

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Heterogeneous reasoning in learning to model

Permalink

<https://escholarship.org/uc/item/2xd7p24c>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 22(22)

Authors

Stenning, Keith
Sommerfeld, Melissa

Publication Date

2000

Peer reviewed

Heterogeneous reasoning in learning to model

Keith Stenning (K.Stenning@ed.ac.uk)

Human Communication Research Centre, Division of Informatics, University of Edinburgh
2 Buccleuch Place, Edinburgh EH8 9LW, UK

Melissa Sommerfeld (msomme@leland.stanford.edu)

Cognition and Learning Lab, CERAS 105, Stanford Ca. 94305-3084, USA

Abstract

Conceptual learning in maths and science involves learning to coordinate multiple representation systems into smoothly functioning heterogeneous reasoning systems composed of sub-languages, graphics, mathematical representations, etc.. In these heterogeneous systems information can be transformed from one representation to another by inference rules, and learning coordination is learning how and when to apply these rules. Heterogeneous reasoning has a particularly important role to play in teaching students how to apply formalisms to real world problems, rather than merely teaching formalism-internal calculation.

This paper analyses three learning incidents which happened in groups of students engaged in learning the mathematics and biology involved in modelling biological populations, from the perspective of the heterogeneous reasoning involved. Greeno, Sommerfeld & Weibe (2000) and Hall (2000) analyse incidents from the same curriculum intervention from other points of view, in this volume.

We observe both learning successes and failures that cannot be understood without understanding the seams joining the representation systems involved, and the inference rules and operations required to get from one to another. One conclusion is that even apparently homogeneous natural language has to be seen as heterogeneous in its fully contextualised application.

Introduction

The coordination of multiple representation systems is frequently instrumental in conceptual learning (see e.g. van Someren et al. 1999; Barwise & Etchemendy 1994; Stenning, Cox & Oberlander 1995). In particular, learning mathematics and science concepts involves learning to coordinate multiple formalisms (numerical, graphical, algebraic, terminological), but it also involves learning how to apply formalisms in contexts. It is all too possible for students to succeed at the first and to fail at the second—to learn the internal operation of some formalism without learning how to apply it to new problems. Barwise and Etchemendy have used the term ‘heterogeneous reasoning’ for reasoning using multiple coordinated representations, and have applied heterogeneity of representation in order to improve students’ grasp of the application of formalisms in computer environments such as Hyperproof. Another related curriculum response to this problem of teaching the application of formalisms to real world problems has been a move toward project-based approaches which teach formalisms in close relation to their context of application—in particular teaching scientific concepts along with the mathematics that goes with them.

The purpose of this paper is to use some example episodes from project-based group learning to illustrate how the concepts of heterogeneous reasoning present themselves in the classroom setting in less formal contexts than Hyperproof. This investigation underscores how local the semantic interpretation of representations is in context. Words change meaning frequently and systematically, and the information they carry is moved into and out of other representations. The investigation also provokes examination of the relation between heterogeneity and localness of interpretation (e.g. Moravcsik 1998). With diagrams, it is usually quite evident to users that the diagram has a local interpretation and that the naive user needs to learn this local interpretation, even though there are regular features of such diagrammatic systems from use to use. With natural language, we are often so practiced at making the contextual interpretation of its local semantics that it is easy to fail to realise that this is what we do. Examining learning discourse in context raises the question whether heterogeneity should also be extended to cases where the linguistic part of the discourse has to be treated as multiply interpreted.

Our longer term aim is to coordinate this approach with others which focus on discourse practices and students’ recruitment of material from their diverse experiential worlds. Greeno, Sommerfeld & Weibe (2000) and Hall (2000) take these respective perspectives on material from the same group learning curriculum intervention. The advantage of pursuing several parallel analyses of the same data for cognitive theory may share something with the advantages of the project-based approach for the students. Applying several kinds of the theory to the same episodes turns up new questions about how the theories relate to each other, and thus may induce conceptual learning and improved ability to apply the theories in novel circumstances.

Heterogeneous reasoning

Theories of human reasoning have begun to pay more attention to how representational systems are selected or constructed, and the variety of systems that may be used in solving a single problem, rather than conceiving of reasoning as a system internal activity. Barwise & Etchemendy have called this use of multiple coupled systems of representation *heterogeneous* reasoning, and have developed several computer environments for teaching heterogeneous reasoning. For example, Hyperproof presents a graphical window containing diagrams of a blocks-world inhabited by regular solids on a chequerboard, and a sentential window containing first order

logic sentences. The proof rules of the heterogeneous system incorporate the inference rules of the conventional sentential calculus, augmented by rules for moving information between diagram and sentences, in both directions. For example, the user can *observe* a feature of the diagram as the basis for inferring a sentence; or may *apply* information from a sentence by inferring (and constructing) a new feature of the diagram. *Observe* and *apply* are two (of about a half dozen) of the heterogeneous inference rules which coordinate the diagrammatic and sentential representations into a heterogeneous system.

Fundamental semantic distinctions between how diagrammatic and sentential representation systems express abstraction have been shown to play an essential part in analysing the learning that occurs as students master the construction of proofs in Hyperproof's heterogeneous environment, allowing the learning to be characterised as learning strategies of representation selection and use (Stenning, et al. 1999). Whether students benefit from the diagrammatic facilities of Hyperproof is determined to a great extent by their facility at grasping useful strategies for using Hyperproof's expressions of graphical abstraction.

Hyperproof reveals an important property of representational systems in use. Its semantics, both of its sentences and its diagrams are *partially interpreted*. That is, the system has some of its meaning fixed while other parts are defined in episodes of reasoning. This contrasts with the usual presentation of first order logic as an entirely uninterpreted language. In Hyperproof, even the sentences of the first order calculus have to be given a partial local interpretation because the predicates and relations have to coincide with those of the diagrams.

Partly diagrammatic systems of representation like Hyperproof reveal the need for coordinating diagrammatic and sentential representation systems, but lead to the further realisation that in situations of real language use, the apparently homogeneous languages in play are in fact often heterogeneous in the fundamental sense that many schemes of interpretation are in play at once. Even when natural language is the only modality, the reasoning systems in operation must be thought of as heterogeneous because the apparently single language can only be understood in terms of overlapping language fragments, each constituting a distinct system of representation.

To illustrate this point, this paper takes some classroom interactions of a group of students learning to model biological populations in terms of mathematical functions, and analyses the multiple partially interpreted representation systems which are in play. The students' representational resources and activities include at least the following: worksheet filling, graph drawing, computer operation, calculator use, group speech and gesture, reference material, and teacher interventions.

The educational issue in focus is the learning about modelling, and particularly learning about the process of formalisation and interpretation. A recurring theme is the struggle to coordinate formalism internal operations (calculation) and formalism external correspondences (semantics). We will analyse both successes and failures of coordination.

The educational setting

The data we analyzed comes from an 8th grade Middle-school Mathematics through Applications (MMAP) classroom in the San Francisco Bay Area. The purpose of MMAP is to have students use math to address real-world problems, often with the assistance of computer applications. In the approximately 30 day unit we will discuss, called Guppies, students created mathematical models of biological population growth. As part of this unit students were to learn both about how to construct mathematical models of population growth and about the exponential functions that underlie them. Our analyses focus on a group of students, Manuel, Lisa, Kera & Nick whose improvement on pre/post assessments placed them about midway in learning of the half a dozen focus groups videotaped by Rogers Hall and his colleagues (Hall, 1999) during this unit in a variety of classrooms. These students are chosen to reflect roughly average performance for the class. For more information about the design of the study, please see Hall (2000, this volume).

Three learning incidents

The three following incidents were chosen from videotapes because they illustrate both successful and problematic learning episodes. The initial incident from the pre-test phase sees the students make at least part of one of the fundamental conceptual discoveries of this field—that population models have a recursive characteristic that leads to exponential growth if unchecked—Malthus' equation.

The second incident, from the body of the course, is of interest because it contains an attempt to creatively diverge from the structure of the assigned worksheet by taking a short cut in the calculation. On the one hand this divergence reveals the germ of another important insight—that functions can be composed. But in the circumstances, the insight is not fully worked out and leads to error and confusion.

The third incident is chosen to illustrate that the confusion that is not resolved in the previous incident appears to persist into the much later post-test phase of the course. It consists of another attempt to calculate a birthrate for a new modelling problem.

In all of these incidents, the students struggle to coordinate multiple representations. We examine some of the coordinations in detail seeking to reveal how some episodes are successful and others not. For this short paper, the transcriptions are compressed by leaving out material which does not relate to our analysis.

Pre-test insight—'babies have babies'

When the group discovers the recursive nature of population growth, they are engaged in constructing a model of a mouse population. They have obtained an initial number of 20 adults from the worksheet, and estimated a birthrate of four per couple. They are now calculating what the population will be after eight breeding seasons. The group initially adopt a linear model implicit in multiplication of a fixed birthrate. Only when they turn to the graphing activity dictated by the worksheet do they begin to think of the process which the calculation is intended to reflect.

60: M. so there's ... equals 40 babies each season

65: M. it's three hundred and twenty

- 66: K. (inaudible) is that including adults?
 67: M. no, three hundred and twenty plus twenty
 69: M. by the end of the winter
 70: M. three hundred and forty mouse ...mice ... mices.
 OK.
 73: M. Now we need to make a graph of it
 ...
 182: M. so let's see ... the first season is over here
 (making a mark on the graph)
 183: L. xxxxxx wait a minute
 186: M. and then sixty plus is going to be a hundred
 189: L. wait a minute its forty (gestures a triangular
 shape) OK its forty right?
 190: L. and then you have to pair those up (brings hands
 together) and then they have kids (spreads hands apart,
 while K and M look at her confused)
 192: M. oh yeaah (embarrassed, laughing at himself)
 we were doing it ...
 194: L. That's a lot of mice
 195: K. gosh that's a lot of nasty mice

The interchange on lines 65/66 is an example of the frequent need to coordinate numbers with their semantics—adults still have to be included in the population, and “three hundred and twenty” is the number of babies in eight seasons just calculated. Similarly line 69 is a further reiteration of the semantics of the number “three hundred and twenty plus twenty”—the number represents a population at a time. Line 73 is an appeal to the authority of the worksheet for what has to be done next. What is interesting about this introduction of a new representation (the graph) is that it appears to be what triggers the new thinking that reveals the error (adopting the linear model) that they have all made. M. makes a mark of sixty on the Y axis at the origin representing the starting population. But L has realised that something is wrong (line 183). M continues calculating the next graph point. But L persists. She starts by reiterating the number and asking for acknowledgement of it (line 189). The number is the number of first season babies. She then states that these have to be paired up, and themselves reproduce (line 190). The gesture is intuitively an important part of her communication that she has a new insight, both for herself and for the group. M fairly rapidly sees their mistake too. They all realise that this is going to make the growth of the population much more rapid though they don't have any number for it. They immediately refer back to the experiential world of ‘nasty mice’. Perhaps the reality of reproduction lies behind the affective tone of the incident. It wasn't just a mathematical mistake, but a failure to apply the ‘facts of life’?

The original adoption of the linear model arises within the ‘mathematical world’. It is, in some sense, the obvious calculation to do—forty babies a season for eight seasons is going to give 320 babies. After all, multiplication is something we learn so as to avoid having to do multiple additions. It is not until the graphing activity makes them break this calculation down into a series of calculations that L sees the error. She thinks about what happens in the world of mice—about the semantics. Her insight is adopted rather rapidly.

An attempt at creative construction—‘discovering function composition’

When the group brushes up against function composition, they are constructing one of their early models of a population. They have a worksheet entitled *Building the Birthrate* which gives them a procedure for calculating, or recording from reference sources, the various parameters of the situation (brood size for different ages, birthrate, survival rate). Parts of this worksheet and the computer interface are condensed into Figure 1.

The worksheet has its own sequence of activities, though it should be noted that this is not the sequence in which this group performs them. The worksheet's (see Figure 1) first table implements the calculation of the total population births in a season from data from reference sources. At Step 2 the percentage *survival rate* is entered from a reference source and, at Step 3, applied to the total from the table to give a number surviving. The lines represent page breaks in the work booklet. Step 4 then converts the total surviving fry into a percent *birthrate* for the computer. The relevant part of the interface appears next. The bottom table of the figure (over the page on the worksheet) keeps track of the model, and will hold several trial models later on.

The group's sequence of work is actually to start by fulfilling steps 1, 2 and 3, followed by entering the result into the computer and recording the model. Step 4 is circumvented initially and is only filled in retrospectively the next day.

The incident opens with M proposing to take a shortcut in the calculation. This is at first taken by L to be a mistake. She requests an explanation and receives one that she finds satisfying. However, she appears to appreciate that there is consequent bookkeeping which needs to be taken care of, but fails to deflect the group from continuing to the entry of data into the computer model.

444: M. hey wait wait wait ... no but listen. If 4% of the fry survive why don't we just forget about the fry survival and just put that amount for the, for how much are born ...

445: L. because the number born are not how much survived

446: M. yes. yes, the ones who survive are the ones we count, not the ones who are dead because we don't make room for the ones that are dead

453: M. OK you know how 4% the whoooooo fry who were born survive so why don't we just put 4% on the guppies birth because that's how many are going to survive

454: L. I get what you're saying because why put however many more guppies in when they're just going to die anyway?

455: K. so why not just put 4% because that's how many are surviving/ that's how many we're going to count

497: L. but what's that 4% ?

498: K. the ones that survive

499: M. The ones that actually survive fryhood

501: L. Yeah, I know, but how many of the guppies are 4% ?

502: M. we don't know, we'll let that mechanical thing

Building the Birthrate

	age	# males	# females	# fry	total
Step 1	young	2	1	4	4
	mature	4	2	50	100
	old	0	1	0	0
	total	6	4		104


Step 2 What percent of fry born survive? What happens to the ones who don't make it?

5% of fry survive. They are eaten

Step 3 Use this survival percent and the total number born to calculate the number that survive.

5.16


Step 4 So what's the birthrate? Now that you have calculated an assumed number of fry that survive past birth, you need to convert this into something that Habitech can use as a birth rate. As you know, Habitech works with percents or constant numbers. You will be using a percent birth rate.



BASED ON YOUR ASSUMPTIONS YOUR BIRTHRATE IS

Congratulations! Now take this birth rate and the death rate you will use and head to Habitech to make your model. Remember this birth rate is based on certain assumptions. If you change an assumption, it will affect your model.

Step 5 Entering numbers into the Habitech interface:



Recording of Models

Initial #	Birth rate %	Death rate %	Years	Descr.
10	4%	4%	2 year	< 13

Figure 1: Parts of the worksheet and computer interface. The numbers in the tables, equation and the italicised answers were entered by the students

work and tell us

At 444, M opens with a proposal to collapse two stages of calculation into one. In fact, this proposal is perhaps something akin to what is embodied implicitly in the worksheet, and is potentially a creative proposal embodying a concept rather close to one of the core aims of this curriculum—the understanding of mathematical functions. M is proposing to compose two functions into a single function taking the argument of the first and the value of the last. L objects to this proposal and justifies her objection by pointing out that ‘the number born are not how much survived’. In fact we will see that in the terminology of the worksheet, the number of fry surviving expressed as a percentage of the whole population *is* the birthrate, which plays its part in this confusion. M appears to understand the objection and explains his proposal’s departure from the worksheet with some success. L accepts the sense of the innovation even though she expresses reservations about its coordination with the worksheet. The activity is turned over to the superior calculating powers of ‘that mechanical thing’—the computer program Habitech.

Unfortunately, the ‘mechanical thing’ does not understand the creative proposal—L’s reservations are well motivated, but, lacking a clear understanding herself, her intervention does not deflect the group (see Greeno, Sommerfeld & Weibe (2000) for further analysis). There are numerous problems of coordination between the representations in Figure 1. The survival rate of 5% at Step 2 gets copied into the model table as 4% (possibly a memory error, or a correction later). But the serious error is in shortcutting the calculation at Step 4 and entering the 4% rate directly into the birthrate box at the end. The algebraic ratio part at Step 4 is returned to only later next day when trying to comply with having the whole thing filled in.

What has gone wrong as the group struggles with the welter of representations and numbers? It is hard to give a crisp interpretation of a murky confusion, but we can suggest some of the contributing factors. An important source may be a divergence of the ordering of biological events and the calculation events that refer to them; another is the terminology. In the fish world, fry are born, and then the vast majority are eaten, and then at the end of the season they are counted. In the calculation world, first the number of births are calculated; then a survival rate is applied; and a census number of surviving fry results. So far so good. But turning the page after Step 3, and after recording model parameters on the next page, the students arrive at a further calculation of the ‘birthrate’, where ‘birthrate’ now means ‘birth-and-survival-to-year’s-end rate’.

So, at Step 1, the birthrate is a set of numbers representing the brood size of the average guppy at different ages (namely the numerals 4, 50, 0); at Step 2, the birthrate is the number (namely the numeral 104) of fry born to the whole population. In steps 3 and 4 *birthrate* is the birth and survival to end of season rate expressed as a percentage of the whole population (namely the numeral 4). The same idea, a very tangible idea, is represented each time by a number, but each time the number counts different kinds of thing, and complex calculations constitute the inference rules which ‘move the number from box to box’.

Unfortunately, M's insight that two functions can be composed requires attendant housekeeping to keep the ontology straight. Perhaps a contributing factor is that because the pre-survival birthrate in Step 1 is never put into the form of a percentage (1040%), M does not appreciate that, after Step 3, it already has been implicitly composed with the survival rate, and the calculation at Step 4 is intended only to get back to a percentage form. The terminology unfortunately exacerbates this problem of 'backward causality'—first calculating a survival rate (using births) and then calculating a birthrate from that figure.

Post-test—the persistence of a confusion

We now present an incident from the post-test in which the group displays evidence that the episode of confusion just described has not been fully resolved. Although in the intervening couple of weeks the group has made good progress in understanding population models, as is illustrated in Hall (submitted to this volume), it is of some concern that the particular confusions surrounding the derivation of birthrate from raw data appear to persist.

The group is working on the post-test problem of constructing a model for a mouse population preyed on by cats. This episode is from fairly early on when they are settling on a birth rate for mice and have not yet considered predation:

- 76: M. four, five or six? per adult?
77: K. If we're going to go four, five or six, let's go four.
78: L. actually lets use five. Its four through six. Let's use five.
82: M. OK how do we find out the birthrate? (grabs a piece of paper) We do the ... five is what we decided on. How many did we start out with (looks at the computer)
83: L. Twenty
86: M. I'm not sure that this is right (as he writes $5/20 = X/10$)
87: M. What's 500 divided by twenty?
88: L. What are you doing?
89: M. Finding out the birth rate
90: L. Oh yeah.
91: M. What's 500 divided by 20? (K hands him the calculator and M starts punching in numbers)
92: M. 25% I could have figured that out myself (K laughs; M goes back to the computer) 25% right? (enters it into the birthrate) and how many die?

Segment 82 illustrates the pervasive struggle with the semantics of numbers. M accepts that they will use 5 (babies per litter per season) which one might think *is* a birthrate, but in this context, 'birthrate' is a specific number that can be entered into certain boxes on worksheet and computer screen. The birthrate, in this sense, they correctly appreciate they do not have, and this is precisely where they had problems before. The number they seek is a percentage. At 87, M has implicitly multiplied the 5 by 100 and is now explicitly going about dividing by twenty (the number in the initial population). L not surprisingly doesn't understand where the 500 came from and asks for clarification, but receives only the description at the completely unhelpful level "finding out the birthrate". The problem is then accepted as a calculation problem, and the semantics is left unaccessed. Why should

the number of babies in one litter divided by the total number adults in the population multiplied by 100 yield a percentage birthrate? The answer would appear to be that the based on some dim memory of a ratio formula (Step 4, Figure 1).

The group is content to continue to the next stage of the problem and does not question the reasonableness of the figure of 25%. This is testimony to the insulation of the numbers from what they mean. If each couple has 5 babies, the actual number is 250%. But the group do not discuss finding this number or acknowledge that adults have to be paired up. The group does not even apply the qualitative reasoning that since the parents are outnumbered by their babies, the birthrate must be more than 100%. Such qualitative inferences are only available if the numbers are treated as standing for something other than themselves—numbers. Even when the model actually turns out to extinguish the mice in short order, the problem is not traced to the low birthrate. It is all too easy for a problem to hide in a complex model. The whole point of models is that many parameters contribute to their outcome. But this means that there are many possible culprits when the outcome is unacceptable.

Discussion

Nothing by way of inferences about the causalities or even correlations between the kinds of events observed here can emerge from an analysis of these few isolated examples. Nor is redesign of a curriculum usefully based on analyses of single incidents. It is clear from other studies of this curriculum that it is highly effective. Indeed, this very group of students shows a considerable mastery of modelling at the post-test phase. The group repeatedly alters parameters of complex models (including not just birth and survival rates but also predation) in the qualitatively correct direction in response to over- or under-shooting of the desired population outcome.

But we believe that these analyses do make clear just what a sea of semantic complexities the group swims in. They are awash with numbers, and those numbers have to travel from one representational system to another to achieve the problem solving task at hand. As they travel, they change their meanings and their names, and their values. Birthrate is rarely the same thing on two occurrences. The whole system cannot be understood as anything other than heterogeneous, and the interpretations as anything other than highly local. If we were to go through the transcript spelling out after each occurrence of a numeral, the type of the entities it enumerates, we would wind up with some splendid and totally incomprehensible sentences. Nor are numerals the only problem. Simply spelling everything out is *not* to be recommended other than as a way of exposing complexity. But we cannot understand the students' problems until this complexity is exposed.

From a theoretical perspective, this may seem either banal or outrageous. Once we are fluent at the skills of transformation required for coordinating the sub-systems of representation, the whole system appears to take on a transparency and homogeneity which is completely illusory. We cease to notice how the very same number means something quite different from occurrence to occurrence, as do many of the other words. We therefore can either forget that the system is heterogeneous (and respond with outrage to the claim), or we can, as theoreticians, claim that there is nothing deep in the

coordinations that are required (and respond with a yawn).

The students do not have the luxury of mastery. For example, one of the banal consequences of the instability of the meanings of the numerals is that there is a huge memory load as evidenced in the repeated mis-recalls of numbers from sheet to sheet of their workings. We do not believe that there is any way out of the heterogeneity. Learning mastery of the coordination of representation systems is a requirement of learning mathematics and science (and probably most other things). But what we can strive to do is to educate both teachers and students into the quirks of the representational furniture they find themselves surrounded by.

Our research experience in classrooms indicates that teachers are rather wary of taking an explicitly metalinguistic stance. They do not often point out the dangers of shifts in meaning of words during an argument. The critical thinking lecturer warns students about equivocation—the same term being used with different meanings in different occurrences in an argument—but only at college. Prevarication is treated as a fallacy, usually assumed to be eradicable, and therefore is perhaps thought to be eliminable from well-kept classrooms. Our analysis in terms of heterogeneity and localness of interpretation strongly suggests that prevarication is not eliminable. We cannot use unique terms for every meaning, and should not if we could. The use of the same term is often essential to anchor the term to the shared concept as the details shift through its various guises. Perhaps signalling when this is likely to be a problem would help? And perhaps teaching teachers to detect the seams that have become transparent for them between systems is an important aim?

But these observations from the classroom are just as important for theories of the semantics of representations. The conventional response to the kind of observations of language we have made here is that everyone knows that natural language is ambiguous. It is easy to acknowledge heterogeneity if a system contains language and diagrams—here the heterogeneity is on the surface. But the idea that natural language consists of many heterogeneous sub-systems is generally resisted, and explained away as polysemy at the lexical level. There are at least two problems with this explanation. The number of polysemous readings required is essentially infinite, and the meaning of one word is systematically related to that of others. Words in these discourses do not function atomistically—they are part of subsystems. If ‘birthrate’ is construed one way, then its contrasting terms such as ‘deathrate’ and ‘survivalrate’ will also be construed in related ways—at least until there is a shift to a different subsystem. Recently, (e.g., Moravcsik, 1998) theories of lexical meaning have paid more attention to the considerable distance between the generalities of the lexicon and the details of contextualised language use. These stratified theories are much more conducive to understanding real language use and the heterogeneous nature of most reasoning.

In learning to get from a real world problem into a formalism, and back out from the formal results of calculation to an implication about the real world, students must cross many experiential worlds and, when working in groups, negotiate complex patterns of authority for knowledge which determine what the group actually does. In these tapes, we again and again see transitions between the world of numbers and the

world of fish and mice. At one point the discourse is entirely numerical and insulated from the real world consequences, as witnessed by the acceptance of completely implausible values. At others, there is a sudden jumping out of the mathematical world to references to the death of a pet fish, or gee, that’s a lot of nasty meeces! Although formalisms distance proceedings from affective states, when we reason about the world, our reasoning should still be animated by affect. We will not understand conceptual learning until we can give an account of how representations, the social arrangements for authority in discourse, and our experiential worlds are all coordinated.

Acknowledgements

We would like to acknowledge invaluable comment and discussion with Randi Engel, Muffie Weibe, Jim Greeno and Rogers Hall. We are grateful for their generous support in allowing access to their data gathered under an NSF grant to Hall. We also acknowledge fellowship support from grant GR #R000271074 from the Economic and Social Science Research Council (UK), and CSLI’s support for the first author’s research at Stanford.

References

- Barwise, J. & Etchemendy, J. (1994). *Hyperproof*. CSLI Publications, Stanford.
- Greeno, Sommerfeld & Weibe (2000). Practices of questioning and explaining in learning to model. Proceedings of the 22nd Meeting of the Cognitive Science Society, Philadelphia.
- Hall, R. (2000). Working the interface between representing and represented worlds in middle school math design projects. Proceedings of the 22nd Meeting of the Cognitive Science Society, Philadelphia.
- Hall, Rogers (1999). Case studies of math at work: Exploring design-oriented mathematical practices in school and work settings. Final report to the National Science Foundation.
- Moravcsik, J. M. (1998). *Meaning, creativity, and the partial inscrutability of the human mind*. CSLI Publications: Stanford, California.
- Oberlander J., Monaghan, P., Cox R., Stenning K. & Tobin, R. (1999). Unnatural language discourse: an empirical study of multimodal proof styles. *Journal of Logic, Language and Information*, **8**, 363–384.
- Stenning, K. Cox, R. & Oberlander, J. (1995). Contrasting the cognitive effects of graphical and sentential logic teaching: reasoning, representation and individual differences. *Language and Cognitive Processes*, **10** (3/4), 333–354.
- van Someren, M., Reimann, P., Boshuizen, E. & de Jong, T. (1999). (eds.) *Learning with multiple representations*. Kluwer.