THE STRUCTURE OF KNOWLEDGE IN COMPLEX DOMAINS

Edwina L. Rissland
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003

TR # 80-07

This work supported in part by NIE grant G-80-0096.
Any opinions, findings, conclusions or recommendations expressed in this report are those of the author and do not necessarily reflect the views of the U.S. government.

Table of Contents

I. INTRODUCTION

Objectives

Learning, teaching and understanding are complex tasks in complex domains like mathematics. Any teacher or student knows that there is rich complexity to the subject of conic sections or continuity. Knowing such domains is much more than knowing formal statements of definitions, theorems and proofs: there are examples; methods; exercises; pictures, graphs and diagrams; rules of thumb; metaphors and analogies; folksy or informal formulations of ideas; a sense of what is important; a sense of what should follow what. This complexity exists whether the purpose is to learn or to teach.

One objective here is to present a conceptual framework in which to describe and understand knowledge in complex domains. Another is to show how such a framework can aid teaching and learning in them.

There are several underlying themes. First, understanding a domain requires a great familiarity with its connections. Polya and Szego [1972] said eloquently:

> There is a similarity between knowing one's way about town and mastering a field of knowledge: from any given point one should be able to reach any other point. One is even better informed if one can immediately take the most convenient and quickest path from the one point to the other. If one is very well informed indeed, one can even execute special feats, for example, to carry out a journey by systematically avoiding certain paths which are

customary...

> There is an analogy between the task of constructing a well-integrated body of knowledge from acquaintance with isolated truths and the building of a wall out of unhewn stones. One must turn each new insight and each new stone over and over, view it from all sides, attempt to join it on to the edifice at all possible points, until the new finds its suitable place in the already established, in such a way that the areas of contact will be as large as possible and the gaps as small as possible, until the whole forms one firm structure.

Second, being aware of what one knows and how one comes to know it are important ingredients of expertise. Disambiguation and explication of the process of knowing enable one to be a smarter student or teacher, as Papert and others have often pointed out [Papert 1971, 1980][Buchanan and Headrick 1970].

Thus, we are engaged in epistemology, the study of knowledge. The epistemology of complex domains includes the study of their content, structure, representation, and use. There are many dimensions and levels to epistemology. There is study of the domain itself, for example, through analysis of its content and structure. There is the epistemology of one domain with respect to different conceptual frameworks; and the epistemology of one compared with another with respect to the same framework. There is the epistemology of epistemologies, that is, the study and comparison of frameworks.

In this paper, we will discuss not only the epistemology of particular domains, but, more importantly, also a general epistemological framework. Our approach reflects concerns and methodologies from cognitive science and artificial intelligence (AI). Important contributions of AI are the tools of information processing and knowledge representation [Feigenbaum and Barr 1981] [Winston 1977]. To be able to use them means that not only should we know the knowledge, but also what we want to do with it.

Our analysis is based upon the experience of the author and others as teachers and learners. The structural notions presented here seem natural to students. Several students, e.g., in calculus, have found that using such notions helps to give them understanding that is different from that of acquiring a large bag of tricks. Our conceptual framework is currently being applied to investigations in learning and understanding probability [Myers 1980], discrete mathematics for computer science [Wegner 1981], and high school mathematics [Davis 1979].

One of the themes running through our work is the importance and interaction of examples with other knowledge such as concepts and results. The role played by examples in the evolution of mathematical concepts is beautifully described by Lakatos [1976]. Examples play a critical role in learning and concept formation by machines [Winston 1975] [Lenat 1976] as well as people [Polya 1965][Collines 1979]. Examples can also be a very effective means of presenting new technical ideas [Wegner 1980].

Of course, attention to and understanding of the topic of heuristics and its role in discovery are due to the Polya. His well-known books are full of insight, e.g., [Polya 1965]; they also provide many examples of the kind of knowledge we wish to explicate further.

Structure and Representation

In this paper, we distinguish between "structure" and "representation". Structure is something inherent in the knowledge domain which can be teased out by examination of the domain and how it is used. Representation is the way in which we encode the knowledge and structure.

We can see structure, or different aspects of it, when we examine the domain. The structure discerned is clearly influenced by our purposes, expectations and point of view [Kuhn 1970], [Bruner 1973]. For instance, if we work with situation-action rules, such rules will jump out as we examine a domain (e.g., [Greeno 1978]). With different purposes, we emphasize different aspects. For instance, if our purpose is to study logical development, we attend to deductive aspects like definitions, theorems and proofs; if our purpose is to study expository style, we attend to rhetorical structures. In this paper, our purpose is to understand the knowledge better, so as to improve our skills as teachers and learners, and to enhance our understanding of understanding.

Often there are many ways to represent the same knowledge; the one chosen also depends on our perferences and purposes -- on what we are going to use such knowledge for. For instance, researchers in logic and theorem proving often use the predicate calculus. We choose the representation scheme that best suits our purposes; "best" in the sense of conceptual clarity, ease of encoding or use. For instance, representation schemes like "frames" [Minsky

1975], "scripts" [Schank and Colby 1973], "schemas" [Rumelhart and Ortony 1977] and "semantic nets" [Quillian 1968], [Brachman 1979] make it easy to encode clustering, typing, and connectivity.

An underlying viewpoint in our approach is that chunking of knowledge is very important. In fact, what might distinguish an expert from a novice is the kind and degree of chunking used in thinking about a domain. As one becomes more expert one is able to use bigger chunks, more encompassing and general. For instance, the expert mathematician sees a proof as a coherent whole, perhaps with substructure [Rissland 1978a], while the novice may see only the succession of individual proof steps. The novice's view is very <u>local</u>; the expert's can be <u>global</u> as well. Experts can switch the level of description easily and naturally.

This pattern is also found developmentally: young children often cannot summarize a story; they retell it line-by-line; older children can present a synopsis. Adults often use more global organizations like "scripts" [Schank and Abelson 1977].

II. A CONCEPTUAL FRAMEWORK

Structure: Items, Relations and Spaces

There are several kinds of structure in the knowledge of complex domains like mathematics:

(1) strongly bound clusters of information: for example, the statement of a theorem, its name, its proof, a diagram, an evaluation of its importance, remarks on its limitations and generality. In this paper, we call such clusters items.

(2) relations between items: for example, the logical connections between results, such as predecessor results on which a result depends logically and successor results which depend on it.

(3) sets of similar types of items related in similar ways: for example proved results and their logical dependencies. Such a set of items and relations constitute a space, for instance, Results-space.

What distinguishes a "space" from a "set" is the prominence of the relations. The structure of a complex domain like mathematics contains not just one but many spaces, each of which describes a different aspect of knowledge. In mathematics, for instance, as found in lectures and textbooks, e.g., [Rudin 1964], there are spaces of:

results: lemmas, propositions, theorems, corollaries, etc., related by logical dependency, i.e., how one result is deduced from others;

examples: illustrative situations and cases related by constructional derivation, i.e., how one is built from others;

concepts: formal ideas like definitions, and informal ones like heuristic principles related by conceptual dependency, i.e., how one concept is defined or presented in terms of others;

exercises: exercises and unresolved conjectures related by how one exercise is generated from others, e.g., as a subproblem that would contribute to solution of a larger problem;

procedures: procedures and methods related by procedural dependency, i.e., how one procedure is composed of or depends on others, as in program subroutines;

strategies: control methods, doctrines or heuristics related by how one strategy engenders or contributes to the practice of others;

goals: goals and purposes related through dependencies of super- and sub-goals [Newell and Simon 1972], i.e., how one goal engenders or contributes to the attainment of others.

In our freedom to use many spaces, we have departed somewhat from the usual semantic net approach in AI since we use many nets, i.e., spaces, simultaneously. Each space represents a different "cut" through the domain. In summary, if the same type of relation can be seen in a set of items, then we can describe it as a space.

Mathematical Knowledge

While all of the above mentioned spaces, and probably others, would be needed to construct an exhaustive epistemological framework for describing complex domains in general, a significant part of the knowledge in domains like mathematics and computer science can be captured in the following spaces: Results, Examples, Concepts, Exercises, and Procedures.

Results

The set of proved results and their logical dependencies constitute Results-space. For instance, in plane geometry [Jacobs 1974], the "Isosceles Triangle Theorem", which states that if two sides of a triangle are equal, then so are the angles opposite them, depends logically upon the side-angle-side postulate which is used in its proof; a corollary result depending on this theorem states that an equilateral triangle is equiangular.

Examples

An important aspect of examples is that one builds examples from other examples. The relation of constructional derivation emphasizes this and allows the collection of examples to be coherently organized. For instance, in order to show that not all second degree equations represent conic sections, one can modify the example of the equation of a circle

x**2 + y**2 = 4

to x**2 + y**2 = -4. Thus, the second example item would point back to the first, and the first would point forward to the second.

Another instance of the construction of examples is the derivation of the example $y = x\sin\frac{1}{x}$ near the origin from the reference examples y = -x, y = x and y = sinx.

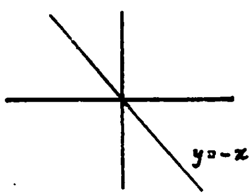

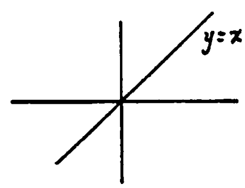

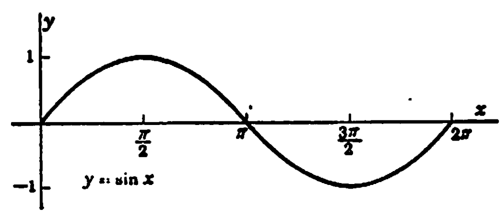

by squeezing the oscillations to occur between the lines y=x and y= -x.

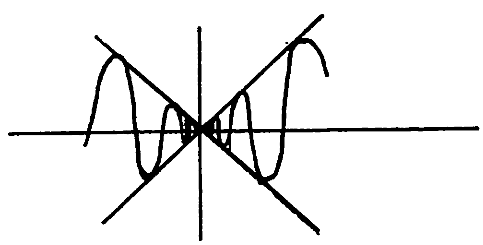

The generation of examples is an interesting study in itself [Rissland, 1980], [Gelbaum and Olmstead 1964], [Steen and Seebach 1970].

Concepts

Concepts include formal and informal ideas. A formal idea is a precisely stated notion like a definition, e.g., the familiar "epsilon-delta" definition of continuity given in calculus. An example of an informal idea is the paraphrase of continuity as "not lifting the pencil from the paper". It is not a bona fide mathematical definition but it is an idea that is frequently used. Informal ideas also include individual heuristics of two types: mega-principles and counter-principles. A mega-principle is a "big" idea that either is an informal truth, such as "Continuous functions have no gaps or breaks", or suggests a good line of attack, such as "Try factoring the polynomial into linear factors". A counter-principle, on the other hand, summarizes troublesome points or offers a warning, like "Be careful about dividing by quantities that might be zero" or "Don't forget to calculate the new dx".

Concepts can be related by dependencies of definition or pedagogy. A definition can be definitionally dependent on another if it uses the other in its formulation; for instance, the limit definition of continuity uses the concept of "limit". Pedagogical dependency is more subjective: it reflects our feeling that one concept should be learned or introduced before another. Some authors map out the conceptual dependencies of topics and chapters in their books e.g., [Dunford and Schwartz 1958], [Royden 1963].

Exercises

Exercises play an important role in helping a student. They can be related in several ways, as in the use of the solution or method of one exercise to pose or solve another. For instance, solving the cubic equation

x**3 - 1 = 0

involves solving the quadratic equation

x**2 + x + 1 = 0

More complicated, and often subtle, relations between exercises are found in higher level mathematics courses like real analysis.

Procedures

In many domains, like computer science and applied mathematics, there are large collections of methods and procedures. Procedures can be related by the dependency of one procedure on another. This relation is important in analyses of skills like arithmetic, where for instance addition is a predecessor skill for multiplication [Davis and McKnight 1979]. The ordering is often related to student "bugs" [Seely-Brown and Burton 1977].

The dependencies can be used to organize procedures hierarchically. For instance, the procedure to find the roots of a cubic equation like the one above is composed of the procedures to factor it as

$$(x - 1)(x^{2} + x + 1)$$

and then to pick off one root, 1, by observation, and, finally to compute the other two by using the quadratic formula. For cubic equations with not so nice a form, the solution procedure might include trying to guess a root, which itself involves procedures, and then proceeding as before.

Representation-graphs

The dependency relations in the spaces allow us to treat the spaces as directed graphs where the nodes represent items and the arrows, the relations between them; the direction shows the inherent predecessor-successor ordering. In this way, we have have several graphs: results-graph, concepts-graph, examples-graph, etc.

Thinking of the spaces as directed graphs is a useful conceptual aid since it allows us to picture them easily and from the picture to discern connections and the overall complexity of the domain. In fact, the use of the word "space" is itself an aid since it calls to mind a mathematician's idea of space with its notions of coherence and distance. The graphs are a way to picture the representation spaces that represent the structure of the domain; the graphs themselves are not a part of the domain.

Structure: The Notion of Dual

An item is related to other items in its representation space through the space's dependency relations. In addition, an item is related to items in the other spaces. The notion of dual concerns these inter-space relations.

As an illustration, consider results, examples, and concepts, which together constitute a large part of the knowledge in many mathematical domains. Results are connected to examples and concepts as well as to other results. The dual items of a result include: (a) the examples motivating and illustrating it; (b) the concepts needed to formulate it and those derived from it.

Similarly examples and concepts are connected to items outside of their spaces. The dual items of an example include: (a) the results illustrated by it and proving things about it; (b) the concepts illustrated and suggested by it. The dual items of a concept include: (a) results leading to its formulation and those proving things about it; (b) examples motivating and illustrating it.

Examples, results and concepts also have dual items from the spaces of procedures, goals, exercises, etc. Exercises are closely tied to the procedures, examples, concepts and results used in their solution and those that the exercise illustrates and affords practice with. In many high school texts, the relations between exercises and worked out examples are particularly close.

Strategies, goals and procedures are often strongly connected because the strategies are dual to the goals that they are useful for achieving; the procedures are ways to implement or carry out the strategies. In programming, this is often called "pragmatics" as compared with the syntax and semantics.

Thus each item is related to certain items from the other spaces: we call the set of those items, its dual. Duals themselves have structure: some dual items tend to precede the item in teaching and learning, others, succeed it. For instance, motivating examples for a result come before the result; other examples, like those showing the limitations of the result (e.g., that the converse is not true) come after it. Thus, we can distinguish pre and post duals. These of course may depend on pedagogical and cognitive style.

The power of the dual notion is that it allows us to associate items that might not be closely related with regards to in-space relations. For instance, two concepts might be widely separate in Concepts-space (e.g., because they are taught many chapters apart), but are closely related in the dual sense in that they share many examples. An important part of understanding is building up such associations; assigning exercises which draw on seemingly distant concepts is a way to help students establish such dual relations.

Taxonomies

Another important component of knowledge is the knowledge that not all items serve the same function in learning and understanding. For instance, experts (teachers and learners) know that certain perspicuous ("start-up") examples provide easy access to a new topic, that some ("reference") examples are quite standard and are always exhibited as illustrations, and that some examples are anomalous and don't seem to fit into one's understanding.

We can develop a taxonomy of items based upon how we use them to learn, understand and teach a domain. In discussing Concepts-space, we have already mentioned different kinds of concepts. Briefly, some important taxonomic classes of items are:

(1) Results

(a) basic results: basic, first proved results;

(b) key results: major, frequently used results;

(c) culminating results: goal results;

(d) transitional results: intermediate logical stepping-stones;

(e) technical results: results establishing technical details;

(2) Examples

(a) <u>start-up</u> examples: perspicuous, easily understood and presented cases;

(b) <u>reference</u> examples: standard, ubiquitous cases;

(c) <u>counter</u> examples: limiting, falsifying cases;

(d) <u>model</u> examples: general, paradigmatic cases;

(e) <u>anomalous</u> examples: exceptions and pathological cases;

(3) Concepts

(a) <u>definitions</u>: formal definitions and specifications;

(b) <u>informal paraphrases</u>: informal formulations;

(c) <u>mega-principles</u>: "big" ideas like certain heuristics;

(d) <u>counter-principles</u>: heuristic warnings;

One can also distinguish kinds of procedures, exercises, and goals. Such a taxonomy is not an exclusive classification since an item can serve more than one role; for instance, some start-up examples become reference examples as one learns more in the domain. It is also surely not exhaustive.

Worth Ratings

Not all items are of equal importance. One can use a worth rating, such as a scheme of *'s (as in the Michelin guidebooks):

- * an interesting item, worth noticing;
- ** an important item, worth a "stop",
- *** a very important item, worth a "detour";
- **** an extremely important item, worth a "journey" in itself;

The worth rating and the taxonomy are related in that certain classes tend to be worth more than others; culminating results tend to have three and four-star ratings, whereas technical or transitional results might not even rate one star.

Just as items can be given a worth rating, so too can relations. For instance, the dual relations of examples to a given result probably have varying degrees of worth and importance: a telling counter-example, say to the converse of a theorem, is probably much more important as a dual example than an example which is just another instance of the theorem. Tagging the relations provides a way of describing the worth of an item relative to another item. Thus in addition to the "global" worth of an item with respect to the entire knowledge domain, items can have "local" relative worth. In this way, for instance, a 1-star example which is usually not very important can be recognized as very important in relation for a particular item.

Context

In describing a domain, we must also include contextual information such as assumptions about what domains precede it and which items can be taken as known or axiomatic; these might be culminating results of a predecessor domain. Also important is the state of knowledge of the learner in studying the domain, e.g., prerequisite skills and material.

In mathematics the "same" set of items can be studied in several mathematical contexts or settings: one can study operators in finite dimensional vector spaces (i.e., matrices) or in Hilbert space or in normed linear spaces. The settings themselves can be organized into a space with a generality relation or what AI researchers refer to as "is-a" or "ako" ("a kind of") hierarchies.

Since our purpose in this paper is to concentrate on a domain within a given context, we shall not dwell upon the influence of setting other than to say that it is important. We will keep our discussion localized to one context.

Representation Frames

In our representation of an individual item we include the kinds of information described above. We use the "frame" format [Minsky 1975] to represent an item each of whose facets or subparts is represented in a "slot" of the frame. Slots can contain declarative, procedural or relational information. In the following discussion, slot names occur in capital letters.

Items have a NAME, like "continuity". The CLASS of an item is its taxonomic class. The WORTH slot contains an evaluation of worth, perhaps in terms of zero to four stars. Its SETTING is the context in which it is valid or presented.

An item can be presented in more than one way: as a declarative STATEMENT, in terms of its derivation or PROCEDURE, or by a PICTURE or diagram. For instance, the declarative statement of an example states what the example exemplifies; its procedural aspect describes its construction; its picture is a schematic diagram or plot or kinetic sequence of pictures. The declarative aspect of a result is a statement of its hypotheses and conclusions (e.g., in "if-then" form); its procedural aspect is its proof. A concept can be stated declaratively or in terms of a procedure that implements it or that provides a way to test whether an item is an instance of it. For instance, the "derivative" can be defined in terms of limits or in terms of differentiation procedures.

The declarative aspect of a procedure is a statement that says what it does; its procedural aspect is its code. The declarative aspect of an exercise is the problem statement; the procedure could be its solution.

Relational information is encoded in slots containing pointers to other items. IN-SPACE pointers point BACK to predecessor items and FORWARD to successors. DUAL-SPACE pointers point to dual items.

The pointers, or sets of pointers, themselves have structure. We have already mentioned the pre- and post- distinction for dual pointers and that pointers can also have worth ratings. Pointers can also be partioned into clusters of dual items that can be thought of as a group bound to the item (e.g., a cluster of examples varying slightly only in their numerical values) or sets of predecessor items that are sufficient to derive the item (e.g., when there may be more than one way to derive a proof from other proved results).

An item can have REMARKS such as NOTA BENE's (NB) and CAVEATS which point out particularly noteworthy or critically limiting things about the item. An item might have PRAGMATIC data telling what the item is good for. In addition an item has PEDAGOGICAL information regarding how and when to present it, like materials needed for a demonstration of it, or leading questions to ask about it; BIBLIOGRAPHIC data like references to books, articles, films; APPLICATIONS data pointing to real-world applications. OTHER data might include historical remarks.

III. AN EXAMPLE

In this section, we map out and discuss a section of a high school algebra text introducing the important notion of "function". This extended example will illustrate how one can distil the ingredients of knowledge from a standard mathematics text and represent them in the conceptual framework developed in the previous sections.

It should be said that in extracting this topic for illustration we lose some of its richness; it will seem shallower and less connected than it would in the context of the entire text or course. However, even in the few pages of text examined, there is a surprising amount of material.

The examples will be taken from a widely used text Dolciani, Wooton, Beckenbach and Sharron [1968]. Our example deals with the topic of functions and relations (Sections 4-1 and 4-2). It assumes background about sequences of real numbers, which was covered in the previous chapter.

Section 4-1 begins with a discussion of the example of the negative, odd integers, which can be thought of as constructed from the important reference example of the odd integers:

$$-1, -3, -5, \ldots, 1-2n, \ldots$$

Using this example, the authors introduce the notion of "pairing" and "ordered pair"; the third number in the sequence is paired with -5, through the formula 1-2n, which evaluates to -5 for n=3, and, of course, by simple counting.

The authors then define the concepts of "coordinates", "components" and "equality of ordered pairs". Equality of ordered pairs is then illustrated with

$$(3, -5) = (3, 1-2*3).$$

Inequality is illustrated with

$$(3,-5) =/= (-5,3).$$

Thus the first pairs serve as an illustration (post-dual example) to the "equality of ordered pair" concept; the second, a derivate of the first, as a limiting counter-example.

The discussion returns to the primary topic of functions:

> "Sequences are special examples of an important mathematical concept, called a function. A function is any pairing of the members of one set (the domain) with the members of another set (the range) so that each member has exactly one partner in the range. Thus, the infinite arithmetic sequence
>
> -1, -3, -5,...
>
> is a function whose domain is {the positive integers} and whose range is {the negative odd integers}."

Thus, using the negative odd integers as a start-up example, the authors have introduced the definition of function. If we were to fill in representation frames for the negative odd integers example and the function concept (as presented so far), they would be:

NAME: Odd negative integers
CLASS: start-up
SETTING:

DECLARATIVE STATEMENT: This example introduces the concept of function through the concept of ordered pair.

PROCEDURAL STATEMENT: $a_n = 1-2n$

IN-SPACE POINTERS: BACK: Odd integers
FORWARD:

DUAL-POINTERS: PRE-CONCEPTS: Sequence
POST-CONCEPTS: Ordered pair, Function

(Note that into the frame we are entering only the knowledge presented in the text. Putting in additional knowledge that reflects our own understanding is quite a different exercise.)

NAME function
CLASS definition
WORTH ***
SETTING

DECLARATIVE STATEMENT: A function is any pairing of the members of one set (the domain) with members of another set (the range) so that each member has exactly one partner in the range.

IN-SPACE POINTERS: BACK: Ordered pair
FORWARD:

DUAL-POINTERS: PRE-EXAMPLES: Odd negative integers
POST-EXAMPLES:

NB: Each member of the domain has exactly one partner in the range.

(We have set the worth at three stars because Dolciani et al have used the word "important".)

To build a complete representation, we would also need to create frames for the concepts of "ordered pair", "coordinate/component", and "equality of ordered pairs". The frame for "ordered pair" would have a PRE-EXAMPLES DUAL pointer to the "negative odd integers" example which was used to motivate it; it would have FORWARD IN-SPACE pointer to the concept "coordinate/component" which in turn would have a FORWARD pointer to the concept "equality of ordered pairs", which would point BACK to "coordinate/component" which in turn would point back to "ordered pair".

Some of the interconnections among the examples and concepts introduced thus far is shown in the following fragments of the examples- and concepts-graphs:

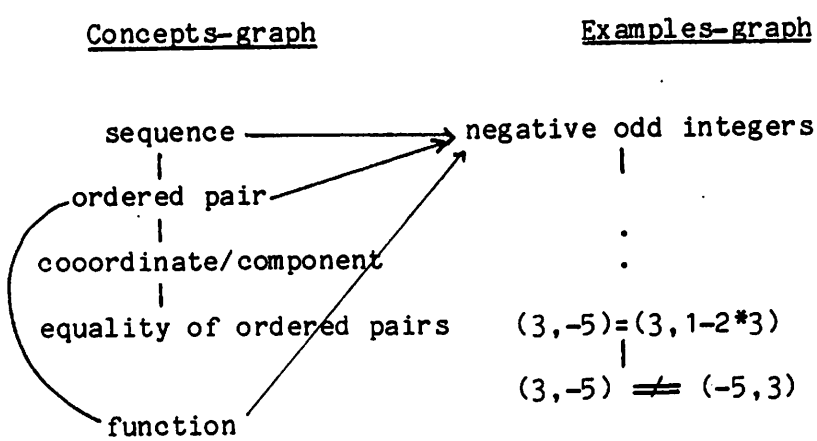

Next, the terminology of value and belong to is explicated. We have the following statement which the authors italicize for emphasis:

> "Thus, you can think of a function as a set of ordered pairs in which different ordered pairs have different first coordinates."

This paraphrase can be thought of as the first variation on the theme of function; we could represent it as a concept which follows from the "function" concept and inherits its slot values (unless otherwise noted) by default from its predecessor, the "function" concept. Alternatively, in our representation frame, we could introduce a slot for PARAPHRASE.

Next, we begin to acquire some procedural knowledge about functions:

> "One way to specify a function is to show its ordered pairs in a table with the members of the domain named in one column (or row) and the corresponding values of the function in another column (or row)."

A simple example illustrates the function table idea:

Domain Element	Function Value
-1	-1
-2	-1
3	1
4	1

Note that the authors are combining the first few integers, particularly +1 and -1, with a sprinkling of minus signs, to generate an example. One could say that this is an instance of the "meta"-heuristic of "using plus and minus one's" for generating examples.

Next comes a nota bene remark on the preceding example and the function concept:

> "Notice that this function does not have a different value corresponding to each different member of its domain."

This NB deserves comment and encoding as a counter-principle, which we could name the "not necessarily one-to-one" counter-principle or CP(not necessarily one-to-one). It is closely connected with the "exactly one" NB of the "function" concept, which also is acquiring the status of a counter-principle, which we name CP(exactly one). As the authors state in their teacher`s manual:

> "This is a good place to make sure that students understand that although each element of the domain of a function is paired with exactly one element in the range, there is no restriction as regards ... an element in the range."

The preceeding function table example is then re-presented as an example of specifying a function by a set of ordered pairs:

$$\{(-1,-1)\ (-2,-1),\ (3,1),\ (4,1)\}$$

The section's start-up example (negative odd integers) is further used to illustrate the table procedure. It is also used to introduce specification of a function in terms of a formula and the use of a function symbol:

> "The formula $a_n = 1-2n$, $n \in$ {the positive integers} provides a compact way to show pairings in the sequence and suggests a convenient way to indicate function pairings in general. However, let us replace the symbol "a<n>" with the function symbol
>
> a(n), read 'a of n'..."

At this point there is another warning to the teacher:

> "Be sure that students do not call the formula a function. The formula does, however, specify or define a function over (in this case) {the positive integers}."

This is really another counter-principle, which we name CP(formula vs function).

To update our function frame, we would now add the function table and formula information to the PROCEDURAL STATEMENT slot, another example pointer, and pointers to the CP(not necessarily one-to-one), CP(exactly one), and CP(function vs formula) counter-principles.

We now come to "Example 1", the first explicitly labelled example, the standard reference example of the absolute value function over the domain of the real numbers. A discussion of the implicit definition of the domain, i.e., the domain of definition, is presented briefly.

From these remarks, we also learn that the tacit setting for all of the discussion is the real numbers, R.

"Also, it is agreed that the domain, unless otherwise specified, consists of those real values of x for which the formula provides a unique real value of f."

This "agreement" is illustrated with the function

f(x) = 1/x**2

Note that the authors have begun to use the letter "f". This seems unremarkable to anyone who has studied functions; it is part of the tacitly held knowledge that "f's", "g's" and "h's" are often used for functions, and "a's" for sequences. Furthermore, the second function (1/x**2), while not explicitly derived from the first ($|x|$), is not unrelated to it, because the absolute value function can be defined in terms of the square root:

$$\sqrt{x**2}$$

Thus, the authors are really using examples coming from a cluster of examples around the x**2 example. This knowledge is not explicitly present in the section of text under examination, but is present in the larger fabric of algebra and mathematics in general. The connections probably help to make the discussion more coherent, at least to the authors and the teachers.

Example 2 is, in fact, the important reference example x**2, which is ubiquitous in discussions about functions, polynomials, parabolas, etc. The example is presented as:

$$h = \{(x, h(x)) : h(x) = x**2 \}$$

It is used as an illustration of the concepts of function, domain, range, and domain of definition; thus its frame would have pre-dual pointers to these concepts and their frames would have post-dual pointers to it.

The next example, Example 3, is a counter-example serving to differentiate between the concept of "set of ordered pair" and "function". It can also be thought of as a counter-example to highlight the "exactly one" counter-principle or nota bene remark (NB) of the function concept:

> "Let g be the of all ordered pairs of real numbers such that g = {(x,y) : xy**2 = 1}."

(Notice the use of second power again.)

This counter-example for "function" serves as a start-up example for "relation", which is then defined in terms of the "ordered pair" concept. This concludes Section 4-1.

Thus at the conclusion of the text of this first section on functions, we have several elements in Examples- and Concepts-spaces, and none in the others (although we have begun to build up knowledge of procedures, like graphing, which could be represented as procedure-items when they acquire "critical mass".) We have spent most of the discussion on the "function" concept and a large part on the "negative odd integers" example. Thus, at in a very few pages (three, to be exact), we have established the "function" concept.

In a high school text, exercises play an important role. The exercises for Section 4-1 are used to provide the student practice with the new concepts. Thus they serve mostly a dual role. A few harder exercises (in group "C") serve as nice introductions to other topics by encouraging the student to discover new knowledge.

The exercises are organized according to the concepts which they treat: i.e., by dual relations to concepts. In other sections, the exercises are often dual to examples in the sense that they ask the student to work through an example or procedure.

In the second section on functions (4-2), the authors begin by discussing the notion of "mapping", introduced with an example of mapping in the everyday sense.

This notion is then pulled into the function discussion:

> "You can think of any function as a mapping of its domain into its range."

Thus the authors have introduced the second important variation of the function concept, which we could call the "function as mapping" mega-principle.

The important schematic of the "mapping diagram" quickly follows:

Domain	Range
5	1
10	2
20	2
30	3

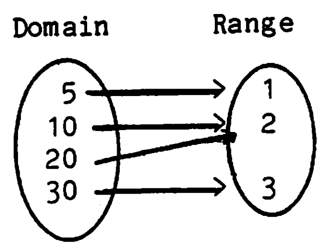

Thus the procedures of making a function table and a mapping diagram are linked. Next comes the graphing procedure:

> "Another way of picturing a function depends on the fact that ordered pairs of real numbers can be graphed in a plane."

There is next a digression about graphing and coordinate systems. Before giving examples, the authors state the graphing procedure for a function and the definition of "graph":

"To picture a function in the coordinate plane, plot the graphs of all the ordered pairs of the function. The set of points obtained is called the graph of the function."

Note that the "graph" concept has been defined procedurally as the result of the "graphing procedure".

Drawing a graph is illustrated with Example 1 of Section 4-2

$$1 - |x|,$$

which is closely related to Example 1 of Section 4-1, the absolute value function

$$f(x) = |x|.$$

Before Example 2,

$$g: x \longrightarrow 3x-5, \quad x \in \{1,2,3,4\}$$

the authors introduce another way of representing a function as a "mapping statement" that involves the use of the arrow. Example 2 is not particularly noteworthy, except that once again it is built around the integers 1,2,3,4.

Example 3, $L = \{(x,y) \mid x=4\}$, is used as a counter-example to distinguish between the "function" and "graph" concepts. It is dual to "exactly one" counter-principle. It is also an example for the "relation" concept. Its graph, which lies on a vertical line, gives a pictorial representation for the "exactly one" CP.

The oral exercises for this section are quite nice, because they exercise both the mapping diagram and graph representations for functions. Included are examples that are not functions. There are several "favorite" reference examples for the concepts of one-to-one, onto, bijection, continuity, differentiability, etc. For instance, the graphical exercises (nos. 5-16) contain: (1) a function composed of other functions (#9); (2) a function with a "jump" discontinuity (#10); (3) a graph which is a well-known figure, but not a function, i.e., a circle (#11); (4) a step function (#13); (5) a non-function of x, but a function of "y" (#15); (6) an absolute value function, -¦x¦ (#12); and the absolutely standard parabola y=x**2 (n 14).

Thus in two short sections there has been a tremendous amount of knowledge clustered about the function concept.

While it is not reasonable to expect anyone to analyze or map out all of a mathematics text in this detailed a fashion, it is useful to examine a section or two in this way to highlight the amount of knowledge that is presented or tacitly assumed, and the degree of interconnectedness of it. Such an excercise also serves to highlight the complexity of the task of learning and understanding new knowledge.

IV. IMPLICATIONS FOR TEACHING AND LEARNING

In addition to representing the knowledge contained in textbooks, the conceptual framework can be used to organize what we know about a subject. The process of filling in the frames prompts us to recall and disambiguate knowledge. It can also help structure and represent lessons, which for instance, could be thought of as a sequence of frames.

The framework can be used explicitly to help students understand or explore a domain, for instance, by asking them to map it out or fill in frames. The frame implicitly contains questions to be asked about an item (e.g., "How important is it, and why?"). Students, at least at the college level, enjoy the mapping process and find it helpful to their acquisition of understanding (e.g., when reviewing the subject). In some ways the mapping process is more important than the representation finally produced. Exercising their knowledge by directing their attention to important items and relations gives focus to their efforts at synthesis. Using specific questions that probe and prompt understanding, such as found in [Rissland 1978], gives the students a way to actively pursue understanding, a process often expected to happen magically.

Thus the framework suggests new homework and test questions, like those to probe dual connections (e.g., "Give two examples of this concept: a standard case and a not so "nice" case."). Students are not asked to give enough examples, much to their detriment since generating examples forces them to understand the involved concepts, results, etc. One can encourage them to explore duals in another way by asking why two items, e.g., theorems, are related (even thought they are found in different chapters). Such questions will encourage students to become active investigators.

For teachers, the conceptual framework offers guidance in choosing items for presentation. For instance, knowing that some examples are better than others for introducing a concept encourages one to look for such a "start-up" example to introduce a new concept. Knowing about the knowledge encourages one to tell others about it. It is easier to share one's knowledge if one has a way of talking about it. It also reminds one to tell students important learning and understanding heuristics that one often forgets to mention because they are so thoroughly assimilated: for instance, instead of hoping that the students will come to recognize an example as a standard, why not tell them it is a "reference", which they should think of as an old friend to be used to test out conjectures and new definitions.

Thus, a conceptual framework for knowledge in complex domains can help organize known knowledge and direct exploration of new. By knowing about knowing, one can be more expert in learning and teaching.