Cognitive Structures and Processes in the Interpretation of Mathematical Proofs

Adrian Robert

SHORT TITLE: Cognitive Structures in Mathematical Proofs

Key words: cognitive semantics, mathematics, reasoning, conceptual structure, knowledge representation

Adrian Robert

2394 Bishop St. #4

Ann Arbor, MI 48105

ph/fx: 734-764-6195

email: arobert@cogsci.ucsd.edu

Abstract

We propose that mathematical concepts are cognitively represented as blended (Fauconnier & Turner, 1998) combinations of schemas acquired through everyday experience; the properties of these source schemas generate the relational characteristics of the concepts that are explored in mathematical proofs. Proofs are understood by constructing novel blends of the concepts under consideration, and they are verified through paying attention to the presence or absence of dissonance in the associated mappings. We analyze these mappings through examples and compare the analysis with the results of an empirical survey of students' mental processes in studying the proofs. The results suggest explanations for the trends of increasing rigor and conceptual 'fine-grainedness' in the historical development of mathematics, and they potentially shed light on processes of natural language understanding because the mechanisms at work in the two cases appear qualitatively similar.

1 Introduction

Mathematical reasoning is an interesting subject for study within cognitive science for several reasons. Firstly, mathematics represents an apparent extreme of human thinking in that it involves more abstract concepts and greater apparent precision than any other area, and the study of a process under extreme conditions can often provide useful information as to its nature. We shall argue in particular that the same processes are active in understanding mathematics as understanding natural language, but that they are more clearly manifest in the former. Secondly, mathematics has one of the longest recorded histories of any human intellectual endeavor, and in the West and Middle East this history is continuous and cumulative, each succeeding generation having built upon the work of the past. This history represents an incomparably rich corpus of material on the development of a human conceptual system. Finally, mathematics is applied as a tool in many other domains of human endeavor: science, engineering, business, and even religion have all employed it to varying degrees of advantage. Insights into cognitive processes in these domains might therefore be gained through an understanding of them in mathematics.

In this paper, we suggest that tools from the field of cognitive semantics, in particular, the concepts of *schema* and *cognitive mapping*, may be profitably applied to further this understanding. The idea of cognitive mapping, a process in which correspondences are drawn between two different conceptual domains, has been applied to analyze metaphor understanding (Lakoff & Johnson, 1980; Lakoff, 1993) as well as analogy-making (Gentner, 1983; Holyoak & Thagard, 1995). The correspondences are conceived of as being drawn from or between cognitive structures summarizing knowledge of classes of objects or relations – often conceptualized as schemas.

Lakoff (1987) and Johnson (1987) proposed that schema-based conceptual mapping and combination processes play a role even in reasoning processes not ordinarily thought of as analogical – in particular, in mathematical reasoning. This proposal suggests that abstract mathematical concepts are grounded in experientially-based schemas, and reasoning with the concepts proceeds according to the correspondences we make between the properties of the grounding schemas and the specific configurations arising in mathematical situations. Here, we carry this proposal forward by examining in detail how the concepts specified in mathematical definitions are grounded and how they are manipulated in the understanding of mathematical proofs.

Grounding occurs through mathematical definitions themselves, which make reference to nonmathematical concepts. We propose that people draw on a cognitive capacity to align and relate different concepts to form a composite, blended structure whose properties are understood, at least at first, by connection to the sources. We will show through analysis that mathematical proofs can be construed as sets of instructions directing the interpreter to form correspondences between conceptual structures involved in the definitions. It is the pattern of mentally observed consonance and/or dissonance arising from the resulting mapping of structural information that leads the interpreter to be convinced of the validity of a proof's conclusion. This stands in contrast to the more conventional conception of mathematical proofs as sequences of logical propositions that follow in necessary fashion from one another by rule application (although this *is* true of a small subset of mathematics), but it receives support both from student reports we collected and statements made by practicing mathematicians.

Our view generates insight into what makes some proofs and proof steps more difficult to understand than others, and it also suggests explanations, related to the need to reduce ambiguity in communication, for certain trends in the historical development of mathematical practice. In particular, both increasing 'granularity' in concepts and increasing standards for rigor in proofs can be understood as enhancing reliability in the processes we describe. Finally, our analysis clarifies interesting similarities between mathematics and other human intellectual activities as well as critical points of difference.

Previous approaches to mathematical reasoning within cognitive science outside of cognitive semantics (Anderson et al., 1981; Koedinger & Anderson, 1990) have focused on problem-solving and the acquisition and application of strategies and special-purpose knowledge structures to this end. They have not addressed the issue of how proofs or definitions are *interpreted* or the nature of the connection between mathematical concepts and structures derived from other domains of experience. The present treatment is therefore complementary to this previous work; however, both the general form of the knowledge structures we propose (schemas) and some of the processes we describe make contact with it, in largely a non-contradictory way.

This paper is organized as follows. In **Section 2**, we marshal the concepts from cognitive semantics to be employed in our analysis. In particular, we clarify our notion of *schema*, and we introduce the *blended spaces* framework developed by Fauconnier & Turner (1994, 1998), a powerful new conception of the cognitive mapping process. **Section 3**, on different modes of mathematical practice, clears some ground for the proposal in **Section 4** on the grounding of mathematical concepts. Our proposal elaborates ideas presented in Lakoff (1987: chap. 20) and Lakoff & Nuñez (1997, 2000) using blended spaces.

In **Section 5**, we apply blended spaces a second time in conjunction with this grounding to analyze the mental process involved in interpreting a simple mathematical proof, and in **Section 6** we describe the categories of mental operations and proof syntax that were developed in performing the analysis. **Section 7** presents the results of an empirical study of mathematics students' reasoning processes in understanding proofs that supports the analysis in certain respects. The parallels and differences between interpretation in mathematics and natural language are discussed in **Section 8**, and some insights into each domain generated by the mapping from the other are presented. Finally, **Section 9** outlines prospects for future research.

2 Concepts from Cognitive Semantics

2.1 Introduction: Metaphor and Mapping

The theory of cognitive mappings grew out of linguistic research on metaphor. In a seminal work, Lakoff & Johnson (1980) discussed a number of different types of metaphors, making the point that in addition to subserving communication, they affect thinking – by determining likely inferences and associations and helping to support focus on subtle aspects of abstract concepts. They suggested that the use of a metaphor in speech reflects an underlying mental correspondence between some aspects of a well-understood source domain and a target domain of interest. For example, we use and understand the sentence *"This discussion is getting off track,"* by drawing a correspondence between the typical situation of a traveling person or vehicle – in which leaving a 'track' may entail making slower progress or getting lost – and the situation of speakers engaged in a conversation with some goal in mind, such as exploring some particular set of ideas.

The word 'typical' represents the intuition that knowledge is mapped at a relatively general level: the source domain is conceived of as made up of one or more *schemas*, involving general relations between categories, not specific relations between instances. Schemas represent abstractions from experience with entities in the world and their interactions (Rumelhart, 1980; Quinn & Holland, 1987, D'Andrade, 1995: c. 6, Mandler, 1984)[1] and are rich, heterarchical structures which may relate to both physical and more abstract domains. For example, a schema for "long thin object" (Lakoff, 1987) involves a visual category as well as a set of potentially available motor manipulations, a schema for "cane" involves many of these same aspects but also adds a number of other properties such as 'curved at one end', 'typically used by elderly people', and so on.

Below, we first describe more precisely the conception of schema we use, then we discuss the more com-

[1] The terms *frame* (Minsky, 1975; Fillmore, 1985) and *script* (Schank, 1982) refer to similar notions.

plex composites such as "typical situation of traveling person or vehicle" which act as sources or targets, and finally we describe the mapping process itself in greater detail.

2.2 Schemas

Langacker (1987) provides a useful classification of schemas[2]: he divides them into *Things* (roughly, categories represented by prototypical members) and *Relations*. The latter consist of a set of one or more Things related to each other in a particular way – for example, the schema underlying the word 'above' consists of two Things joined by a particular spatial relationship. The fact that the arguments in a Relation are not objects but Things (i.e., categories of objects) captures the fact that a variety of different objects can play their roles. The relation within a Relation is similarly composite – for instance, spatial Relations like 'above' and 'over' apply to a number of subcases (distinguished by contact between the objects and other properties) that are sometimes divided in other languages (see Lakoff, 1987: c. 24; Choi & Bowerman, 1991; Bowerman, 1996).

For our purposes, we will define a schema as either a Thing concept (of objects, actions, events, or anything else) or a set of concepts (which we will sometimes refer to as 'slots') connected by a relation-category. The details of the structure of concepts and how they are represented are under debate in cognitive science (e.g., Komatsu, 1992), but we will assume that they:

1. are *acquired through experience*, either directly, through learning, or indirectly, through the effects of past natural selection on brain structure;

2. contain some representation of the *range of permissible values* for attributes of the class in question; and

3. contain some representation of the *possible, necessary, and impossible* interrelations between subsets of attribute values.

The first of these assumptions is not controversial. Regarding the second, we assume that representations contain this information (Homa & Vosburgh, 1976; Komatsu, 1992), but we will not concern ourselves with such issues as whether the range is represented through storage of instances or complex prototypes

[2]He considered only schemas underlying word meanings, however from the viewpoint of his theory these structures are basic elements fundamental to cognition in general. In particular, concepts and all of their relational properties are assumed to be represented by networks of schemas.

of some kind (Hintzman, 1986; Anderson, 1990). The third assumption falls out of similarity-based categorization models sensitive to local correlations between subclusters of attributes (Medin et al., 1982; Malt & Smith, 1984; Anderson, 1990) and is a feature of more sophisticated relation-based models (Lakoff, 1987; Goldstone et al., 1991). In addition to helping account for certain aspects of concept learning and category judgment data, this structure can act as a constraint on concept combination and cognitive mapping more generally. For instance, to explain people's intuitions about the extensions of (possibly novel) concept combinations on the basis of the components, a representation of necessity and impossibility relations among attributes has been helpful (Hampton, 1987). We will discuss the role of this structure more fully below when we introduce blending.

To avoid a long side discussion, we will not attempt to state explicitly how attribute relations are represented or which aspects are derived in working representations versus being recalled from static ones. The assumption that detection of and reaction to necessity and impossibility occurs in using categories is sufficient for our purposes. We do, however, find it useful to distinguish between two types of impossibility, which we refer to as *incompatibility* and *incommensurability*. An incompatible combination does not occur in the real world but could conceivably occur. For example, the attribute "purple fur" is not compatible with the other attributes associated with the category 'cats', but we can nevertheless imagine purple cats as existing in some possible world. On the other hand, an incommensurable combination could not occur even conceivably. For instance, the attribute "green" (as a color) is not commensurable with any of the attributes associated with the category 'ideas', and the attribute "has a handle" is not commensurable with anything in the category "wind". The difference between knowledge of compatibility and that of commensurability has been discussed by Keil (1979), however we do not wish to claim this is a simple dichotomy, for commensurability shows some characteristics of a continuous parameter. This is clear from the two examples of incommensurability we have just given, for it is easier to think of metaphorical applications of "green" to 'ideas' (e.g., ideas which are new and lacking in support from other quarters, or those which appear to be weak – "Under pressure from the professor's criticism, the student's ideas began to look a little green.") than of "handle" to 'wind'. We will suggest that incompatibility and incommensurability play a role in determining the structure of viable mathematical concepts.

2.3 Composite Structures

Typically, as briefly described in the introduction to this section, the mappings observed in metaphors operate between complex configurations. Schemas are the reusable parts making up these configurations. For example, the source domain for "This discussion is getting off track," involves a traveler and a path, while that for "Vanity is the quicksand of reason," involves a traveler and an obstruction. It is widely believed that schemas contain at least some and possibly all of the information about how they can combine in configurations such as these, but the configurations themselves may not be stored as static structures but rather constructed in working memory to suit the conditions of a particular situation (Schank, 1982; Fauconnier, 1985, 1997; Barsalou, 1983, 1987). The level of complexity at which this transition occurs is not well agreed-upon (and is likely to be domain- and individual-dependent), nor is the division of storage of composability information between schemas and some less modularized store of "background knowledge", however we will not be concerned with these issues.

We shall use the term *schema* to stand for a relatively general-purpose cognitive entity at roughly the level of the single lexical item, the terms *property* and *attribute* to refer to components at a lower level, and the term *frame* to refer to multi-schema configurations, which in at least some cases are likely to be represented as transient constructs rather than permanent structures.

2.4 Conceptual Blends

2.4.1 Overview

The theory of *conceptual blending* (Fauconnier & Turner, 1994, 1998) was developed to encompass a broad set of cognitive phenomena including metaphor which involve cognitive mapping. Other instances include analogy, categorization, grammatical combination (of verb frames and their arguments), evolution of socio-cultural models, and creative thought – and, in many if not most cases of these, the result of the mapping process itself begins in working memory but can potentially assume properties of a separate cognitive structure in its own right. That is, it can be stored for later access, as with the entrenchment of meaning of novel compounds (Hampton, 1987), and it can play a role as source in subsequent mappings, as with elaborative applications of scientific metaphors such as the 'mind is a computer'. Within the blending framework, this phenomenon is accomodated by postulating the construction of a *blend* incorporating aspects of each input structure.

The description of blending builds on the concept of a *mental space* (Fauconnier, 1985), a working memory construct containing one or more interrelated schemas and, possibly, instances. Through a process of mapping, elements in two or more mental spaces become aligned and a new space is constructed deriving aspects of its structure from each of the inputs. The criteria and mechanisms for alignment and construction are not well understood, but the main idea in the simplest cases is that instances in one space that fit the category specifications for slots in relations within the other are aligned with them – and the resulting blend is an impression of the relations holding between these instances (Fauconnier & Turner, 1998).

For example, in the sentence "Sally is the mother of Paul," one input space contains the relation 'mother of', complete with slots for the two elements thereby related, while the other contains the specific instances Sally and Paul. The resultant interpretation – and the structure in the blended space – is of the 'mother' relation holding between Sally and Paul. Here, the conventions of grammar actually guide the choice in mapping Sally and Paul to the "mother" and "mothered" slots respectively, but in the construction of the blend there is more occurring at a finer level that is not so specified, as is revealed by considering the sentence "Necessity is the mother of invention." Here, 'necessity' and 'invention' are again aligned with the same slots of "mother" and "mothered", but only certain aspects of the relation "mothering" are carried over into the blend, namely, those concerned with nurturing and fostering and not those concerned with biological relationship or provision of food and other support during childhood (see Figure 1). In fact this carrying over of only part of an input is true even of cases like "Sally is the mother of Paul," for depending on context this could have the meaning that Sally is the biological mother of Paul (but perhaps did not raise him) or that she adopted him, or that she is a warm, caring person who provides him with a lot of encouragement (see Lakoff, 1987: c. 6-8).

Note, Figure 1, for practical reasons, leaves out much of the detail that is assumedly part of the 'mother' relation and the entities taking part in it, leaving only enough to illustrate the point at hand. Furthermore, the figure should not be taken to imply we are claiming the existence of a definite 'MOTHER schema' with a particular structure. The left-hand input could be some dynamically-constructed average of previously experienced instances or anything else that contains the relevant structural information. Our later depictions of mathematical concepts and relationships will have a similar nature.

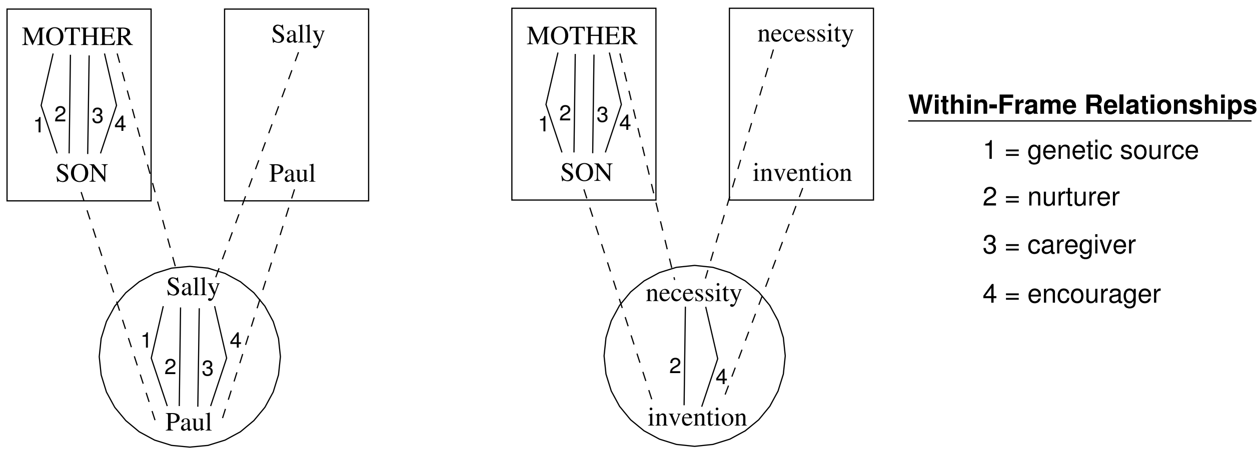

Figure 1: Schematic illustration of blending for the cases "Sally is the mother of Paul" and "Necessity is the mother of invention". In each diagram, the left rectangle represents an input containing a schema for MOTHER, consisting of two role slots and 4 relations between them, shown in legend. The right rectangle represents the contents that are given in the particular sentence. The blends (circles) consist of a representation of the schema with the contents mapped into its slots. In the case of "necessity...", only part of the schema's relational structure is mapped into the blend (connecting lines not shown), whereas in that of "Sally..." all of it is (but see text).

2.4.2 Mechanisms

In Fauconnier & Turner's (1998) treatment, the process of alignment of elements and slots is referred to as *cross-mapping* while the partial carrying-over of aspects into the blend is referred to as *selective projection*. The latter is driven by the structure of the spaces themselves, and in fact, in many cases, the former may be as well. For instance, in the sentence "Vanity is the quicksand of reason," nothing specifies that the reasoner's situation with respect to his vanity is to be likened to that of a *traveler confronted by* quicksand, yet this is the mapping by which it is most readily understood (Turner, 1991). Here, both the specific structure of travelers, reasoners, and confrontation and the alignment between them emerge driven by the elements – given by the sentences – originally in the input spaces. The pulling in of previously unspecified structures is termed *completion*. It may be provisionally viewed as deriving from a process in which schemas commonly associated with an element (e.g., 'quicksand') are partially activated by its presence and then those that match some structure in the other space are stimulated further, until they pass some threshold to activate completely, similarly to the "two-thirds rule" of adaptive resonance networks (Grossberg, 1987). For instance, in the present example, the potentially deletory effect that getting caught up in vain self-inflating thoughts has on the attention-requiring, effortful advancement of slippery logical chains shares certain components with the scenario of a traveler being held back by quicksand; these components convey excitation between the two structures so that their activation grows, and grows faster than any competing organization. Hence, in a real sense, the final structure of the mapping drives its construction.

In terms of the discussion of a hierarchy of cognitive structures above, the shared components that drive the mapping correspond to properties or clusters of properties smaller than a schema. Researchers have constructed models demonstrating the functioning of mapping and completion mechanisms as we have described using hierarchical structures (Gentner, 1983; Holyoak & Thagard, 1995; Hofstadter et al., 1995; Hummel & Holyoak, 1997). These models differ somewhat in their aims and organization, but they share the property that the represented necessary, possible, and impossible relations between components discussed above play roles in specifying strongly or moderately excitatory or inhibitory connections between units representing activation of various elements and relations. A mapping then grows by spreading activation.

To facilitate our treatment of mathematical concepts and proof understanding, we will find it useful to distinguish four types of conflictive property sub-mappings. The first three, direct conflict, indirect conflict, and incommensurability, are fatal in the sense that the sub-mapping in question is not viable and can only contribute to the final blend if other elements change or are removed. The last one, incompatibility, is avoided if possible but can remain in a blend if not.

Direct conflict occurs when two opposing features or values on a dimension are mapped onto the same slot. For example, in "colorless green ideas", 'colorless' and 'green' both map to the 'color' slot in the blend that would result. A direct conflict can be resolved either by ignoring one or the other input, or by reinterpreting one or both inputs so that the conflict is eliminated. For instance, 'colorless' can be seen as meaning "pale" or "translucent" in the example just given.

Indirect conflict results when, for example, attribute A only occurs experientially in conjunction with attribute B, which itself is in direct conflict with something else present in the proto-blend. In this case attribute A will not occur in the final blend unless the conflict is removed. For example, 'saturated' in "*saturated* colorless green ideas" indirectly conflicts with 'colorless' if it is taken in its color-related sense, since it requires a color (in this case 'green') to be present with it. An indirect conflict can be resolved by getting rid of both the directly and indirectly conflicting attributes, or by getting rid of the element(s) generating the direct conflict, or by reinterpretation as with a direct conflict (e.g., we take a non-color-related sense of 'saturated' in our example).

Incommensurability is more subtle. We say attribute A is incommensurable with attribute B if A never occurs (experientially) together in any categories with B. An incommensurable attribute cannot participate in a blend not because it reacts in a negative way with other elements, but simply because, past experience having supplied no correlations, it does not react at all. An intuitive justification for this conception is that a blend should be a self-reinforcing, resonant conceptual structure, much like a cell assembly in Hebb's (1949) sense, else it cannot maintain integral existence as a cognitive entity, however transient. "Green idea", disregarding, for the moment, the metaphorical interpretation, is an example of an incommensurable combination. The combination does not stimulate any implication or association that "green" and "idea" do not stimulate separately; the two do not act as a whole because there is no experiential basis for understanding or predicting their interaction, and the upshot is a failed blend. Incommensurability is resolved simply by excluding whichever attribute or set of attributes is least reinforced by other structures in the proto-blend. Reinterpretation is also possible, as in the case of our example, where the metaphorical interpretation of an untried idea can be (more or less effortfully) formed by associating 'green' with certain culturally associated attributes relating to inexperience, for which parallels can be found with attributes relating to the role ideas play in discourse.

Finally, *incompatibility* occurs when the value an attribute takes in a blend is different from that which it ordinarily takes in the presence of the other elements in the blend but does not lead to conflict (and incommensurability is impossible by definition). Our "purple cat" example is an instance, since purple does not occur in the context of the other attributes characterizing 'cat'.

Below, we shall argue that mathematical concepts are mentally represented through blending of previously-acquired schematic structures specified by definitions; the mappings in these blends usually involve incompatibility. Steps in proofs, on the other hand, lead recipients to attempt to form blends and examine their suitability vis-a-vis conflicts potentially resulting directly or indirectly from correspondences specified by the statements. We term these conflicts *dissonance* in this context. First, we must clear some ground regarding the nature of mathematical practice.

3 Formal and Informal Mathematics

To forestall misinterpretation we precede the main presentation by drawing a distinction between *formal* and *informal* mathematical practice (see Lakatos, 1976: p. 1–5), even though it will only become completely clear subsequent to the discussion itself.

3.1 Formal Mathematics

Under formal mathematics we group three approaches that are united by the fact that they involve a separation between the statement-generating activity of mathematical development and the original ideas which motivate it and connect it to other realms of experience. They all developed around the turn of the century as attempts to resolve questions about mathematical foundations, though we will argue in the Discussion that they were culminations of more general trends. The essence of formal mathematics is most clearly evident in the (historically latest) *formalist* approach, which received its most complete development in the work of Hilbert & Bernays (1934/1968)[3].

Here, mathematics is done within a *formal system* made up of a collection of symbols, a set of specifications of what constitutes a well-formed (not necessarily true) string, a set of axioms or strings of the symbols assumed true, and a set of rules specifying how new true strings may be derived from old. A *theorem* is a string that has been derived through a sequence of rule applications starting from one of the axioms, and its *proof* is simply the sequence of strings produced by the applications. A theorem can provide the starting point (in place of an axiom) for new proofs. To verify a proof, it is merely necessary to check that the preconditions for the given rule's application are satisfied at each step, and original theorems can be produced simply by the selecting applicable rules and applying them in sequence starting with an axiom or theorem.

Notice that the production of proofs and theorems is essentially internal to the system: valid chains of applications depend only on the form of the rules and axioms, not on what the symbols stand for. As formal mathematics is typically employed, the symbols *are* tied to external entities or concepts, as specified by a special mapping to a different system, termed a 'model'; this is the route by which formal mathematics may be applied to understand and predict phenomena in the experiential world. But the properties of these experiential concepts play a role only in initially determining the forms of the axioms and rules, and subsequently no reference is made to them in judging whether a given proof is valid. This latter task is easily automated, and in fact all present machine-based theorem proving systems (see review by Bundy (1983) and Farmer et al. (1995) for a recent example) operate within the formalist program[4].

[3]The *Formulaire de Mathematiques* of Peano (1895-1908) and colleagues was an important precursor. See Wilder (1952) for further details on historical development.

[4]Gödel's incompleteness theorem invalidated the formalist program as a foundation for all of mathematics (or any substantial part of it), but for more practical and restricted applications formal systems are useful.

Importantly, to use a formal system in this way, there must be a set of prescriptions and conventions as to how the strings of symbols are manipulated, and this, together with mechanisms for proving properties about the system like consistency are the province of *metamathematics*. Overall, the situation may be summarized as in Figure 2, left.

In the *logicist* approach, most fully developed by Whitehead & Russell (1925)[5], the situation is similar, with symbolic logic playing the role of the formal system. Instead of arbitrary symbolic rules, therefore, the inferential principles within logic operate, and the symbols themselves do have a meaning within this context. Their function with regard to mathematics, however, is that appropriate combinations of them stand for mathematical concepts, as did individual symbols before. The ordinary meanings of the concepts are relegated to the outside – derivation of consequences within the system proceeds only according to logical rules of inference independently of any interpretation given them (Figure 2, left).

The *axiomatic* approach shares the quality of relegating any prior meaning[6] of mathematical concepts to an indirect role, but in doing so leaves more substance within the derivational system itself than the previous two methods. Although its roots go back to Euclid's *Elements* (trans. Heath, 1956), the axiomatic method achieved its modern form in Hilbert (1899/1959)[7]. The method begins with a set of primitive terms which are left undefined, aside from what is said in a set of axioms concerning relations between them. These axioms generally make reference to terminology from logic, set theory, and arithmetic (unless it is one of these that is being axiomatized), as well as possibly concepts from ordinary experience, and it is the understanding of this terminology that drives the derivation of further statements within the system beyond the axioms. The primitive terms themselves and their interrelations, however, may be related to concepts outside the system only through a model relation as for the formal system or logic-based approaches.

Because the axiomatic method allows a wider scope for inference within the system, formulations within it are less cumbersome than the previous two methods, and it has been widely employed within mathematics as a means of ensuring consistency and integrity of development (Kramer, 1970: c. 3). However, in the teaching of mathematics at the precollege and even, for the most part, at the college level, a different approach is taken, an approach which also plays a role in the development of mathematics in new areas today, just as it did in all areas before axiomatization[8]. We term it, following Lakatos (1976), *informal mathematics*,

[5]Frege's (1884/1953) treatment of arithmetic was a major precursor.

[6]By this word we intend "connection to other concepts and aspects of experience".

[7]Pasch and Peano both published treatments in the 1880's that were precursors (see Wilder (1952: c. 1)).

[8]Axiomatic treatments for traditional areas of study were developed largely during the second half of the 19th and first half of the 20th century; see Kramer (1970: c. 22) and Bell (1945: c. 8-11).

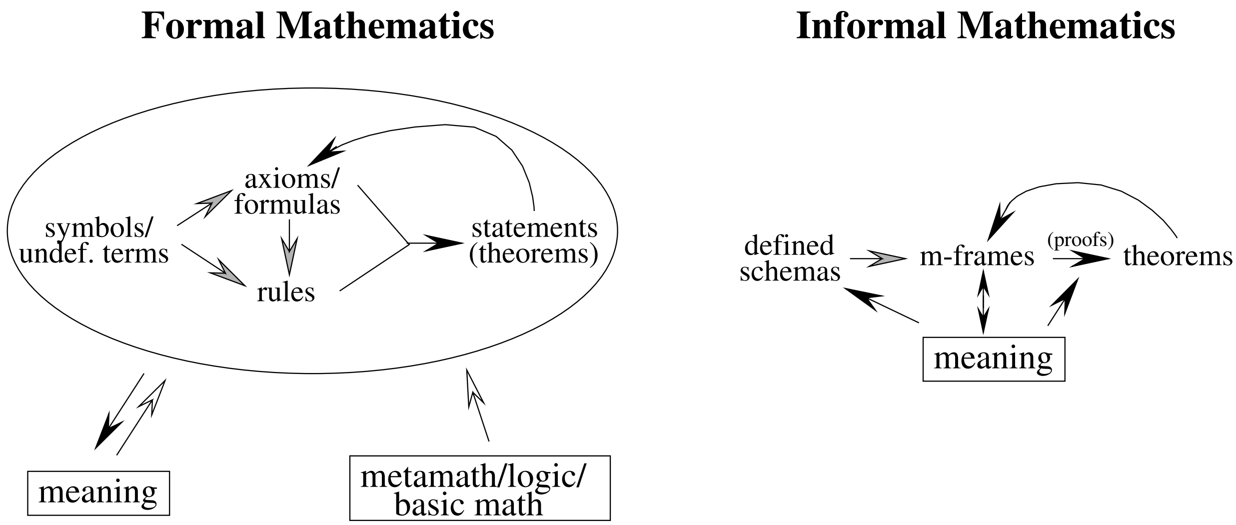

Figure 2: Black arrows represent mappings or determination of structure performed repeatedly as a system is used, outline arrows represent mappings/structure determinations that occur only during the formulation of the system, and grey arrows represent the provision of input to one component of the system from another.

although it is only so in the technical sense of not employing an encapsulated derivational system.

3.2 Informal Mathematics

In informal mathematics, there is no separation between meanings of mathematical concepts and their manipulation. Concepts are defined making reference via natural language to categories and intuitions deriving from everyday experience and then reasoned with in proofs according to their resulting properties. A theorem is a (possibly complex) statement of relationship between defined concepts, and a proof is a sequence of simpler statements of relationship between defined concepts. The validity of each step in a proof is verified by reference to the properties of the concepts. The situation is as illustrated in Figure 2, right.

We will describe the informal framework in further detail as we present our treatment of definitions and proofs in the following two sections. For now, we note that the scenario we describe for informal mathematics also underlies two important aspects of formal mathematical practice. First, there is the issue of how the metamathematical component is understood and applied. Whether this is as simple as ideas about string and symbol manipulation (formalism) or as complex as what is involved in the axiomatic method, it must be understood, ultimately, in terms of intuitions about how these work which have their basis in experience (Lakoff, 1987: c. 20). Second, the naive interpretations of the symbols or primitive terms provide guidance in the manipulation of the system, indicating which results are interesting and what routes might work to attain them (Borasi, 1993; Kline, 1966; Wilder, 1952: c. 1–2). We will return to these issues later in the analysis and in the Discussion.

4 The Grounding of Mathematical Concepts

We begin by reviewing the antecedents to our approach to grounding, then we describe it using first simple then more complex examples. A clarification of a couple of the implications of our proposal follows, and the section concludes with a brief description of ontogenetic (synchronic) concept development within this framework.

4.1 Previous Approaches

To review the history of ideas on the grounding of mathematical knowledge would require much space and lead us far afield; we limit ourselves to briefly describing two recent accounts that relate most closely to our own. They deal essentially with informal mathematics.

The philosopher of science Philip Kitcher (1984) provides a systematic and extensive development of what he terms a *empiricist* position on mathematical knowledge, which distinguishes itself from the "apriorist" positions in holding that this knowledge is grounded in sensory-motor experience. He suggests, for example, that our knowledge of arithmetic is grounded in experiences with collecting, shifting objects between groups, and drawing correspondences, and our knowledge of the concept of congruence in geometry is based on experiences with matching figures with each other, either perceptually (by scanning from object to object and comparing), or manipulatively (by physically juxtaposing one object with the other). These experiences are then generalized mentally by imagining the performance of such operations by an "ideal agent" operating in "supertime", meaning that in mathematics we conceive of these operations as they would be performed without human limits on scale.

Kitcher proposes that mathematics beyond what is perceptually grounded in this way has developed by a series of "rational transitions" between sets of mathematical beliefs and practices over history; rational transitions are portrayed as being similar to transitions in scientific practice, driven by such factors as the desire to answer questions, unify observations, and explore consequences. This provides no explicit account of the development of mathematical knowledge within individuals, and he seems to imply that as long as the historical development can be explained in this way, it is conceivable that the individual may follow the same route (or a shorter one) to abstract mathematical knowledge.

Lakoff & Nuñez (1997, 2000) provide a more detailed account of empirically-grounded mathematics, building on suggestions in Lakoff (1987: c. 20). They propose that we understand mathematical concepts via two kinds of metaphorical mappings: *grounding metaphors*, which map structures from everyday experience onto mathematical objects, and *linking metaphors*, which map mathematical structures onto other mathematical structures. The proposal of grounding metaphors makes the same claims that Kitcher made about the conception of operations by an ideal agent in a more explicitly structured way: rather than saying only that we understand arithmetic in terms of collection operations, the metaphorical account says that this understanding consists of mapping schemas generalized from manipulative experience onto the mathematical entities, for example, "Addition **Is** putting collections together with other collections to form larger collections," and "Equations **Are** scales weighing collections that balance." With each of these statements of metaphorical correspondence comes the implication that at some level the cognizer is able to draw a detailed correspondence between the source and target as illustrated above in Section 2.4, and it is via these correspondence links that certain inferences regarding the mathematical objects are rendered permissible while others are not.

Phrasing things in terms of grounding metaphors may appear to add little besides window-dressing to an account like Kitcher's, however it is valuable not only for its greater concreteness, which allows the implications of empirical grounding to be understood more clearly, but also because it relates mathematics nontrivially to a host of other phenomena that have also been analyzed in terms of cognitive mappings (see Section 2.4).

Two examples of linking metaphors are the conceptions "A Line **Is** a set of points (fulfilling certain conditions)," and "Exponentiation **Is** repeated addition." The proposal is that we conceptualize abstract mathematical objects by means of mappings from more concrete ones, just as it has been claimed that we understand abstract *concepts* in terms of concrete ones (e.g., Lakoff & Johnson, 1980). This provides a simpler and more direct account of the more abstract knowledge of an individual mathematician than Kitcher's history of rational transitions.

The account we put forward below is similar to Lakoff & Nuñez's, except that we describe the metaphorical mappings associated with a given mathematical concept as part of a blended space network, in which the concept itself is the blend proper. This makes possible a clear description of the process of acquisition of a mathematical concept, from first contact with a definition, through familiarization with its behavior in various situations, to a direct grasp of its characteristics without reference back to the source domains.

4.2 Simple Examples

Consider the following definitions from Euclid's *Elements*[9] (trans. Heath, 1956, p. 153):

1. A **point** is that which has no part.
2. A **line** is [that which is] breadthless length.
5. A **surface** is that which has length and breadth only.
8. A **plane angle** is the inclination to one another of two lines in a plane which meet one another and do not lie on a straight line.

Regarding the objects of the first two definitions, POINTs and LINEs, the account of Kitcher claims that we understand them by reference to our experiences with objects in the world that have locations and lengths, and the ways in which we observe them to interrelate. The account of Lakoff & Nuñez holds that we understand them by metaphorical mappings from schematic cognitive structures for location in space and physical length. Both of these accounts refer to *understanding* these concepts, but neither deals explicitly with the matter of what they actually are, aside from targets of mapped structure – or, to put it another way, what referents they have. We suggest the concepts start out as *blends* of different "metaphorical" sources invoked by the text of the definitions. Like our earlier "necessity" example but unlike the "Sally/Paul" one, the result is a generality rather than a specific instance. The overall effect of the definition is similar to what might be imagined if 'necessity' were *defined as* "the mother of invention".

Figure 3 illustrates the structure of these blends for the definitions of POINT and LINE using a similar 'entity and relation' approach to Figure 1. In the first case, the blended concept is essentially "an object with no size". It has the properties of having identity (schema A) and location (schema B), which are triggered by the words "that which has" in the definition as well as the property of having no size (schema C), summarizing what the cognizer interprets from "no part" given by the text. Of course, there is really nothing in previous experience that corresponds to an object like this; nevertheless, we suggest we are able to conceive of it by the same means as we do a purple cat (Section 2.4) – by mapping a nonconflicting value

[9] This text, although an important precursor to and inspiration for the axiomatic approach, falls within the realm of informal mathematics as we described it (Wilder, 1952: c. 1). Axiomatic approaches to geometry (e.g., Gemignani, 1971) leave the central concepts undefined.

for a commensurable attribute onto a category. In the blend, all such mutually allowable attributes are combined, and directly conflicting ones are resolved in favor of one or the other. A large number of other attributes, indirectly conflicting (e.g., 'mass' and 'color') or simply irrelevant (e.g., 'velocity') are excluded. There are no incommensurable elements in the inputs to this blend.

Objects experientially have the property of having a definite physical size, however projection of this property to the blend is ruled out because it would directly contradict schema C, 'has no part', which is called for explicitly by the definition. 'Mass' and 'color' are excluded by indirect conflict, since they require physical extension. Objects have other properties, such as velocity, that are not mapped into the blend even though they do not conflict with anything that is specified. This is partly due to the fact that the general context (geometry, the study of static spatial relations) prevents them from being mapped because they are irrelevant, but it also owes strongly to the influence of how the concept is subsequently *used* – the types of situations it is placed in and the examples of it that are put forward. The point that definitions themselves do not suffice to *specify* a concept will become clearer as we examine more cases, and we will see in some detail how usage affects conception when we examine proofs later.

In the case of LINE, one input (A) is the schema for a long thin object, one (C) is a conception of absence of breadth (similar to 'has no extension'), or a dimension with a zero entry. The 'line of sight' schema (B) represents something that satisfies the 'no breadth' requirement and has many of the qualities we understand to be true of a geometric line, but it lacks the qualities of identity and persistence in time possessed by the 'long thin object' input. In the blend, the result is a conception of an object that has length but no other dimension.

Thus, we propose that mathematical concepts are blends, the inputs of which may themselves be blends but eventually ground out in cognitive schemas. We term the final blends resulting from definitions mathematical schemas, or *m-schemas*, to reflect the fact that they are not instances, but rather schematics, specified by reference to other schematic structures.

When a given m-schema is employed in some context, we are able to understand how it relates to other elements, even though it is a novel structure with which we have no experience, by virtue of the connections which are maintained from the blend back to its inputs. When we read definition 8 (ANGLE) for instance, it is clear from experience with the input sources to the LINE blend that two directed lengths, when they meet, will do so with a certain quantitative relationship between their directions. Subsequent development in Euclid's *Elements*, as well as in any other mathematical text, relies on our ability to understand in this

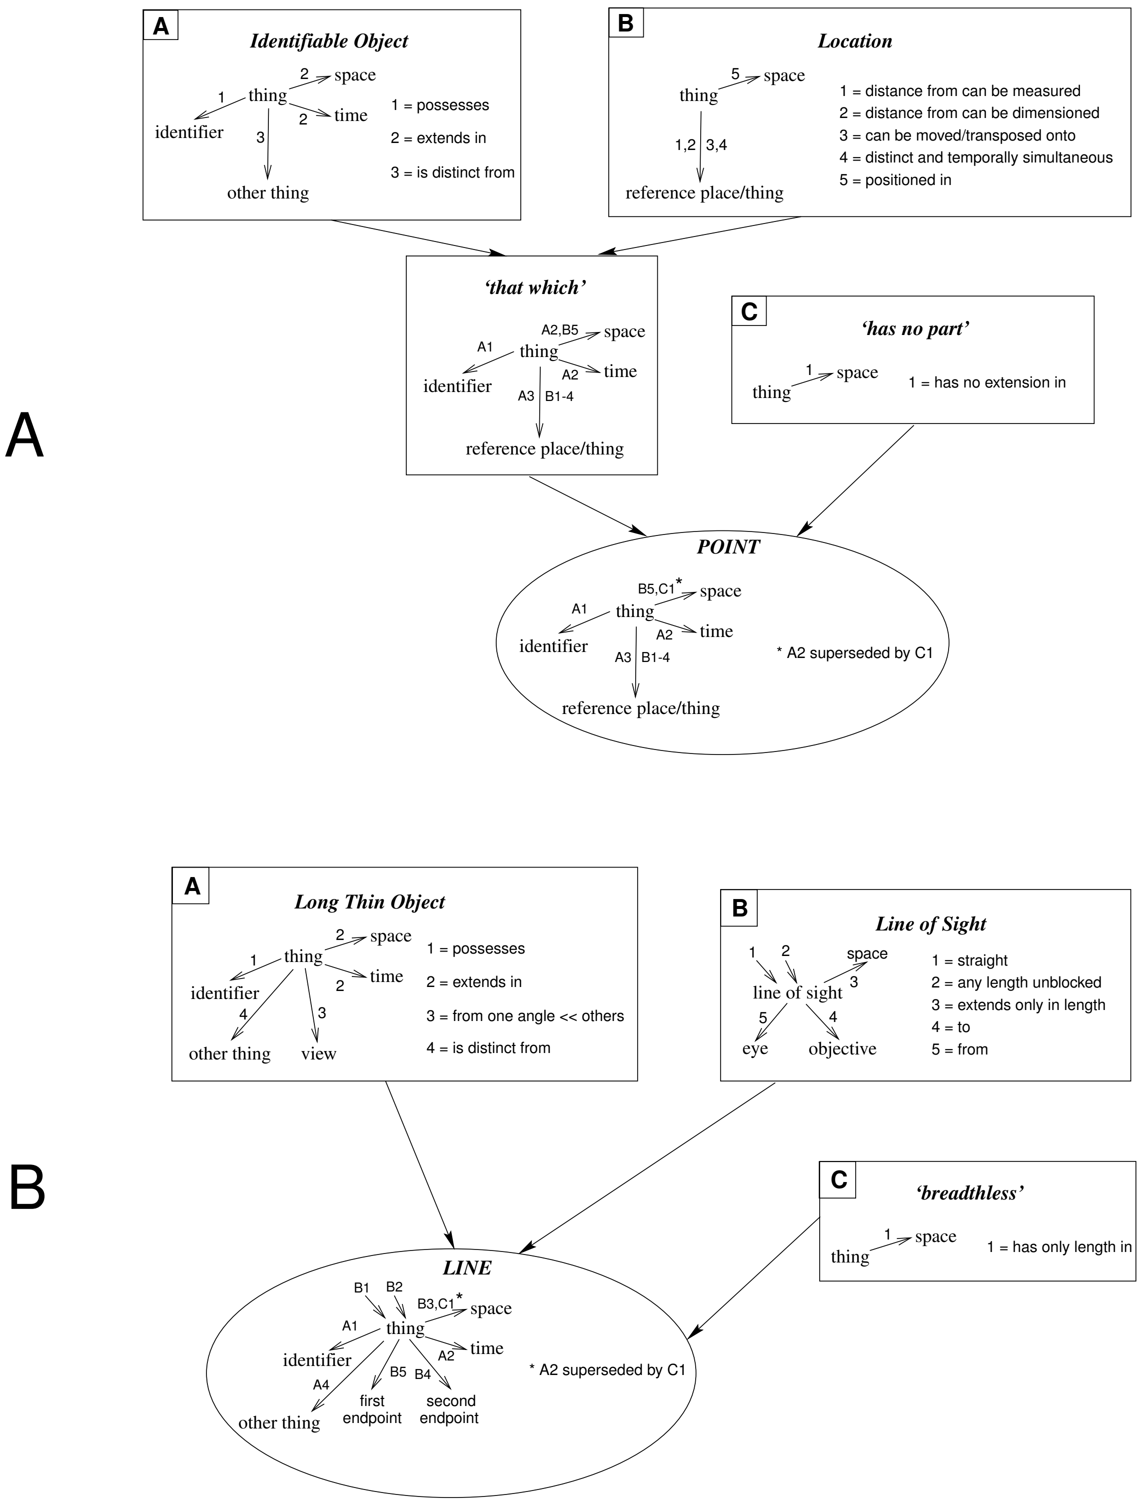

Figure 3: Illustration of the construction of mathematical concepts (m-schemas) POINT and LINE as a blend with image-schematic inputs. The blend is on the bottom in each case, with its inputs just above, all represented by a set of entities and some of their interrelations (cf. Figure 1). In the blends, entities and relations are drawn from each input as indicated (using the letter 'codes' for the relations). As with Figure 1, this and the following illustrations are only intended as rough characterizations of the true detailed cognitive structures and correspondences.

way, as well as the ability to form and hold in mind specific instantiations of the m-schemas in question. We shall examine how this works more closely in the section on proofs.

One question that arises is what determines which blended mathematical objects can be formed and which cannot? What makes a viable mathematical concept? We suggest that this is governed by one hard criterion and two soft ones. The hard criterion is that the definition cannot directly specify a mapping generating (direct or indirect) conflict or incommensurability. For example, a definition like "a **point** is that which has both no extension and infinite extension" is not viable, nor is "a **line** is one-to-one breadthless length."

The two soft criteria are interrelated and more difficult to define precisely. The first is that the source schemas should be relatively general. That is, only categories picking out very broad classes of instances – in terminology originated by Rosch et al. (1976), superordinate or higher – should be used. Classes like 'object' and 'surface' are good, those like 'cat' and 'tool' are bad. Highly general schemas are not only powerful in having a wide range of analogical applicability, but also make for simpler, more reliable definitions since there are fewer features and alignments that must be gotten straight in the definition. The second soft criterion is that the definitions should be relatively precise in specifying what sources are used and how they are combined. For a given length, definitions which involve only explicitly mapping one or a few properties onto a category will be favored over those which involve larger numbers of properties and/or elements. For example, a hypothetical definition of "vanity" by saying it is the "quicksand of reason" (see above) would not be favored, since it is not overtly specified what is supposed to be mapped between the 'quicksand' domain and the 'reason' domain, and in fact not every interpreter may form the same blend. A longer definition could eliminate this problem by spelling out most of the intended parallels, but such definitions are little favored in mathematics or in, for that matter, ordinary dictionaries, probably because they lead to concepts that are difficult to learn or remember. Typically short definitions are the rule, and the ambiguity problem is partially alleviated by the use of very general, less property-laden sources in definitions that leave less room for alternate mappings. That it remains nevertheless is illustrated by the definitions in the next subsection and will be discussed further following the proof analysis.

4.3 Complex Examples

Other mathematical definitions are more complex, involving source image schemas that are not simply "perceptual" properties such as length or objecthood. Consider the following definitions of SET, SUBSET, and NEIGHBORHOOD[10]:

1. ...a set should be a well defined collection of objects.

   ...we are specifying a *universe*, or *universe of discourse*, which is usually denoted by U.

   We then select only elements from U to form our sets.
2. ...we say that C is a *subset* of D and write $C \subseteq D$ or $D \supseteq C$, if every element of C is an element of D.
3. A *neighborhood* of a point p is a set $N_r(p)$ consisting of all points q such that

   $d(p, q) < r$ [distance between p and $q <$ r].

Figure 4 sketches definitions 1 and 2. In the case of SET, the consideration of 'collection abilities' – that some sets fitting the description of a "collection of objects" might nevertheless not be viable because of conceptual difficulties with assembling the collection – is not overtly specified by the definition, but the issue does come up in certain situations. This was not widely noticed until Russell (1903/1937) published his discovery of the paradox involving sets containing themselves, but it has since led to more specific and involved definitions of SET – which are, however, only employed in situations where the problem is likely to come up. The *axiom of choice* (e.g., Hoffman & Kunze, 1971: p. 399; Kramer, 1970: p. 595–7) relates to a similar concern over the ability to form infinite collections in certain circumstances. This property that schemas or relations not specified or highlighted in a definition may come to have an important bearing on the existence or nature of the defined entities in certain situations sharply distinguishes informal from formal mathematics, for, as we shall see below, it is the (possibly changing) nature of m-schemas that determines their relational properties in proofs, whereas in formal mathematics, the relational properties of symbols are entirely fixed beforehand in the specification of the system.

An important second point is that the relations in the result are not necessarily copied from the inputs, but may themselves undergo alteration by contextual effects. The property of occupying a restricted portion of space in the 'Collection' input is transformed to possession of an identifier in the blend, because, via the "multiplex-mass transformation" (Lakoff, 1987), a localized group can be seen as a single distinct (hence identity-possessing) thing. Because 'space' is irrelevant to the other elements of SET in the blend, the other aspects of this property disappear. Similarly, the distinctiveness of objects in the 'Object from U' source is transformed into the requirement that there be no duplication of elements in a set. In the context of the impoverished types of objects usually dealt with in mathematics, this often ends up needing to be pointed

[10] Definitions 1 & 2 from Grimaldi, 1989: p. 98; definition 3 from Rudin, 1976: p. 32, italics original.

out explicitly, even though in the physical world it is trivially true – every instance of 'two objects', for example, is distinct. Thus, the salience of various properties of the sources can change dramatically in the context of other factors in the blended definitions or in the usage of the concept.

The case of SUBSET illustrates three additional points. First, a definition may build upon an earlier m-schema, in which case it is incorporated as an input into the blend. Second, a defined entity may only be 'perceivable' through a temporally extended process of verification. In the SUBSET blend, the 'exhaustive selection' schema (triggered by "every element" in the definition) implies that verifying whether something is an instance of SUBSET involves a kind of sequential testing (for element membership) that is an interplay between specific perceptual and motor processes. This interplay arises naturally in the blend from the combination of the containment relation (the perceptual part) and the exhaustive selection schema (the motor part).

The third point regarding SUBSET is that despite this 'temporally-extended' definition, it *can* be interpreted in an instantaneous, 'global' fashion as reflected by relation 2 in the blend, which may have its source in experience with containers within containers. Such a mapping is not overtly specified by the definition, and in fact in order to rigorously prove that something is a subset of something else it is necessary to demonstrate that every element is contained, as per the 'letter' of the definition. The later containment mapping simply serves as a kind of 'mental shorthand' allowing a person to think more quickly and globally about a situation involving SUBSET. The notion of NEIGHBORHOOD (not illustrated) is similar, involving an exhaustive selection process that can be summarized by thinking of a more global circle-containment relation.

Other mathematical definitions draw on different schemas of verification that also involve both perceptual and motor components; for example, the notion of CONGRUENCE is based on the imagined possibility of placing figures over one another and noting their correspondences, and the notion of NUMBER is based on the idea that certain regular correspondences are maintained in imaginary augmenting and reducing operations on collections of abstract entities (see Kitcher (1984) for further discussion).

4.4 Mathematics and Reality

4.4.1 Universality and Specificity

What are the implications of this conception of mathematical grounding from a philosophical standpoint? The schools of thought on the reality of mathematical objects have traditionally fallen into two main classes,

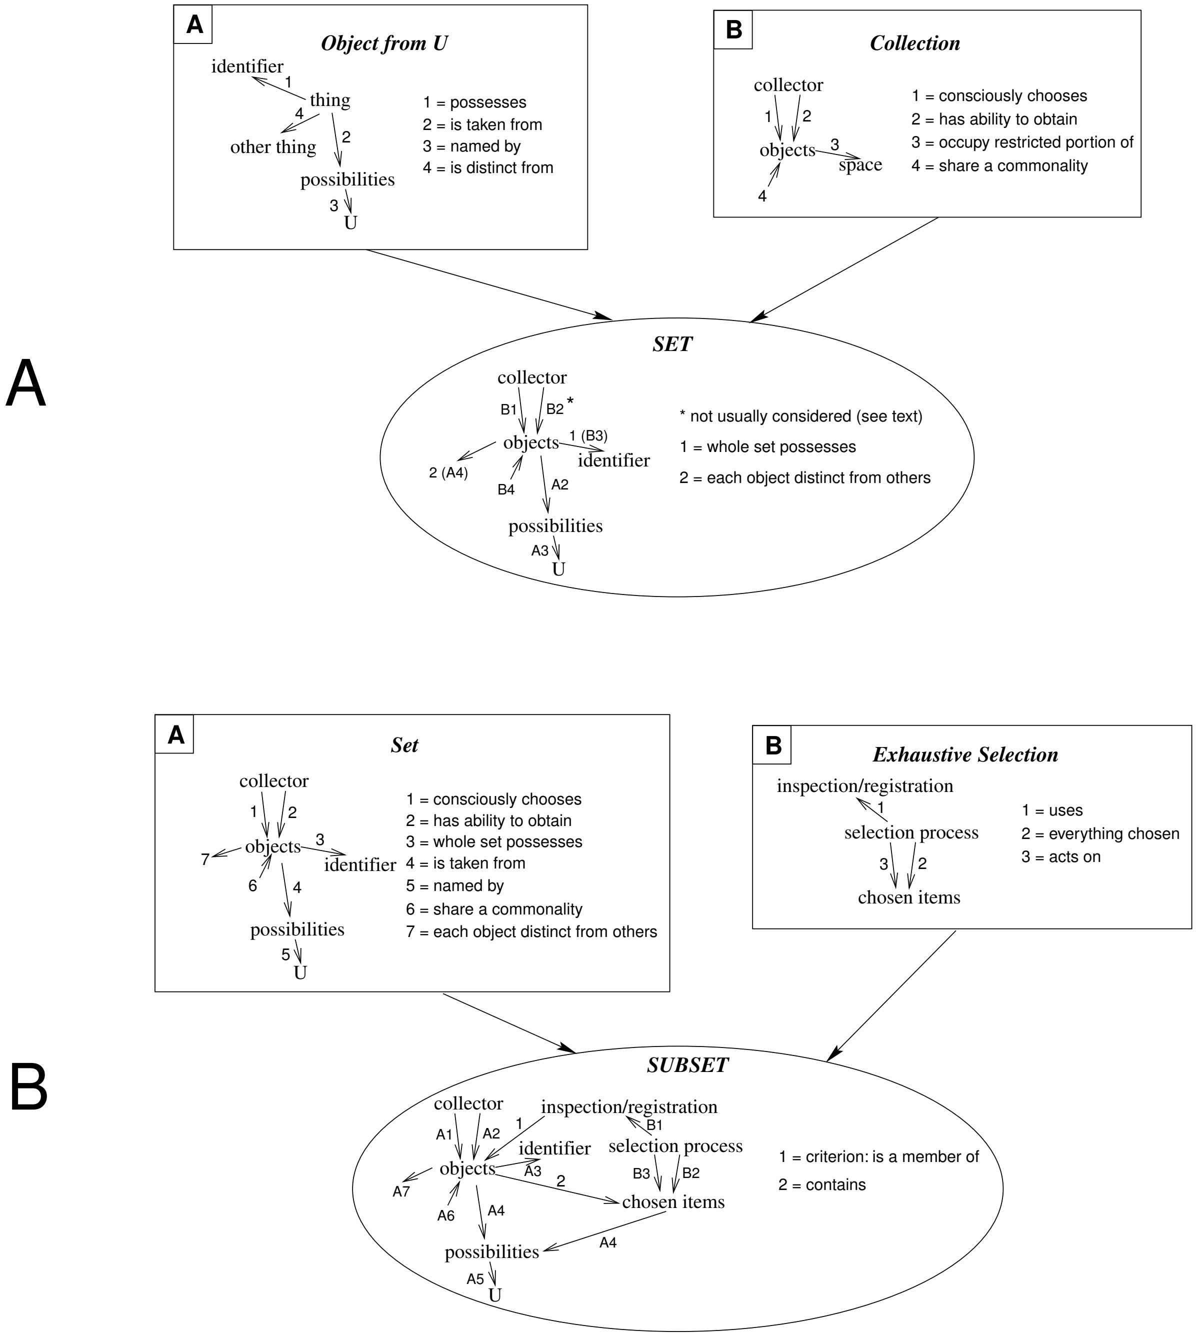

Figure 4: Illustration of construction of m-schemas for SET and SUBSET (conventions as in Figure 3). In the case of SET, a schema for an entity chosen out of a particular kind of things is combined with a schema for 'collection' (i.e., a cluster of objects brought together for a purpose) to result in the blended conception. The 'collection abilities' property (B2) was initially ignored historically. In the case of SUBSET, the blended m-schema involves the satisfaction of a particular process within the context of a SET m-schema (property 1 in the blend).

which may be termed *realist* and *constructivist* (Davis & Hersh, 1981). The realist position holds that mathematical objects exist independently of human thought, while the constructivist claims they do not. Realists typically point to either the combination of the indubitability of mathematical knowledge and its unpredictableness – e.g., the distribution of prime numbers, or the outline of the Mandelbrot set, as implying that mathematics could not be a human creation since humans have no knowledge of these structures until they are 'discovered' and 'explored', or they point to the effectiveness of mathematics for explaining aspects of the physical world. Constructivists, on the other hand, point to such things as the existence of the non-Euclidean geometries, the evolution of mathematical concepts, and the changing standards for proof acceptance as indicating that mathematics is a living human creation.

Our account suggests they are both right. Mathematical concepts depend on regularities in perceptions and interactions with the world. Because the character of these regularities depends on both the part of the universe one lives in and the characteristics of sensory and motor systems with which one interacts with it, mathematics has the properties of something constructed. If we assume that the universe operates according to the same laws everywhere, then there will be some reason to expect the mathematics developed in different places to share commonalities. The laws, together with all possible sensory/motor systems that they allow to exist, determine a basic stock of mathematical structures that in some sense are 'real', or 'out there', and different races and groups of mathematics-using creatures discover subsets of these. There is also, however, another factor determining what mathematics is actually *observed*, which is the choices that are made as to which properties to focus on and include in mathematical definitions. These are not arbitrary (see end of Section 4.2), but there is ample room for variation, as a historical survey of mathematical concepts readily reveals.

Consider, as an informative extreme, a race of creatures such as Arthur C. Clarke once wrote about living in a star[11], made up of organized regions of plasma. Suppose they have no means of self-locomotion, and their only sensory mechanism is to sense the temperature within their bodies as they are blown about by currents within the star. Assume for the sake of argument that they are intelligent and develop mathematics. This mathematics is quite unlikely to contain anything related to our geometry based on shape and space, and even natural numbers will likely be an unknown concept without experience with discrete objects to underlie it. However they will quite likely have some kind of a highly developed theory of the structure and evolution of continuous distributions and stochastic processes that builds from and helps describe the

[11] "Out of the Sun", available in *The Nine Billion Names of God* (New York: Harcourt Brace & World, 1967).

patterns of rises and falls in temperature they experience, and from this basis they might eventually come to, with much effort, something like our idea of space as a theoretical construct explaining regularities in this experience. Whether their approach to concepts like CIRCLE or ANGLE (if they had them at all) would be intelligible to us is an entirely different question.

The point is that although there *may be* a basic stock of 'real' mathematical structures (and this is a metaphysical question in the end), it is very large, and the mathematics that actually is developed by a given civilization is a small subset of this highly dependent on their sensory and motor capabilities. We advocate studying the nature of this dependence using the tools of cognitive science.

4.4.2 Mathematics in Physics

The "unreasonable effectiveness of mathematics" in application to the physical world famously pointed out by Wigner (1960) becomes less paradoxical in the context of the above considerations, for it is the same schemas which mathematical concepts are derived from that are used in perceiving and categorizing the world to which the mathematics is applied. Wigner in fact acknowledged this for special cases, noting that "the concepts of elementary mathematics and particularly elementary geometry were formulated to describe entities which are directly suggested by the actual world," but goes on to say that "the same does not seem to be true of the more advanced concepts, in particular the concepts [group theory, function spaces, etc.] which play such an important role in physics."

Responses to this question (e.g., Davis & Hersh, 1981: p. 68-76) have tended to focus on the selectivity in the application of mathematics to the world – i.e., if the mathematics *does not* fit, we either switch to different mathematics or change the physical features we attend to until it does. While this is undoubtedly an accurate characterization of the application process, it says little about why matches can be found at all. Our cognitive semantics based account suggests that the reason is that all of the "more advanced concepts" are constructed from and construed in terms of schematic components grounded in experience – which are the same as those which organize the pre-mathematical perceptions and conceptions of the physical phenomena. Moreover, the latter have often themselves undergone considerable abstraction and recombination (consider, for example, momentum, force fields, or waves). Thus both the mathematical structures and the physical conceptions are likely based on a limited subset of the human conceptual stock – those structures which can be grasped, applied, and communicated most readily, and are simple enough that mapping-based composites can be reliably constructed (also see Discussion). Finally, physics and mathematics have heavily influenced each others' development (Bell, 1945: c. 16). These factors raise the likelihood that systematic correspondences between the two domains can be found.

4.4.3 Non-Euclidean Geometries; Paradoxes

Given this account of grounding, certain facts of mathematics itself, however, might appear puzzling. For example, how can there be geometries other than Euclidean in which different theorems are true if both are grounded in real-world concepts which are presumably consistent with each other? The reason is that certain choices are embodied in the selection and combinations of properties in mathematical definitions. For instance, the various plane geometries, differing in their version of Euclid's 5th (parallel line) postulate, can be construed as referring to situations on a planar, hemispherical, or other surface (see, e.g., Kramer, 1970: p. 52–4). Inconsistency arises only in attempts to map the different systems onto the same set of relations in the world.

Regarding paradoxes, only a few cases are known (Kramer, 1970; Wilder, 1952), and in each case careful examination shows that an absurdity is passed over unnoticed prior to the paradox. For example, Russell's paradox requires one to entertain the notion of sets containing themselves, something which simply violates the properties of the 'collection' input to the definition (Figure 4). We will discuss the nature of such violations and their detection in the following section, but for now one might ask, why were such cases considered at all? For the case of a self-containing set, we suggest that the relevant constraints from the collection schema input are simply not noticed since attention is focused on the prospective member (the set itself) viewed as (simply) a member of U. As shown above, the entire set of constraints from a definition input is in general *not* used, but those that *are* used are not likely to lead to paradox, since the selection is made on the basis of experience with the world as far as what subsets of properties have proven separable from others during the development of abstracted nonmathematical concepts and modes of perception. It is only on occasions when a property that is not completely separable is ignored in application of a mathematical concept that a problem arises.

4.5 Concept Development

In our account, we have described the defined mathematical concept as an independent cognitive unit resulting from a blend. One of the roles of this unit is to serve as an organizing nexus for the complex of metaphorical sources that are drawn upon when the concept is used. However, we hypothesize that it can also attain status as a schema in its own right, through a three-stage process of entrenchment.

In the first stage, the concept is understood entirely in terms of the metaphorical sources as they are given in the definition. In the second stage, the student becomes familiar with the application of the concept in different contexts, the range of its examples, and the shifts in metaphorical construal it can undergo. However, reference is still made to the metaphorical sources in application and manipulation situations. (Indeed, as mentioned, these situations themselves play a role in determining which of the array of possible sources are accessed). In the third stage, the inferential properties of the concept in relation to other mathematical concepts become internalized directly, so that it can be reasoned with in purely mathematical contexts without reference to the original sources. This process is sketched in Figure 5. The analysis presented in Section 5 is relevant principally to the second stage; we shall elaborate on the nature of the third stage and its relationship to the second in Section 7.

This account of development is similar in certain respects to notions in cognitive psychology of the development of cognitive expertise. The latter process is assumed to reflect a transition from slow application of general-purpose algorithms and knowledge to rapid application of special purpose rules or memorized correspondences – though opinions differ as to the details. Logan (1988), for instance, proposes increasing reliance on pure memory storage of stimulus-response pairs, while Anderson (1993) suggests a more highly structured process in which oft-used complexes of declarative and procedural structures are compiled into new procedural structures. Our proposal follows the spirit of the latter in that the connections back to input spaces can be thought of as a kind of "declarative" (application-general) knowledge (which is often employed via analogical mapping in Anderson's model) while the inferential properties that are eventually internalized are a kind of production-like "procedural" (application-specific) knowledge. (Anderson et al. (1981) discusses development within this framework for expertise in geometry.)

5 Blending in the Interpretation of Mathematical Proofs

In this section, we analyze the cognitive process of interpreting and verifying mathematical proofs in terms of the mapping of components of defined concepts (m-schemas) onto the specific elements of an imagined configurational situation. We shall describe this mapping using the framework of blended spaces a second time, but first we begin by clarifying the practical nature of proof in informal mathematics. We then present the analysis followed by a discussion of the structural regularities in cognitive processes it reveals

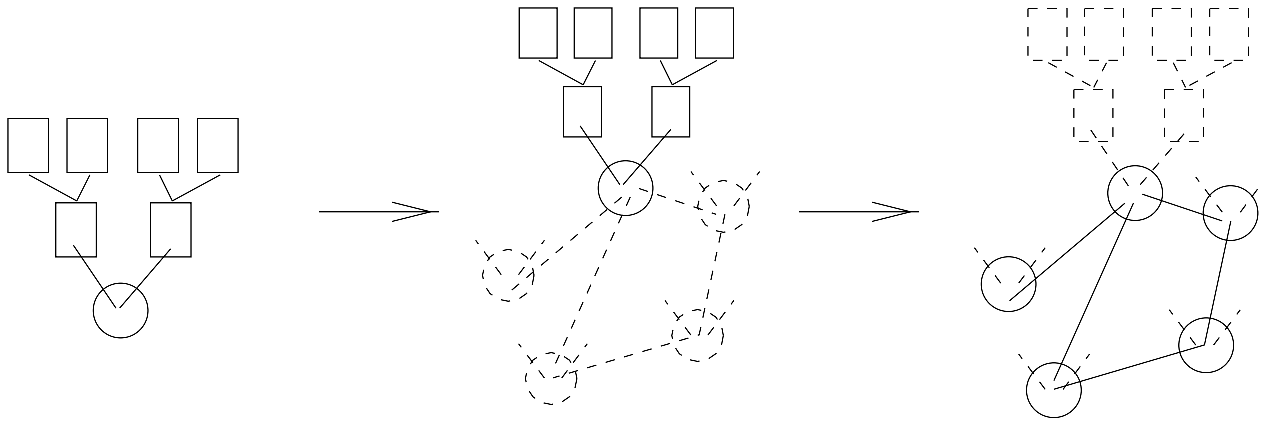

1 **2** **3**

Figure 5: Sketch of the development of a mathematical concept (see text). In the first stage, the concept is understood entirely by reference to the inputs in its definitional blend. In the second stage, the cognizer becomes familiar with the way this concept relates to other mathematical concepts, and in the third stage, this familiarization has resulted in the network of relations becoming entrenched in its own right, so that inferences can be drawn within it without reference to the original sources.

in Section 6. In Section 7 below, we present evidence from an empirical study supporting our analysis.

5.1 The Nature of Proof in Informal Mathematics

A mathematical proof is a sequence of statements which ultimately lead the reader to perceive a valid connection or relationship between a theorem's conditions and its consequences. This is accomplished by breaking down the relationship, which may be quite complex, into simpler relationships and guiding the reader in such a way that he may convince himself that each of them is true. This process of guided observation has been likened by the mathematician G.H. Hardy (1928) to pointing out the features of a distant mountain range for observation, except that the observed objects are internally imagined instantiations viewed by the mind's eye (see also Wilder, 1944).

Insight into the character of this process can be gained by considering the following example. Suppose you and a friend are looking up at a rock cliff face and the friend remarks that he can see a face in the wall. You confess that you can't see what he is talking about, but then after he points out which features in the rock correspond to the eyes, the nose, and the mouth, it becomes clear. Unable to perceive a complex configuration in the stimulus, you were led to it by a sequence of simpler configurations that you knew fit together to make it. The pointings out of features were 'steps' in the 'proof' of the existence of a reasonably

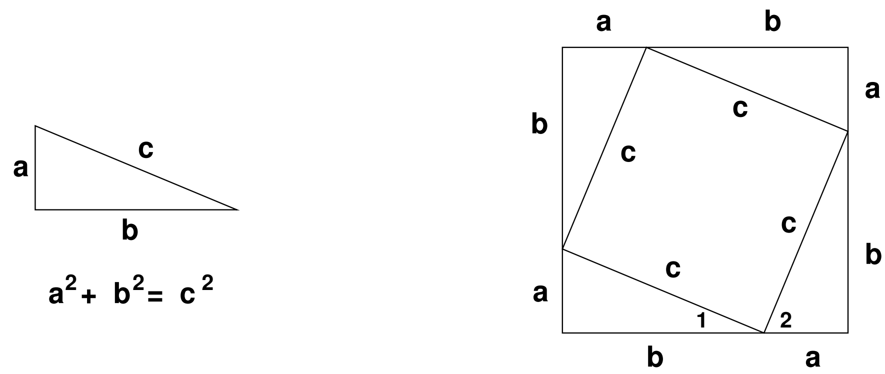

Figure 6: A simple proof of the Pythagorean Theorem (left), illustrating how a proof breaks a complex relationship down into simpler, easily perceivable ones. Assuming triangle *abc* is a right triangle implies that angles 1 and 2 sum to 90, and the interior figure on the right is a square. The area of the outside square minus the area of the triangles is the area of the small square: $(a+b)^2 - 4(1/2 \cdot ab) = c^2$, implying that $a^2 + 2ab + b^2 - 2ab = c^2$, or $a^2 + b^2 = c^2$.

face-like subconfiguration in the stimulus.

What differentiates a 'complex' relationship requiring proof from a 'simple' relationship able to serve as a proof statement (i.e., assumed intuitively obvious)? The rock cliff example provides some limited insight here. The schema of 'face' involves more elements that must be found and relations which must be verified between them than the simpler schemas for 'nose', etc., so it presumably requires more cognitive effort to verify its presence in the stimulus. Certainly other factors would affect this amount of cognitive effort, however, such as familiarity and perhaps 'goodness of structure' in the Gestalt sense. Interestingly, intuition is not always a reliable guide; in fact, one of the subtler skills learned by mathematics students at the undergraduate level is how to determine when statements require proof. A simple example of a real proof is illustrated in Figure 6.

5.2 Statements in Proofs

Mathematical proofs are made up of two principal types of statement, *introduction* statements, which induce the interpreter to construct an internal model for examination, and *inference* statements, in which observations are made.

Introduction and inference statements are explicitly marked in proof discourse by particular key words or phrases (marked in bold in the analysis below in Table 1), which we term *directives*; these terms actually specify additional characteristics of statements as well: the complexity of the inference required, and whether reference to other parts of the proof are required to draw it. Analysis of these functions will be the subject of a future paper.

5.2.1 Introduction

Introduction statements typically take the form '**Let** x be a y', as in "Let E be a set" or "Let x be an element of E", and their interpretation involves mapping a label or an imagined entity onto part of a given m-schema; the resulting "meaning" is an *imagined instantiation* of the schematic entity, an object held in mind with certain properties. Given the first such statement, a mental space is set up containing the instantiation, and subsequent statements add structure to this same space. When a configuration has been built up in a space in this way by multiple statements (see analysis below), we term the result an *m-frame*. The schema-frame distinction is somewhat fuzzy cognitively, because as we saw above, definitions can involve other (and, in fact, sometimes multiple) m-schemas, yet we still call these m-schemas. However it is clear in mathematical discourse: an m-schema has been singled out as sufficiently tightly integrated to be given a name, whereas an m-frame has not.

We make a further terminological distinction between instantiations of m-schemas and m-frames and the schemas and frames themselves. An instance of an m-schema introduced in a proof statement is an *instantiated m-schema*, or an im-schema for short. A configuration of more than one interrelated im-schemas introduced by a series of proof statements is an *instantiated m-frame*, or im-frame. Theorems themselves are examples of non-instantiated m-frames. Finally, we shall use the term *structure* depending on context to refer to instantiated or noninstantiated m-schemas and m-frames as well as subparts thereof.

5.2.2 Inference

Inference statements state a relationship between m-schemas and parts of instantiations held in mind which it is up to the interpreter to verify for himself. As we shall discuss, the interpretation of these statements generally involves constructing a mapping between the slots of the schema(s) and the elements of the instantiation(s) and observing that no conflict or mismatch arises.

5.3 A Proof in Mathematical Analysis

Table 1 contains the analysis of the first half of a proof in elementary point-set topology taken from Rudin (1976: p. 34). This book is a standard text used in introductory college courses in *mathematical analysis*, the modern theoretical underpinnings of calculus. Typically analysis is only studied by mathematics majors and a few others with an interest in advanced pure mathematics. (It is not treated axiomatically in these courses, though it could be.) *Point-set topology* is concerned with the properties of sets of points in spaces that are generalizations of the coordinate spaces on which mathematical functions are defined.

In the table, the text of each step of the proof is given in the left-hand column, with the directive in bold type. The second column lists the associated statement type, either introduction or inference. In the last column, descriptions of the mental operations required to interpret the step (and accept its validity if it is an inference) are given, with all mathematical schemas marked by capitalization. The third column lists the operation types under a classification to be presented following the analysis and may be ignored on the initial reading.

The complete text of the proof as it appears in the book is given below. Note that there are no illustrations or diagrams associated with this proof, nor with any other in Rudin's text. (The theorem is about sets of points drawn from spatially-structured universes such as the Cartesian plane. Understanding of the terms *limit* and *interior point*, and *open* and *closed set* is not necessary to follow our analysis, but their definitions are included in Appendix A.

Thm. 2.23: *A set is open if and only if its complement is closed.*

> First, suppose E^c is closed. Choose $x \in E$. Then $x \notin E^c$, and x is not a limit point of E^c. Hence there exists a neighborhood N of x such that $E^c \cap N$ is empty, that is, $N \subset E$. Thus, x is an interior point of E, and E is open.
>
> Next, suppose E is open. Let x be a limit point of E^c. Then every neighborhood of x contains a point of E^c, so that x is not an interior point of E. Since E is open, this means that $x \in E^c$. It follows that E^c is closed.

In the remainder of this subsection, we explain the interpretation of the first four steps of the proof in detail (see Table 1).

5.3.1 Introduction

In the first mental "substep" of step 1, substep 1.1, we call to mind the schema for SET and combine it with the label 'E'. This can be analyzed using the blended spaces framework as in Figure 7A. This blend is similar to the "Sally is the mother of Paul" blend described in Section 2.4 in that the left-hand side contains a general schema and the right-hand side contains specific content items, in this case a single label and an attached entity standing schematically for the thing to be labeled (we follow this convention of placing

Table 1: Analysis of a proof.

#	STEP	FUNC.	OP	DESCRIPTION
1.	First, **suppose** E^c is closed.	introduce	I1	instantiate an entity (labeled 'E') in SET schema
			I2	embed this in COMPLEMENT schema
			I2	embed this part in CLOSED schema
2.	**Choose** $x \in E$.	introduce	F1e	return focus to 1a: E in SET schema
			I2	imagine instantiation ('x')
3.	**Then** $x \notin E^c$,	infer	F1e	recall previous embedding of E into COMPLEMENT schema
			O1e	blend $\notin$ schema with arguments x, E^c
			N1	see that no conflict/dissonance arises
			O1e′	counterfactually blend $\in$ schema with arguments x, E^c
			N1	note that dissonance does arise
4.	**and** x is not a limit point of E^c.	infer	O1e′	blend LIMIT POINT schema with x, E^c
			F1i	recall E^c CLOSED blend from 1d
			N2	draw inference $x \in E^c$ (schematic implication)
			F1i	recall result of previous step ($x \notin E^c$)
			N1	note dissonance
5.	**Hence** there exists a neighborhood N of x such that $E^c \cap N$ is empty,	infer	O1e	blend NEIGHBORHOOD schema with x
			O1e	blend INTERSECTION schema with neighborhood and E^c
			F1e	embed into LIMIT POINT schema
			O1e′	counterfactually assume INTERSECTION nonempty
			N2	draw inference that x is a LIMIT POINT
			O3	blend result of previous step (x not a LIMIT POINT of E^c) with this
			N1	note dissonance
6.	**that is**, $N \subset E$.	infer	F1e	from prior focus on N, E^c (as disjoint), recall COMPLEMENT schema
			F2	use schema to shift focus to E
			I2	blend SUBSET schema with args E, N
			N1	note absence of dissonance
			O1e′	counterfactually assume N extends out (not SUBSET) of E
			N2	draw inference (using COMPLEMENT schema) that $N \cap E^c$ nonempty
			O3	blend part of previous step ($E^c \cap N$ is empty) with this
			N1	note dissonance
7.	**Thus**, x is an interior point of E,	infer	O1e	blend INTERIOR POINT schema with x, E
			O3	bring N (with context $\ni x, \subset E$) in from previous discourse
			N1	note precondition satisfaction (absence of dissonance)
8.	**and** E is open.	infer	O1e	blend OPEN schema with x, E
			O3	bring in x INTERIOR POINT from previous step
			N1	note precondition satisfaction (dissonance absence)

more abstract inputs on the left in all our illustrations). The result is that the interpreter constructs a mental space containing a *specific instantiation* of the entity (called for by the right side) described by the schema (structure mapped from the left side)[12].

In substep 1.2, we blend the COMPLEMENT schema with two arguments (Figure 7B). The first of these is 'E^c', another label, and the second is the extant instantiation E, an entity compatible with one of the slots in the COMPLEMENT schema. The result is that we have two instantiations held in mind, with a particular relationship between them specified by the schema. In substep 1.3 (not illustrated, but similar to 1.2), additional structure is given to the instantiation E^c by mapping it into a compatible slot in the CLOSED schema.

In substep 2.1 (see Figure 7C), focus is shifted back to the SET instantiation E, triggered by the reference in the text. This shift may be understood as a blend in which the combination of the label 'E' with the mental structure previously built up results in a new version of this structure in which the set E is highlighted for attention. We single this bringing into prominence out as a separate substep because the introduction statements cause a complex structure to be built up within a single mental space, not all of which is easily held in mind at one time. Since changes of focus are not always explicitly marked when they are needed, they can therefore constitute a significant part of the challenge of interpreting a proof. An example illustrating this will be discussed under step 4.

In substep 2.2, a blend similar to that in substep 1.2 is made associating the label 'x' with a part of the set im-schema just focused on (Figure 8A).

5.3.2 Inference

Thus far, the operations have involved generating instantiations held in mind and bound to certain properties. The result is that an imagined configuration has been set up and in the next step (step 3) an observation is made on it. The text of the step simply states that a certain relationship that has not been specified in the preceding construction holds. The interpreter must look for evidence of the relationship in what is held in mind.

[12]This blend and the next are not called for explicitly (there is no statement "**let** E be a set" or "**let** E^c be the complement of E"); they are understood because E is a conventionally-used label for a SET in this text and the entities referred to by the theorem statement are the most salient and only possible referents for the two m-schemas. In fact, some texts do include statements of this type, particularly in early chapters, for clarity.

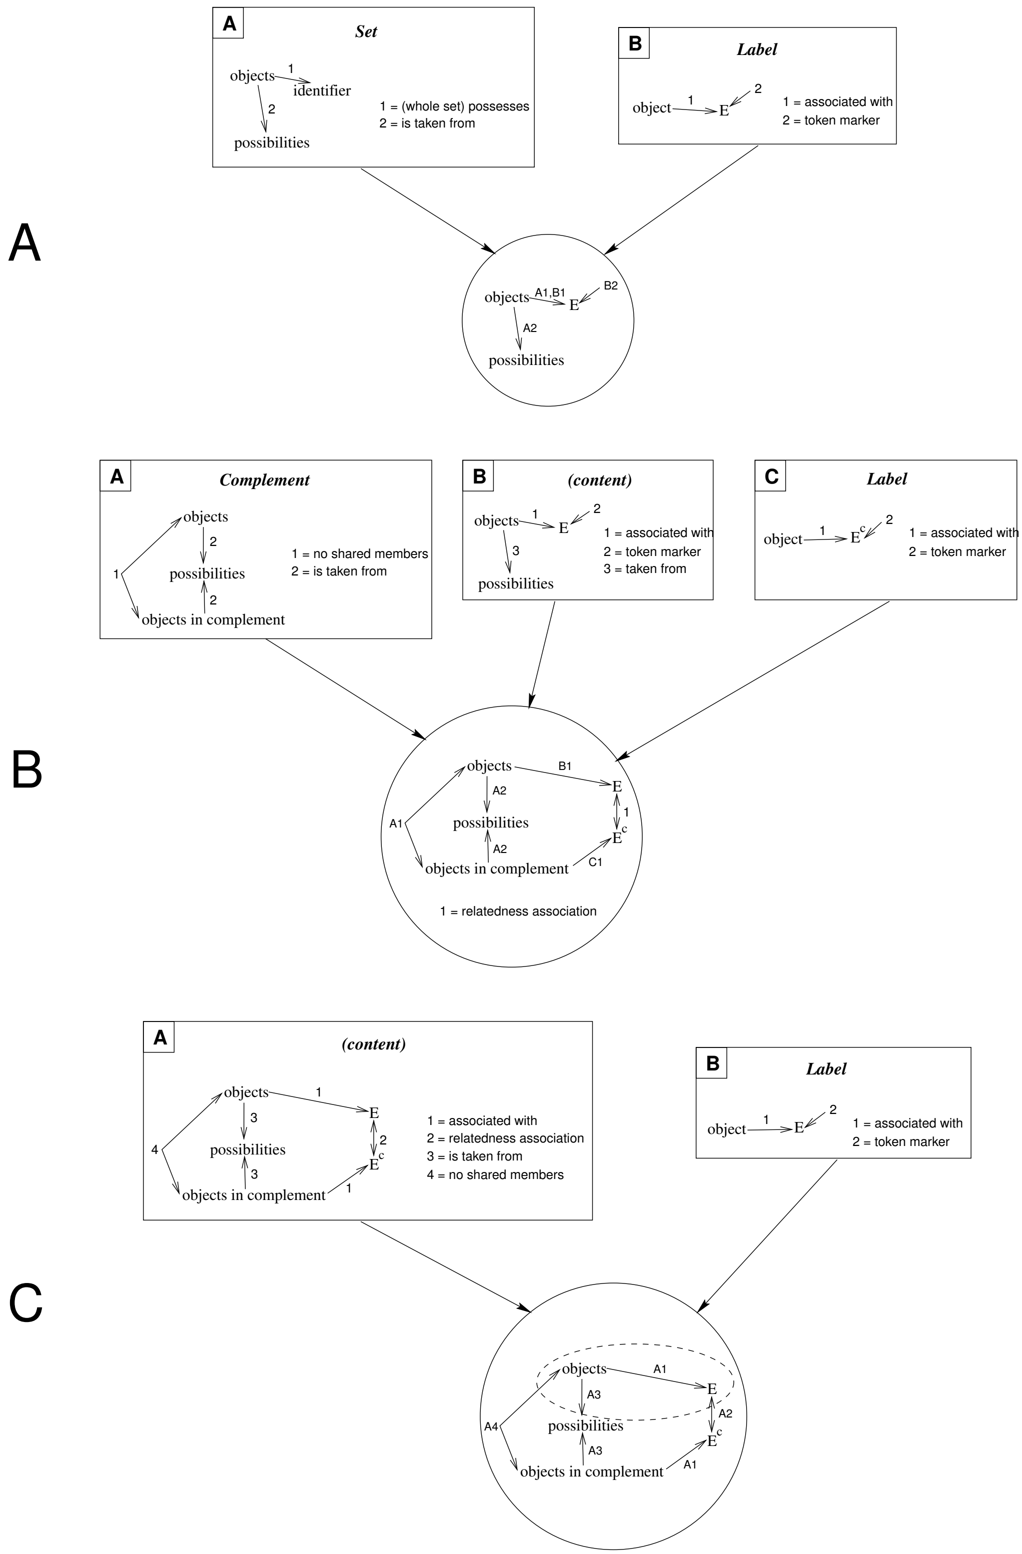
A
Set
objects
identifier
possibilities
1 = (whole set) possesses
2 = is taken from
B
Label
object
E
1 = associated with
2 = token marker
A
objects
A1,B1
E
B2
A2
possibilities
A
Complement
objects
possibilities
objects in complement
1 = no shared members
2 = is taken from
B
(content)
objects
E
possibilities
1 = associated with
2 = token marker
3 = taken from
C
Label
object
E^c
1 = associated with
2 = token marker
B
objects
B1
E
A2
A1
possibilities
E^c
objects in complement
C1
1 = relatedness association
A
(content)
objects
E
possibilities
E^c
objects in complement
1 = associated with
2 = relatedness association
3 = is taken from
4 = no shared members
B
Label
object
E
1 = associated with
2 = token marker
C
objects
A1
E
A3
A4
possibilities
A2
E^c
objects in complement

Figure 7:

[Caption for Figure 7]: Blends in steps 1–2 of proof. As in Figures 1, 3, and 4, cognitive structures are depicted by 'entity and relation' diagrams that provide only a coarse level of detail. In this case, we leave out even some of what was in Figure 4 for SET, because we suggest less-relevant aspects of concepts are cognitively backgrounded. In (A) and (B), the input on the left contains an m-schema. The input on the right in (A), substep 1.1, consists of a 'labeled instantiation' schema – a label ('E') and a slot for an instantiated entity. In (B), substep 1.2, the right input again contains a 'labeled instantiation' schema, while the middle input contains the structure from the previous blend. In (C), substep 2.1, the label 'E' is combined with the structure built thus far to result in a version of that structure in which the entity associated with the label is made prominent.

In substep 3.1, focus is shifted back to the embedding of E and E^c into the COMPLEMENT schema, triggered by mention of the label E^c. In substep 3.2 (not illustrated), a blend is constructed using the $\notin$ schema as one input and the instantiations x and E^c as the other. The difference between this type of blend (type O1) and the earlier ones in the proof is that there are no new labels or entities being introduced, and *all* content items already engage in relationships, so that the situation regarding potential mappings of the schema is more restricted. In this case, the constraints do not generate a conflict with the specifications for the slots being mapped to in the $\notin$ schema, and this lack of dissonance is what is noted in substep 3.3.

In most cases when similar slot-filler blends are made in natural language (as with our old friend, "Sally is the mother of Paul"), such an absence of dissonance suffices to accept the constructed meaning without further (or any) thought. However in mathematics, experience proves that dissonance may sometimes exist in such cases without being noticed (see following section), and one important check to eliminate this possibility is to construct the counterfactual blend and see that dissonance *does* exist in this case. Thus, a mathematician verifying the given proof would perhaps also carry out substeps 3.4 and 3.5, attempting to construct a blend (in a temporary mental space incorporating the structure built up prior to this step) using a schema contrary to the specified one ($\in$ instead of $\notin$) and observing that conflict arises (from the structure of the COMPLEMENT schema) which prevents the mapping from being formed (see Figure 8B).

Step 4 illustrates both an implicitly triggered focusing on background structure constructed earlier and the drawing of an inference that is not directly specified in the text. To understand the mathematics in this step, it is only necessary to know that a CLOSED SET contains all of its LIMIT POINTS. The way the inference it states is verified is by following this reasoning: "If x was a LIMIT POINT of E^c, it would be in E^c since E^c is closed. But this would violate the previous step." In terms of the mental operations, first, a

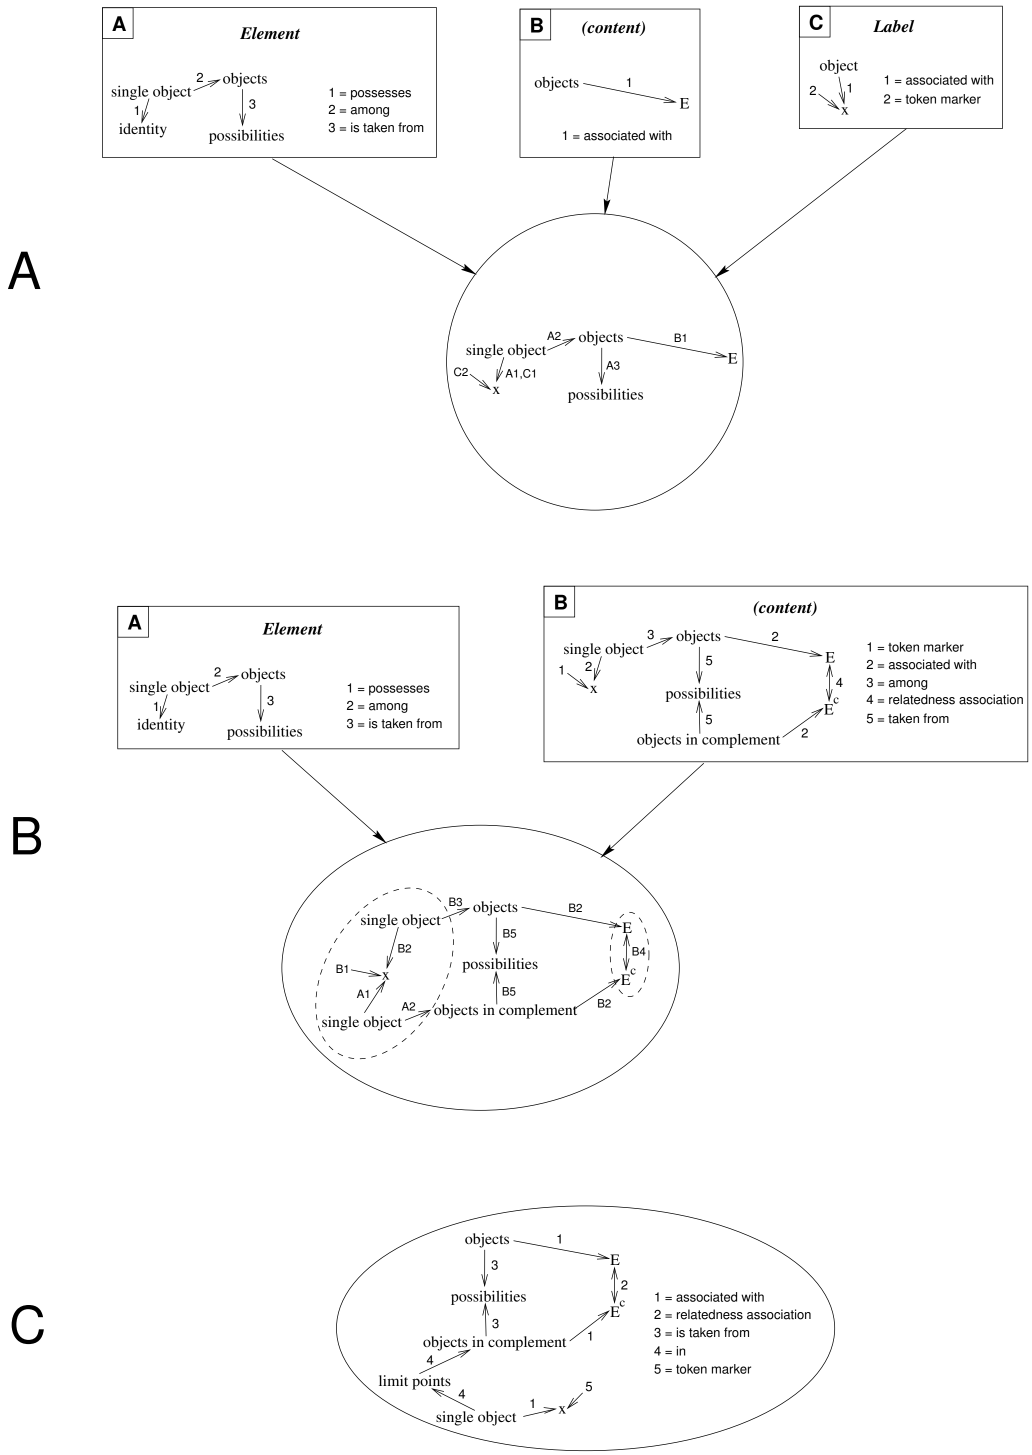

Figure 8: Blends in steps 2-4 of proof. In (A), substep 2.2, the instantiation E highlighted just previously has new structure added to it, labeled by 'x' and specified by the ELEMENT substructure of the SET schema. In (B), substep 3.4, the ELEMENT subschema with its SET component corresponding to E^c is blended with the previously built structure in which x is in E not E^c. This generates a conflict because of the properties of the COMPLEMENT schema, and the blend fails. (Content elements involved in the conflict are circled.) In (C), substep 4.3, the assumption that x is a LIMIT POINT of E^c is seen to imply x is in E^c, because the CLOSED schema (true of E^c) contains all its limit points.

counterfactual space is constructed in which x is a LIMIT POINT of E^c. Then focus is returned to the fact that E^c is CLOSED (substep 1.3). Since the LIMIT POINTS of a set play a particular role in the CLOSED schema (they are contained in the set), this leads to an observation – that x must be an element of E^c (Figure 8C). Shifting focus back to the inference previous step ($x \notin E^c$), a conflict is noted as in substep 3.5, and the counterfactual situation is rejected.

6 Issues Raised by the Analysis

6.1 Simple and Complex Concepts

The inferences we focused on in the analysis pertained to elementary relations in set theory and were very simple. We did this to make the presentation clear and emphasize, therefore, that the same mechanisms described are at work with more complex concepts. Although, as discussed earlier in connection with grounding, it is likely that the inferences relating to sets and containment are actually directly retrieved from memory by anyone with the prerequisite mathematical maturity to even encounter proofs such as the one above (but, see next section), our analysis elucidates the basis from which such inferences can be drawn when no memory exists. Appendix B contains an analysis of a different proof involving continuous functions that contains more difficult inference steps.

6.2 The Nature of Dissonance

The process we have analyzed consists of assembling particular blended configurations of m-schemas, mapping further m-schemas onto these, and detecting the presence of any dissonance in the mappings. Dissonance is constituted either by one of the three types of inviable sub-mappings described earlier, in which an entity does not fit the category constraining the role or "slot" it is mapped into in a blend, or by a structural mismatch , as when an element should clearly be mapped to multiple distinct roles within one input schema. Although in cases like step 3 above the dissonance is obvious and noted immediately by interpreters, in other cases it is not, and the situation must be considered attentively and at length to detect it (see next section for examples). This highlights the fact that neither m-frames nor even m-schemas are held completely active in the mind at one time. Instead, they serve as sources of connected background structure from which relevant substructure is drawn as attention is directed to various elements and aspects of instantiated structures.

This last point should clarify why it is *not* the case that being able to read through a proof and verify the appropriate dissonances and consonances implies that it is well-understood. It is possible to satisfy oneself as to the validity of each step without being able to hold the entire structure in mind at a single time. Nevertheless, most, if not all, mathematicians consider grasp of the entire structure essential to understanding the proof. Hadamard (1945), for example, writes, "I do not feel that I have understood it [a mathematical argument] as long as I do not succeed in grasping it in one global idea and ... this often requires a more or less painful exertion of thought." We will discuss how this process might be understood within our framework after describing our empirical study.

6.3 Structural Classification of Mental Operations

6.4 Mental Operations in Interpreting Proof Statements

The types of elementary operations found thus far sufficient for the analysis of statements from proofs in analysis and algebra are listed in Table 2, and the seven among them that involve the construction of blends are illustrated structurally in Figure 9. We have classified these based on two factors: the character of the result (whether it is an instantiation, focusing, or observation operation), and the number of inputs and the degree of specificity of their contents (whether they refer to schemas or frames, or particular instances). Note that for the inputs in this figure we follow the same convention as throughout the paper of placing the most general structures on the left and the most specific and focal structures on the right. Thus, m-schemas and m-frames go to the left of i-structures go to the left of labels for specific i-structures or parts of them.

The classification intentionally takes no account of the complexity of the various components in terms of the number of roles in a schema, the number of distinct elements in a frame, and so on, because there is no clear foundation provided for by the theory of conceptual structure we assume for drawing discrete boundaries along these dimensions. The intent is to identify discrete classes of mental operations that, component complexity being equal, are likely to be distinct in terms of level of difficulty from a mathematical perspective, and also to characterize dimensions of variation that may be useful in distinguishing the thought processes that go on during mathematical proof interpretation from those in other, nonmathematical communication. For this purpose, we rely on the distinctions between schema and instance and between schema and frame as outlined above in Section 2. Labels, which provide transient names to temporary constructs in a parallel fashion to the names definitions attach to entrenched concepts, are distinguished specially rather than being considered simply another schema-instance blending operation because this

Table 2: Elementary Mental Operations for Proof Interpretation

CLASS	#	OPERATION
instantiation	I1	imagine entity exemplifying or satisfying part or all of an m-schema
	I2	imagine entity involving blended combination of m-schemas (usually embedding of extant, focused entity into a new frame)
focus	F1	shift focus to structure present (focused on) in previous discourse – triggered by (e) text, either explicitly or semiexplicitly (e.g., symbol in text is part of substructure of a previously given schema), or (i) implicit cue
	F2	focus on an element/substructure in an imagined structure, usually specified by a subschema
observation	O1	blend an m-schema from background knowledge with current focused structure, triggered by (e) text, explicitly or semiexplicitly, or (i) implicit cue; O1′: use a schema contrary to that specified from background knowledge
	O2	blend an m-frame from background knowledge or previous discourse (e.g., a lemma), triggered by (e) text, explicitly or semiexplicitly, or (i) implicit cue; O2′: use a schema contrary to that specified from background knowledge
	O3	blend structure from previous discourse with current focused structure
inspection	N1	note consonance/dissonance in immediate mapping situation
	N2	draw relevant schematic inference/notice relevant schematic implication

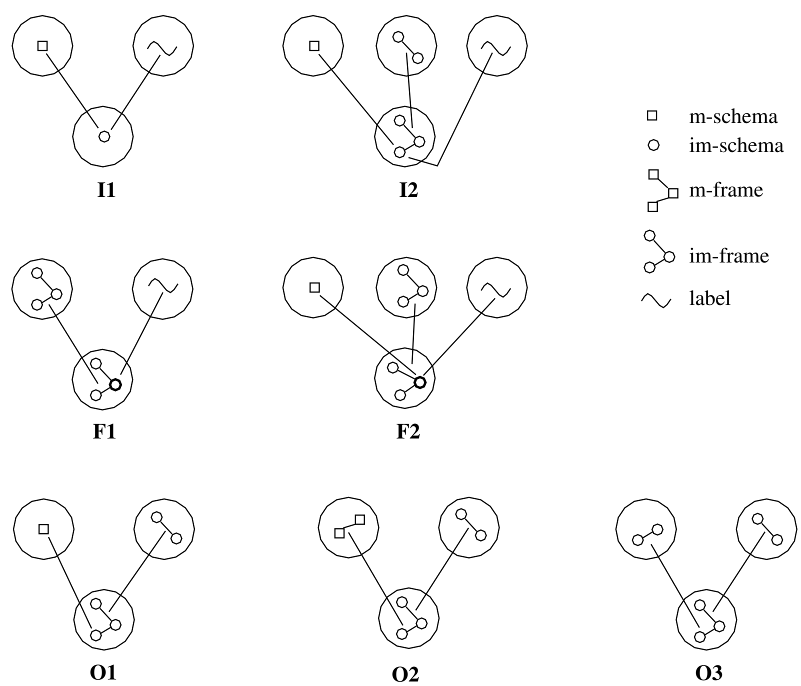

Figure 9: Types of blends corresponding to the classes of elementary mental operations given in Table 2.

function assumes particular importance in mathematical proofs as a signal for instantiation.

As far as differences between mathematical and nonmathematical communication processes, it appears likely that 1) the usage patterns of instantiation, focus, and observation differ, and 2) the complexity of the components for comparable structural blends differs. We will return to this issue in the Discussion. Whether there are differences along the dimension of input types to a given class of blend is a question left for another work.

Despite the attempt to rely on obvious theoretical distinctions, the classification left some latitude for judgment, namely, relating to how fine the classes should be. For example, there seems a major cognitive difference between trying to combine an entire m-frame with the present i-structure and only combining an m-schema (O1 vs. O2), but the distinction between refocusing on an i-structure with or without some new substructure (F1) seems less important. One alternative to the proposed classification would involve combining O1-O3 on the basis of their common global structure (only one input changes between them), while another would involve dividing O2 into two types based on the fact that the left-hand input comes from background knowledge in one case and local constructions in the other. In an earlier version of the classification, we included an extra focusing blend type similar to F1 but differing in that between the previous and current explicit references to a given input im-schema new substructure had been added to it.

7 Empirical Validation

As mentioned, the above analysis was developed on the basis of the author's introspection, based on his knowledge and experience as an undergraduate mathematics major. It was desirable, therefore, to obtain some confirmation that these introspective impressions were not primarily due to idiosyncracy or theory-generated preconceptions. To this end, an empirical study was performed in which undergraduate and graduate mathematics students were asked to report on their mental processes during their interpretation of proofs that were previously analyzed.

7.1 Methods

Students were presented with two proofs in introductory mathematical analysis (one at a time) and requested to describe their mental processes in reading and understanding them, as prompted by the following instructions (which were presented in written form):

> Using the paper provided to you by the experimenter, please describe, using words and/or diagrams, your mental process as you read this proof and convince yourself that it is true. Please take as much time as you want and provide as much detail as you can. We are not interested so much in, for example, how a teacher would describe the proof, but rather what is going through YOUR OWN MIND that leads you to BELIEVE it. Imagine that, for example, you are explaining to someone unfamiliar with mathematics (your mother, or a student in philosophy or English literature) what exactly it is you do when reading your mathematics text. You may read through the entire proof and think about it if you wish before writing anything. You may refer to the packet of definitions provided by the experimenter at any time.

These instructions were designed guided by considerations drawn from critical reviews of what conclusions can and cannot be drawn from subject self-report data (e.g., Ericsson & Simon, 1993; Goguen & Linde, 1993). The aim was to elicit descriptions of mental processes that are highly precise and as uniform as possible despite involving primarily nontechnical language. For half of the subjects, some additional guidance as to the structure of the report was provided by splitting the proofs up into steps and asking them to report for each step separately (the "split" condition). The other half could split the proof up in any manner they chose (the "block" condition).

Students were also asked to state what they thought the "main idea or crux" of each proof was, and to reconstruct the entire proof from memory. The former provided an indication as to whether the students formed a conception of the proof as a whole, rather than its individual steps, and the latter provided information on how confident students were in their understanding (students were also asked to rate their grasp of the proofs directly). Three groups of students at the University of California, San Diego served as subjects: third year undergraduate mathematics majors (group U), pure mathematics graduate students in their fourth year or later (group G), and an intermediate group consisting of two third year graduate students in mathematical physics and one beginning graduate student in pure mathematics (group I). Each group contained 3 students, and the total included 2 females. While all groups had previously taken courses in mathematical analysis covering the theorems employed in the study, group U students were still in the process of familiarizing themselves with both the concepts and their applications as well as the process of constructing proofs, group G students were highly experienced with proofs in general and had long mastered the particular concepts used in our examples, and group I students were intermediate in experience. One of the two proofs used was the one subjected to analysis above; the second one is presented and analyzed in Appendix B.

7.2 Results

7.2.1 General

Students varied widely in both the amount and the nature of what they wrote, but for the most part variations were uncorrelated with group (U, G, or I), and we ignore it in our report. (This may have been a ceiling effect relating to the short, simple proofs we used, and investigations with more subjects and a wider range of difficulties are called for.) Regarding the division into steps, three of five students in the block condition broke their reports down by the steps of the proof in almost exactly the same manner as the split condition (determined from the analysis), while the other two employed larger units made up of several of the smaller steps. This suggests that the steps the proofs were broken down into in the theoretical analysis are cognitively salient.

As to the steps themselves, the most typical descriptions involved a small sketch and a few words clarifying how it related to the structure(s) specified by the text. Others just verbally broke down the step in a little more detail, one student commenting that, "Notice they are implicitly assuming you will fill in these 'obvious' properties...". Some described mental procedures like "Translate the whole thing into the X space via f...". Table 4 presents the responses of a typical student for the first half of the proof analyzed previously in Section 5. Below we focus on a few aspects of the report data relating most directly to the issues raised by our analysis.

7.2.2 Pictures

Two thirds of the students drew pictures for at least some of the steps, and 90% of the steps had pictures drawn for them by from 1–4 of the subjects. Students often emphasized that the pictures they drew were only partial representations of much richer mental images. In particular, the images contained dynamic elements. For example, one student elaborates, "I have a picture like I drew in my mind, but again it is more of a dynamical process: I see myself taking open sets on $f(X)$, and making them smaller and smaller so that I could fit infinitely many over $f(X)$, trying to contradict that it is compact." (G44) Often students would indicate dynamic elements verbally in the drawing, as, for example, an arrow pointing towards a series of points with the caption "infinite sequence of points in the set getting closer and closer to the limit point." (G78)

Because pictures (with or without dynamical elements) are a particularly natural representation for schemas originating in visuo-motor experience, these aspects of the reports lend support to the idea that

Table 4: Reports of Math Student on Proof Steps

#	STEP	MENTAL PROCESS
1.	First, **suppose** E^c is closed.	E^c is closed means that if there is a limit point of E^c then it is in E^c. A limit point of a set is a point such that there is a sequence of points in the set getting closer and closer to it. [picture] Is the limit point in E^c? Yes because its closed. Want to show E is open.
2.	**Choose** $x \in E$.	If $x \in E$ then the picture looks like this [picture]
3.	**Then** $x \notin E^c$,	Yup. [picture]
4.	**and** x is not a limit point of E^c.	If x were really close to the edge of E then there could be an infinite sequence of points in E^c that approach x and so x could be a limit point of E^c but if it was then x would be in E^c (and not in E) because E^c is closed.
5.	**Hence** there exists a neighborhood N of x such that $E^c \cap N$ is empty,	Since x is not a limit point of E there exists a neighborhood around x such that the neighborhood does not contain any point of E^c (the neighborhood might have to be really small but if it doesn't exist then x is a limit point of E^c and we said this wasn't possible).
6.	**that is**, $N \subset E$.	If N does not contain any points of E^c then $N \subset E$.
7.	**Thus**, x is an interior point of E,	An interior point of E is a point that contains a neighborhood that is entirely contained in E. Thats what we just showed x was.
8.	**and** E is open.	Since x was *any* point of E then that means that every point of E is an interior point and this is the definition of open.

the students are employing structures similar to those we proposed in their understanding of the proofs, at least in the respect of involving links to experientially-based schemas. However, they do not prove that such schemas underlie their primary representation of the concepts, for the same reason that direct or indirect evidence of the use of mental imagery in a task does not imply *it* is fundamental: it may be an epiphenomenal appendage to operations in a different underlying medium (Pylyshyn, 1984). We will return to this issue below, but first we examine the individual differences that were found.

7.2.3 Modes of Thinking

As mentioned, not all students drew or mentioned using pictures, and students who did use them did so for some steps and not others. This suggests that conscious appeal to schemas is not a necessity but an option in understanding proofs, which is in fact in line with what our hypothesis on concept development (Section 4.5) suggests. In particular, we proposed that the experientially-grounded source schemas play a greater role in earlier stages of familiarity. The experience of some of the students seemed to support this idea. For example, one student, a sixth year graduate, elaborated on this point with great clarity.

> One point I should make, is that I can read this proof at different levels. I can read each step, without the mental picture, if I am familiar with the theorems, techniques, or notation involved. But if I get stuck at some point, I might have to go back and think through the ideas or pictures that enable me to understand the statements. Indeed, this is what I do in trying to understand new material. In fact, I draw pictures or special examples as I go along, to make things more clear, or I think about them in my head, if it is not too hard." (G44)

During informal questioning after the experiment, this student clarified that when he read the steps without mental pictures, he was aware of their meaning by direct knowledge of their implicational structure. He could see that the steps followed from one another, and, if required to produce a proof, he could do so simply by remembering that if one thing was the case, then certain other things were also, without any conscious reference to the auxiliary aid of an image.

We suggest that this corresponds to the third stage of conceptual development we outlined. Knowledge of the concept's relations with other mathematical concepts are stored directly in some way. We do not know whether in this case non-mathematical schemas are still active subconsciously, or if the implicational structure is simply memorized, or compiled into a separate structure independent of its foundational image schemas. We also do not know whether other individuals employ these two levels of thinking in the same way (i.e., images more for more complex or unfamiliar situations).

In fact, the opinions of some students seem to contradict this. For example:

> Using visualizations of these concepts ([e.g.,] a set with boundaries [picture] as closed and without them [picture] as open) can be very misleading when you are first reading the theorem – visualization helps me only when I'm sure that the theorem is self-consistent, because I can try to see which visualizations work and which don't. (G39)

However, that this student is in fact referring to imagery in particular rather than knowledge additional to lists of inferential relationships more generally is indicated by something he wrote down later:

> Probably I tend to use my background knowledge of the problem, even if I am trying to derive everything from first principles [i.e., logical implications from axioms]. I don't visualize concepts, but my intuition is guided by ideas that I grasp very intuitively (like a continuous function). (G39)

This student also indicated that he *did* use images to help him *remember* proofs and theorems, something difficult to explain if the underlying representations of the concepts do not have important aspects in common with images.

7.2.4 Counterfactuals

Some students (4/9), including two out of the three who did not draw pictures, reported constructing *counterfactuals* not specified by the text to convince themselves of the truth of three of the steps. These included, in particular, steps 4 & 5 of the proof analyzed earlier (see Table 4 for examples). These constructions strongly suggest that the proof steps are not simply interpreted by checking against an internal list of potential inferences for the current relational situation.

7.2.5 Directives

In some cases, the reports demonstrated that students were sensitive to the complexity and referential meta-information conveyed by directive keywords. For example, under a step with text "Since f is continuous", one student writes, "OK, this is a key fact. Be alert for important consequences to follow." (U22) Another wrote for the same step, "I know that we will be using the property that if f is continous and X is open then $f^{-1}(X)$ is open." (G78)

7.3 Conclusions

Overall, the results of the empirical study support the idea that our analysis *is* getting at some aspects of mathematical thinking that play a role in students' understanding of mathematical proofs. The results are not, however, precise enough to provide confirmation of the details of the analysis – i.e., that blends are constructed online exactly as proposed. If future studies are to address this, it is clear that the possible differences in modes of thinking must be taken into account, and before this can be done a clearer theoretical conception of what these differences are (as we are developing) would be useful.

8 Implications

8.1 Parallels and Differences Between Mathematics and Language

We have argued that while derivation of theorems within formal mathematics is driven by a relatively spare set of operators acting on structures disconnected from the rest of experience, proofs in informal mathematics in fact rely on a much richer set of interpretive mechanisms. Our analysis showed that both grounding and proof interpretation can be understood as the interaction of cognitive mapping machinery with extensive conceptual structures, sometimes brought in without explicit specification from discourse context or background knowledge, deriving ultimately from experience. Conventions of mathematical expression such as the use of directive terms and the frequent use of certain letters to stand for certain concepts serve to guide this process in a general way. Overall, this view of mathematical practice shares much in common with the view of natural language developed within cognitive linguistics, in which sentences are understood through mappings directed (but not specified) by syntactic structure (Fauconnier & Sweetser, 1996; Fauconnier, 1997; Langacker, 1987).

We suggest that an important difference, however, is that conscious awareness of underlying conceptual operations is greater in mathematics than language. One reason for this is that the schemas employed in blends in the former case are novel composite structures that have been consciously constructed by the users rather than subconsciously acquired from regularities in linguistic and extralinguistic experience. Another reason is that whereas in language it is generally only the results of blending that are attended to, in proofs conscious attention must be paid to whether or not inferential blends actually *work* – that is, whether or not dissonance arises as a result of the mapping.

This high degree of conscious awareness enabled us to provide a detailed account of the structure of a class of concepts (those constructed by mathematical definitions) and to identify and classify the cognitive components of the interpretive process for mathematical proofs. It is possible that both of these accounts apply to natural language processes as well: that novel concepts are acquired and manipulated as blended configurations of more familiar concepts or schemas, and that declarative sentences are understood by constructing their meanings via blending operations falling into the classes mentioned. This is a set of issues for consideration in future work.

One point worth mentioning in connection with the parallels we *have* provided evidence for is that the 'formal' view that the operation of mathematics *is* an encapsulated system operating according to regular rules has been largely rejected in mathematical practice. It exists mainly as a subject of academic interest to those engaged in foundations research or automated theorem proving. Axiomatizability *is* seen as the hallmark of consistent, systematic development, but both unspecialized teaching and advanced research often take place outside of axiomatic systems, and even within them role of extra-formal processes is pervasive. The mathematician Morris Kline summarizes that the axiomatic approach serves as "a check on the thinking and systematic organization of the knowledge gained ... to supply a logical basis for what we already know," but it "limits the creative process, is often contrived, and is obscure," (Kline, 1966). Other mathematicians have echoed these sentiments (Hadamard, 1945; Hardy, 1928; Rota, 1993). This situation in mathematics suggests that it will be difficult to do completely without a non-Fregean approach to meaning in natural language either.

8.2 Understanding Trends in the Development of Mathematics

Owing to the parallels in grounding and manipulation just discussed, interpretation in informal mathematics is subject to some of the same user difficulties inherent to language. In particular, there are two potential sources of *miscommunication* leading to disagreement over the validity of theorems and proofs, depending on whether the problem lies with m-schemas or their manipulation.

The problem with m-schemas lies in their specification via natural language definitions. If the wording of a given definition is not sufficiently precise, there may be more than one definitional blend that is compatible with it. This can arise from uncertainty in how source schemas are combined in the definition, from the fuzzy, prototype-centered nature of the source schemas themselves (see Lakoff, 1987: c.2), or from differing ideas as to which sources are relevant. In any case the result is potentially that different mathematicians use the same terminology to refer to different m-schemas with different cognitive properties, leading to confusion and disagreement over which relationships hold.

A particular example of such a situation, for the definition of polyhedron, is described and examined in detail by Lakatos (1976). An initially rough definition was progressively honed over a period of almost 150 years as the role of the concept in different situations (making different source schemas seem relevant) was explored through mathematical discourse. Lakatos was skeptical that ambiguity can ever be completely resolved in any case but showed how discrepancies in properties between different mathematicians' versions of a concept can be made arbitrarily small in particular situations of interest, by means of a kind of dialogue he termed the 'method of proofs and refutations'.

The second kind of difficulty arises in the manipulation of m-schemas that occurs in proofs. As shown earlier, instantiations must be mapped into slots and their eligibility for the slots verified. Often several factors (required properties) bear on such eligibility, and, as the situation grows more complex, it becomes more difficult to hold everything clearly in mind without making any oversights. The more complex the schemas and their slot requirements, the more difficult and error-prone these kinds of operations will be. Evidence from the self-reports of proof interpretation supported this idea, for more complex steps were more likely to elicit counterfactual constructions or to be reported as unclear by students (steps 5 & 6 of proof given in Appendix B).

During the 18th and particularly the 19th centuries, awareness of these difficulties led to modifications of both mathematical concepts and the complexity of the inference steps used in their manipulation (Bell, 1945[13]). Mathematical concepts were defined more precisely and by appeal to simpler schemas from ordinary experience. One major development of this type was the so-called arithmetization of analysis (e.g., Kramer, 1970: c. 22), in which many abstract notions in calculus such as limit and continuity were tied definitively to specific constructions involving the natural numbers. Because the natural numbers and their properties are (for reasons beyond the scope of this paper) generally held to be more concrete and definitely apprehensible by mathematicians, the resulting formulations of concepts are less subject to problems of disagreement between mathematicians.

These simplifications to concepts in turn affected reasoning in proofs, it being easier to handle concepts and examine the interactions of their properties when they are simpler. And when the concepts remain complex, mathematicians have been increasingly careful to split reasoning into smaller steps with less needing

[13]This book, though inaccurate in some biographical and attributive details, provides an excellent survey of the development of mathematics. See in particular chapters 13 and 23.

to be verified in each step. This was the substance of the trend towards increasing rigor in mathematics.

The culmination of these trends of increasing precision and rigor came at the beginning of the 20th century, with the instigation of formal mathematics as we described in Section 3. Although these efforts are largely considered to have been directed at resolving the problem of foundations, precision and rigor are intimately connected with this since they are required for a solid connection to whatever base is settled upon (see Bell, 1945: c. 15, 23).

9 Conclusions and Future Directions

In this paper, the blended spaces framework of Fauconnier & Turner has been used to provide a description of the grounding of mathematical concepts in everyday experience, and some proposals have been made as to how these conceptual structures interact in the interpretation of mathematical proofs. The analysis resulting from this proposal highlights three important observations.

Firstly, the cognitive operations involved in interpreting proofs do not seem qualitatively different from those involved in ordinary natural language, but their structure is generally more consciously accessible in mathematics. This suggests that one direction to extend the present work would be to utilize the analysis of mathematical discourse as a *window* on processes in natural language interpretation. For example, the fine detail of patterns of reference to discourse context is accessible to introspection in mathematical proof, and it was even possible to classify focus-shifts and blending operations on this basis. Another area which might potentially benefit is the understanding of acceptability constraints for conceptual mappings (see Turner (1991) and Lakoff (1993)). Such constraints are accessible to introspection in mathematical reasoning because it is often the question of whether a given mapping is acceptable or not that is focused on in verifying a proof step.

Secondly, pending further development, aspects of the utility of artificially defined concepts for human reasoning may be effectively studied within cognitive semantics. In particular, better understanding of the nature of conflict generation in the mapping process may afford predictions about the *viability* of potential concepts, while their *usability* might be clarified through characterizing the effects of working memory limitations on the construction of cognitive mappings.

Finally, our account of proof interpretation expands the notion of declarative-procedural knowledge compilation as discussed in some production-system frameworks for cognition (e.g., Anderson (1993)). Interpreting proof steps to the point they can be seen to be true at first requires reference to the conceptual sources of the mathematical concepts, and the resulting implicatures may be memorized if they occur frequently, eventually obviating the need for such referring back.

Appendix A: Definitions of Concepts Used in Analyzed Proofs

A *neighborhood* of a point p is a set $N_r(p)$ consisting of all points q such that $d(p, q) < r$ [distance between p and q < r].

A point p is a *limit point* of the set E if *every* neighborhood of p contains a point $q \neq p$ such that $q \in E$.

[A set] E is *closed* if every limit point of E is a point of E.

A point p is an *interior point* of E if there is a neighborhood N of p such that $N \subset E$.

[A set] E is *open* if every point of E is an interior point of E.

By an *open cover* of a set E in a metric space X we mean a collection $\{G_\alpha\}$ of open subsets of X such that $E \subset \cup_\alpha G_\alpha$.

A subset K of a metric space X is said to be *compact* if every open cover of K contains a *finite* subcover.

Suppose X and Y are metric spaces, $E \subset X$, $p \in E$, and f maps E into Y. Then f is said to be *continuous at* p if for every $\epsilon > 0$ there exists a $\delta > 0$ such that

$$d_Y\left(f(x), f(p)\right) < \epsilon$$

for all points $x \in E$ for which $d_X(x, p) < \delta$.

Appendix B: Analysis of a Second Proof

Thm. 4.14[14]: *Suppose f is a continuous mapping of a compact metric space X into a metric space Y. Then $f(X)$ is compact.*

Let $\{V_\alpha\}$ be an open cover of $f(X)$. Since f is continuous, Theorem 4.8 shows each of the sets $f^{-1}(V_\alpha)$ is open. Since X is compact, there are finitely many indices, say $\alpha_1, ..., \alpha_n$, such that

$$X \subset f^{-1}(V_{\alpha_1}) \cup \cdots \cup f^{-1}(V_{\alpha_n}). \tag{1}$$

Since $f(f^{-1}(E)) \subset E$ for every $E \subset Y$, (1) implies that

$$f(X) \subset V_{\alpha_1} \cup \cdots \cup V_{\alpha_n}. \tag{2}$$

This completes the proof.

[14] Rudin, 1976, p. 89.

#	STEP	FUNCTION	OP #	OP DESCRIPTION
1.	**Let** $\{V_\alpha\}$ be an open cover of $f(X)$.	introduce	I1	instantiate an entity labeled X in METRIC SPACE schema
			I1	instantiate an entity labeled Y in METRIC SPACE schema
			I2	instantiate an entity labeled f in FUNCTION schema with domain and range labeled X and Y
			I2	assign label $\{V_\alpha\}$ to OPEN COVER schema with argument $f(X)$
2.	**Since** f is continuous,	introduce	F1e	shift focus to a different aspect of $f(X)$ (from context)
3.	Theorem 4.8 **shows** each of the sets $f^{-1}(V_\alpha)$ is open.	infer	O3e	blend $f(X)$ and V_α (with OPEN property) into the theorem's schematic structure
			N2	observe consonance
4.	**Since** X is compact,	introduce	F1i	shift focus to X as blended with METRIC SPACE schema in theorem preconditions
			I2	embed X in COMPACT schema, also in preconditions
5.	**there are** finitely many indices, say $\alpha_1, \ldots, \alpha_n$, **such that** $X \subset f^{-1}(V_{\alpha_1}) \cup \cdots \cup f^{-1}(V_{\alpha_n})$ [this equation labeled (12) in text].	infer	F1i	recall blend from step 1 ($\{V_\alpha\}$ an open cover of $f(X)$)
			O1e	blend inverse function schema with elements of $f(X)$ as arguments
			N2	note implication that all of X is covered since all of $f(X)$ was
			F1i	recall X is COMPACT
			I2	associate sets $f^{-1}(V_{\alpha_i})$ with COVER subschema of COMPACT schema
			N2	infer according to COMPACT schema that a finite subcollection of these sets is also a COVER
6.	**Since** $f(f^{-1}(E)) \subset E$ for every $E \subset Y$,	infer	I2	instantiate E as a set of Y, given in the preconditions
			I2	blend E with f into INVERSE schema
			F2	shift focus to $f^{-1}(E)$
			O1e	blend function application schema with this and f
			N2	observe result is certainly within E
7.	(12) **implies** that $f(X) \subset V_{\alpha_1} \cup \cdots \cup V_{\alpha_n}$.	infer	F1	recall (12) as a more detailed m-frame (i.e., to level of detail requiring reference to written representation)
			O1e	blend function application schema with this equation and f, possibly using written representation to keep track (result is $f(X) \subset f(f^{-1}(V_{\alpha_1})) \cup \cdots \cup f(f^{-1}(V_{\alpha_n})))$
			O3	blend result from previous step (recalled because of similarity) with each term of the right-hand side
			N2	perhaps using the written representation of (12) to keep track, draw the implication of that previous result that each term may be replaced by the entity on the other side of the subset, leaving the result given
8.	This completes the proof.	infer	F1i	recall that V_α was a COVER, and that indices $\alpha_1, \ldots, \alpha_n$ represented a finite subcollection
			O1i	blend COMPACT schema with $f(X)$ and the finite covering subcollection $V_{\alpha_1} \cup \cdots \cup V_{\alpha_n}$
			N1	observe consonance/satisfaction

References

[Anderson, 1990] Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Lawrence Erlbaum Associates. chapter 3.

[Anderson, 1993] Anderson, J. R. (1993). *Rules of the Mind*. Hillsdale, NJ: L. Erlbaum Associates.

[Anderson et al., 1981] Anderson, J. R., Greeno, J. G., and Kline, P. J. (1981). Acquisition of problem-solving skill. In Anderson, J. R., editor, *Cognitive Skills and their Acquisition*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[Barsalou, 1983] Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition*, 11:211–227.

[Barsalou, 1987] Barsalou, L. W. (1987). The instability of graded structure: Implications for the nature of concepts. In Neisser, U., editor, *Concepts and Conceptual Development*, pages 101–40. Cambridge: Cambridge University Press.

[Bell, 1945] Bell, E. T. (1945). *The Development of Mathematics*. New York: McGraw-Hill.

[Borasi, 1993] Borasi, R. (1993). Appreciating the humanistic elements within mathematical content: The case of definitions. In White, A. H., editor, *Essays in Humanistic Mathematics*, pages 123–39. Washington, D. C.: Mathematical Association of America.

[Bowerman, 1996] Bowerman, M. (1996). Learning how to structure space for language: A crosslinguistic perspective. In Bloom, P., Peterson, M. A., Nadel, L., and Garrett, M. F., editors, *Language and Space*. Cambridge, Mass.: MIT Press.

[Bundy, 1983] Bundy, A. (1983). *Computer Modelling of Mathematical Reasoning*. New York: Academic Press.

[Choi and Bowerman, 1991] Choi, S. and Bowerman, M. (1991). Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition*, 41:83–121.

[D'Andrade, 1995] D'Andrade, R. G. (1995). *The Development of Cognitive Anthropology*. Cambridge: Cambridge University Press.

[Davis and Hersh, 1981] Davis, P. J. and Hersh, R. (1981). *The Mathematical Experience*. Boston: Birkhauser.

[Ericsson and Simon, 1993] Ericsson, K. A. and Simon, H. A. (1993). *Protocol Analysis: Verbal Reports as Data, revised edition*. Cambridge, MA: MIT Press.

[Farmer et al., 1995] Farmer, W. M., Guttman, J. D., and Thayer, F. J. (1995). Contexts in mathematical reasoning and computation. *J. Symbolic Computation*, 19:201–16.

[Fauconnier, 1985] Fauconnier, G. (1985). *Mental Spaces*. Cambridge, Mass.: MIT Press.

[Fauconnier, 1997] Fauconnier, G. (1997). *Mappings in Thought and Language*. Cambridge: Cambridge University Press.

[Fauconnier and Sweetser, 1996] Fauconnier, G. and Sweetser, E. (1996). *Spaces, Worlds, and Grammar*. Chicago: University of Chicago Press.

[Fauconnier and Turner, 1994] Fauconnier, G. and Turner, M. (1994). Conceptual projection and middle spaces. Technical Report 9401, Department of Cognitive Science, University of California, San Diego. Available from `http://cogsci.ucsd.edu`.

[Fauconnier and Turner, 1998] Fauconnier, G. and Turner, M. (1998). Conceptual integration networks. *Cognitive Science*, 22:133–87.

[Fillmore, 1985] Fillmore, C. (1985). Frames and the semantics of understanding. *Quaderni di Semantica*, 6.2:222–53.

[Frege, 1953] Frege, G. (1953). *The Foundations of Arithmatic*. New York: Philosophical Library. Originally published 1884 Die Grundlagen der Arithmetik.

[Gemignani, 1971] Gemignani, M. C. (1971). *Axiomatic Geometry*. Reading, Mass., Addison-Wesley Pub. Co.

[Gentner, 1983] Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7:155–70.

[Goguen and Linde, 1993] Goguen, J. A. and Linde, C. (1993). Techniques for requirements elicitation. In IEEE, editor, *Proceedings of the IEEE International Symposium on Requirements Engineering*. Los Alamitos, CA: IEEE Computer Society Press.

[Goldstone et al., 1991] Goldstone, R. L., Medin, D. L., and Gentner, D. (1991). Relational similarity and the nonindependence of features in similarity judgments. *Cognit. Psychol.*, 23:222–62.

[Grimaldi, 1989] Grimaldi, R. P. (1989). *Discrete and Combinatorial Mathematics*. Reading, MA: Addison-Wesley.

[Grossberg, 1987] Grossberg, S. (1987). *The Adaptive Brain, volume I: Cognition, Learning, Reinforcement, and Rhythm*. Amsterdam: North-Holland.

[Hadamard, 1945] Hadamard, J. (1945). *An Essay on the Psychology of Invention in the Mathematical Field*. New York: Dover.

[Hampton, 1987] Hampton, J. A. (1987). Inheritance of attributes in natural concept combinations. *Memory and Cognition*, 15:55–71.

[Hardy, 1928] Hardy, G. H. (1928). Mathematical proof. *Mind*, 38:1–25.

[Heath, 1956] Heath, T. L., editor (1956). *The Thirteen Books of Euclid's Elements*. New York: Dover Publications.

[Hebb, 1949] Hebb, D. O. (1949). *The Organization of Behavior*. New York: Wiley.

[Hilbert, 1959] Hilbert, D. (1959). *The Foundations of Geometry*. La Salle, Ill.: Open Court Pub. Co. originally published Grundlagen der Geometrie, 1899.

[Hilbert and Bernays, 1968] Hilbert, D. and Bernays, P. (1968). *Grundlagen der Mathematik*. Berlin: Springer. Originally published (2 vol.s) 1934, 1939.

[Hintzman, 1986] Hintzman, D. L. (1986). 'Schema abstraction' in a multiple-trace memory model. *Psychol. Rev.*, 93:411–28.

[Hoffman and Kunze, 1971] Hoffman, K. and Kunze, R. (1971). *Linear Algebra, 2nd Edition*. Englewood Cliffs, N. J.: Prentice Hall.

[Hofstadter and Group, 1995] Hofstadter, D. and Group, F. A. R. (1995). *Fluid Concepts and Creative Analogies*. New York: Basic Books.

[Holyoak and Thagard, 1995] Holyoak, K. J. and Thagard, P. (1995). *Mental Leaps: Analogy in Creative Thought*. Cambridge, Mass.: MIT Press.

[Homa and Vosburgh, 1976] Homa, D. and Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. *J. Exper. Psychol.: Human Learning and Memory*, 2:322–30.

[Hummel and Holyoak, 1997] Hummel, J. E. and Holyoak, K. J. (1997). Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104:427–466.

[Johnson, 1987] Johnson, M. (1987). *The Body in the Mind*. Chicago: University of Chicago Press.

[Keil, 1979] Keil, F. C. (1979). *Semantic and Conceptual Development: An Ontological Perspective*. Cambridge, Mass.: Harvard University Press.

[Kitcher, 1984] Kitcher, P. (1984). *The Nature of Mathematical Knowledge*. Oxford: Oxford University Press.

[Kline, 1966] Kline, M. (1966). Mathematics and axiomatics. In of the Mathematical Sciences, T. C. B., editor, *The Role of Axiomatics and Problem Solving in Mathematics*, pages 57–62. Washington, D. C.: Ginn and Co.

[Koedinger and Anderson, 1990] Koedinger, K. R. and Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14:511–50.

[Komatsu, 1992] Komatsu, L. K. (1992). Recent views of conceptual structure. *Psychological Bulletin*, 112:500–26.

[Kramer, 1970] Kramer, E. E. (1970). *The Nature and Growth of Modern Mathematics*. New York: Hawthorn Books.

[Lakatos, 1976] Lakatos, I. (1976). *Proofs and Refutations*. Cambridge: Cambridge University Press.

[Lakoff, 1987] Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.

[Lakoff, 1993] Lakoff, G. (1993). The contemporary theory of metaphor. In Ortony, A., editor, *Metaphor and Thought, 2nd Edition*, pages 202–51. Cambridge University Press.

[Lakoff and Johnson, 1980] Lakoff, G. and Johnson, M. (1980). *Metaphors We Live By*. Chicago: University of Chicago Press.

[Lakoff and Nuñez, 1997] Lakoff, G. and Nuñez, R. (1997). The metaphorical structure of mathematics: Sketching out cognitive foundations for a mind-based mathematics. In English, L., editor, *Mathematical Reasoning: Analogies, Metaphors, and Images*, pages 21–89. Hillsdale, NJ: Erlbaum.

[Lakoff and Nuñez, 2000] Lakoff, G. and Nuñez, R. (2000). *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. New York: Basic Books.

[Langacker, 1987] Langacker, R. W. (1987). *Foundations of Cognitive Grammar, Volume I: Theoretical Prerequisites*. Stanford, Calif.: Stanford University Press.

[Logan, 1988] Logan, G. D. (1988). Toward an instance theory of automatization. *Psychol. Review*, 95:492–527.

[Malt and Smith, 1984] Malt, B. C. and Smith, E. E. (1984). Correlated properties in natural categories. *J. Verbal Learning and Verbal Behavior*, 23:250–69.

[Mandler, 1984] Mandler, J. (1984). *Stories, Scripts, and Scenes: Aspects of Schema Theory*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[Medin et al., 1982] Medin, D. L., Altom, M. W., Edelson, S. M., and Freko, D. (1982). Correlated symptoms and simulated medical classification. *J. Exper. Psychol.: Learning, Memory, and Cognition*, 8:37–50.

[Minsky, 1975] Minsky, M. (1975). A framework for representing knowledge. In Winston, P., editor, *The Psychology of Computer Vision; New York: McGraw-Hill*. New York: McGraw-Hill.

[Peano, 1908] Peano, G. (1895-1908). *Formulaire de Mathematiques*. Turin: Bocca Freres.

[Pylyshyn, 1984] Pylyshyn, Z. (1984). Mental imagery and functional architecture. In Pylysyhn, Z., editor, *Computation and Cognition*, pages 225–83. Cambridge, Mass.: MIT Press.

[Quinn and Holland, 1987] Quinn, N. and Holland, D. (1987). Culture and cognition. In Holland, D. and Quinn, N., editors, *Cultural Models in Language and Thought*. Cambridge: Cambridge University Press.

[Rosch et al., 1976] Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., and Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognit. Psychol.*, 3:382–439.

[Rota, 1993] Rota, G. (1993). The concept of mathematical truth. In White, A. H., editor, *Essays in Humanistic Mathematics*. Washington, D. C.: Mathematical Association of America.

[Rudin, 1976] Rudin, W. (1976). *Principles of Mathematical Analysis, 3rd Edition*. New York: McGraw-Hill.

[Rumelhart, 1980] Rumelhart, D. E. (1980). Schemata: The building blocks of cognition. In Spiro, R., Bruce, B., and Brewer, W., editors, *Theoretical Issues in Reading Comprehension*, pages 33–58. Hillsdale, NJ: L. Erlbaum Associates.

[Russell, 1937] Russell, B. (1937). *Principles of Mathematics, 2nd edition*. New York, W. W. Norton & company. First edition published 1903.

[Schank, 1982] Schank, R. (1982). *Dynamic Memory: A Theory of Reminding and Learning in Computers and People*. Cambridge: Cambridge University Press.

[Turner, 1991] Turner, M. (1991). *Reading Minds*. Princeton, N. J.: Princeton University Press.

[Whitehead and Russell, 1925] Whitehead, A. N. and Russell, B. (1925). *Principia Mathematica, 2nd edition*. Cambridge: Cambridge University Press. First edition published 1910.

[Wigner, 1960] Wigner, E. P. (1960). The unreasonable effectiveness of mathematics in the natural sciences. *Comm. on Pure and Applied Math.*, 13:1–14.

[Wilder, 1944] Wilder, R. L. (1944). The nature of mathematical proof. *The American Mathematical Monthly*, 51:309–23.

[Wilder, 1952] Wilder, R. L. (1952). *Introduction to the Foundations of Mathematics*. New York: John Wiley and Sons.