Grounding Symbol Structures in Space: Formal Notations as Diagrams

David Landy (dlandy@indiana.edu)
Department of Computer Science
Indiana University, Bloomington, IN 47408 USA

Robert L. Goldstone (rgoldsto@indiana.edu)
Department of Psychological and Brain Sciences
Indiana University, Bloomington, IN 47408 USA

Abstract

Although a general sense of the magnitude, quantity, or numerosity is common both in untrained people and animals, the abilities to deal exactly with large quantities and to reason precisely in complex but well-specified situations—to behave formally, that is—are skills unique to people trained in symbolic notations. These symbolic notations employ typically complex, hierarchically embedded structures, which all extant analyses assume are the product of concatenative, rule-based processes. The primary goal of this article is to establish, using behavioral measures on naturalistic tasks, that the some of the same cognitive resources involved in representing spatial relations and proximities are also involved in representing symbolic notations. In short, formal notations are used as a kind of diagram. We examine self-generated productions in the domains of handwritten arithmetic expressions and typewritten statements in a formal logic. In both tasks, we find substantial evidence for spatial processes even in these highly symbolic domains.

Keywords: symbolic processing, mathematics, embodied cognition, relational reasoning

Introduction

It is clear that mathematical equations written in modern notation are, in general, visual forms, and furthermore that they share some properties with diagrammatic or imagistic displays. Equations and mathematical expressions are often set off from main text, use non-standard characters and shapes, and deviate substantially from the linear placement of text. Furthermore, evidence indicates that at least some mathematical processing is sensitive to the particular visual form of its presentation notation (Campbell, 1999, McNeil & Alibali, 2004). Nevertheless, notational mathematical representation is typically considered 'sentential,' and placed in opposition with diagrammatic representations in fields as diverse as education (Zazkis et al, 1996; Stylianou, 2002), philosophy of science (Galison, 1997), computer science (Iverson, 1980), and cognitive studies of problem-solving (Anderson, 2005; Stenning, 2002).

The standard conception of mathematical notation is best understood via Stephen Palmer's (1978) classic distinction between intrinsic and extrinsic representational schemes. A representation is intrinsic "whenever a representing relation has the same inherent constraints as its represented relation" (p. 271). Line A being shorter than Line B can be intrinsically represented by the representational element that corresponds to A being shorter, taller, brighter, or larger than the element representing B – any relation that is inherently asymmetric and transitive. Representations are extrinsic when their inherent structure is arbitrary. They model the represented world by explicitly building the necessary structure so as to conform to the world. Palmer argues that analog representations are intrinsic; correspondences and inferences between represented and representing worlds come for free because of their shared intrinsic structure. Propositional representations, including language, logic, and mathematics, are extrinsic, and hence come to represent objects by explicitly establishing relations with whatever structure is needed. The only intrinsic relation necessary to propositions is the left-right concatenation of basic symbols. Although traditionally understood as extrinsic, it is possible that representations in mathematics and logics nonetheless possess intrinsic and analog properties, and it is this possibility that we empirically pursue here.

Specifically, we propose that formal notations are diagrammatic as well as sentential and the property conventionally described as syntactic structure is cognitively mediated, *in part*, by spatial information. Elements of expressions are "bound" together through perceptual grouping—often induced by simple spatial proximity. Thus, our claim is that mathematical formalizations of syntax are not themselves the direct cognitive mechanisms typically employed in processing that structure. The former really are concatenative, but we propose that people use space and spatial relationships in representational schemas to facilitate the processing of syntax. To be clear: we are not claiming here that the execution of each individual step in a proof or computation is inherently spatial or processed exclusively using sensori-motor mechanisms. We do suggest that spatial reasoning over the physical layout of notational forms is common in reasoning with formal languages, and that spacing practices play a significant role in human reasoning using notations.

We have argued previously that a broadly similar interference of metric (non-order-related) spatial properties on syntactic judgments provides evidence that syntax processing typically involves spatial and perceptual grouping processes (Landy & Goldstone, in press; Kirshner & Awtry, 2004). These works conclude that mathematical reasoners are sensitive to non-formal properties of presented equations, and in particular align close proximity with high operator precedence. For instance, in Landy & Goldstone (in press), participants were much more likely to attribute equality to “n + w*y + b = y + b*n + w” than to the formally equivalent “m+p * e+g = e+g * m+p”.

The fact that people are sensitive to small changes in the physical spacing of formal expressions suggests that symbolic interpretation is processed, in part, through spatial reasoning. If so, then people might also respond spatially to the syntax of internally represented expressions, leading symbolic productions to reflect syntactic structure, e.g., the more tightly two mental terms are bound syntactically, the closer they should be written physically. A relationship between spatial and syntactic proximity would be at best unexplained if space is not part of how we represent formal syntax. Study 1 directly examines our proposal by measuring physical inter-operand spacing in handwritten equations constructed by participants from presented word equations.

Study 1

In this study, participants were asked to write out simple equations by hand. If, as we propose, formal notations automatically encode spatial relations corresponding to structural relations, then spacing in handwritten equations should reflect the formal structure of the equation. In particular, equality spacing should be very large, since equality signs denote, in all cases, the broadest partition of the sentence. Within the 2-operator side of each equation, spacing should depend on the structure of the expression. In mixed expressions, the middle term is syntactically “bound” to the higher-order multiplication sign, and so the spacing around that sign should be compressed (or that around the lower-order addition sign expanded) relative to its unmixed spacing. Thus, we expect a dependency of operator spacing on the interaction between operator and context.

Although traditional sentential accounts of notational reasoning provide no reason to expect operator spacing to ever be non-uniform, experience with typeset or handwritten equations might drive some kinds of spacing regularities. Typeset equations are generally not fixed-width, and multiplication symbols are generally narrower than addition signs. Experience with typeset equations could lead equation writers to generally space multiplicands more narrowly than addends. However, no prominent equation typesetter adjusts the spacing of terms based on syntax (and if one did, we would regard this as evidence favoring our view!): therefore, while either the spatial encoding or the amodal perspective might predict a main effect of operator, neither the width of the operators themselves nor experience with typeset equations could lead to the predicted interaction. Thus, the principle theoretical measure is the interaction between operator sign and mixed/unmixed structure.

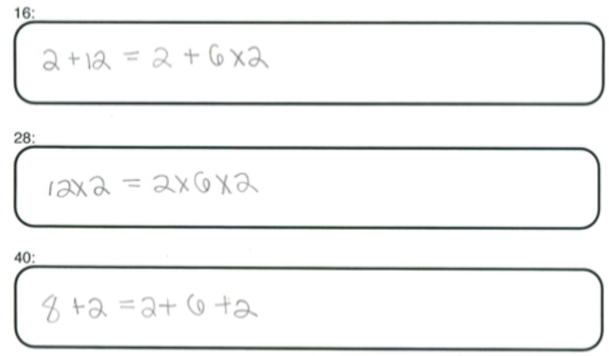

Figure 1: Sample responses in Study 1.

Method

24 Indiana University undergraduates participated in the experiment, which fulfilled a partial course requirement. This experiment lasted about 25 minutes.

Word equations were presented one at a time on a computer; participants wrote out corresponding equations using standard mathematical symbols. Participants were instructed to use only standard Arabic numerals and formal operator symbols (‘+’, ‘x’, ‘=’), and were explicitly asked not to use any parentheses. Participants were not asked to solve or evaluate the correctness of any equation, nor were they reminded of the correct order of operations.

For each participant, 10 triples of numbers between 2 and 9 were randomly generated (because 1 is much narrower than other numbers, it was excluded from this experiment)

Syntactic binding between terms was systematically manipulated by altering the operator of equations. Each triple appeared in 4 different equations, one with each of the operator structures plus-plus, plus-times, times-plus, and times-times, making 40 translations in all. The other side of the equation contained the same expression, but with the first operation completed. Thus, if the number triple was {2,4,9}, the equations would be 6 + 9 = 2 + 4 + 9, 2 + 36 = 2 + 4 x 9, 8 + 9 = 2 x 4 + 9, and 8 x 9 = 2 x 4 x 9. The middle two equations are labeled *mixed* operator conditions, while the first and last are termed *unmixed* (since there is no particular hierarchical structure on either side of the equation). In 5 of the 10 equation sets, the triple appeared as the right-hand side of the equation; in the other 5, the triple appeared on the left. This procedure eliminates any interference of particular number choices, since each production is compared to productions that are identical except for operator context.

Each participant received a different, randomly generated stimulus set. Word equations were presented on a computer screen, one at a time, and remained on the screen while participants wrote the corresponding symbolic equation in a printed box (1.1cm high by 10.4cm wide, see Figure 1) on a piece of paper. Word equations employed number words along with the words “times,” “plus,” and “equals”. For instance, if the word equation probe was “six plus five times four equals two plus nine times three,” participants would respond by writing “6 + 5 x 4 = 2 + 9 x 3”. Each participant viewed 40 equations in total. Responses that were left blank, contained parentheses or other extraneous marks, or contained crossed out values or other errors were dropped from the analysis. The measure was the distance between the innermost points of each pair of adjacent operands.

Results

For each participant, spacing was averaged across the stimuli in each condition. The mean distances across participants for each context are shown in Table 1. These mean values were analyzed using a 2-way within-participants ANOVA, using distance as a dependent measure, and operator and expression structure as independent categorical variables. As predicted by the typesetting hypothesis, the ANOVA revealed a main effect of operator type: multiplicands were spaced more closely than addends, ($F(2, 46)=7.9$, $MSE=3.35$, $p<0.01$), and equality signs were spaced substantially more widely than either ($F(2,46)=105.7$, $MSE=249.2$, $p<0.001$). The interaction between operator type and expression structure was also significant: participants' compression of multiplicands relative to addends increased in the mixed condition ($F(1,23)=4.726$, $MSE=1.28$, $p<0.05$).

Table 1: Mean spacing (with standard error) by measurement condition (mm).

Operator	Unmixed	Mixed	Overall
Addition	9.52±.48	9.77±.48	9.65±.48
Multiplication	9.38±.50	9.16±.53	9.27±.51
Equality	12.38±.59	12.13±.57	12.25±.58

Discussion

Study 1 indicates that syntax in arithmetic equations, at some level, is processed automatically. Although the simple transcription task requires no consideration of syntax at all, the results show a modulation of productions in virtue of syntactic structure. Moreover, writers are not swayed arbitrarily by syntax—they construct spatial properties that match their psychological groupings. Terms were spaced more narrowly when they were grouped more closely. Historical interactions with typeset equations do not predict these effects, nor do traditional symbolic accounts of mathematical competence. So of the hypotheses considered earlier, this result is compatible only with the suggestion that participants systematically vary spacing according to the particular syntactic structure of the current equation. Because this behavior presumably generalizes to the population at large, including the teachers and parents of our particular participants, historical interactions with other hand-written equations could account for the results—participants in our task might reasonably be sensitive to syntax because their teachers and parents were, and therefore sensitivity of spacing to syntax forms part of the participants' training. That is, our participants may have received more training with mathematical expressions in which the spacing is consistent with the syntactic structure. However, this explanation does not provide any additional insight into why this spacing convention has been adopted in the first place. For that, the most parsimonious account for the environmental regularity is, once again, that spatial processes are involved in the representation of mathematical syntax in the normal course of algebraic reasoning.

There is a plausible alternative to the spatial information hypothesis: it might be that syntax processing (somehow) mediates access to the lexicographic forms for numbers and symbols. If syntactically bound items are chunked in memory, for instance, then access to terms within a chunk may be quicker than access to terms across chunks (Cheng & Rojas-Anaya, 2006). If so, and if horizontal pen movement between characters correlates with access time (if for instance the pen is moved at some more-or-less fixed velocity while the lexical form is being accessed), then a memory delay could produce increased spacing. Study 2 addresses this possibility by exploring spacing behavior on a typed input task.

Study 2

A limitation with Study 1 is that the formal system used was a small (though important) one: pre-algebra using equality, addition and multiplication. While this system is convenient in that it is widely known and studied, nevertheless generalizing from such a small system is difficult. Study 2 broadens the scope of our examinations by exploring a very different notational system: formal propositional (quantified and unquantified propositional) logic. Instead of asking participants to write unused and useless pseudo-equations in a laboratory setting, Study 2 involves a corpus analysis of self-generated expressions by participants interacting with a Web-based teaching tool designed and maintained by Colin Allen and Chris Menzel, and based on an accompanying textbook (Allen & Hand, 2001).

Method

The current analysis is based on the "Logic Daemon and Quizmaster" (Allen & Menzel, 2006, http://logic.tamu.edu/) which is an interactive Website designed for use with the Logic Primer textbook by Allen & Hand (2001). Although the site is publicly accessible, the primary users are probably students in introductory logic courses, who can use this Website to tackle exercises found in the textbook as well as additional problems of the same type (because the data we analyze was anonymous and already collected, informed consent was not obtained. This research was approved as exempt by Indiana University IRB (approval number 06-11025). For our analysis, we chose to focus on the 'translation' exercises which require students to render sentences of English into the formal system described in chapters one (propositional) and three (predicate logic) of Logic Primer. In these exercises, students are presented with up to five English sentences and under each sentence is a standard Web form single-line text input field. Students freely type a response for one or more of the sentences and click a submit button. For instance, problem 9 from problem section 1.3 (in chapter 1) states that “If Mary dances although John is not happy, Bill will dance”. Instructions state that “Q”, “S”, and “R” are to be used to denote the atomic sentences “Mary dances”, “John is happy”, and “Bill dances”, respectively. Participants enter a formal sentence corresponding to this sentiment (one correct response would be “((Q & ~S) -> R)”).

Each string of characters submitted in this way is checked first to see whether it represents a well-formed formula (wff) according to the specifications of the formal system. Although the textbook uses non-ASCII characters, these are mapped to ASCII strings for keyboard input; specifically, the single arrow is represented as '->' (dash-greater-than, double arrow with '<->' (less-than-dash-greater-than), the upside-down 'A' of universal quantification with '@', and the backwards 'E' of existential quantification with '$'. The wff formation rules specify the use of parentheses around binary sentential connectives: '&' (and), 'v' (or), '->' (if...then), and '<->' (if and only if). Some of these parentheses may also be omitted following a formal convention that is defined in chapter one of the text. The parenthesis-dropping conventions follow the specified order of operations: & and v precede ->, and -> precedes <->. Any string that passes the wff test is next checked for correctness with respect to the particular translation problem attempted (logical equivalency to a stored answer. Both the wff check and the correctness check is indifferent to any white space introduced by the student, and when problems are returned to students with feedback, any spaces are removed.

It is worth noting that although interface, formal system, physical situation, and participant pool and motivation are different from Study 1, the task is quite similar: in both cases, participants are asked to take a natural language statement and translate it into a formal system.

Analysis & Predictions

129,526 submissions to the translation verification interface (exercises 1.3 and 3.2) submitted between May 5, 2005 and April 4, 2006 were collected and analyzed. Junk entries, entries which used incorrect symbols (e.g., “=>” instead of “->”), and repeated entries coming from a single IP address on a particular day were removed. After all of these reductions, 48,131 statements from 595 unique IP address/time stamp combinations remained. The same verification system used to provide submission feedback was used to categorize submissions for our analysis.

We distinguish three physical spacing conditions: spacing *consistent* with the operator structure, spacing *inconsistent* with operator structure, and no spacing at all (*unspaced*). An expression is considered consistent when the space around every operator in the expression is appropriate: spaces around conjunction, disjunction, conditional, and bi-conditional signs should be even, spaces should only be inserted to the left of negation signs and quantifiers. If any spacing violated these constraints, then the expression was flagged as inconsistent.

Our predictions are as follows: first, we predict that because representations of space play a role in the way reasoners process syntax, participants using the site will at least occasionally insert spaces. Although random insertion of spaces would be far more likely to produce inconsistent than consistent spacings, we predict that .spacing will be primarily consistent, and that only consistent spacing will improve performance.

Study 1 indicated that spacing is modulated in the presence of hierarchical syntax, participants should be more likely to produce spaces in responses containing more than one operator. Also, we predict that accuracy will be higher on consistently spaced statements than on other statements, whenever structure matters (i.e., whenever there are two or more operators).

The problems studied came from two sections of the book; one on propositional logic (chapter 1), and one on predicate logic (chapter 3). Because the latter is from a later section of the book, submissions on this section can be assumed to come from more advanced reasoners; one might wonder whether such reasoners will space more or less frequently and regularly than beginners. Both possibilities are compatible with out hypothesis. One might think that advanced reasoners should have a better mastery of spatial systems of reasoning, and consequently space more. On the other hand, one might think that advanced reasoners have internalized the appropriate perceptual structures, and do not need physical cues to indicate them. Because more experienced reasoners are less likely to be dependent on perceptual support (Chi, Feltovich & Glaser, 1981), we also predict that more advanced participants will be less likely to space expressions at all. Our prediction then is that spacing will be more common in the first, more elementary section, suggesting that advanced reasoners have internalized the appropriate perceptual segmentation.

In analyzing these data, we do not attempt to evaluate the statistical significance of our results for two reasons: first, the breakdown of submissions by unique IP address/date does not adequately divide submission into independent samples. Since we have no way to determine unique individuals, and furthermore no way to determine the relationship between individuals, statistical tests based on the assumption of independent samples are inappropriate. Second, the large size of the sample guarantees that standard statistical measures will indicate significance (all of the contrasts considered here are highly significant by standard measures), regardless of the underlying mechanisms. For these reasons, we report frequencies without invalid statistical measures.

Results

Table 2 presents the frequencies of submission broken down by spacing, number of operators, and logic type. As expected, participants frequently spaced expressions. 10.8% of all expressions submitted contained some spacing. When

Table 2: The use of spacing in typed formal translations.

Type	Logic Type					
	Propositional			*First Order*		
	U	C	I	U	C	I
Single Op						
Correct	1,344	160	11	1,187	12	0
Incorrect	1,592	145	15	735	9	0
Multiple Ops						
Correct	6,175	1,058	159	12,899	1,044	165
Incorrect	5,182	844	213	13,817	1,011	339

Results from Study 2. Entries specify number of submissions of each category. U denotes unspaced submissions, C those in which spacing and operations were consistent, and I submissions with at least one inconsistently spaced operation.

expressions were spaced, moreover, they were predominately consistently spaced: 82.6% of all spaced equations were consistent with operator syntax. Consistently spaced submissions were also more likely than either inconsistent or unspaced expressions to be correct; 53% of consistent equations were correct, compared with only 37% of inconsistently spaced and 50% of unspaced equations.

In order to test the structure sensitivity of consistent spacing, we divided the dataset according to whether a problem required syntax resolution (that is, whether it had two or more connectives). Participants did indeed space more frequently on multi-operator problems (9.2% of multi-operator problems were consistently spaced against 6.8% of few-operator problems). Furthermore, accuracy was highest (53%) when expressions were consistently spaced, and lowest on those inconsistently spaced (37%; 50.1% of all unspaced expressions were correct), principally on multi-operator expressions. On few-operator problems, 52% of consistent, 42% of inconsistent, and 52% of unspaced submissions were correct. However, the number of problems of this type is very small (there were only 26 inconsistently spaced few-operator submissions in all).

We also tested the theory that more training would reduce the need for formally extraneous spacing. Translation problems appear in two sections of the textbook: Chapter 1 (propositional logic) and Chapter 3 (predicate logic). We divided the full dataset into these two categories, and measured spacing frequency across these two categories; both consistent and inconsistent spacings were more frequent on propositional problems (13.1% and 2.4%, respectively) than on predicate logic problems (6.7% and 1.6%).

Discussion

Despite being formally unnecessary and informally discouraged, spaces were frequently inserted into typed sentences of formal logic. These spacings were nearly always consistent with the operations they abutted; submissions with consistent spacing were also more likely to be correct than unspaced submissions. Together with Study 1, Study 2 establishes that people working in two very different domains systematically space formal systems that do not require differential spacing.

Because the participants in this study were typing on a keyboard, a chunking account that predicts differential spacing on handwritten equations as a result of the time-course of memory retrieval within and across chunks cannot account for spacing here. Other accounts of the results of Study 2 are possible, however. For instance, other formal systems, such as programming languages, are often taught with explicit instructions to space logical terms; participants might be transferring this practice from programming experience. This possibility cannot be definitively eliminated (though why spacing is common in programming is *still* mysterious), but the fact that more experienced reasoners space less, not more, seems incompatible with the idea that appropriately spacing is an acquired skill.

General Discussion

In both typed logic and handwritten arithmetic translation tasks, participants created formally irrelevant spatial relationships in stimuli. In both cases, these relationships aligned with the syntactic structure of the formal statement being expressed. The kinds of regularities produced in Study 1 have been shown to benefit correct syntactic interpretation (Landy & Goldstone, in press). People seem to spontaneously create alignments of space and syntax that help them reason formally.

That spacing is connected to syntax is important for our understanding of mathematics and mathematical learning, but it is also important for education and cognitive psychology more generally. For education, our results suggest increased sensitivity to the physical features of how mathematics is presented to students and how they present mathematics to their teachers. Physical properties such as spacing may be used to give students a perceptual scaffold for the rules underlying algebra. Further research is necessary to know whether these scaffolds, when removed, help students to continue to obey the appropriate mathematical rules or if they act as crutches that thwart rule development. Reciprocally, by examining students' spacings of their own productions, we may be able to diagnose their misunderstandings. In the same way that manual gestures are sensitive indicators of inchoate explicit mathematical understandings (Alibali & Goldin-Meadow, 1993; Goldin-Meadow, Wein & Chang, 1992), production spacing may indicate the beginnings, or lack thereof, of knowledge for order of precedence.

Most fundamentally, our results challenge conceptions of symbols as amodal and divorced from analog, spatial information. In this respect, we offer a reinterpretation of Newell and Simon's (1963, 1976) influential "Physical Symbol System Hypothesis." Their hypothesis was that physical symbol systems had the necessary and sufficient means for producing intelligent action. A symbol system includes both physical symbols such as marks on paper or punches on a computer tape, and the explicit rules for manipulating these tokens. In action, all of their physical symbols were distantly related to their worldly referents, and were digital and discrete entities such as the strings "P ⊃Q" and "GOAL 7 TRANSFORM L3 INTO LO." The arbitrary nature of these entities was by design because they wanted symbols to be able to designate any expression whatsoever without any *a priori* prescriptions or limitations. We concur with Newell and Simon's emphasis on *physical* symbols, and believe in paying even more attention to symbols' physical attributes involving space, shape, and perceptual grouping. Accordingly, our revised physical symbol systems hypothesis is that symbols are not arbitrary, unconstrained tokens, but rather are represented and processed using space and perceptually organized groups. This conception of physical symbols makes them far more constrained than those underlying Newell and Simon's General Problem Solver, but these constraints are not only limiters, but permitters as well. For Specific Problem Solvers that are humans, it is good policy to design symbols that can be processed efficiently given what we know about perceptual and cognitive mechanisms. From this perspective, it is hardly surprising if the symbols we write look a lot like those that we are good at reading, and if the symbols we think with are a lot like those we are good at thinking.

Acknowledgments

This research was funded by Department of Education, Institute of Education Sciences grant R305H050116, and National Science Foundation ROLE grant 0527920. The authors thank Colin Allen for his support, both material and conceptual.

References

Allen, C. & Hand, M. 2001 *Logic Primer, 2nd edition*. Cambridge, MA: The MIT Press.

Allen, C. and Menzel, C. (2006) Logic Daemon and Quizmaster website (http://logic.tamu.edu/). Accessed Sep 28, 2006.

Alibali, M. W. & Goldin-Meadow, S. (1993). Gesture-speech mismatch and mechanisms of learning: What the hands reveal about a child's state of mind. *Cognitive Psychology 25*, 468-523.

Anderson, J.R. (2005). Human symbol manipulation within an Integrated Cognitive Architecture. *Cognitive Science 29*:313-341.

Campbell, J.I.D. (1999). The surface form x problem size interaction in cognitive arithmetic: Evidence against an encoding locus. *Cognition* 70, B25-B33.

Cheng, P. Rojas-Anaya, H. (2006). A temporal signal reveals chunk structure in the writing of word phrases. *Proceedings of the 28th Annual Conference of the Cognitive Science Society*. Vancouver, B.C.

Chi, M. T. H, Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science, 5*, 121-152.

Galison, P. (1997). *Image and Logic*. Chicago, IL: University of Chicago Press.

Goldin-Meadow, S., Wein, D. & Chang, C. (1992). Assessing knowledge through gesture: Using children's hands to read their minds. *Cognition and Instruction, 9*(3), 201-219.

Iverson, K. E.(1980). Notation as a tool of thought. *Communications of the ACM, 23*(8), 444-465.

Kirshner, D., & Awtry, T. (2004). Visual salience of algebraic transformations. *Journal for Research in Mathematics Education, 35*(4), 224-257.

Landy, D. & Goldstone, R. (in press). How abstract is symbolic thought?

McNeil, N. M., & Alibali, M. W. (2004). You'll see what you mean: Students encode equations based on their knowledge of arithmetic. *Cognitive Science*, *28*(3), 451-466.

Newell, A., and Simon, H.A. (1963). GPS: A Program that Simulates Human Thought. In E.A. Feigenbaum and J. Feldman. (eds.), *Computers and Thought.* McGraw-Hill, New York.

Newell, A., and Simon, H.A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM 19*(3), 113--126.

Palmer, S. E. (1978). Fundamental aspects of cognitive representation. In E. Rosch & B. B. LLoyd (Eds.), *Cognition and categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates. (pp. 259-303)

Stenning, K. (2002). *Seeing Reason: Image and Language in Learning to Think*. Oxford University Press, Oxford.

Stylianou, D. A. (2002). On the interaction of visualization and analysis: the negotiation of a visual representation in expert problem solving. *Journal of Mathematical Behavior*, 21, 303-317.

Zazkis, R., Dubinsky, E., & Dautermann, J. (1996). Coordinating visual and analytic strategies: A study of the students' understanding of the group d4. *Journal for Research in Mathematics Education, 27* (4), 435-456.