locks and sluices along with pumps and control mechanisms are added, so that the use of the whole canal system is both tailored to how the world works and also capable of rapid global reorganisation to meet changing needs and opportunities. (We leave unanswered for the purposes of this contrast where the needs come from!) Of course these are only crude metaphors, partly inspired by Waddington's notion of an epigenetic landscape [50], which we have extended

to allow for different patterns of development.

4.1 Physical differences at birth

A notable feature of altricial species is that they tend to be physically under-developed for prolonged periods after birth or hatching. Young mammals such as canids or felids are hairless, with very uncoordinated motor activity and underdeveloped senses. Kittens' eyes are closed for about 9 days, and are not fully clear until 32 days after birth, and their external auditory canal opens after 12 days [49]. In comparison to a two day old caribou calf, a kitten of the same age looks like a foetus. Human babies are even more extreme in the length of their period of helplessness. While we have approximately the same gestation as other apes, our infants' post-natal development is grossly retarded compared to them [29]. Chimpanzees achieve independence at around 6 years old, and while the age of independence varies a great deal cross-culturally in humans, we generally take at least twice as long to become independent from our parents.

It is hard to escape the conclusion that for many altricial species, and particularly for humans, this prolonged period of development while protected and cared for by adults, and while the body and brain are growing and developing, is crucial for cognitive development. Since it is possible for evolution to produce a competent mammal or bird without such a prolonged developmental period, and since the production of incompetent offspring requires parents to invest a great deal in the protracted period of care, it is very likely that this pattern of development gives altricial species some advantage not available to (or required by) precocial species. Our speculations in section 3.3 about the benefits of cascaded development may be part of the answer. Another part may be that during the earliest stages where control is still at a primitive level it may be best for the initial learning to happen when limbs are relatively small and light, and muscles are relatively weak, so that control errors are unlikely to cause damage to the individual or members of its family.

5 ALTRICIAL, SELF-BOOTSTRAPPING ARCHITECTURES

How is the variety of developmental modes found in biology relevant to artificial self-organising systems? Inspired by the life history strategies found in nature, we can divide the types of development required in robots into two extreme categories, with mixtures in between, as we proposed in [38]. Application domains where tasks and environments are fairly static and machines need to be reliably functional quickly, require preconfigured skills (possibly including some adaptation and self-calibration, leading to minor and gradual changes within a largely fixed behavioural repertoire).

In contrast, meta-configured skills would be more appropriate where tasks and environments have widely varying details that are unknown to the designers of the machines, and where they also change in complex ways over time while the machines are in use, so that machines need to learn how to cope without being sent for re-programming.

Examples could include domestic robots to be deployed in a wide variety of cultures where there may be fast technological and cultural change. Another type of example might be a robot designed for underwater exploration or exploration on a new planet, where most of the details of the environment cannot be determined by the designers in advance, and where remote control by humans is impossible. Such robots would need fairly sophisticated default precocial (preconfigured competences) in addition to powerful and rapid learning mechanisms for generating and testing both theories about the environment and new altricial (meta-configured) competences appropriate to the environment, e.g. if a local life form is present whose members threaten to disrupt the operation of the robot it would be desirable to develop techniques for frightening them off without harming them.

In many such applications (e.g. domestic robots) the machines are likely to need an extended 'dependent' period, during which a human caretaker guides and supports development and the machine is not fully capable of its eventual range of 'adult' skills. Indeed, this period of dependency, active development, exploration and the resultant construction of new cognitive skills might be an important determinant of the eventual sophistication of the system, in the same way that enriching a rat's environment during development positively affects cognitive sophistication of the adult rat [48].

Interestingly, the effects of enrichment only seem to work if the rat is an active participant in the environment, not merely an observer of it [48], which is consistent with the hypothesis that such altricial development depends on the animal learning about the consequences of its own actions on the world. This would be the case if many such actions were specifically generated as part of the process of testing and debugging newly constructed competences and theories about the environment. Observing actions produced by others would be far less likely to produce information pertinent to the observer's current cognitive needs.

More precisely, if you have formed a hypothesis about why something worked or failed, then it is useful to be able to test the hypothesis by acting in the environment. If you cannot do that, you have to wait until a suitable test is provided the actions of others, or some other environmental effect – like a theoretical scientist who has no influence over experimental scientists, and can merely read their publications. This has two disadvantages: first it may require an arbitrarily long wait, and second it requires memory mechanisms that can relate all new observations to stored hypotheses awaiting tests. In contrast, being able to include testing of hypotheses during exploration and play while developing hypotheses has obvious advantages. This close coupling of autonomous actions to the construction of rich cognitive structures which are adapted to complex features of the environment may be very important both for species that develop meta-configured competences, and for some future AI systems.

Architectures, mechanisms, forms of representation and types of learning may differ sharply between the precocial and altricial extremes. Moreover the end results of altricial learning by the same initial architecture may differ widely in different environments

(discussed further in Section 6). Note that neither extreme involves an organism that is a 'blank slate', and which learns by incremental reinforcement learning only. Such a system would be unlikely to be able to produce systems even of comparable complexity to a day-old chick or a flying insect as it emerges from its pupa. As McCarthy wrote in [27] 'Evolution solved a different problem than that of starting a baby with no *a priori* assumptions.'

We noted previously that even altricial biological species such as humans have a *mixture* of preconfigured (precocial) and meta-configured competences. Robots may also need a mixture, for example where some features of the environment are predictable and therefore need not be learnt by the individual, or where there is no carer to provide support for a helpless learner in the initial stages. In cases like the underwater explorer and interplanetary robot, where humans are not available to provide initial care, the innate competences will have to play the role of a surrogate parent, and designing them may require designers to pre-program some coarse-grained knowledge of the features of the remote environment, even if many details are left to the robot to discover.

6   ONTOLOGIES FOR USE IN ALTRICIAL ARCHITECTURES

We have conjectured that some meta-configured competences arise because certain innate bootstrapping mechanisms, spontaneously discover discrete, re-usable and (recursively) recombinable chunks of information, using, for example, successively higher level 'macro operators' as described in section 3.3. We now consider the content of those operators in more detail, specifically the ontologies they use.

We wish to challenge two assumptions that have recently become commonplace among certain sorts of AI researchers, neuropsychologists and philosophers, namely (a) that all concepts acquired by an organism are derived from experience of instances (as claimed in the centuries old philosophical theory of *concept empiricism*, recently renamed 'symbol-grounding' theory) and (b) that all knowledge expresses observed correlations between instances of such concepts, often referred to as 'sensorimotor contingencies', and expressed in terms of continuously varying dynamical systems. An example of such a theory, in [23]) was given in Section 2.

What is rarely noticed is that this implies that all knowledge is about patterns and relationships between patterns, of processes that occur *within the body of the animal or robot*. Thus it is commonly assumed (at least implicitly) that all learning is about 'somatic' sensorimotor contingencies (concerned with relations between states, events and processes in sensors and motor controllers within an animal or machine). We wish to contrast that with learning about 'exosomatic' objective condition-consequence contingencies (concerned with relations between states, events and processes in the environment). The latter requires the learner to use an ontology referring to things and processes in the environment whose existence is independent of whether and how they are perceived or produced.

19

The last kind of learning is important for the sorts of altricial species we are considering, because it allows chunks of knowledge acquired in the context of particular percepts or action to be re-used in different context. For example, young human infants reach to grasp an object with their mouths if their hands are not available [29], and once they have discovered that they can pull a toy resting on a blanket towards them by pulling the blanket, they can transfer the same action to perceptually very different, but functionally similar materials (such as a sheet of paper). In other words, what is learnt is something *objective* about the world (namely, that objects can be grasped by inserting them between two surfaces and then bringing the surfaces together, or that flexible materials can be scrunched and pulled). If this information is not tied closely to particular sensory and motor signals (within the body) but uses an *exosomatic* ontology referring to objects and relationships in the environment, then generalisations learnt about those objects and relationships can be applied to new contexts, where the perceptual and motor details are very different. (If this is correct, so-called 'mirror neurons' should probably have been described as 'abstraction neurons'.)

Second, we conjecture that genetically-programmed mechanisms might determine which kinds of actions and percepts the animal learns, but also that this learning — rather than being driven by direct reward and punishment — might instead be driven by an exploratory drive, where a notion of 'interestingness' determines which actions, effects and percepts are explored and stored as new units of knowledge, and the environment contributes substantial information. More specifically, altricial learning may be based on genetically determined mechanisms with:

- implicit meta-level knowledge⋆ about kinds of information chunks that might be learnt, including
  – perceptual chunks (using concepts of space and time)
  – action chunks (using a concept of causation)

- implicit meta-level knowledge about kinds of associations that might be perceived, and knowledge about how to investigate which are causal and which spurious

- mechanisms for combining old chunks into more complex wholes (e.g. complex goals, or action sequences)

- mechanisms for discovering new complex wholes that occur in the environment, including
  – enduring structured objects that have persistent features, parts and patterns of behaviour,
  – processes extended in time in which objects endure even when not perceived
  – more and more complex actions produced and controlled by the individual

---

⋆ Our use of the word 'knowledge' here does not carry philosophical implications of 'justified true belief', since we are simply referring to information acquired through evolution and learning and then later used.

- mechanisms for creating and manipulating hypothetical structures which might describe unobserved portions of reality or possible future complex actions

- mechanisms for deriving consequences from complex information structures and for comparing and selecting between complex structures with different consequences

- mechanisms (in a small subset of species) for discovering useful types of information about mental phenomena, in oneself and in others.

The history of science, including the development of theoretical concepts like 'electromagnetic radiation', 'neutrino', 'valency', 'gene' shows that besides the ability to extend an ontology by defining new concepts as useful abbreviations for combinations of previously understood concepts, humans also have the capacity of *substantive ontology extension*. We suggest that this is not only true of scientists, but also of young children discovering what sort of environment they have been born into. But how that is done, and how the potentially astronomical search for *useful* conceptual extensions is controlled remains unsolved. It may be partially constrained by the requirement to develop coherent causal explanations. In [38] we suggested, on the basis of recent philosophy of science, that some concepts referring to unobservables are largely defined by their role in a theory along with some links (sometimes called 'meaning postulates', following [6]) between the theory and observation and experiment. We called that idea 'symbol attachment', but now prefer 'symbol tethering'.

As illustrated in [39] and in the work of Gopnik, [15] and Rochat [29] there seem to be many hundreds of small experiments a young child spontaneously conducts during the first few years of life, as a result of which a wide variety of competences are acquired that are to some extent independent, but can be combined in various ways, like using the initial part of an action acquired for the purpose of transferring food to one's mouth and combining it with another sort of action to transfer the food to a roughly horizontal surface. This seems to require mechanisms that spontaneously generate goals that have nothing to do with biological needs such as food, comfort, warmth, avoiding injury, but instead serve the processes of cognitive development, especially when the goals are not achieved, predictions fail, and current concepts and theories about the environment have to be debugged. Eventually such a child creates appropriate concepts and theories using those concepts, regarding the relationships that prevent or enable the occurrence of various processes.

Our suggestion is that this growing web of concepts and theories refers to entities, relationships and processes in an *external* environment not to patterns in sensor and motor signals. This requires far more sophisticated and varied forms of representation than can be found in control theory where everything is expressed in terms of vectors of numerical values and something like differential equations relating the changes of values of those variables and their derivatives. A video of an 11 month old baby playing with a spoon

and a tub of yogurt is available at http://www.cs.bham.ac.uk/~axs/fig/yog.mpg. The child seems to attempt to transfer yogurt to his leg, and later to the carpet, but fails to understand the role of the bowl in preventing the transfer if the spoon is not inverted. In order to represent the role of the bowl as both supporting transfer of yogurt to the mouth and preventing transfer of yogurt to a leg or a carpet, an ontology is needed in which there are relationships between three dimensional surfaces and volumes, that depend on kinds of material that are rigid and impervious. Further discussion of these points can be found in [35], [36], [37], and [39].

6.1 Pre-linguistic development

The kinds of exploration and forms of learning discussed here do not depend on the use of an external human language, for they occur in pre-linguistic children and to some extent in animals that never learn to talk. So although the mechanisms require some form of syntax for expressing goals, hypotheses, plans and predictions, that formalism probably evolved before human language, as argued in [32]. Variants of such meta-configured competences might have evolved later, tailored to communication, supporting the development of external languages with combinatorial syntax and semantics.

Moreover, in humans, and perhaps some other species, meta-configured capabilities that were originally *outwardly* directed (e.g. perceiving and acting on external objects and processes) might, after suitable architectural extensions, also be *inwardly* directed, allowing individuals to develop more and more complex chunks of information not only about the environment, but also about their own internal processes of perception, reasoning, learning, problem solving, motivation, choosing, etc. Some of the benefits for a robot of formalisms and architectures allowing self-understanding are discussed by McCarthy in [25].

Ontologies used for such internal 'meta-management' could also be used in mechanisms for perceiving, reasoning about and behaving towards others (e.g. conspecifics, prey and predators). Both the inward-directed and outward-directed cases require *meta-semantic* competence: the ability to represent and reason about entities which themselves process information. Animals and machines with such mechanisms can, for example, try to produce, change or prevent beliefs, plans or desires in others.

In humans, growth of the architecture seems to be multi-faceted, involving acquisition of many new sub-ontologies, new forms of representation, new collections of skills required for particular domains, e.g. learning a new language, learning to read music and play an instrument, learning programming, learning academic disciplines, learning athletic or dancing skills, learning mathematics, or quantum physics. Later growth enriches the architecture by growing new links between such domains – including using some as 'metaphors' for others.

As explained in section 3.2, searching for a combination of large units that solves a complex problem may be very much faster than if the search either had to use more primitive units or had to use gradual modification of existing units. Of course, if the

22

animal is restricted to searching in a space of combinations of large chunks, that may prevent solutions being found to some problems. So it is also useful to have clues that some searching in a lower-level space is needed (like designers of computing systems sometimes having to invent new programming languages, or new types of computer hardware). This is one of the ways in which an individual can learn that a lower level ontology included richness that has not yet been exploited.

7 EVIDENCE FOR META-CONFIGURED COMPETENCES AND ARCHITECTURES

How can we provide evidence for and examine the kinds of re-usable, re-combinable ontologies described above? Adult human behaviour is so rich, flexible and quickly adaptable to complex and novel features of the environment that it is hardly credible to attribute it to learning of somatic patterns in sensor and motor signals alone. Furthermore, because humans have language, introspection and self-reports about cognitive processes can provide some information, even though they can sometimes be misleading. But how can we investigate architectures in pre-verbal humans and non-verbal animals? Superficially, learning of somatic sensorimotor contingencies and the kinds of cognitive altricial architectures we are discussing here can both result in complex patterns of behaviour, so determining which behaviours result from which kind of learning process is difficult. One approach is to examine the details of the behaviour under experimental manipulations to see how flexible the behaviour is, and in particular, whether the animal appears to be abstracting knowledge gained in one context to use in another [19].

Despite these difficulties, we do have some evidence from both pre-verbal humans and non-human animals. From very early in life, human infants seem to have certain expectations about the way that the world works, as reported by Spelke [44] and Rochat [29]. These expectations may help to guide and organise infants' learning and exploration of their environment and form re-usable chunks of knowledge about what happens when, for example, two moving objects make contact with one another.

In the physical domain, a tool-making New Caledonian crow (*Corvus moneduloides*) showed spontaneous, novel modification of a tool for a specific task [51], apparently applying a new manufacturing technique to an unfamiliar material. Chimpanzees seem to be able to use causal cues (such as the sound of food being shaken in a cup) to infer the location of food [5].

Furthermore, some animals seem capable of certain aspects of forward planning (implying that their actions are not solely controlled by the immediately available stimuli). Bonobos (*Pan paniscus*) and orangutans (*Pongo pygmaeus*) selected, saved and transported appropriate tools between rooms, and used them to obtain a goal up to 14 hours later [28], even when the apparatus was not visible when they were making their selection. Western scrub jays (*Aphelocoma californica*) remember which individual scrub jay was watching when they stored a food cache, and adjust their behaviour accordingly

when re-caching the food [9]. Thus, they act as if they are prospectively planning for the likelihood that their caches will be pilfered by another individual.

For most species, the appropriate experiments have not been carried out, so it is likely that more evidence of meta-configured architectures will be found in non-human animals as time goes on. Moreover, if we can produce working models of the sorts sketched here that will be evidence of feasibility, and will probably generate new predictions and questions that can be tested empirically.

8   SUMMARY SO FAR

We can summarise the key ideas presented here as follows, many of which are, of course, old and familiar ideas. After starting with very simple self-replicating machines and biological evolutionary processes gradually added several different kinds of complexity, including the following (though not necessarily in exactly this order):

(1) making machines that, instead of replicating themselves directly, produce seeds or eggs that *grow* physical replicas partly under the control of a genome (represented chemically) and partly under the influence of the environment;

(2) producing more and more complex physical machines that use that method of reproduction, requiring increasingly complex resources such as food, water, shelter, mates, protection from predators — which in turn led to a need for increasingly complex control mechanisms;

(3) adding more and more complex systems for controlling behaviour of those machines, in the form of information-processing (virtual) machines implemented in a variety of physical mechanisms, including both chemical and neural mechanisms;

(4) finding ways of representing behaviour control systems in the genome so as to support evolution of behaviours through mechanisms of natural selection;

(5) allowing some of the behaviour control mechanisms to be determined by adaptive learning processes during the period of growth and development of the individual instead of being preconfigured in the genome;

(6) extending those control systems so that instead of being limited to on-line, dynamic, control of behaviour some of them were able to acquire, store and manipulate information for future use, including subsequently supporting prediction and planning several steps into the future, using discrete, structured forms of representation with compositional semantics;

(7) overcoming slowness of *general-purpose* learning mechanisms, through the operation of genetically determined meta-competences tailored to learning about particular *types* of environment and *types* of task, through action (including play and exploration), often requiring ontologies to be extended by introducing semantic extensions not definable in terms of previous semantic contents;

These ideas seem to be close to some aspects of Piaget's theories about a child's construction of reality, or at least the modified versions of Piaget's theories presented

24

by Rochat. Like Piaget we have assumed that what the child learns as it develops has a layered structure, and like Piaget we do not postulate a rigid sequence of development or rigid boundaries between the layers. Exactly which chunks of information, and which meta-configured competences a particular altricial individual learns will be influenced by many factors, including the physical actions possible for its body, the environment and its affordances, and the individual's history. These factors could produce different kinds of understanding and representation of space, time, motion, causality and social relations in different species, or in similar individuals in different environments.

Nevertheless we should not ignore the fact that humans born blind, or deaf, or with missing limbs, or with control deficiencies as in cerebral palsy, can develop a similar understanding of the environment and communicate in rich and deep ways by the time they are adults. So although our bodies are relevant to how we learn about the environment, and what we learn, the fact that we can develop an exosomatic ontology seems to allow human minds to become more and more independent of their bodies so that people with very different sensorimotor information end up talking and thinking about the same things.

9 CONCLUSION

We have attempted to provide evidence suggesting that by attending closely to the diversity of biological phenomena we can gain new insights into (a) how evolution happens, (b) what sorts of mechanisms, forms of representation, types of learning and development and types of architectures have evolved, (c) new ways to think about the relations between genetic and environmental influences on individual development (instead of regarding them as in competition we need to see how they cooperate), (d) how to explain ill-understood aspects of human and animal intelligence, and (3) new useful mechanisms for artificial systems. We have analysed some of the tradeoffs between alternative designs for organisms or machines, and identified a need for a kind of architecture that grows itself, using powerful symbolic mechanisms to achieve discontinuous learning, including development of an exosomatic ontology, referring to things in the environment, in contrast with mechanisms that learn only sensorimotor patterns and associations between them and adaptive mechanisms that make only continuous modifications within a fixed architecture.

The ideas presented here are not established in any detail: rather they define a research programme that will have to be evaluated mainly by its results in the long term (as described by Lakatos in [21]), which could include new contributions to theoretical biology, developmental psychology and the set of designs worth using in AI, including designs for human-like robots. Some of the new contributions to science will be new questions, some of which are already asked in this paper.

Whether computers as we know them can provide the infrastructure for the types of systems we have described, or whether new low-level information processing mechanisms

are needed, remains to be seen. If new mechanisms are needed, advances in brain science will presumably provide clues as to their nature. One of the ways to answer the question is to start from detailed descriptions of collections of animal competences and patterns of development, use them to define precise requirements for integrated working systems, such as robots with perceptual and manipulative capabilities, going beyond the current generation of mobile robots that can do little more than move around, and see whether the techniques developed in different branches of AI, including both symbolic and neural mechanisms, can be combined in ways that will demonstrate the sorts of competences we have described. If that does not work, it could be simply because our designs need to be improved, or it could be because the available low-level mechanisms are inadequate. It is too early to predict the outcome.

10   ACKNOWLEDGEMENTS

REFERENCES

[1] The Genetic Programming Notebook. Website with information about Genetic programming http://www.geneticprogramming.com/index.html.

[2] Patrick Bateson and Gabriel Horn. (1994). Imprinting and recognition memory: A neural net model. *Animal Behaviour*, 48(3):695–715.

[3] M. Berdoy, J. P. Webster, and D. W. Macdonald. (2000). Fatal Attraction in Rats Infected with Toxoplasma gondii. *Proceedings: Biological Sciences*, 267(1452):1591 – 1594.

[4] Alain Berthoz. (2000). *The Brain's sense of movement*. Perspectives in Cognitive Science. Harvard University Press, London, UK.

[5] Juliane Brauer, Juliane Kaminski, Julia Riedel, Josep Call, and Michael Tomasello. (2006). Making inferences about the location of hidden food: Social dog, causal ape. *Journal of Comparative Psychology*, 120(1):38–47.

[6] R. Carnap. (1947). *Meaning and necessity: a study in semantics and modal logic*. Chicago University Press, Chicago.

[7] L. Cosmides and J. Tooby. (1992). Cognitive adaptations for social exchange. In J. Barkow, L. Cosmides, and J. Tooby, editors, *The adapted mind*. Oxford University Press., New York.

[8] J. E. Cushing Jr. (1944). The relation of non-heritable food habits to evolution. *Condor*, 46(6):265–271. Online at http://elibrary.unm.edu/sora/Condor/files/issues/v046n06/p0265-p0271.pdf.

[9] Joanna M Dally, Nathan J Emery, and Nicola S Clayton. (10.1126/science.1126539). Food-caching western scrub-jays keep track of who was watching when. *Science*, 18 May 2006:1–4. Online at http://dx.doi.org/10.1126/science.1126539.

[10] R Dawkins and J R Krebs. (1979). Arms races between and within species. *Proceedings of the Royal Society of London B*, 205:489–511.

[11] Margaret Anne Defeyter and Tim P German. (2003). Acquiring an understanding of design: evidence from children's insight problem solving. *Cognition*, 89:133–155.

[12] D.C. Dennett. (1996). *Kinds of minds: towards an understanding of consciousness*. Weidenfeld and Nicholson, London.

[13] S. E. Fahlman and C. Lebiere. (1990). The cascade-correlation learning architecture. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems*, volume 2, pages 524–532, Denver 1989. Morgan Kaufmann, San Mateo. Online at citeseer.ist.psu.edu/article/fahlman91cascadecorrelation.html.

[14] J.J. Gibson. (1986). *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates, Hillsdale, NJ. (originally published in 1979).

[15] Alison Gopnik and Laura Schulz. (August 2004). Mechanisms of theory formation in young children. *TRENDS in Cognitive Sciences*, 8(8):371–377. Online at http://ihd.berkeley.edu/gopnik_tics.pdf.

[16] Richard J. Haler. (1990). EDITORIAL: The End of Intelligence Research. *Intelligence*, 14:371–374.

[17] Masahiko Haruno, Daniel M. Wolpert, and Mitsuo Kawato. (2003). Hierarchical mosaic for movement generation. *Excepta Medica International Congress Series*, 1250:575–590.

[18] M Hauser. (2001). *Wild Minds: What Animals Really Think*. Penguin, London, UK.

[19] A Kacelnik, J Chappell, A A S Weir, and B Kenward. (2006). Cognitive adaptations for tool-related behaviour in New Caledonian crows. In E A Wasserman and T R Zentall, editors, *Comparative Cognition: Experimental Explorations of Animal Intelligence*, pages 515–528. Oxford University Press, Oxford.

[20] I. Kant. (1781). *Critique of Pure Reason*. Macmillan, London. Translated (1929) by Norman Kemp Smith.

[21] I. Lakatos. (1980). The methodology of scientific research programmes. In J. Worrall and G. Currie, editors, *Philosophical papers, volume I*. Cambridge University Press.

[22] J Le Dividich, J A Rooke, and P Herpin. (2005). Nutritional and immunological importance of colostrum for the new-born pig. *The Journal of Agricultural Science*, 143:469–485.

[23] M. Lungarella and O. Sporns. (October 2006). Mapping Information Flow in Sensorimotor Networks. *PLoS Computational Biology*, 2(10:e144):1301–1312. Online at http://dx.doi.org/10.1371

[24] P. Marler. (1990). Innate learning preferences: Signals for communication. *Developmental Psychobiology*, 23(7):557–568. Online at http://dx.doi.org/10.1002/dev.420230703.

[25] J. McCarthy. (1995). Making robots conscious of their mental states. In *AAAI Spring Symposium on Representing Mental States and Mechanisms*. Accessible via http://www-formal.stanford.edu/jmc/consciousness.html.

[26] J McCarthy. (1996). From here to human-level intelligence. In *KR '96*.

[27] J. McCarthy, (1996). The Well Designed Child. Discussion paper: http://www-formal.stanford.edu/jmc/child1.html.

[28] Nicholas J Mulcahy and Josep Call. (2006). Apes save tools for future use. *Science*, 312:1038–1040. Online at http://dx.doi.org/10.1126/science.1125456.

[29] Philippe Rochat. (2004). *The Infant's World*. The Developing Child. Harvard University Press.

[30] O. G. Selfridge. (1993). The Gardens of Learning A Vision for AI. *AI Magazine in*, 14(2):36–48. Online at http://www.dcs.fmph.uniba.sk/ hegedus/su/AIMag14-02-005.pdf.

[31] N. Shefferly and K. Joly. (2000). Rangifer tarandus (On-line). In *Animal Diversity Web*. University of Michigan. Accessed July 26, 2006 at http://animaldiversity.ummz.umich.edu/site/accounts/information/Rangifer_tarandus.html.

[32] A. Sloman. (1979). The primacy of non-communicative language. In M. MacCafferty and K. Gray, editors, *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979*, pages 1–15, London. Aslib. http://www.cs.bham.ac.uk/research/cogaff/.

[33] A. Sloman. (2001). Beyond shallow models of emotion. *Cognitive Processing: International Quarterly of Cognitive Science*, 2(1):177–198.

[34] A. Sloman. (2006). Introduction to Symposium GC5: Architecture of Brain and Mind Integrating high level cognitive processes with brain mechanisms and functions in a working robot. In *Proceedings of the AISB '06 Adaptation in Artificial and Biological Systems*, Bristol. Online at http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0602.

[35] A. Sloman. (2006). Orthogonal recombinable competences acquired by altricial species. Technical Report COSY-DP-0601, School of Computer Science, University of Birmingham, UK. http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601.

[36] A. Sloman. (2006). Sensorimotor vs objective contingencies. Technical Report COSY-DP-0603, School of Computer Science, Birmingham, UK. Online at http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603.

[37] A. Sloman. (2006). Two views of child as scientist: Humean and Kantian. Technical Report COSY-PR-0506, School of Computer Science, University of Birmingham, UK. Online at http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506.

[38] A. Sloman and J. Chappell. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings IJCAI'05*, pages 1187–1192, Edinburgh. Online at http://www.cs.bham.ac.uk/research/cogaff/05.html#200502.

[39] A. Sloman, J. Chappell, and Cosy-Project-Team, (2006). How an animal or robot with 3-D manipulation skills experiences the world. Poster for 10th annual meeting of the Association for the Scientific Study of Consciousness, Oxford, Online at http://eprints.assc.caltech.edu/46/.

[40] A. Sloman and R. L. Chrisley. (June 2005). More things than are dreamt of in your biology: Information-processing in biologically-inspired robots. *Cognitive Systems Research*, 6(2):145–174.

[41] A. Sloman, R.L. Chrisley, and M. Scheutz. (2005). The architectural basis of affective states and processes. In M. Arbib and J-M. Fellous, editors, *Who Needs Emotions?: The Brain Meets the Robot*, pages 203–244. Oxford University Press, Oxford, New York. http://www.cs.bham.ac.uk/research/cogaff/0-INDEX03.html#200305.

[42] A. Sloman and J.L. Cunningham, (1983). TEACH FINGER. Pop-11 Tutorial file, part of the Poplog programming environment available here http://www.cs.bham.ac.uk/research/projects/poplog/teach/finger; Originally inspired by O.G. Selfridge.

[43] A. Sloman, J.L. Wyatt, N.A Hawes, J.M. Chappell, and G-J. M. Kruiff. (2006). Long Term Requirements for Cognitive Robotics. In *Proc. Cognitive Robotics O6 Workshop, AAAI'06*. AAAI.

[44] E Spelke. (1994). Initial knowledge: six suggestions. *Cognition*, 50:431–445.

[45] S. Stepney, (2006). A Grand Challenge for Computing Research: Journeys in Non-Classical Computation. http://www.cs.york.ac.uk/nature/gc7/GC-7_2000_words.pdf.

[46] G.J. Sussman. (1975). *A computational model of skill acquisition*. American Elsevier.

[47] N. Tinbergen. (1951). *The Study of Instinct*. Oxford University Press, London.

[48] Henriette van Praag, Gerd Kempermann, and Fred H Gage. (2000). Neural consequences of environmental enrichment. *Nature Reviews Neuroscience*, 1:191–198.

[49] J R Villablanca and C E Olmstead. (1979). Neurological development of kittens. *Developmental Psychobiology*, 12(2):101–127.

[50] C.H. Waddington. (1957). *The strategy of the genes*. MacMillan, New York.

[51] A A S Weir, J Chappell, and A Kacelnik. (2002). Shaping of hooks in New Caledonian crows. *Science*, 297:981.

[52] P. H. Winston. (1975). Learning structural descriptions from examples. In P. H. Winston, editor, *The Psychology of Computer Vision*, pages 157–209. McGraw-Hill, New York.