

If Learning Maths Requires a Teacher, Where did the First Teachers Come From?

Aaron Sloman¹

<http://www.cs.bham.ac.uk/~axs>

Abstract. This is the latest progress report on a long term quest to defend Kant's philosophy of mathematics. In humans, and other species with competences that evolved to support interactions with a complex, varied and changing 3-D world, some competences go beyond discovered correlations linking sensory and motor signals. Dealing with novel situations or problems requires abilities to *work out* what can, cannot, or must happen in the environment, under certain conditions. I conjecture that in humans these products of evolution form the basis of mathematical competences. Mathematics grows out of the ability to use, reflect on, characterise, and systematise both the discoveries that arise from such competences and the competences themselves. So a "baby" human-like robot, with similar initial competences and meta-competences, could also develop mathematical knowledge and understanding, acquiring what Kant called synthetic, non-empirical knowledge. I attempt to characterise the design task and some ways of making progress, in part by analysing transitions in child or animal intelligence from empirical learning to being able to "work things out". This may turn out to include a very general phenomenon involved in so-called "U-shaped" learning, including the language learning that evolved later. Current techniques in AI/Robotics are nowhere near this. A long term collaborative project investigating the evolution and development of such competences may contribute to robot design, to developmental psychology, to mathematics education and to philosophy of mathematics. There is still much to do.

1 Mathematics as a Biological Phenomenon

1.1 Evolved competences and meta-competences.

I once met a software engineer who poured scorn on evolution as a designer because it had produced no animals with wheels. I have not met one who poured scorn on evolution for designing human neonates to be helpless while many other species have self-sufficient neonates: e.g. chicks get themselves out of the egg and soon peck for food, while a new foal or calf soon gets up on four legs and finds its way to the mother's nipple. It is commonly supposed that humans are born so immature because larger skulls would kill their mothers, but evolution produced some mammals with larger heads, e.g. elephants. Another common suggestion is that in an environment that is continually changing, learning capabilities are more important than innate knowledge: but that does not explain why evolution does not produce both in the same species, namely a collection of innate modules that control various aspects of behaviour along with a general purpose learning module (one version of the so-called "modularity" thesis).

An alternative answer is that in addition to those competences required from birth (e.g. breathing, sucking, digesting food, doing various things that attract attention of parents) there are special-purpose, more abstract, *meta-competences* for acquiring competences (e.g. for acquiring competences involved in use of independently movable manipulators, such as hands, claws, or trunk), where the perceptual and motor processes are so complex and varied that they cannot be encoded in the genome (possibly for reasons suggested in [33]). On this alternative, some behaviours are produced by sophisticated run-time control competences that have to be learnt, using specialised meta-competences (for learning) that also develop through the influence of the environment on meta-meta-competences: throughout life, humans go on learning how to learn many different things.

Sloman and Chappell [5] conjectured that humans are born with, among other things, sophisticated, evolved, *meta-knowledge about how to learn in a complex 3-D world containing more or less complex structures and processes, opportunities and obstacles, and also other intelligent individuals*. This may include learning how to grow layers of meta-competence suited to increasingly complex and/or abstract features of the environment, through play and exploration. The innate/learned dichotomy does not do justice to such diversity.

This paper extends those ideas by providing some (informal) empirical observations and theoretical conjectures about layered development and learning, that may lead to further progress in understanding how mathematics fits in with the rest of animal cognition, including precursors in human infants and toddlers, and in other animals, unlike more widely studied, allegedly innate, precursors concerned with recognition or estimation of objecthood, numerosity, size, etc. [41].

1.2 Philosophy of mathematics.

This proposal is related to the philosophy of mathematics of Immanuel Kant [15], according to whom mathematics is neither empirical (as suggested by Mill [22]), nor merely about "relations between ideas" (as suggested by Hume). But that does not mean mathematical knowledge is innate, or that mathematical competences are infallible. Neither is the nature of mathematical knowledge usefully described in terms of "feelings of compulsion" [1], e.g. since mathematical discoveries may be tentative conjectures at first. Rather the structure, kinematics and dynamics of the 3-D world generally requires generalisations and predictions to be tested *empirically*, but includes features that allow powerful *short-cuts* (using reasoning) to be discovered and used in creative problem-solving [6]. This makes possible biologically useful *productive laziness* which reduces or eliminates the need for empirical trial and error, or symbolic searching, and sometimes provides relatively quick solutions for

¹ University of Birmingham, UK, email: a.sloman@cs.bham.ac.uk

novel practical problems, avoiding empirical experiments that could be dangerous, or even fatal. However, the mechanisms are not infallible: testing and debugging may be required, as Lakatos [17] demonstrated.

In humans, these products of evolution develop into what we understand as mathematical competences. (Examples are given below.) Many of species-specific details are of psychological interest, but that does not mean that the subject matter of mathematics is psychological: it is primarily concerned with the structure of the world, and also how that structure can be discovered, and used, leading to the study of alternative possible structures.

1.3 Russell vs Feynman on Mathematics.

Russell and Whitehead [43] attempted to demonstrate that all of mathematics could be reduced to logic (Frege had attempted this only for Arithmetic). Despite the logical paradoxes, Russell thought the goals of the project could be or had been achieved, and concluded that mathematics was just the *investigation of implications that are valid in virtue of their logical form, independently of any non-logical subject matter*. He wrote: “Mathematics may be defined as the subject in which we never know what we are talking about, nor whether what we are saying is true” (*Mysticism and Logic* [25]).

In contrast, Feynman described mathematics as “the language nature speaks in”. He wrote: “If you want to learn about nature, to appreciate nature, it is necessary to understand the language that she speaks in” [7]. I believe that Feynman’s description is closely related to generally unnoticed facts about how a child (and perhaps future robots) can develop powerful, reusable concepts and techniques related to patterns of perception and patterns of thinking that are learnt through interacting with a complex environment, part of which is the information-processing system within the learner [28, Ch 6]. Despite the role of experience in such learning, the *results* of such learning are not empirical generalisations. Kant wrote, in [15], “though all our knowledge begins with experience, it by no means follows that all arises out of experience.” Feynman seems to agree.

1.4 Learning Mathematics and Learning Language.

I suspect that some of the re-conceptualisation involved in developing mathematical insight uses mechanisms produced by evolution for unifying disparate empirical discoveries in a deeper, more general, explanatory theory. A special case of this, which evolved later, and had to be modified to allow exceptions to explanatory theories, seems to be the mechanism for language learning that allows empirical learning of useful patterns of verbal communication after a time to trigger the development of a rule-based syntactic theory that is far more general in its predictive and explanatory power than the previously learnt patterns. In children this leads to “U-shaped” learning² because at first the new system cannot handle exceptions, so incorrect utterances are produced (e.g. “she speaked”, “he hitted me”) by a child who earlier used the correct forms (“she spoke”, “he hit me”). Later the architecture is extended to cope with exceptions (a non-trivial alteration to an information-processing system) and both previously known and newly learnt exceptions to the rules are dealt with correctly. Further reorganisation may occur as the language changes.

This ability to develop a powerful (implicit) generative theory of the language in use, replacing previously acquired empirical generalisations, may be a specialised version of mechanisms that evolved earlier because they support productive laziness in coping with a physical environment. Language is not the only aspect of

the world that has general laws with exceptions that have to be accommodated. But some aspects are exceptionless, including things learnt about topological features of space and time and some of the metrical features. There are also exceptionless generalisations about numbers and about logical forms, which can be learnt later and embedded in a powerful theory.

The possibility of empirical learning being supplemented by more powerful mechanisms supporting “productive laziness” does not seem to be widely recognised. For example, many examples of learning about affordances in young children are described in [8]. The authors, like many others, assume there are only three ways for a child to have knowledge about what can and cannot occur, namely either it is innate, or it is acquired by learning associations by observation and trial and error, or the information is acquired from someone else, by imitation or being instructed. They, like many others, ignore a fourth way: learning by *working things out*. Having an appropriate theory about some aspect of reality allows a child (or animal, or physicist, or engineer) to *work out* what must happen even in some situations that have never previously been encountered (e.g. dealing with new spatial configurations) whereas if the learning were purely empirical, every novel situation would have to be tested before predictions could be relied on. Examples are given in the next section. Systematic collection and testing of examples in young children and other species could form a major interdisciplinary research project. Finding explanations, producing working models of the processes, and explaining how neural mechanisms can support both the processes of development and the processes of discovery is a bigger challenge. Meeting that challenge is a requirement for explicating and defending Kant’s philosophy of mathematics – though that will not be obvious to most philosophers.

From this standpoint, someone composing a poem or story is working out what can be thought or said, using mechanisms that are partly similar to mechanisms used by someone working out what can be done in the environment, or working out what structures and processes can and cannot exist in the environment, or what relationships can and cannot hold between numbers. All of these require learning through interaction between meta-meta-(meta-...) competences produced during evolution of the species, and information provided by the individual’s environment.

2 Examples of “Working Things Out”

The ability to learn by creatively working something out is being actively studied in animal behaviour research, a spectacular example being Betty the hook-making crow (search for “betty+hook+crow” for videos). Paradoxically the phenomenon has escaped the notice of many researchers investigating human children, mainly because of the wide-spread assumption that their learning is mainly Bayesian. My examples are mainly from human learners – young and old.

2.1 Examples of “Productive Laziness”

(1) A deceptively obvious, though quite sophisticated, example concerns invariant results of sequential matching operations. A child may discover empirically (and with surprise and delight) that counting the same set in different orders (e.g. counting fingers on a hand from left to right and from right to left) always produces the same result – if no items are added/removed or counted twice, omitted from the count etc. This can start off as an empirical discovery and later become obvious (at least for some learners). How? Is it merely that after finding much evidence to support the generalisation learners associate a very high *probability* with it?

² <http://www.cis.udel.edu/~case/slides/nugget-ushape.pdf>

That would be rash, since the evidence would be limited to only a tiny subset of possible cases: what justifies the extrapolation to all the very many sets of objects much larger than any the learner has encountered? Can you give a proof? Some philosophers believe that the result is not proved – merely stipulated as a criterion for doing the counting properly. (They have presumably never experienced the process of making a mathematical discovery.) Few mathematical results have only one proof. This result could be understood on the basis of transitivity of one-to-one correspondences, shifting the problem – an exercise for the reader.

(2) Another example can occur after the child has learnt that there are different kinds of matter, some rigid, some not. It is possible to discover empirically that if a part of a rigid object is grasped and rotated all the other parts of the object will move, and their movements will be greater the further they are from the rotated portion. Later on comes understanding that rigidity (which is a local property of resistance to deformation) can have global effects, so that movement of part of object can cause all other parts to move, and that rotary movements have amplified remote effects. The reasoning can be applied to objects of many shapes and sizes in many contexts, and makes it possible to predict consequences of rotation and other movements, and also constraints on such movements (assuming rigidity and impenetrability of the materials). This sort of discovery can, of course, have practical consequences that are important.

(3) A child who can lift a cut-out picture from its recess may know which recess it should go back into but fail to insert it without a lot of random movement. Later, the child discovers the need to *align the boundary of the picture with the boundary of the recess*. That may require a non-trivial extension of the child's ontology to include concepts of boundary and alignment. If the boundaries are aligned and the picture piece is slightly smaller than the recess then all parts of the piece can simultaneously avoid obstruction from the surrounding material. Any slightly different position causes overlap that prevents insertion. Unless the recess and picture piece are both circular, most *rotations* will cause obstruction. However if the piece is symmetrical there may be several rotations that allow insertion, but most rotations will cause obstruction: why?

A child who can reliably get such pictures back into their recesses, may at first merely attach a very low probability of success to insertion without boundary alignment then later come to see that the insertion *must* fail, but without being able to explain that this impossibility depends on the materials being rigid and impenetrable. What changes between noticing an empirical regularity and seeing that it cannot have exceptions?

(4) There is much to learn about strings. If an inelastic but flexible string is attached to a remote movable object, then if the end is pulled away from the object a process can result with two distinct phases: (1) curves in the string are gradually eliminated (as long as there are no knots), and (2) when the string is fully straightened the remote object will start moving in the direction of the pulled end. However, if the string is looped round a fixed pillar, the first sub-process does not produce a single straight string but two straight portions and a portion going round the pillar, and in the second phase the attached object moves toward the pillar, not toward the pulled end; and so on. Many more complex variants can be reasoned about.

Many more examples including string, elastic, pins, stones, mazes, and marks on paper are described in [26]. Some of their experiments can be performed on *simulated* physical objects, e.g. simulated rubber bands, using a graphical package like 'dia'. E.g. if a *star* is defined to be a polygon with alternating concave and convex vertices, what's the minimum number of pins required to hold a

stretched rubber band in a star shape? This and other problems are discussed in my slides on "toddler theorems" [36]. As discussed in [27] proofs generally work on *representations* of the subject matter rather than *instances* of the subject matter, which would not be possible for empirical knowledge.³ The paper also refutes the common assumption that analogical/pictorial representations are always *isomorphic* with what they represent.

(5) J.J. Gibson [9] drew attention to the biological importance of perceiving not only *structures*, but also *processes* (e.g. optical flow patterns), and being able to detect possibilities for action, and constraints on actions, in the environment: positive and negative affordances for the perceiver. That requires the ability to perceive the *possibility* of some processes and configurations and the *impossibility* of others [30]. Gibson considered only possibilities involving the perceiver's actions. However those are only a subset of a broader range of possibilities and constraints that can be perceived: they can be called "proto-affordances", since they involve possibilities and constraints on processes that need not involve actions, e.g. seeing that an object is too large to fit through a gap, whereas another could get through, as long as it is rotated during the process (e.g. if it is L-shaped). These are *potential* affordances, not *actual* affordances if you don't care about moving the object.



Figure 1. How many different action sequences can do the transformation?

There is a huge variety of ways in which using your right hand, left hand, both hands, teeth, a pair of pliers, tweezers, etc., you could rearrange the objects on the left to the configuration on the right. To do this a robot would need forms of representation that: (a) can be derived from sensory input (even low-resolution noisy input, as here); (b) have sufficient definiteness to allow goals to be expressed; (c) allow representation of effective multi-step plans to achieve goals; (d) are sufficiently abstract to allow effective reasoning about future possibilities and impossibilities (e.g. grasping the saucer in the middle, using finger and thumb is impossible), without considering all the continuously variable configurations and processes; and (e) allow plans to be used to control actions (using visual servoing, not necessarily ballistically). An infant cannot do all this but a normal child will eventually develop the ability to do it and later the ability to describe ways to do it. Byrne & Russon [4] showed that some other species seem to have partly similar abilities.

Figure 1 illustrates a vast array of possible processes and constraints on processes involving a small number of objects which can be moved from one configuration to another via a very large collection of different trajectories, using many possible modes of manipulation. Such transformations of 3-D configurations are achievable by children before they can talk, and also by other animals, as discussed in [37]. So the ability to perform them cannot depend on use of human language, though there must be an *internal* form of information processing (a "generalised language") supporting not only perception of what exists but also reasoning about what is and is not *possible* including representation of goals (non-existent states to be achieved) so that action sequences can be thought about and selected before they are performed, and their consequences *worked out*, using general capabilities, even if the task is a new one [29]. Physical trial-and-error is not always required, for toddlers, or other animals solving novel problems: they seem to be able to *work things out*: they are productively lazy. (That's not easy!) (6) A familiar example is stacking nested cups (facing up). At first, as

³ Instances can also function as representations, however!

Piagetians found, children succeed only by a procedure that appears to involve randomly selecting the next cup, and later, when a cup is found too big to insert, removing some of the previously inserted cups and trying again with the omitted cup. Eventually a child may notice that all the back-tracking involves wasted effort and time and go on to “debug” the procedure, until she works out that wasted moves can be avoided if the cups are selected in order of decreasing size: inserting the largest remaining cup at each stage.

A child who cannot *describe* that successful procedure, may nevertheless discover it. Perhaps she has *understood* but is unable to *say*: (a) ‘smaller than’ is transitive; (b) a cup cannot be inserted into a smaller cup; (c) if at some stage a cup X is inserted that is smaller than some other available cup Y, then (d) all the cups inserted after X will be smaller than X (from (a)) and therefore (e) all the cups inserted after X will be smaller than Y, and (f) Y cannot therefore be inserted after X – so that back-tracking will be required. A child who has somehow understood all that may be able to *work out* that always inserting the largest remaining cup will prevent the blockage occurring. A further step is understanding that preventing that blockage will suffice to get all the cups stacked. Yet another step is generalising the result to nesting the cups upside down – requiring a partly similar, but importantly different, strategy.

I doubt that most adults who understand the strategy can *explain* why the strategy works, unless they have had training in computer science or mathematical reasoning. *How* such understanding is implemented is another matter, requiring further research. It seems unlikely that a young child, or even most non-mathematical adults, can use reasoning in predicate calculus, or a computational process-algebra. Rather, I suspect the reasoning uses a form of representation that is more pictorial than logic, but more abstract than normal pictures – features required also for planning multi-step future actions. Compare [27, 10, 14, 45, 36]. Spelling out details remains a hard task for the future.

(7) There are very many examples involving counting and numbers. Much philosophy of mathematics assumes numbers are abstract entities named by numerals, but for a young child, as pointed out (independently) in [44] and [28, Ch 8], number words are primarily used for *doing* things, including reciting them, matching them and using them to perform tasks involving collections of objects (which later can include collections of number names). In addition to example (1) there are many more numerical discoveries that can be made empirically, then later understood as necessarily exceptionless, including the discovery that adding 3 more objects to a set of 4 objects gives the same “How many?” result as adding 4 more objects to a set of 3 objects; then later generalising this so as to understand that adding N objects to a set of M objects gives the same result as adding M objects to a set of N objects. One way to see this is to make the (non-trivial) discovery that if the objects are indistinguishable, and addition is done one object at a time, then adding the larger set to the smaller set necessarily goes through a stage which is the initial state of adding the smaller set to the larger set.⁴ (Contrast the analysis in [12].) Moving from the empirical view to the non-empirical view of these results requires a fairly sophisticated ability to attend to abstract features of *processes* as well as structures, namely parametrisable patterns that can be instantiated in different ways, while (necessarily) preserving some of their structure. This mathematical use of *abstraction* to bring different cases under the same specification, is often confused with use of *metaphor*. Metaphors cannot explain why certain things are possible and others

impossible, though they may suggest explanations.

(8) Links between counting abilities and abilities to manipulate and rearrange objects can be discovered when playing with blocks, e.g. the discovery that some sets of blocks can be arranged in rectangles, or in cubes, while others cannot. E.g. 5 or 7 blocks cannot be arranged as a rectangle, except a rectangle with only one row, whereas 8 or 9 blocks can be arranged in a rectangle (the latter a square). In contrast, 8, but not 9 blocks can be arranged to form a cube, and so on. In very unusual children this exploration might lead to invention of the concept of a prime number or even conjecturing the unique-factorisation theorem for natural numbers.

Additional examples were investigated by Piaget, though as far as I know Piagetians did not ask the questions posed here from an AI standpoint, including: What forms of representation does the learner need? What perceptual and reasoning mechanisms are required? What is the information processing architecture within which such processes function? What meta-cognitive mechanisms are needed? What ontology is required, and how are ontologies extended? Are all the components, and the architecture containing them, present from birth (which seems unlikely) or do they grow? If so, what mechanisms (forms of representation, architectures, etc.) are required to produce that growth and what are the roles of the genome and the environment in making such growth possible? The work of Karmiloff-Smith on representational redescription, e.g. [16], seems to be related, but as far as I can see it does not propose mechanisms.

2.2 Toddler Theorems

The preceding examples are a small subset from a potentially much larger collection of examples, based on informal observations of infants, toddlers and older children, as well as observations reported in developmental psychology and animal behaviour literature. One way to make progress in this field is to go on collecting, analysing, and organising, examples of what could be called “toddler theorems”, supported, where possible, by videos. Analysing the behaviours, from the viewpoint of a designer of robots that need to function and learn in 3-D environments of the sorts that humans and other animals can cope with, would include partially ordering the examples in terms of the information processing mechanisms that are required. That could be tested by building and demonstrating such mechanisms producing similar achievements, and analysing dependency relationships.

I suspect there are very many examples of the transition, based on “productive laziness”, from empirical discovery to knowledge about necessity of structural relations (in spatial configurations, in processes, etc.). Many of the examples will be about proto-affordances (possibilities for and constraints on change), vicarious affordances (possibilities and constraints for *other* agents) and epistemic affordances (how changes in the environment alter the information available to particular perceivers). Building a shared repository of examples may require overcoming the misguided prejudice among many psychologists that if your observations cannot be fed into a statistics package you are not doing science.

A large and varied collection of examples could be the basis of systematic investigation into requirements for robots able to go through similar (partially ordered) transitions – including requirements that have not been noticed by most robot designers (for instance researchers who over-emphasise the importance of embodiment, morphology, and dynamic coupling with the environment, criticised in [39]). Since animals with different morphologies seem to learn complex abilities to manipulate 3-D objects, and some seem to be able to solve non-trivial planning

⁴ An anonymous referee helpfully noted a flaw in an earlier formulation.

problems (as illustrated in [4]) it may be that what is *common* to animals with *different* morphologies, namely information-processing capabilities evolved in a 3-D environment, is more important for their intelligence (productive laziness!) than their morphology. Compare humans born limb-less, or blind, etc.

Much of human (and animal) intelligence is concerned not with real-time dynamic interaction with the environment (which is obviously important some of the time) but with *avoiding* unnecessary, time-wasting, dangerous, or inefficient interaction. It is no accident that much mathematical thinking about physical events and processes is done without actually producing those events and processes. As Kant noticed, development of those competences is triggered by experience, e.g. producing and observing physical processes in play and exploration, but what results is not a conclusion *derived* from or based on those observations but a new theoretical understanding, providing *explanations* of those observations. (Compare Popper's view of science, in [23].)

3 The Core Conjecture

A core conjecture, partially expounded in [5] is that the individual information processing architecture builds itself by adding new "layers" of competence, partly as a scientific community does.

3.1 Growing layers of competence

Diverse mathematical competences seem to be based on expanding layers of meta-cognitive competences that grow out of more common biological competences that originally evolved to support interactions with a complex, varied and changing 3-D world. These competences include: (a) the ability to detect the need to extend the current ontology substantively (as happens in deep science, e.g. adding "gene", and "valence" without defining them in terms of previous concepts [34]), (b) the ability to create new explanatory and predictive theories and to extend or modify those theories (e.g. by modifying the ontology, basic postulates and reasoning methods they use), (c) the ability to deploy theories, including noticing new opportunities for deployment, (d) the ability to detect and debug flaws in those theories, and (e) the ability to create new forms of representation, including new ways of manipulating representations that are usable in reasoning, explaining, predicting, and inventing.

So far, robots are nowhere near matching these competences. We don't yet understand the requirements well enough to specify the resources needed to support them: e.g. the forms of representation, information-processing mechanisms, and self-modifying multi-functional architectures. There has been much AI work on attempting to model aspects of mathematical competence, e.g. by Seely-Brown, Lenat, Bundy, Jamnik, Winterstein, Colton, Pease, and very recently [12], but that work starts with systems that manipulate abstract symbolic structures (e.g. logical formulae) rather than objects in the environment. Instead, we need to investigate how empirical knowledge is acquired through interactions with the environment, and later transformed to a non-empirical status, as suggested by Kant.

A good way to test Kant's view, is to try to design and build a robot that develops mathematical understanding as humans seem to do, initially by building on their understanding of structures and processes in 2-D and 3-D space using more general biological competences, and also reflecting on their own thinking about structures and processes, and then later reflecting on and talking about the results of their earlier reflection. Formal work in AI on theorem proving does not address these issues, because it starts at a much later stage of sophistication. Contrast [36].

3.2 Psychological Theories About Number Concepts

It is often supposed that the visual or auditory ability to distinguish groups with different numbers of elements (subitizing) displays an understanding of number. However this is simply a perceptual capability. A deeper understanding of numbers requires a much wider range of abilities. Rips *et al.* [24] rightly criticise theories that treat number concepts as abstracted from perception of groups of objects, and discuss alternative requirements for a child to have a concept of number, concluding that having a concept of number involves grasping (not necessarily consciously) a logical schema something like Peano's five axioms. They claim that that is what enables a child to work out various properties of numbers, e.g. the commutativity of addition, and the existence of indefinitely larger numbers. This requires logical capabilities to draw conclusions from the axioms, possibly unconsciously. How can children acquire such competences? They conclude that somehow the Peano schema and the logical competences are innately *built into* the child's "background architecture" (but do not specify how that could work).

Instead, I conjecture such competences are "meta-configured", i.e. not wholly determined in the genome, but produced through interactions with the environment that generate layers of meta-competences, competences that enable new competences to be acquired ([5, 40]. Some hints about how that might occur are presented below.

Many psychologists and researchers in animal cognition misguidedly search for experimental tests for whether a child or animal does or does not understand what numbers are.⁵ Rips *et al.* are not so committed to specifying an experimental test, but they do require a definition that makes a clear distinction between understanding and not understanding what numbers are.

3.3 An alternative approach.

As far as I know, no developmental psychologists have considered the alternative view, presented over 30 years ago in [28], chapter 8, that there is no single distinction between having and not having a concept of number, because learning about numbers involves a never-ending process that starts from relatively primitive and general competences that are not specifically mathematical and gradually adds more and more sophistication, in parallel with the development of other competences. Independently developed ideas in [44] also suggest that learning about numbers involves developing capabilities of the following sorts:

1. memorising an ordered sequence of arbitrary names;
2. performing a repetitive action;
3. performing two repetitive actions together, in synchrony;
4. initiating such processes, with different stopping conditions, depending on the task;
5. doing so when one process is uttering a learnt sequence of names;
6. learning rules for extending the sequence of names indefinitely;
7. observing various patterns in such processes and storing information about them, e.g. information about successors and predecessors of numerals, or results of counting onwards various amounts from particular numerals ([28], Chap 8);
8. noticing commonalities between static mappings and process mappings (e.g. paired objects vs paired events);
9. finding mappings between components of static structures as well as the temporal mappings between process-elements;

⁵ Compare the mistake of striving for a definitive test for whether animals of some species understand causation, criticised here in presentation 3: <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac>

10. noticing that such mappings have features that are independent of their order (e.g. counting a set of objects in two different orders must give the same result);
11. noticing that numbers themselves can be counted, e.g. the numbers between two specified numbers;
12. noticing possibilities of and constraints on rearrangements of groups of objects – e.g. some can be arranged as rectangular arrays, but not all;
13. learning to compare continuous quantities by dividing them into small components of a standard size and counting.
14. grasping (or inventing) the notion of a *limiting case*.

Such competences and knowledge can be extended indefinitely. Some can be internalised, e.g. counting silently. Documenting all the things that can be discovered about such structures and processes in the first few years of life could fill many pages. (Compare [26, 18, 36, 32].) The sub-abilities are useful in achieving practical goals by manipulating objects in the environment and learning good ways to plan and control such achievements, for example: fetching enough cups to give one each to a group of people, or matching heights of two columns made of bricks, to support a horizontal beam, or ensuring that enough water is in a big jug to fill all the glasses on the table.

Gifted teachers understand that any deep mathematical domain has to be explored from multiple directions, gaining structural insights and developing a variety of perceptual and thinking skills of ever increasing power: learning new constructs, new reasoning procedures, learning to detect deficiencies in understanding, and finding out how to remedy them. Our work and Wiese's [44] provide conjectures about some of the information-processing mechanisms involved. As far as I know, nobody has tried giving a robot such capabilities. Simplified versions should be attempted. It seems that no other known species has an information-processing architecture able to support all this – and neither does a newborn human infant. We need to understand not “core concepts” [41] so much as *core architecture-building mechanisms*.

3.4 Internal Construction Competences

The processes described above require the ability to create (a) new internal information structures, including, for example, structures recording predecessors of numbers, so that it is not necessary always to count up to N to find the predecessor of N , and (b) new algorithms for operating on those structures. As these internal information-structures grow, and algorithms for manipulating them are developed, there are increasingly many opportunities to discover more properties of numbers. The more you know, the more you can learn.

Moreover those constructions do not happen instantaneously or in an error-free process. Many steps are required including much self-debugging, as illustrated by Sussman's HACKER [42]. This depends on self-observation during performance of various tasks, including observations of both external actions and thinking. One form of debugging is what Sussman called detecting the need to create new “critics” that run in parallel with other activities and interrupt if some pattern is matched, for instance if disguised division by zero occurs.

The ongoing discovery of new invariant patterns in structures and processes produced when counting, arranging, sorting, or aligning sets of objects, leads to successive extensions of the learner's understanding of numbers. Initially this is just empirical exploration, but later a child may realise that the result of counting a fixed set of objects cannot depend on the order of counting. (Why not?) That invariance (a kind of transitivity) is intrinsic to the nature of one-to-one mappings and does not depend on properties of the things

being counted, or on how fast or how loud one counts, etc. However, some learners may never notice this non-empirical character of mathematical discoveries until they take a philosophy class!

3.5 Approaches to infinity

One of the non-empirical discoveries is that the sequence of natural numbers is never ending. Kant suggested that this requires grasping that a rule can go on being applied indefinitely, which fits the outline theory summarised above in 3.3. A child able to produce repetitive processes and simultaneously monitor and reflect on what is happening, could learn that there are repetitive processes of two kinds: those that start off with a determinate stopping condition that limits the number of repetitions and those that do not, though they can be stopped by an external process. Tapping a surface, walking, making the same noise repeatedly, swaying from side to side, repeatedly lifting an object and dropping it, are all examples of the latter type. (The idea of approaching a limiting value might grow out of such of meta-cognition.) This contrasts with the suggestion by Rips *et al.* [24] that a child somehow acquires logical axioms which state that every natural number has exactly one successor and at most one predecessor, and that the first number has no predecessor, from which it follows logically that there is no final number and the sequence of numbers never loops.

The general notion of something not occurring is clearly required for intelligent action in an environment. E.g. failure of an action to achieve its goal needs to be detectable. So if the learner has the concept of a repetitive process leading to an event that terminates the process, then the general notion of something not happening can be applied to that to generate the notion of something going on indefinitely. From there, depending on the information processing architecture and the forms of representation available, it may be a small step to the representation of two synchronised processes going on indefinitely, one of which is a counting process.

What is more sophisticated is acquiring a notion of a sequence of sounds or marks that can be generated indefinitely without ever repeating a previous mark. An obvious way to do that is to think of marks made up of one or more dots or strokes. Then the sequence could start with a single stroke, followed by two strokes, followed by three strokes, etc., e.g. | || ||| |||| ||||| etc. Such patterns grow large very quickly. That can motivate far more compact notations, like arabic numerals, though any infinitely generative notation will ultimately become physically unmanageable.

3.6 Backward to zero, and beyond

A different sort of extension is involved in adding *zero* to the natural numbers, which introduces “anomalies”, such as that there is no difference between *adding* zero apples and *subtracting* zero apples from a set of apples. A special meaning has to be invented for “multiplying by zero” – or it could be deemed meaningless.

Negative integers add further confusions. This extension is rarely taught properly, and as a result most people cannot give a coherent explanation of why multiplying two negative numbers should give a positive number. It cannot be *proved* on the basis of previous knowledge because what multiplying by a negative number means is undefined initially. For a mathematician, the *simplest* way to extend multiplication rules to include negative numbers without disruption of previous generalisations, is to stipulate that multiplying two negatives produces a positive. (Compare defining what 3^{-1} and 3^0 should mean.)

Some teachers use demonstrations based on the so-called “number line” to introduce notions of negative integers, but this can lead to serious muddles (e.g. about multiplication). Pamela Liebeck [19]

developed a game called “scores and forfeits” where players have two sets of tokens: addition of a red token is treated as equivalent to removal of a black token, and vice versa. Playing and discussing the game seemed to give children a deeper understanding of negative numbers and subtraction than standard ways of teaching, presumably because the set of pairs of natural numbers can be used to model accurately the set of positive and negative integers. Productive laziness can stimulate a search for good notations and procedures for manipulating them: a difficult task with high payoff.

3.7 Beyond discrete sets

Cardinality and orderings are properties of *discrete* sets. Extending the notion of number to include *measures* that are continuously variable, e.g. lengths, areas, volumes and time intervals, requires sophisticated extensions to the learner’s ontology and forms of representation – leading to deep mathematical and philosophical problems (e.g. Zeno’s paradoxes). I suspect that the development of such notions in children requires more intermediate stages than most investigators have noticed, including stages where there are only *partial orderings* of “amount” or “quantity” rather than global measures. For example, if P1 and P2 are polygons, and P2 can fit inside P1, then P2 is clearly smaller than P1. If P1 can fit inside P2, then clearly P2 is larger than P1. But if P1 is a circle and P2 is a square, and the diagonal of the square is slightly more than the diameter of the circle, then it can turn out that neither can fit inside the other. So deciding which is bigger (or has more area) requires invention of a new mode of comparison. Thinking of doing that by dividing both into many small squares of the same size and counting is both a major achievement and a source of problems because areas with curved boundaries, or with sides that have irrational ratios cannot be exactly tiled using squares. I expect few primary school teachers understand the intellectual challenges facing the child. Psychologists who investigate whether children have or have not grasped some notion of area, or volume, or conservation of amount, do not always probe for the intermediate stages. There is no requirement for all learning trajectories to be the same, when the terrain is so rich and complex: making rigidly sequenced curricula deeply flawed.

4 Problems Faced by Biological Evolution

4.1 Affordances, Visual Servoing, and Beyond

Analysis of biological requirements for vision (including human vision) enlarges our view of the functions of vision, requiring goals of AI vision researchers to be substantially expanded. An example is the role of vision in servo-control, including control of continuous motion as well as discrete condition-checking.

Gibson’s work on affordances in his [9] showed that animal vision provides information not merely about geometrical and physical features of entities in the environment, as in [2, 21], nor merely about recognising or categorising objects (the focus of much recent AI ‘vision’ research), but about what the perceiver can and cannot do, given its physical capabilities and its goals. I.e. visual information includes what processes *can and cannot occur in the environment* [30]. This requires an “exosomatic” ontology that is more concerned with what exists in the environment than with contents of retinal arrays. Gibson did not go far enough, as we have seen in discussing proto-affordances, vicarious affordances, etc., in 2.1 and 2.2.

4.2 Perception of Actual and Possible Processes

An EU FP6 cognitive robotics project CoSy,⁶ included analysis of

⁶ Described in <http://www.cs.bham.ac.uk/research/projects/cosy/>

requirements for a robot capable of perceiving and manipulating 3-D objects, including: (a) representing 3-D objects with parts and relationships; (b) multiple ontological layers in scene structures (as in chapter 9 of [28]); (c) representing “multi-strand relationships”, since both objects and parts of objects are related within and across object boundaries (e.g. faces, edges and vertices of a block); (d) representing “multi-strand processes” when multi-strand relationships change, e.g. with metrical, topological, causal, functional, continuous, and discrete changes occurring concurrently; (e) representing new *possible* multi-strand processes and constraints on such possibilities – the positive and negative “proto-affordances” that underlie affordances; (f) predicting outcomes of possible processes that are not currently occurring (e.g. when planning); (g) constructing possible explanations of how a perceived situation came about. Some representations of processes may use a partial simulation of the processes [6, 11]. Those biological mechanisms, not yet matched in AI, appear to be the basis of mechanisms capable of use in geometrical/topological reasoning.


As explained in [38, 31], that requires an *amodal, exosomatic* form of representation of processes; one that is not tied to the agent’s sensorimotor processes. That possibility is ignored by researchers who focus only on sensorimotor learning and representation, and assume all semantic content requires “symbol-grounding” [34].

4.3 The Importance of Kinds of Matter

A child, or robot, learning about kinds of process that can occur in the environment, needs to be able to extend the ontology she uses indefinitely, and not merely by defining new concepts in terms of old ones: there are also *substantive* ontology extensions (as in the history of physics and other sciences). For example, concepts of different kinds of matter are required to explain different sorts of deviations from rigidity: string and wire are flexible, but wire retains its shape after being deformed; an elastic band that returns to its original length after being stretched, need not restore its shape after bending. Some kinds of stuff easily separate into chunks in various ways, if pulled, e.g. mud, porridge, plasticine and paper. A subset of those allow restoration to a single object if separated parts are pressed together. There are also objects that are marks on other objects, like lines on paper, and there are some objects that can be used to produce such marks, like pencils and crayons. Marks produced in different ways and on different materials can have similar structures. Compare [26]. However, it is important to remember ([27]) that 2-D pictures/diagrams of 3-D structures/processes *cannot* be isomorphic with what they represent, as commonly assumed.

4.4 Perception and Mathematical Discovery

Many mathematical discoveries involve noticing an invariant in a class of processes. Mary Ensor, a mathematics teacher, once told me she had found a good way to teach children that the internal angles of a triangle add up to a straight line, as demonstrated in the figure.



An arrow starting at one corner, pointing along one side, can be rotated three times as indicated, going through positions shown in the three successive figures. The successive rotations *a*, *b* and *c* go through the interior angles of the triangle, and reverse the arrow, so they must add up to a straight line. This discovery may initially be made through empirical exploration with physical objects, but the

pattern involved does not depend on what the objects are made of, the colours used, lengths of lines, particular angles in the triangle, temperature, strength magnetic field, atmospheric pressure, etc. How such invariants are discovered, represented, and used, and which animals are capable of this is still unknown. But that capability lies at the heart of Kant's claims about the nature of mathematical knowledge.

4.6 Non-empirical Does Not Imply Infallible or Innate.

The processes are not infallible: mathematical discoveries can have “bugs” as Lakatos [17] demonstrated using the history of Euler's theorem about polyhedra. Further examples are in [36]. This does not imply that mathematical knowledge is empirical in the same way as knowledge about the physical properties of matter. Lakatos used the label “quasi-empirical”, but it needs to be explicated in terms of mechanisms required. Even logic is quasi empirical, insofar as logicians can produce buggy proofs and generate paradoxes requiring “debugging”. The kind of self-debugging demonstrated in [42] is very different from finding out more about what is in the physical environment. The discovery of bugs in proofs and good ways to deal with them is an important feature of mathematical learning as well as learning to program. Ensor's proof (in 4.4) would break down on a spherical surface, for example, since the triangle could have three 90 degree angles (among other possibilities inconsistent with the original result). Noticing this might lead a learner to investigate properties of triangles that distinguish planar and non-planar surfaces. But that exploration does not *require* experiments in a physical laboratory, though it may benefit from them. Kant claimed that such discoveries are about the perceiver's forms of perception, but they are not restricted to any particular perceivers. Making Kant's claim more precise, in the form of a working design for a robot learner remains a challenge.

4.7 Humean and Kantian Causation

The forms of representation required for counting and matching groups of entities or parts of processes, are different from those required for understanding topology and Euclidean geometry. The latter seem to build on reasoning/planning competences based on visual competences, illustrated in [26]. Some of these competences are apparently shared with some other animals – those that are capable of planning and executing novel spatial actions [4]. Adding properties of matter, such as rigidity and impenetrability, to representations of shape and topology allows additional reasoning about and prediction of results of processes. An example is the ability to use the fact that two meshed gear wheels are rigid and impenetrable to *work out*, in advance of experimenting, how rotating one will cause the other to rotate. That kind of reasoning about processes in the environment is not always available.

If the wheels are not meshed, but there are inaccessible, hidden mechanisms, then the only basis for predicting the consequence of rotating the wheels is to use a Humean (statistical, Bayesian) notion of causation: basing predictions of results of actions or events solely on observed correlations. In contrast, where the relevant structure and constraints are visible, or already known, then mathematical reasoning (using geometry and topology) may be used. I believe this underlies Kant's anti-Humean conception of structure-based, deterministic, causation.⁷ A significant subset of the causal understanding of the environment that a child acquires (though not all) is Kantian because it allows the consequences of novel processes to be *worked out* on the basis of geometric and

topological relationships, and kinds of matter involved. But for this, much engineering design would be impossible.

4.5 Doing Philosophy by Doing AI

I suspect that many philosophical disputes about the nature of mathematics will be transformed if we can build not a standard theorem proving machine but a robot that works out how to achieve “productive laziness” by doing spatial/geometrical reasoning. Many researchers (e.g. those in [10]) have pointed out the need to provide intelligent machines with spatial forms of representation and reasoning, but progress in replicating animal visual/spatial abilities discussed above has been very slow. In part, this is because the requirements for human-like (or animal-like) visual systems have not been analysed in sufficient depth (as illustrated in [31, 38]). There is a vast amount of research on “object recognition” that contributes nothing to our understanding of how 3-D spatial structures and processes are seen or how information about spatial structures and processes is used, for instance in reasoning and acting. Differences between perception of 3-D structures and proto-affordances, and mere recognition are discussed in [35]. The deeper competences appear to be required to support Kant's philosophy of mathematics.

5 Individual and social processes

5.1 Is Human Mathematical Competence Unique?

I have tried to present a view of products of biological evolution used for interacting with a complex spatio-temporal environment whose structure allows some things to be *worked out* by an individual instead of having to be learnt empirically, even though the working out is sometimes preceded by empirical discovery of the result, followed by a re-interpretation at a later stage. Further work is required to understand both the evolutionary pressures and what exactly those pressures have produced in humans and other animals (e.g. corvids, elephants, dolphins, octopuses, dogs, and others).

The evidence considered here suggests that the *long term potential* to learn about even fairly advanced mathematics is in most newborn human infants (though there may also be genetic individual differences which can only be understood in the light of future, more detailed theories of the mechanisms involved). Moreover, the existence of the potential does not mean that that potential can be realised in a normal lifetime without additional help. One product of evolution that is not unique to humans, but is supported by especially powerful mechanisms in humans, is the ability of individuals to speed up their learning through guidance and scaffolding based on learning done by predecessors. This makes cultural evolution, and cultural learning possible. It is facilitated by toys, games, development of tools, shelters, clothing, collaborative experiences with acquiring and consuming food, etc. (The explanatory power of imitation can be over-rated if the mechanisms required for imitation are ignored.)

Many animals, e.g. dogs, parrots, various primates, elephants, and others, are able to acquire unusual cognitive competences as a result of domestication and training by humans, but are not able to pass on what they have learnt to others, although there are other things that they teach their young. This means that humans in groups can share the process of acquiring new knowledge and competences and passing them on with much faster progress between generations than other species. A well-known spectacular example was the spontaneous development of a novel sign language by a group of deaf children in Nicaragua, but human history is full of examples of collaborative advance, usually spread out over longer time periods and with a large age spread in the communication from experts to learners. That applies also to mathematical advances. So although the

⁷ See also <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac>

products of biological evolution discussed in previous sections are necessary for development of mathematics, they do not suffice for the development of advanced and systematised mathematical knowledge in a human lifetime: that requires cultural development. It also seems to be needed for the growth of meta-mathematical and philosophical reflection on what has been learnt by doing mathematics. Perhaps future robots, or a future animal species, will not need that cultural support for rapid mathematical or philosophical development – if they can internalise the feedback loops.

6 Conclusion

I have sketched a research programme motivated (for me, though not necessarily for others) by the aim of explaining the nature of mathematical discovery, and supporting major features of Kant's philosophy of mathematics. We need further investigation of architectures and forms of representation that allow playful exploration by a robot, leading to development of "framework theories" that can support the discovery of patterns and invariants in structures and processes that are needed for development of deep mathematical concepts and mathematical forms of reasoning. The robot should be able to go through the following stages:

1. Acquiring familiarity with some domain, e.g. through playful exploration;
2. Noticing (empirically) certain generalisations;
3. Discovering a way of thinking about them that shows they are not empirical;
4. Generalising, diversifying, debugging, deploying, that knowledge;
5. Formalising the knowledge, possibly in more than one way.

This will be based on mechanisms for building powerful explanatory theories about the environment that seem to have evolved in other species, and are used (unwittingly) by infants and toddlers before they start discovering what I have called "toddler theorems".

Although this work emphasises embodiment, because spatial embodiment presents both the problems that need to be solved (including problems of manipulating, route-finding, building things, eating food, avoiding being eaten, finding shelter, etc.), it does not emphasise embodiment in the way that much current research in neuroscience and AI/Robotics does e.g. [20, 3], namely by focusing on the detailed morphology of a species or the *real-time dynamic interaction with the environment*. That is certainly a problem for all organisms. What is special about the organisms we regard as most intelligent (including humans, primates, some birds, octopuses) seems to be their relatively unusual ability to use an understanding of the structures and potential in the environment for purposes of deliberation rather than action, e.g. selecting goals, and creating plans, as well as carrying out plans. They can control their behaviour in a productively lazy way: not driven simply by the physics of the environment and their own morphology (like a ball going down a helter-skelter), but in part by using information acquired by perception of structures to reason (mathematically) about structures and future possible and impossible processes [39].

I have so far described only a subset of the requirements for working designs. Some more detailed requirements are in cited papers. But I don't claim that we understand all the requirements yet. It is clear that AI still has a long way to go before the visual and cognitive mechanisms of robots can match the development of a typical human child going through infant and toddler learning, but it is not clear exactly what the gaps in our understanding are.

A possible (fanciful?) outcome of this symposium might be creation of an informal multidisciplinary network for collaboration

on refining the requirements, and collecting more evidence regarding the fine structure of individual development of the kind discussed here. This could feed into a collaborative project to produce a working prototype system as a proof of concept, using a simulated robot, perhaps one that manipulates 2-D shapes in a plane surface, discovering properties of various kinds of interactions, involving objects with different shapes made of substances with various properties that determine the consequences of the interactions, e.g. impenetrability, rigidity, elasticity, etc. A later stage would add reflective capabilities enabling the robot to discover the differences between what it can work out and what it merely discovers about its world empirically. Much later it might become a philosopher of mathematics able, unlike us, to reflect on known features of its own design to explain its capabilities, fulfilling the dream (nightmare for some) expressed in my 1978 book [28] of basing Kant's epistemology and philosophy of mathematics on a working model.

Acknowledgements

I have benefited from discussions with too many people to list here, including colleagues in philosophy, psychology, biology, AI and computer science. Referees for the symposium kindly made several useful comments. Some of these ideas were developed while working on two EU-funded cognitive robotics projects, CoSy and its successor CogX.

REFERENCES

- [1] Jody Azzouni, 'The compulsion to believe: Logical inference and normativity', *ProtoSociology*, **25**, 69–88, (2008).
- [2] H.G. Barrow and J.M. Tenenbaum, 'Recovering intrinsic scene characteristics from images', in *Computer Vision Systems*, eds., A. Hanson and E. Riseman, 3–26, Academic Press, New York, (1978).
- [3] Alain Berthoz, *The Brain's sense of movement*, Perspectives in Cognitive Science, Harvard University Press, London, UK, 2000.
- [4] R W Byrne and A Russon, 'Learning by imitation: a hierarchical approach', *Behavioral and Brain Sciences*, **21**, 667–721, (1998).
- [5] Jackie Chappell and Aaron Sloman, 'Natural and artificial meta-configured altricial information-processing systems', *International Journal of Unconventional Computing*, **3**(3), 211–239, (2007). <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>.
- [6] K. Craik, *The Nature of Explanation*, Cambridge University Press, London, New York, 1943.
- [7] R.P. Feynman, *The Character of Physical Law*, MIT Press, Cambridge MA, 1965. The 1964 Messenger Lectures.
- [8] Eleanor J. Gibson and Anne D. Pick, *An Ecological Approach to Perceptual Learning and Development*, Oxford University Press, New York, 2000.
- [9] J. J. Gibson, *The Ecological Approach to Visual Perception*, Houghton Mifflin, Boston, MA, 1979.
- [10] *Diagrammatic Reasoning: Computational and Cognitive Perspectives*, eds., J. Glasgow, H. Narayanan, and B. Chandrasekaran, MIT Press, Cambridge, MA, 1995.
- [11] R. Grush, 'The emulation theory of representation: Motor control, imagery, and perception', *Behavioral and Brain Sciences*, **27**, 377–442, (2004).
- [12] Markus Guhe, Alison Pease, and Alan Smaill, 'A Cognitive Model of Discovering Commutativity', in *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, eds., N. A. Taatgen and H. van Rijn, pp. 727–732, Austin, TX, (2009). Cognitive Science Society.
- [13] Eva Jablonka and Marion J. Lamb, *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*, MIT Press, Cambridge MA, 2005.
- [14] M. Jamnik, A. Bundy, and I. Green, 'On automating diagrammatic proofs of arithmetic arguments', *Journal of Logic, Language and Information*, **8**(3), 297–321, (1999).
- [15] I. Kant, *Critique of Pure Reason*, Macmillan, London, 1781. Translated (1929) by Norman Kemp Smith.

- [16] A. Karmiloff-Smith, 'Precis of: Beyond Modularity: a Developmental Perspective on Cognitive Science', *Behavioral and Brain Sciences*, **17**(4), 693–706, (1994). <http://www.bbsonline.org/documents/a/00/00/05/33/index.html>.
- [17] I. Lakatos, *Proofs and Refutations*, Cambridge University Press, Cambridge, UK, 1976.
- [18] Pamela Liebeck, *How Children Learn Mathematics: A Guide for Parents and Teachers*, Penguin Books, Harmondsworth, 1984.
- [19] Pamela Liebeck, 'Scores and Forfeits: An Intuitive Model for Integer Arithmetic', *Educational Studies in Mathematics*, **21**(3), 221–239, (Jun 1990).
- [20] M. Lungarella and O. Sporns, 'Mapping information flow in sensorimotor networks', *PLoS Computational Biology*, **2**(10:e144), (2006). DOI: 10.1371/journal.pcbi.0020144.
- [21] D. Marr, *Vision*, W.H. Freeman, San Francisco, 1982.
- [22] J. S. Mill, *A System of Logic, Ratiocinative and Inductive: Being a connected view of the principles of evidence and the methods of scientific investigation*, John W. Parker, London, 1843.
- [23] K.R. Popper, *Objective Knowledge*, Oxford University Press, Oxford, 1972.
- [24] L. J. Rips, A. Bloomfield, and J. Asmuth, 'From Numerical Concepts to Concepts of Number', *The Behavioral and Brain Sciences*, **31**(6), 623–642, (2008). doi:10.1017/S0140525X08005566.
- [25] B. Russell, *Mysticism and Logic and Other Essays*, Allen & Unwin, London, 1917.
- [26] J. Sauvy and S. Suavy, *The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology*, Penguin Education, Harmondsworth, 1974. Translated from the French by Pam Wells.
- [27] A. Sloman, 'Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence', in *Proc 2nd IJCAI*, pp. 209–226, London, (1971). William Kaufmann. <http://www.cs.bham.ac.uk/research/cogaff/04.html#200407>.
- [28] A. Sloman, *The Computer Revolution in Philosophy*, Harvester Press (and Humanities Press), Hassocks, Sussex, 1978. <http://www.cs.bham.ac.uk/research/cogaff/crp>.
- [29] A. Sloman, 'The primacy of non-communicative language', in *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference, Oxford, March 1979*, eds., M. MacCafferty and K. Gray, pp. 1–15, London, (1979). Aslib. <http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43>.
- [30] A. Sloman, 'Actual possibilities', in *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR '96)*, eds., L.C. Aiello and S.C. Shapiro, pp. 627–638, Boston, MA, (1996). Morgan Kaufmann Publishers. <http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15>.
- [31] A. Sloman, 'Putting the Pieces Together Again', in *Cambridge Handbook on Computational Psychology*, ed., Ron Sun, chapter 26, 684–709, Cambridge University Press, New York, (2008). <http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#710>.
- [32] A. Sloman. Ontologies for baby animals and robots. From "baby stuff" to the world of adult science: Developmental AI from a Kantian viewpoint., 2009. Online tutorial presentation: <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#brown>.
- [33] A. Sloman and J. Chappell, 'The Altricial-Precocial Spectrum for Robots', in *Proceedings IJCAI'05*, pp. 1187–1192, Edinburgh, (2005). IJCAI. <http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>.
- [34] Aaron Sloman. Why symbol-grounding is both impossible and unnecessary, and why theory-tethering is more powerful anyway., 2007. <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models>.
- [35] Aaron Sloman. A Multi-picture Challenge for Theories of Vision, 2008. <http://www.cs.bham.ac.uk/research/projects/cogaff/misc/multi-pic-challenge.pdf>.
- [36] Aaron Sloman. A New Approach to Philosophy of Mathematics: Design a young explorer, able to discover "toddler theorems", 2008. Online presentation <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#toddler>.
- [37] Aaron Sloman. Evolution of minds and languages. What evolved first and develops first in children: Languages for communicating, or languages for thinking (Generalised Languages: GLs)?, 2008. <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0702>.
- [38] Aaron Sloman, 'Architectural and representational requirements for seeing processes and affordances', in *Computational Modelling in Behavioural Neuroscience: Closing the gap between neurophysiology and behaviour*, eds., D. Heinke and E. Mavritsaki, 303–331, Psychology Press, London, (2009). <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0801>.
- [39] Aaron Sloman, 'Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress', in *Creating Brain-like Intelligence*, eds., B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, and K. Doya, 248–277, Springer-Verlag, Berlin, (2009). <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804>.
- [40] Aaron Sloman and Jackie Chappell, 'Computational Cognitive Epigenetics (Commentary on [13])', *Behavioral and Brain Sciences*, **30**(4), 375–6, (2007). <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703>.
- [41] Elizabeth Spelke, 'Core Knowledge', *American Psychologist*, 1233–1243, (November 2000).
- [42] G.J. Sussman, *A computational model of skill acquisition*, American Elsevier, San Francisco, CA, 1975. <http://dspace.mit.edu/handle/1721.1/6894>.
- [43] A. N. Whitehead and B. Russell, *Principia Mathematica Vols I – III*, Cambridge University Press, Cambridge, 1910–1913.
- [44] Heike Wiese, 'The Co-Evolution of Number Concepts and Counting Words', *Lingua*, **117**, 758–772, (2007).
- [45] D. Winterstein, *Using Diagrammatic Reasoning for Theorem Proving in a Continuous Domain*, Ph.D. dissertation, University of Edinburgh, School of Informatics., 2005. <http://www.era.lib.ed.ac.uk/handle/1842/642>.