

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

Aaron Sloman¹

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

Contents

1	Mathematics as a Biological Phenomenon	1
2	Russell vs Feynman on Mathematics	2
3	Examples of “Productive Laziness”	2
4	Toddler Theorems	4
5	Psychological Theories About Number Concepts	5
6	Internal Construction Competences	6
6.1	Approaches to infinity	6
7	Doing Philosophy by Doing AI	7
8	Affordances, Visual Servoing, and Beyond	7
9	Perception of Actual and Possible Processes	7
10	The Importance of Kinds of Matter	7
11	Perception and Mathematical Discovery	7
12	Humean and Kantian Causation	8
13	Individual and social processes	8
14	Conclusion	8
	References	9

Abstract. One way to defend Kant’s philosophy of mathematics is to show how a “baby” human-like robot could develop mathematical competences on the basis of biological competences that evolved to support interactions with a complex, varied and changing 3-D world. Current techniques in AI/Robotics are nowhere near that. This paper helps to characterise the 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”. It is conjectured that mathematical competences grow out of those competences. 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. This may turn out to be a special case of a very general phenomenon involved in so-called “U-shaped” learning, including the language learning that comes later than some of the learning discussed here. There is still much hard work to be done.

1 Mathematics as a Biological Phenomenon

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 that ignores the existence of other mammals born with large heads, e.g. elephants. Another 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. One version of that suggestion is that humans are born with a collection of innate modules that control various aspects of behaviour including a general purposes learning module.

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 meta-competences for acquiring competences (e.g. those required for manipulation using independently movable manipulators) that involve such complex and varied perceptual and motor processes that they cannot be encoded in a genome (possibly for reasons suggested in (Sloman & Chappell, 2005)). For instance, some behaviours need to be produced by sophisticated run-time control competences that have to be learnt, using specialised meta-competences that also develop under the influence of the environment: we learn how to learn many different things. (Chappell & Sloman, 2007) 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. (Spelke, 2000). This proposal is related to the philosophy of mathematics of Immanuel Kant (Kant, 1781), according to whom mathematics is neither empirical (as suggested by Mill), nor merely about “relations between ideas” (as suggested

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

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” (Azzouni, 2008). 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* to be discovered and used in creative problem-solving (Craik, 1943). 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 novel practical problems, avoiding empirical experiments that could be dangerous, or even fatal. However, the mechanisms are not infallible, so that testing and debugging may be required.

The central claim is that these products of human evolution can and do develop into what we understand as mathematical competences. (Examples are given below.) 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. Many of species-specific details are of psychological interest, however.

2 Russell vs Feynman on Mathematics

Russell and Whitehead (1910–1913) 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* (Russell, 1917)).

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” (Feynman, 1965). 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 (Slovan, 1978, Ch 6). Despite the role of experience in such learning, the *results* of such learning are not empirical generalisations. Kant wrote, in (Kant, 1781), “though all our knowledge begins with experience, it by no means follows that all arises out of experience.” Feynman seems to agree.

I suspect that some of the re-conceptualisation involved in developing mathematical insight uses mechanisms related to the more specialised mechanisms for language learning that allow 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 set of 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 would have 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.

This possibility does not seem to be widely recognised. For example, many examples of learning about affordances in young children are described in (E. J. Gibson & Pick, 2000). 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. How is it possible to avoid that need?

3 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 becomes 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? That would be very 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.)

(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

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

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 learns 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 two boundaries are aligned and the picture piece is slightly smaller than the recess then all parts of the piece can simultaneously move in without obstruction from the material surrounding the recess. Any translation before insertion will cause overlap that prevents insertion. What is not so obvious is that unless the recess and picture piece are both circular, most *rotations* will also 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.

(5) J.J. Gibson (1979) 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 (Sloman, 1996). 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 for someone who has no interest in moving the object.

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 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 therefore the ability to perform them cannot depend on use of human language, though there must be an *internal* form of information processing 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




Figure 1. *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 (1998) show that some other species seem to have partly similar abilities.*

before they are performed, and their consequences *worked out*, using general capabilities, even if the task is a new one. Physical trial-and-error is not always required.

(6) A familiar example is stacking nested cups (facing up). At first, as 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 (Sloman, 1971; Glasgow, Narayanan, & Chandrasekaran, 1995; Jamnik, Bundy, & Green, 1999; Winterstein, 2005). 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 (Wiese, 2007) and (Sloman, 1978, 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 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. (Contrast the analysis in (Guhe, Pease, & Smaill, 2009).) 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.

(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.

Many more examples including string, elastic, pins, stones, mazes, and marks on paper are described in (Sauvy & Suavy, 1974). Some of their experiments can be performed on *simulated* physical objects, e.g. simulated rubber bands, using a graphical package like ‘dia’. This illustrates a point made in (Sloman, 1971) namely that some proofs work on *representations* of the subject matter rather than *instances* of the subject matter, which would not be possible for empirical knowledge. (However, as pointed out in that paper analogical/pictorial representations are often *not isomorphic* with what they represent, as commonly assumed.)

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. (Karmiloff-Smith, 1994), seems to be related, but as far as I can see it does not propose mechanisms.

4 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 (Sloman, 2009b)). 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 problems (as illustrated in (Byrne & Russon, 1998)) 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.

One of the reasons for this importance, is that 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, summarised in (Popper, 1972).)

Conjecture: diverse mathematical competences are 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, i.e. non-definitionally), (b) the ability to create new explanatory and predictive theories and to extend or modify those theories (e.g. by modifying the ontology they use, as well as basic postulates and reasoning methods), (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, but there is no known reason why they could not do so in future – once we understand the competences well enough to specify the computational resources required to support them: forms of representation, information-processing mechanisms, and self-modifying multi-functional architectures. Although 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 (Guhe et al., 2009), such work starts with systems that manipulate abstract structures rather than objects in the environment. For the current project, we need to investigate how empirical knowledge is acquired by such interactions, and later transformed to a non-empirical status.

I suggest that the best way to test Kant's view of the nature of mathematics, is to try to show how 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 all starts at a much later stage of sophistication, where much knowledge has already been formalised and reasoning methods are (usually) based on logic, though those techniques may later play an important role in this project.³

5 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 wider range of abilities, some of them already mentioned. Rips *et al.* (Rips, Bloomfield, & Asmuth, 2008) 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 having some understanding (not necessarily consciously) of 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 implies that such children have the logical capabilities required to draw conclusions from the axioms, though not necessarily consciously. That immediately raises the question how children can 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).

An alternative possibility was presented in (Chappell & Sloman, 2007; Sloman & Chappell, 2007) according to which such competences are meta-configured, i.e. not 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). 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.⁴ 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.

As far as I know, no developmental psychologists have considered the alternative view, presented over 30 years ago in (Sloman, 1978), 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 (Wiese, 2007) also suggest that learning about numbers involves developing capabilities of the following sorts:

1. performing a repetitive action;
2. memorising an ordered sequence of arbitrary names;
3. performing two repetitive actions together and keeping them in synchrony;
4. initiating such a process and then being able to use different stopping conditions, depending on the task;
5. doing all this when one of the actions 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;
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;
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.

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 (Liebeck, 1984).) The sub-abilities involved in these processes are useful in achieving practical goals by manipulating objects in the environment and learning good ways to plan and control such achievements. An example might be fetching enough cups to give one each to a group of people, or matching heights of two columns made of bricks, to

³ See also this PDF presentation on whether a baby robot can grow up to be a mathematician
<http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#math-robot>

⁴ 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>

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 is something that has to be explored from multiple directions, gaining structural insights and developing a variety of perceptual and thinking skills of ever increasing power. That includes learning new constructs, new reasoning procedures, learning to detect partial or erroneous understanding, and finding out how to remedy such deficiencies. The work of Wiese and Sloman presents some 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 achievable soon.

6 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 in (Sussman, 1975). This depends on self-observation during performance of various tasks, including observations of external actions and of 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!

6.1 Approaches to infinity

One of the non-empirical discoveries is that the natural numbers form an infinite set. Kant suggested that this requires grasping that a rule can go on being applied indefinitely. This contrasts with the suggestion by Rips *et al.* (Rips *et al.*, 2008) 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. Instead, a child 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 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.

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 mathematicians, it is defined by the rules for multiplying negative numbers, and the *simplest* way to extend multiplication rules to 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 (Liebeck, 1990) 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. (Multiplication was not included.) The person with the biggest surplus of black over red wins. Giving a player both a red and a black token, or removing both a red and a black token makes no difference to the total status of the player. Playing, and especially 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.

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). In humans, an understanding of Euclidean geometry and topology seems to build on reasoning/planning competences combined with visual competences, illustrated in (Sauvy & Suavy, 1974). This requires different forms of representation from counting and matching groups of entities. Some of these competences are apparently shared with some other animals – those that are capable of planning and executing novel

spatial actions (Byrne & Russon, 1998).

7 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 discovers ways of achieving “productive laziness”, for instance by doing spatial/geometrical reasoning. Many others (e.g. (Glasgow et al., 1995)) have pointed out the need to provide intelligent machines with spatial forms of representation and reasoning, but progress in replicating animal abilities 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 (Sloman, 2009a, 2008)). E.g. 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.⁵

8 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 (1979) showed that animal vision provides information not merely about geometrical and physical features of entities in the environment, as in (Barrow & Tenenbaum, 1978; Marr, 1982), nor 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. vision needs to provide information not only about actual objects, structures and movements, but also what processes *can and cannot occur in the environment* (Sloman, 1996). In order to do this, the visual system must use an ontology that is only very indirectly related to retinal arrays. Gibson did not go far enough, as we have seen in discussing proto-affordances, vicarious affordances, etc.

9 Perception of Actual and Possible Processes

Previous work in a cognitive robotics project included analysis of requirements for a robot capable of manipulating 3-D objects, e.g. grasping them, moving them, and constructing assemblages, possibly while other things were happening. Analysis of the requirements revealed (a) the need for representing scene objects with parts and relationships (as everyone already knew), (b) the need for several ontological layers in scene structures (as in chapter 9 of (Sloman, 1978)), (c) the need to represent “multi-strand relationships” because not only whole objects but also parts of different objects are related in various ways, (d) the need to represent “multi-strand processes”, because when things move the multi-strand relationships change, e.g. with metrical, topological, causal, functional, continuous, and discrete changes occurring concurrently, and (e) the need to represent *possible* multi-strand processes and constraints on such

⁵ Some of the differences between recognition and perception of 3-D structure are illustrated in <http://www.cs.bham.ac.uk/research/projects/cogaff/challenge.pdf>

possibilities – the positive and negative “proto-affordances” that are the substratum of affordances.

As explained in (Sloman, 2009a, 2008), that requires an *amodal, exosomatic* form of representation of processes; one that is not tied to the agent’s sensorimotor processes. That possibly is ignored by researchers who focus only on sensorimotor learning and representation, and base all semantics on “symbol-grounding”.


The ability to perceive a multi-strand process requires the ability to have internal representations of the various discretely and continuously concurrently changing relationships. Some of the mechanisms used for perceiving multi-strand processes can also be used both to predict outcomes of possible processes that are not currently occurring (e.g. when planning), and to explain how a perceived situation came about. Both may use a partial simulation of the processes (Craik, 1943; Grush, 2004). Our conjecture is that those biological mechanisms, not yet matched in AI, evolved/grew into mechanisms capable of use in geometrical/topological reasoning.

10 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, whereas many perceived processes involve objects that preserve all their metrical relationships, there are also many deviations from such rigidity, and concepts of different kinds of matter are required to explain those deviations: string and wire are flexible, but wire retains its shape after being deformed; an elastic band returns to its original length after being stretched, but does 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 (Sauvy & Suavy, 1974). However, it is important to remember ((Sloman, 1971)) that 2-D pictures/diagrams of 3-D structures/processes *cannot possibly* be isomorphic with what they represent, a common incorrect assumption.

11 Perception and Mathematical Discovery

I have argued that many mathematical discoveries involve noticing an invariant in a class of processes. For example, 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, demonstrated in the figure.



Imagine an arrow starting at one corner, pointing along one side. It 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 clearly does not depend on what the objects are made of and changing conditions such as colours used, lengths of lines, particular angles in the triangle, temperature, strength of gravitational or magnetic field cannot affect the property of the process.

However, such mathematical discoveries can have “bugs” as Lakatos (Lakatos, 1976) demonstrated using the history of Euler’s theorem about polyhedra. That is sometimes wrongly taken to 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. Compare the self-debugging in (Sussman, 1975).) The discovery of bugs in proofs and good ways to deal with them is an important feature of mathematical learning. Ensor’s proof 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.

12 Humean and Kantian Causation

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 is not always available.

If the wheels are not meshed, but there are hidden connections, then the only basis for predicting the consequence of rotating the wheels is to use a Humean (statistical) notion of causation: basing predictions of results of actions or events solely on observed correlations. In contrast, where the relevant structure and constraints are known, and mathematical proofs (using geometry and topology) are possible, Kant’s notion of causation, which is structure-based and deterministic, can be used. 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.⁶ But for this, much engineering design would be impossible.

13 Individual and social processes

I have tried to present a view of products of biological evolution concerned with 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

(as shown by research on primates, corvids, elephants, dolphins, octopuses, dogs, and others).

The evidence considered here suggests that the *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 detailed theories of the mechanisms involved – not yet available). 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 under the guidance and scaffolding of results of learning of those who have come before. 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. (I think the role of imitation is sometimes grossly over-rated. Another common error, which there is no space to discuss here is to confuse processes of creative abstraction with use of metaphorical reasoning. Metaphors cannot produce an grasp of why certain things are possible and others impossible.)

Many animals seem to be 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 from humans to others, e.g. dogs, parrots, various primates, elephants, 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 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. Perhaps future robots, or a future animal species, will not need that cultural support for rapid mathematical development.

14 Conclusion

We need further investigation of architectures and forms of representation that allow playful exploration by a robot to produce discoveries of patterns in structures and processes that are the basic mechanisms 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.

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

⁶ <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac>

not emphasise embodiment in the way that much current research in neuroscience and AI/Robotics does e.g. (Lungarella & Sporns, 2006; Berthoz, 2000), namely by focusing on 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 structure of the environment to select goals, to create plans, to carry out plans, and to control their behaviour in a productively lazy way: not being driven simply by the physics of the environment and their own morphology, but in part by using information acquired by perception of structures to reasoning about structures and future possible and impossible processes (Sloman, 2009b).

This paper merely reports on a subset of the requirements for working designs. Some more detailed requirements are in cited papers. 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. There is still a great deal more to be done, and meeting all the requirements will not be easy.

If others are interested in this project, perhaps it would be useful to set up an informal multidisciplinary network for collaboration on refining the requirements and then producing 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.

Perhaps one day, a team of robot mathematicians, philosophers of mathematics and designers will also be able implement such systems. Don't expect that soon.

Acknowledgements

[To be added]

References

- Azzouni, J. (2008). The compulsion to believe: Logical inference and normativity. *ProtoSociology*, 25, 69–88.
- Barrow, H., & Tenenbaum, J. (1978). Recovering intrinsic scene characteristics from images. In A. Hanson & E. Riseman (Eds.), *Computer vision systems* (pp. 3–26). New York: Academic Press.
- Berthoz, A. (2000). *The brain's sense of movement*. London, UK: Harvard University Press.
- Byrne, R. W., & Russon, A. (1998). Learning by imitation: a hierarchical approach. *Behavioral and Brain Sciences*, 21, 667–721.
- Chappell, J., & Sloman, A. (2007). Natural and artificial meta-configured altricial information-processing systems. *International Journal of Unconventional Computing*, 3(3), 211–239. (<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0609>)
- Craik, K. (1943). *The nature of explanation*. London, New York: Cambridge University Press.
- Feynman, R. (1965). *The Character of Physical Law*. Cambridge MA: MIT Press. (The 1964 Messenger Lectures)
- Gibson, E. J., & Pick, A. D. (2000). *An Ecological Approach to Perceptual Learning and Development*. New York: Oxford University Press.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
- Glasgow, J., Narayanan, H., & Chandrasekaran, B. (Eds.). (1995). *Diagrammatic reasoning: Computational and cognitive perspectives*. Cambridge, MA: MIT Press.
- Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, 27, 377–442.
- Guhe, M., Pease, A., & Smail, A. (2009). A Cognitive Model of Discovering Commutativity. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 727–732). Austin, TX: Cognitive Science Society.
- Jablonka, E., & Lamb, M. J. (2005). *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life*. Cambridge MA: MIT Press.
- Jamnik, M., Bundy, A., & Green, I. (1999). On automating diagrammatic proofs of arithmetic arguments. *Journal of Logic, Language and Information*, 8(3), 297–321.
- Kant, I. (1781). *Critique of pure reason*. London: Macmillan. (Translated (1929) by Norman Kemp Smith)
- Karmiloff-Smith, A. (1994). Precis of: Beyond Modularity: a Developmental Perspective on Cognitive Science. *Behavioral and Brain Sciences*, 17(4), 693–706. (<http://www.bbsonline.org/documents/a/00/00/05/33/index.html>)
- Lakatos, I. (1976). *Proofs and Refutations*. Cambridge, UK: Cambridge University Press.
- Liebeck, P. (1984). *How Children Learn Mathematics: A Guide for Parents and Teachers*. Harmondsworth: Penguin Books.
- Liebeck, P. (1990, Jun). Scores and Forfeits: An Intuitive Model for Integer Arithmetic. *Educational Studies in Mathematics*, 21(3), 221–239.
- Lungarella, M., & Sporns, O. (2006). Mapping information flow in sensorimotor networks. *PLoS Computational Biology*, 2(10:e144). (DOI: 10.1371/journal.pcbi.0020144)
- Marr, D. (1982). *Vision*. San Francisco: W.H.Freeman.
- Mill, J. S. (1843). *A System of Logic, Ratiocinative and Inductive: Being a connected view of the principles of evidence and the methods of scientific investigation*. London: John W. Parker.
- Popper, K. (1972). *Objective Knowledge*. Oxford: Oxford University Press.
- Rips, L. J., Bloomfield, A., & Asmuth, J. (2008). From Numerical Concepts to Concepts of Number. *The Behavioral and Brain Sciences*, 31(6), 623–642. (doi:10.1017/S0140525X08005566)
- Russell, B. (1917). *Mysticism and logic and other essays*. London: Allen & Unwin.
- Sauvy, J., & Sauvy, S. (1974). *The Child's Discovery of Space: From hopscotch to mazes – an introduction to intuitive topology*. Harmondsworth: Penguin Education. (Translated from the French by Pam Wells)
- Sloman, A. (1971). Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In *Proc 2nd ijcai* (pp. 209–226). London: William Kaufmann. (<http://www.cs.bham.ac.uk/research/cogaff/04.html#200407>)
- Sloman, A. (1978). *The computer revolution in philosophy*. Hassocks, Sussex: Harvester Press (and Humanities Press). (<http://www.cs.bham.ac.uk/research/cogaff/crp>)

- Sloman, A. (1996). Actual possibilities. In L. Aiello & S. Shapiro (Eds.), *Principles of knowledge representation and reasoning: Proceedings of the fifth international conference (kr '96)* (pp. 627–638). Boston, MA: Morgan Kaufmann Publishers. (<http://www.cs.bham.ac.uk/research/cogaff/96-99.html#15>)
- Sloman, A. (2008). Putting the Pieces Together Again. In R. Sun (Ed.), *Cambridge Handbook on Computational Psychology* (pp. 684–709). New York: Cambridge University Press. (<http://www.cs.bham.ac.uk/research/projects/cogaff/07.html#710>)
- Sloman, A. (2009a). Architectural and representational requirements for seeing processes and affordances. In D. Heinke & E. Mavritsaki (Eds.), *Computational Modelling in Behavioural Neuroscience: Closing the gap between neurophysiology and behaviour*. (pp. 303–331). London: Psychology Press. (<http://www.cs.bham.ac.uk/research/projects/cosy/papers#tr0801>)
- Sloman, A. (2009b). Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. In B. Sendhoff, E. Koerner, O. Sporns, H. Ritter, & K. Doya (Eds.), *Creating Brain-like Intelligence* (pp. 248–277). Berlin: Springer-Verlag. (<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804>)
- Sloman, A., & Chappell, J. (2005). The Altricial-Precocial Spectrum for Robots. In *Proceedings IJCAI'05* (pp. 1187–1192). Edinburgh: IJCAI. (<http://www.cs.bham.ac.uk/research/cogaff/05.html#200502>)
- Sloman, A., & Chappell, J. (2007). Computational Cognitive Epigenetics (Commentary on (Jablonka & Lamb, 2005)). *Behavioral and Brain Sciences*, 30(4), 375–6. (<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703>)
- Spelke, E. (2000, November). Core Knowledge. *American Psychologist*, 1233–1243.
- Sussman, G. (1975). *A computational model of skill acquisition*. San Francisco, CA: American Elsevier. (<http://dspace.mit.edu/handle/1721.1/6894>)
- Whitehead, A. N., & Russell, B. (1910–1913). *Principia Mathematica Vols I – III*. Cambridge: Cambridge University Press.
- Wiese, H. (2007). The Co-Evolution of Number Concepts and Counting Words. *Lingua*, 117, 758–772.
- Winterstein, D. (2005). *Using Diagrammatic Reasoning for Theorem Proving in a Continuous Domain*. Unpublished doctoral dissertation, University of Edinburgh, School of Informatics. (<http://www.era.lib.ed.ac.uk/handle/1842/642>)