

Definitions: Dealing with Categories Mathematically

LARA ALCOCK, ADRIAN SIMPSON

The Concise Oxford Dictionary defines 'swan' as:

[a] large waterbird of genus *Cygnus* etc., with long flexible neck, webbed feet, and in most species snow-white plumage. (OED, 1982, p. 1077)

This is not a definition in the mathematical sense, since it is neither possible to say with absolute certainty that everything satisfying the definition is a swan, nor that every swan satisfies the definition. Neither is it reasonable to assume that deductions made from this definition would apply to all swans, nor that something which is true of all swans could necessarily be deduced from the definition. Yet people have few problems thinking and reasoning about swans.

By contrast, a mathematical definition does have the property that everything satisfying it belongs to the corresponding category and that everything belonging to the category satisfies the definition. Deductions made from the definitions provide us with theorems that hold for every member of the category and, in the context of the problems provided by those lecturing to first year undergraduates, any theorem a student is asked to prove can be deduced from the definitions. Yet many students have significant difficulties thinking and reasoning about mathematical concepts.

This article explores these differences by contrasting three approaches to mathematical reasoning. Two of these are more and less sophisticated versions of strategies that we claim are employed in everyday situations. The third type is that required of undergraduates taking proof-oriented mathematics courses. By examining the consequences of the use of these approaches in a first course in real analysis we address the following two questions.

1. Why is the transition to university mathematics difficult?
2. Why is analysis particularly hard?

The first section of this article introduces three terms that are significant in our discussions: *specific object*, *category* and *property*. We then illustrate the three reasoning strategies using examples from a substantial study of students' reasoning behaviour in analysis (Alcock, 2002). Finally, the framework built up from study of these examples is used to suggest reasons for student difficulties with the transition to university mathematics in general and with analysis in particular.

Why is the transition to university mathematics difficult?

In the UK context, as in many others, the transition from school to university mathematics could be seen as an amalgam of many transitions: social transitions (from relatively homogeneous home environments to heterogeneous ones);

pedagogical transitions (from a personal teaching relationship to a fairly impersonal one; from immediate feedback to delayed response; from clear authority relationships to unclear ones); content transitions (from more to less contextualised mathematics); philosophical transitions (from utilitarian to systematic viewpoints), and so on. The issues addressed here cross the last two transitions listed: we focus on the nature of the mathematical objects that students are asked to deal with and the ways in which they reason about them.

We might begin to explore the nature of these changes in objects and reasoning by asking what university mathematics requires that school mathematics does not. One well-researched answer is 'formal proof'. Other authors have given careful delineations of proof-related behaviour (e.g. Harel and Sowder, 1998) and have attributed this behaviour to factors such as relationship with authority and students' previous educational experiences. Here, the focus is on the cognitive origins of behaviour that is commonly considered incorrect or unproductive in situations requiring such proof. We address this by drawing together ideas from the cognitive psychology literature on human categorisation and data from task-based interviews with university students taking their first course in real analysis.

The framework developed in the next section describes different student reasoning approaches in terms of how they handle *specific objects*, *categories* and *properties*.

We use the term 'specific object' in the sense of Sfard's (1991) *structural* notions in mathematics: mathematicians often think of mathematical notions as "static structure[s], existing somewhere in space and time" (p. 4). It is difficult to clarify this meaning further, because of the fluidity with which mathematicians move between thinking of different constructs as objects at different times. However, in the restricted context of a first course in analysis, we can be more precise by associating its meaning with the logical structure of the topic. Thus, specific objects in this work include specific sequences, such as $(1/n)$ or $(\cos(n))$, specific series, such as $\sum n$ or $\sum 1/2^n$, and specific real numbers, such as $0.999\dots$.

Specific objects are collected together in *categories*, each of which is usually associated with a mathematical term, such as 'convergent sequence'. This term is used instead of the more usual 'set', since an important distinction in this article is that between formally defined mathematical *sets* and everyday human *categories*. So while, formally, 'strictly increasing sequence' contains all and only those sequences in which each term strictly exceeds its predecessor, we shall see that not all students appear to think about such categories in this way.

An object may have *properties* and, indeed, all the specific objects in a category may have properties in common. For example, a sequence might have the property of monotonicity and every object in the category of convergent sequences definitely has the property of being bounded.

These terms allow us to highlight what university mathematics requires of students that school mathematics does not. School mathematics primarily involves calculations performed upon specific mathematical objects. For example, students are required to integrate a specific function or solve a specific differential equation. Even the few proofs encountered at this level (in the U.K. context) have this property: students are asked to prove by induction that *this* formula gives the sum of the first n terms of *this* series or prove that *this* trigonometric identity is equivalent to *that* one.

Proof at university goes beyond this. Work with specific objects is still required: students are asked to find the limit of a given sequence or to find the rational number that is represented by a given infinite decimal. However, they must now also work with entire categories of objects. This might involve showing that a specific object is an element of a category – for example, that a given sequence is convergent – or showing that a whole category is contained within another category – for example, that all convergent sequences are bounded. It might be argued that at least the first of these is essentially equivalent to a calculation on a specific object. Logically this may be the case, but we will argue that psychologically it is not.

We will build up a framework describing how different strategies for reasoning about whole categories develop and how these differ from mathematical notions of proof and deduction. In particular, we will outline three distinct strategies. Two of these rely on the use of a ‘prototype’: a representation that an individual considers prototypical of the category and which might be seen to correspond to an individual’s schema for an object-concept (Skemp, 1979) or to a frame with default values instantiated (Minsky, 1975). The last strategy, at least formally, abandons the use of prototypes in favour of definitions.

Briefly, these three strategies are given below.

1. *Generalising*. Students begin with a prototype, inspect it in order to evaluate or generate a conjecture and generalise their conclusions to the whole category.
2. *Property abstraction*. Students abstract a salient property from their prototype and make deductions (intended for the whole category) based on this.
3. *Working from definitions*. Students use agreed defining properties, making deductions (for the whole category) based on these.

Generalising

Wendy is typical of students who appear to use the first strategy. In these excerpts, she is working on the following question during week 9 of a real analysis course (described in more detail in Alcock and Simpson, 2001).

When does $\sum \frac{(-x)^n}{n}$ converge?

Justify your answer as fully as possible.

Wendy begins by expanding the given series and identifying a special case of it as a familiar object. This also reminds her of a test that might be useful.

W: Erm, just, if x is bigger than or equal to one then the series will be erm ... minus x , plus x squared over two, minus x cubed over three, plus x to the four over four and then continuing on.

W: If x is one, it’ll be, minus ... we’ve done that, that converges, doesn’t it. Because we did the one plus a half plus a third plus a quarter plus a fifth doesn’t converge yet ... when you have the alternating ... what was that alternating series test?

She and her interview partner Xavier establish similar known results for other special cases and Wendy then considers intermediate values.

W: Erm, if ... if you take x as between nought and one, say x equal to a half, erm ... and put it into the series, you’d get minus a half, plus a half squared over two, minus a half cubed over three, and so on ...

X: So that’s going to be smaller.

W: Erm, the terms are decreasing in size, so ... x_n is bigger than x_{n+1} [*she misreads notation here*]

X: And tending to zero.

Pause (writing).

W: Converges?

In this instance, Wendy is using $\sum \frac{(\frac{1}{2})^n}{n}$ as a generic example or ‘prototype’ in reasoning about a category of objects. She believes her reasoning will generalise and, in this instance, her final answer is correct. Her work shows a solid relational understanding of the material: she is able to switch sensibly between representations of familiar objects, identify potentially useful results, make qualitative comparisons between different objects and organise her results.

In this situation, Wendy’s strategy is efficient and successful. We argue that she is working in exactly the way that one works with human cultural categories. Rosch’s (1978) work indicates that such categories have structure that could not derive from a ‘classical’ interpretation of categories as mathematical sets. For example, some (perhaps most) categories do not have well defined boundaries – there are no criteria for deciding absolutely who belongs to the category ‘tall man’ or to the category ‘chair’. Similarly, some members of the category ‘bird’ have clearer categorical membership than others; in a U.K. context for example, a robin is more clearly a bird than a turkey is.

The latter phenomenon is known as a *prototype effect*. Within a culture, ratings for the extent to which a member of

a category is considered a 'good example' are consistent across various experiments. This, in itself, does not constitute any particular theory of the use of categories in cognitive processing (Lakoff, 1987), but individuals must somehow cope in a society where categories are of this nature. We have claimed elsewhere (Alcock and Simpson, 1999) that a natural way to reason about categories under these circumstances is to use the *general cognitive strategy* (illustrated in Figure 1 below): the individual assesses a conjecture by evaluating its validity for a prototype and draws a conclusion by generalising to the category as a whole. Others have accounted for human deductive capacity in a similar way (e.g. Johnson-Laird and Byrne, 1991).

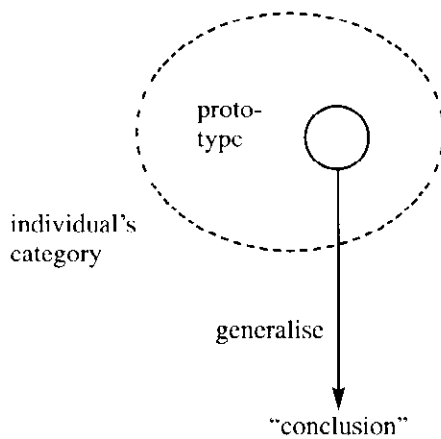


Figure 1

In these terms, Wendy uses $\sum \frac{(\frac{1}{2})^n}{n}$ as a 'good example' of the category she was asked to investigate and makes (correct) statements about the whole of the category by examining this and similar prototypes. While this strategy yields a correct answer in this case, at no point does Wendy verify the fact that her conclusion really does hold for *all* values of x within the range she is considering. The problems associated with this omission can be seen in an extract from her earlier attempt at the following question, posed in an interview in week 7 of the same course.

Consider a sequence (a_n) . Which of the following is true?

- a) (a_n) is bounded $\Rightarrow (a_n)$ is convergent;
- b) (a_n) is convergent $\Rightarrow (a_n)$ is bounded;
- c) (a_n) is convergent $\Leftrightarrow (a_n)$ is bounded;
- d) none of the above.

Justify your answer.

In this case, the difference between Wendy's notion of justification and mathematical proof becomes clear. Wendy begins as follows (I = interviewer).

W: Well if it converges, you get closer and closer ...

Pause (drawing).

W: Is that enough to, like, justify it ... a little diagram, what have you?

I: Well, I'd like you to prove it, if you can.

W: Oh dear! *(laughs)* Oh right, well, if a to the n is increasing ... *(writing)* ... then it's bounded ...

After this prompt for proof, her attempt at this is not much more sophisticated: she describes her picture, but provides no more indication of the logical necessity of her conclusion.

W: It's convergent [*draws a monotonic increasing convergent sequence*] ... yes so if it's convergent it's always ... or ... say it could be the other way round it could be, going down this way [*draws a monotonic decreasing convergent sequence*]. It converges, so it's always above that limit.

In effect, Wendy is offering what Harel and Sowder (1998) would call a perceptual proof, one which is inadequate due to a failure to look beyond features of the particular image she has in mind (see also Presmeg, 1986). In fact, it proves difficult to persuade her to consider convergent sequences that are not monotonic. Whether this occurs because she does not think that there are any such sequences or whether she believes them to be in some way less important is not clear. In any case, she is not willing to move beyond her strategy of examining a prototype and generalising.

Property abstraction

The next, more sophisticated, reasoning strategy still relies on prototypes. However, rather than making a direct generalisation from the prototype to the whole category, we see students attempting to abstract properties from their prototype and work with those properties to justify a conclusion. Cary's initial approach to the week 7 convergence/bounded problem is similar to Wendy's. He establishes an answer using generalization from visual prototypes, making the sketch given in Figure 2 and describing his intentions.




Figure 2: Cary's sketch of convergent sequences

C: I've drawn ... er ... convergent sequences, such that ... I don't know, we have er ... curves ... er ... approaching a limit but never quite reaching it, from above and below, and oscillating either side. I think that's pretty much what I've done. I was trying to think if there's a sequence ... which converges yet is unbounded both sides. But there isn't one. Because that would be ... because then it wouldn't converge. Erm ... so I'll say b) is true.

Though initially he is still working with prototypes, we can see a difference from Wendy's work. Cary considers more examples, making explicit his inclusion of non-monotonic sequences. Following this, he tries to abstract properties from his prototypes in order to formulate arguments, evaluating each candidate property by performing a counterexample check to see whether there are any cases for which this would not be valid.

- C: If it converges ... that has to be ... well I don't suppose you can say bounded. It doesn't have to be monotonic. Erm ...
- C: Yes, I'm trying to think if there's, like ... if you can say the first term is, like, the highest or lowest bound but it's not. Because then you could just make a sequence which happens to go ... to do a loop up, or something like that.

Finding an appropriate property proves difficult, which is not surprising since it took the mathematical community a considerable time to formulate a property that would adequately capture the essence of convergence of sequences.

Tom, who in many ways works like Cary, does find such a property in working on this question, and he attempts a deduction from his choice.

- T: ... If a_n tends to big A , okay ... Then erm ... a_n does not tend to infinity, therefore there is a bound, a , lower than infinity ...

Tom is correct in asserting that since the sequence is convergent it cannot tend to infinity. However, his counterexample checking is not as thorough as Cary's and this means that his deduction is not valid: a sequence that does not tend to infinity may nevertheless fail to have an upper bound (although it seems likely that a student's prototype for this type of sequence *would* have an upper bound). Nonetheless, this gives us an overview of the abstraction-and-deduction strategy, which is contrasted with the generalisation strategy in Figure 3.

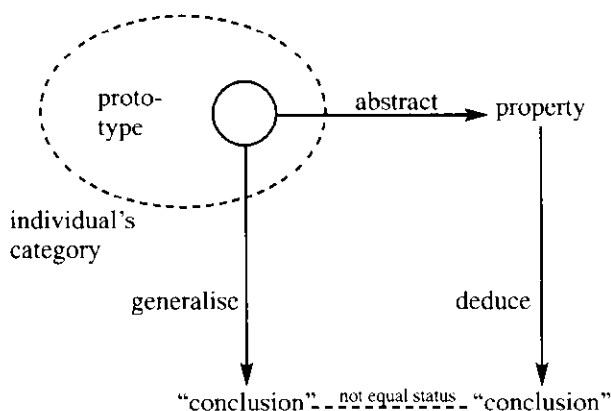


Figure 3: Abstraction and deduction

Our use of the term *abstraction* differs slightly here from that described by Harel and Tall (1991), in which they say

that properties of *objects* are *isolated*. Here, the student's prototype need not be a specific object and a property abstracted from the individual's prototype is not necessarily isolated from it for them, although it may appear to be so for a listener.

In any case, abstraction-and-deduction yields more mathematically acceptable conclusions than direct generalisation, although the latter remains favoured in everyday reasoning. This is no coincidence and the difference is related to the goals of any piece of argumentation. In mathematics, as in any technical field, accurate communication is necessary. Participants must be sure that they intend to indicate the same objects when using a category word. Hence, while they may think using their individual prototypes, they communicate by abstracting verbally formulated properties and arguing in terms of them.

In everyday situations, this type of precision is not usually necessary and may be a hindrance to cognitive efficiency. In their own reasoning, individuals may generalise directly from a prototype without needing to consider explicitly what properties of this prototype make a conclusion valid. In cases where communication is required, the reasoner relies in general on the other party's prototype being similar to their own or on citing specific cases and inviting agreement with the generalisation. In effect, they attempt to communicate by sharing their ideas of prototypes rather than of entire categories. Most of the time this communication will succeed, so the approach is quick and efficient, and is usually valued over guaranteed accuracy in everyday reasoning (Balacheff, 1986).

So we might conclude that Cary is engaging in better mathematical thinking than Wendy. However, as noted above, he finds it difficult to choose a property that will allow him to construct a satisfactory argument. He rejects the assumption of monotonicity and the possibility that the first term of the sequence must be one of the bounds. A mathematician watching this behaviour would recognise that the student needs to introduce the definition of convergence. In fact, with some prompting, Cary is able to do this and to use the definition to outline an appropriate, though incomplete, argument. This is still very much tied to a picture for him - in a similar way to the work of the student Chris discussed in Pinto and Tall (2002). Cary's diagram is given in Figure 4.




Figure 4: Cary's illustration of his definition-based argument

- C: Yes, your n_0 ... that could just be called your n_0 instead, so going back to your definition up there, there exists this point here, such that after that

point, i.e. when n is greater than n_0 , the sequence ... that statement there won't be less than any epsilon which you just happen to pick ... And so it's ... and so the upper bound - so because there's finitely many terms before n_0 , then er ... your upper bound will either be plus or minus epsilon, or it'll be the maximum of those finite terms beforehand.

Working with definitions

While Cary can work with definitions, this is not his first approach; it requires several prompts to persuade him to write down a reasonably complete definition and to use this in preference to his original strategy. By contrast, in answering the same question, Greg shows immediate recourse to the definition and rapidly outlines an argument.

G: Right. It's easiest to use the definition. Just say definition of ... of convergence is that eventually ... Okay! ... The definition of convergence is eventually you'll find epsilon such that ... for all n bigger than big N , epsilon - no a_n minus a is smaller than epsilon. So you've got it bounded between ... a_n is ... a plus epsilon ... a minus epsilon even. And a plus epsilon. And for ... n bigger than big N . And for all n smaller than big N , you know that a_n has a minimum, and maximum because it's finite. Er ...

This approach is radically different from those seen so far. Greg does not generalise directly from a prototype (as Wendy does), nor does he attempt to abstract appropriate properties from a prototype in order to make deductions (as Cary does). His approach does use a property - but it is the *defining* property of convergence of a sequence.

The result is a fundamental difference in the nature of the category the students work with. For both Wendy and Cary, the category is pre-existing (and non-classical) and for Cary the properties of the category *follow* from it. Greg's approach to property use, however, goes beyond Cary's by inverting the property/category relationship: the defining property *determines* the category.

Since a definition is precisely a set of necessary and sufficient conditions for category membership, a mathematically-defined category has a fixed boundary and no 'better' members. So, once a definition is chosen, both the original objects and any individual's prototype constructed through experience with those objects, become secondary in importance to the property itself (cf. Gray *et al.*, 1999). This is illustrated in Figure 5.

The process of choosing such properties is institutionalised within the mathematical community. While definitions such as those used by Greg can often be traced to properties abstracted from an individual's prototype, one function of the community is to debate which of these properties best capture what is common to those objects in the category under discussion (Lakatos, 1976). This is a worthwhile enterprise, because making such decisions facilitates communication on a large scale by making reasoning in the subject systematic (Bell, 1976). Greg's work reflects the results of this debate in analysis and his appropriate use of

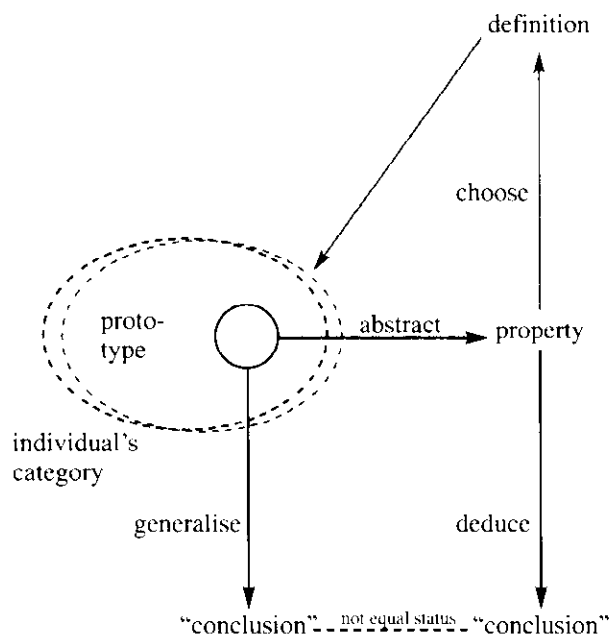


Figure 5: One property is chosen as the definition, precisely determining the mathematical category

the definition means that any correct deductions he makes will be valid for all members of the mathematical category of convergent sequences.

We do not argue that mathematicians think solely in terms of definitions, although our particular example happens to show no evidence of the use of a prototype. Indeed, prototypes can remain vital: Tall's (1991) description of Poincaré's thinking and Thurston's (1990) self-reflection demonstrate the importance of prototypes as a source of both conjecture and guidance for the direction of deduction. There is no inconsistency here: individuals can learn to formulate arguments within the logical structure of mathematics while still reasoning using the same psychological strategies as in other contexts. Indeed, we argue that the inclination to work from definitions is a 'prefix' imposed on general thinking skills (the 'rigour prefix', described in Alcock and Simpson, 1999). Mathematicians might think in terms of prototypes, but they are aware that, in order to ensure universal validity for their arguments, they must eventually formulate these in terms of appropriate definitions.

A communication breakdown

Although students are rarely involved in selecting definitional properties (Harel and Tall, 1991), working from those dictated by their teachers might in principle be expected to make their task easier. Proof tasks in beginning advanced-level courses are usually quite straightforward in structure: in order to show that an object is a member of a category, one checks that it satisfies the definition; and in order to show that a conjecture is true of all objects in a category, one makes deductions from the definition. Hence, while formulating the detail of such deductions may be difficult, the 'top level' (Leron, 1985) or 'proof framework' (Selden and Selden, 1995) is often very simple.

However, as demonstrated by the excerpts from interviews with Wendy and Cary, and as noted elsewhere in the literature (Moore, 1994), the shift from “show that x is X ” to “show that x satisfies the definition of X ” is not one that all students readily make. The status of definitions in mathematics seems to elude them and this is the source of our earlier claim that, psychologically, showing that an object is a member of a category is quite different from performing calculations on this object.

To be more accurate, we might say that what eludes the students is the distinction between a dictionary definition as a *description* of pre-existing objects and a mathematical definition as the chosen basis for deduction, one which serves to *determine* the nature of the objects. The result is that a breakdown in communication occurs when a lecturer gives the student a definition. In a non-technical context, good students do not simply learn a definition, they use further experience with examples to deepen their understanding of the concept.

So when a lecturer provides a mathematical definition, expecting the student to work with it in future, the student may try to build or refine a prototype for use in this non-technical way and perhaps abandon the definition once this is accomplished (Vinner, 1991). If the lecturer also provides examples, intending them to be illustrative, the student may even construct their prototype on the basis of experience with them, ignoring the definition altogether. This contrast between using a dictionary definition to construct a prototype and using a mathematical definition to generate a category is illustrated in Figure 6.

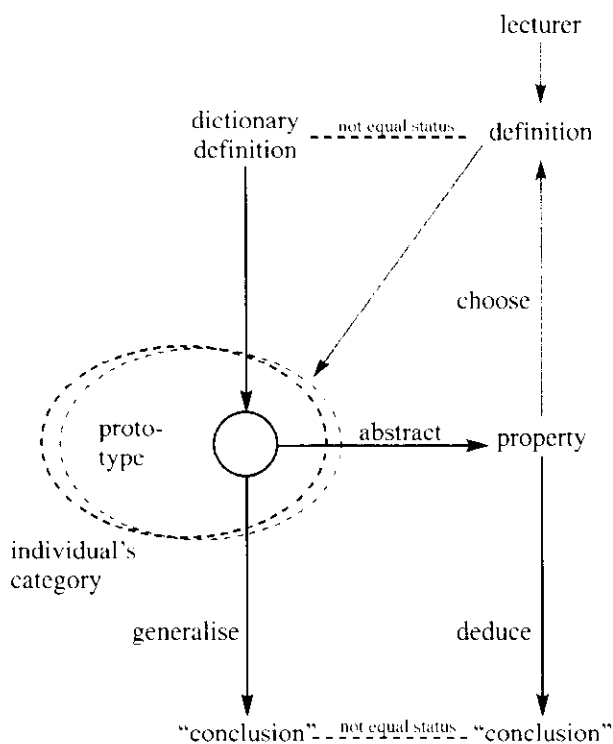


Figure 6: Contrasting the use of dictionary and mathematical definitions

Crucially, it may not be apparent to either party that this communication breakdown has occurred. In cases where the technical meaning of a term is similar to – or perhaps, derived from – the everyday meaning, there will be a large overlap between a student’s idea of what is in the category and the formal version. Hence, there may be very few cases where conflict seems to arise and both student and teacher may feel that they are communicating successfully.

Hence, one reason for the difficulty of the transition to university mathematics is that certain reasoning strategies are inadequate when applied to university mathematics, although they may be efficient and successful in non-technical contexts and in the kind of reasoning with specific objects required by school mathematics. The student must learn to override these strategies with the new approach of working from the dictated definitions, but since the role of mathematical definitions usually remains below the level of consciousness of working mathematicians, this is rarely communicated and may be far from transparent.

Why is analysis hard?

We can now offer an explanation of why real analysis, of the topics studied at the beginning of a mathematics degree, proves particularly difficult. Definitions in analysis are logically complicated; they often involve multiple mixed quantifiers (Dubinsky, Elterman and Gong, 1988). However, teachers of analysis also typically make use of visual representations of objects and results. Such representations are useful in building prototypes, such as those seen in use by Wendy and Cary. These prototypes are in turn conducive to arguing by direct generalisation: it is easy to convince oneself that an increasing sequence which is bounded above must converge, without recourse to abstracting properties or formulating algebraic arguments.

Indeed, considerable exasperation may result when students’ prototypes provide them with a strong feeling of intuitive intrinsic conviction, but they are asked to justify their assertions precisely (Fischbein, 1982). The requirement to use the complex definitions means that they are often in a position from which they must prove something they consider obvious using algebraic formulations which make them feel insecure (Gray *et al.*, 1999).

Comparing analysis with other beginning university subjects such as group theory shows how much more significant the problem is in analysis. Definitions in group theory may be long, but they are logically simpler than those in analysis. Other types of representation are less readily available or, at least, less often taught in the early part of the course. Hence, in group theory, it is likely that more students will produce work that competently makes use of the definitions. This does not mean that they necessarily understand the structure of this topic or the role that definitions play within it, only that these are easier to handle and that they have no obvious other option.

Indeed, it would be interesting to investigate the consequences of the difference for student complaints. Since the visual representations used in analysis occupy an intermediate position between realistic pictures and verbal/symbolic representations (Gibson, 1998), these appear more concrete. We may find that complaints are distributed so that analysis

attracts more comments like: “It’s obvious, but I don’t know how to prove it” and group theory more comments like “It’s just too abstract”.

Thus, in analysis, the availability of visual representations means that more students initially have access to a way of coming to understand the concepts. The understanding gained in this way means that they feel less need to engage seriously with the complex algebraically-expressed properties from which the formal categories are constructed. So, paradoxically, analysis may be difficult not only because the material is complex *per se*, but because it is initially *less* ‘abstract’ than other beginning subjects in advanced mathematics.

References

- Alcock, L. (2002) *Categories, Definitions and Mathematics: Student Reasoning about Objects in Analysis*, unpublished Ph.D. thesis, Coventry, Warwickshire, University of Warwick.
- Alcock, L. and Simpson, A. (1999) ‘The rigour prefix’, in Zaslavsky, O. (ed.), *Proceedings of the 23rd International Conference on the Psychology of Mathematics Education*, Haifa, Israel, 2, pp. 17-24.
- Alcock, L. and Simpson, A. (2001) ‘The Warwick analysis project: practice and theory’, in Holton, D. (ed.), *The Teaching and Learning of Mathematics at University Level*, Dordrecht, Kluwer, pp. 99-111.
- Balacheff, N. (1986) ‘Cognitive versus situational analysis of problem-solving behaviours’, *For the Learning of Mathematics* 6(3), 10-12.
- Bell, A. (1976) ‘A study of pupils’ proof-explanations in mathematical situations’, *Educational Studies in Mathematics* 7(1), 23-40.
- Dubinsky, E., Elterman, F. and Gong, C. (1988) ‘The student’s construction of quantification’, *For the Learning of Mathematics* 8(2), 44-51.
- Fischbein, E. (1982) ‘Intuition and proof’, *For the Learning of Mathematics* 3(2), 9-14.
- Gibson, D. (1998) ‘Students’ use of diagrams to develop proofs in an introductory analysis course’, in Schoenfeld, A., Kaput, J. and Dubinsky, E. (eds), *Research in Collegiate Mathematics Education* 7, pp. 284-307.
- Gray, E., Pinto, M., Pitta, D. and Tall, D. (1999) ‘Knowledge construction and diverging thinking in elementary and advanced mathematics’, *Educational Studies in Mathematics* 38(1-3), 111-133.
- Harel, G. and Sowder, L. (1998) ‘Students’ proof schemes: results from exploratory studies’, in Schoenfeld, A., Kaput, J. and Dubinsky, E. (eds), *Research in Collegiate Mathematics Education* 7, pp. 234-283.
- Harel, G. and Tall, D. (1991) ‘The general, the abstract and the generic in advanced mathematics’, *For the Learning of Mathematics* 11(1), 38-42.
- Johnson-Laird, P. and Byrne, R. (1991) *Deduction*, Hillsdale, NJ, Lawrence Erlbaum Associates.
- Lakatos, I. (1976) *Proofs and Refutations: the Logic of Mathematical Discovery*, Cambridge, Cambridge University Press.
- Lakoff, G. (1987) ‘Cognitive models and prototype theory’, in Neisser, U. (ed.), *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*, Cambridge, Cambridge University Press, pp. 63-100.
- Leron, U. (1985) ‘Heuristic presentations: the role of structuring’, *For the Learning of Mathematics* 5(3), 7-13.
- Minsky, M. (1975) ‘A framework for representing knowledge’, in Winston, P. (ed.), *The Psychology of Computer Vision*, New York, NY, McGraw Hill, pp. 211-277.
- Moore, R. (1994) ‘Making the transition to formal proof’, *Educational Studies in Mathematics* 27(3), 249-266.
- OED (1995) *Concise Oxford Dictionary of Current English*, Oxford, Clarendon Press
- Pinto, M. and Tall, D. (2002) ‘Building formal mathematics on visual imagery’, *For the Learning of Mathematics* 22(1), 2-10.
- Presmeg, N. (1986) ‘Visualisation in high school mathematics’, *For the Learning of Mathematics* 6(3), 42-46.
- Rosch, E. (1978) ‘Principles of categorization’, in Rosch, E. and Lloyd, B. (eds), *Cognition and Categorization*, Hillsdale, NJ, Lawrence Erlbaum Associates, pp. 28-48.
- Selden, J. and Selden, A. (1995) ‘Unpacking the logic of mathematical statements’, *Educational Studies in Mathematics* 29(2), 123-151.
- Sfard, A. (1991) ‘On the dual nature of mathematical conceptions: reflections on processes and objects as different sides of the same coin’, *Educational Studies in Mathematics* 22(1), 1-36.
- Skemp, R. (1979) *Intelligence, Learning and Action: a Foundation for Theory and Practice in Education*, New York, NY, John Wiley and Sons.
- Tall, D. (1991) ‘The psychology of advanced mathematical thinking’, in Tall, D. (ed.), *Advanced Mathematical Thinking*, Dordrecht, Kluwer, pp. 3-21.
- Thurston, W. (1990) *Three Dimensional Geometry and Topology*, Princeton, NJ, Princeton University Press.
- Vinner, S. (1991) ‘The role of definitions in the teaching and learning of mathematics’, in Tall, D. (ed.), *Advanced Mathematical Thinking*, Dordrecht, Kluwer, pp. 65-81.