Ontology, Society, and Ontotheology

Joseph A. Goguen
Department of Computer Science & Engineering
University of California at San Diego, USA

Abstract. This paper surveys a wide range of work currently little known among formal ontologists, but that shows promise for improving the state of the art. Two conclusions are that the embodied, embedded, situated nature of human concepts undermines attempts to reify context, and that category theory provides relevant tools for problems associated with structural and logical heterogeneity. Ethnomethodologists emphasize the negotiable, situated, embodied, emergent character of classification, as of all human activity. Cognitive linguists and psychologists study categorization, conceptual domains, metaphor and blending, and reach similar conclusions. Sociologists of science observe the intensely political and ethical aspects of classification systems, as well as their malleability, evolution, and local interpretation. French post-structuralists consider writerly texts, intertextuality, deconstruction, etc. Heidegger criticized "ontotheology" as the alienating notion of "being" that is the essence of modern technology. Taken together, these results motivate skepticism about extreme claims for ontologies in the technical sense of the Semantic Web, database integration, etc., despite the undoubted applicability of this technology to many specific problems. What can emerge from carefully considering skeptical arguments, hyperbolic claims, technical advances, and logical foundations is a balanced assessment of what seems possible and desirable, versus what seems impossible and undesirable, as well as a plea for greater humility, better ethics, better theory, and more humanity.

1 Introduction

One of the most pressing problems for the technological application of ontology[1] is to understand its limitations. One reason this problem is so pressing is that organizations, managers, and even experienced engineers, often expect too much, and know little or nothing about the technical details. This is partly due to the technology being new and relatively untried, partly due to the hyperbole that so frequently accompanies new technology, especially when it has a comforting reductionist flavor, and partly due to insincere marketing by some researchers and organizations. It should be noted that these reasons are mainly social, with of course an economic background.

Computer ontologies, which consist of logical axioms that relate terms of interest, have genuine promise when restricted to appropriate, well-understood domains, such as B2B transactions in a car manufacturer's supply chain. Though many engineers are no doubt skeptical, philosophy can indeed make significant contributions to understanding the nature and limitations of computer ontologies, and philosophical ontology is particularly relevant, though I do not think the connection as direct as some might, since I have in mind some ideas of Martin Heidegger. Moreover, other areas can also contribute to a more complete picture of

[1] There are two distinct but related senses of this word, its technical computer science sense, associated with the World Wide Web, and its technical philosophical sense; when confusion is possible, we call the first "computer" or "computational" ontology, and the second "philosophical" or "formal" ontology.

the difficulties with computer ontologies, which indeed have been and are being, encountered on a daily basis in many research centers. These areas include sociology of science, post-structuralism (which can be seen as part of philosophy), ethnomethodology, and cognitive linguistics. This paper attempts to assemble results and arguments from a range of fields, in order to discern some limits of computer ontology.

Unfortunately, limitations of time preclude writing detailed expositions of the diverse fields surveyed here, and even with three or four times as much space, it still would not be possible to give adequate expositions, let alone to convince skeptics about these sometimes rather radical views. But I hope that at least some readers will be convinced of the relevance of these areas, and be moved to learn more about some of them. I am afraid it is the nature of such a paper to raise more questions than it can answer, and I will be happy if it stimulates a wider debate about the relevance to computational ontology of disciplines that seriously challenge the current widespread somewhat uncritical acceptance of rather strong forms of philosophical realism.

Acknowledgements: I thank Jenny Wang and Young-Kwang Nam for collaboration on the schema matching tool, Kai Lin and Vitaliy Zavesov for work on its implementation, and Bertram Ludäscher for valuable discussions. This material is based on work partially supported by the National Science Foundation under Grant No. ITR 0225676, the Science Environment for Ecological Knowledge (SEEK) project.

2 Computational Ontology

Data integration is emerging as a major challenge in the early 21st century. The rise of inexpensive storage media, data warehousing, and especially the web, have made available vast amounts of data. But it can be very difficult to find what you want, and to combine it properly to get what you need. Our laboratory has designed and built a tool called SCIA which supports integration and transformation of databases having schemas in DTD and XML Schema format [1, 2, 3]; it will soon be extended to other kinds of schema. Since fully automatic schema mapping generation is infeasible, this tool attempts to minimize total user effort by identifying the critical decision points, where user input can yield the largest reduction of future matching effort.

Difficulties with schema integration include highly variable structure and quality of data and meta-data: science labs and businesses often have data stored in spreadsheets, or even just formatted files, with little or no documentation of format or meaning; moreover, some entries may be incomplete, corrupted, or inconsistent. If all documents had associated schemas (also called data models) to accurately describe their structure, and if fully automatic schema integration were feasible, then some interesting problems could be solved at the syntactic level. But these assumptions are far from true, and format is only a small part of the difficulty.

One proposed solution is ontologies, in the sense of formal terminological systems, items from which can be attached to items in e-documents. These cannot capture real world semantics, but only logical relations between terms, such as that all humans are mammals; the actual meanings of “human” and “mammal” remain unformalized, as do potential exceptions to logical relations (e.g., consider bionic appendages, cyborgs, androids, etc.). Moreover, a given domain may have several ontologies, each in some ways incomplete and/or ambiguous, and possibly written in different ontology languages, which in turn may be based on different logical systems. OWL and RDF are currently most prominent, but others include Ontologic, $\mathcal{ALC}$, KIF, KL-ONE, XSB, Flora, and OIL; specialized ontology languages, e.g., Ecolingua and EML for ecology, tend not to have a formal semantics. It follows that the ontology approach to data integration may require not just schema and ontology integration, but also ontology language integration, and even ontology logic integration, such that semantics is respected throughout the entire "integration chain," from actual datasets or "documents," through schemas and ontologies, up to ontology logics. It also follows that this is not sufficient to deal with low quality data.

A computer ontology is just a theory over a logic, i.e., a set of sentences in that logic. Using ontologies to integrate data raises issues analoguous to those for schema integration. Such issues can be addressed using **institutions** [4], which axiomatize the notion of logical system based on Tarski's idea that the *satisfaction* of a sentence by a model is fundamental. Institutions have been successfully applied to give semantics for powerful module systems [5], and multi-logic specification languages [6], databases [7], behavioral types, and semantics for object oriented programming [8], as well as to generalize many results in classical model theory, such as Craig interpolation [9]. See [10] for details of our approach using Grothendieck institutions [6] to integrate ontologies written in different logics.

It should not be thought that because category theory is very abstract, it is incompatible with an embodied, enacted, situated philosophy; on the contrary, its very abstractness makes it more useful, by freeing it from additional presuppositions, though of course, it remains true that anything said in the language of category theory is a model, an imperfect representation, useful for specific purposes in particular situations, not to be confused with any concrete entities that it might be used to represent.

3 Cognitive Science

This section reviews research from cognitive science that can help evaluate the potential of computer ontology. We first discuss *cognitivism*, a now receding movement which arose as a rebellion against the restrictive worldview of behaviorism, which tried to study behavior without invoking mind. Cognitivism in the broad sense of taking mind seriously, is admirable, but in fact, most cognitivist research takes a much more narrow view, in which cognition is considered computation, so that body, emotion, and society are neglected, and the representation of knowledge emerges as a central problem. In its classic form, now called "good old fashioned AI" or "GOFAI," knowledge is represented in symbolic logic, an approach which the logical positivists of the Vienna Circle would presumably have endorsed. The conspicuous failure of this approach, e.g., in the Japanese Fifth Generation project, has inspired a number of biologically motivated refinements, such as neural nets and so-called artificial life, which do not, however, abandon the computational model, nor do they solve the problem of representation, which can be more precisely formulated as the *symbol grounding problem*, stated (but not solved) by Stevan Harnad [11]: the issue is how the symbolic representations used in a computational model can come to refer to the real world. While the information processing models of cognitivism might be adequate for formalized games like chess, their exclusion (or cursory treatment) of embodiment, emotion and society render them unsatisfactory as a theory of what it means to be human [12, 13].

In the late 1960s, Eleanor Rosch began a systematic experimental study of categorization [14], which overturned then prevalent ideas about their propositional nature. In brief, basic-level categories (like "bird") are determined by similarity to prototypes, and are then expanded radially by analogies. This research, brilliantly summarized in [15], became the foundation for the "conceptual metaphor theory" (abbreviated "CMT") of George Lakoff and others, which has greatly deepened our understanding of metaphor [16]. One result is that many metaphors come in families, called **image schemas**, that share a common pattern. An example is BETTER IS UP, as in "I'm feeling up today," or "He's moving up into management," or "His goals are higher than that." Some image schemas, including this one, are grounded in the human body[2] and are called **basic image schemas**; they tend to yield the most persuasive metaphors.

It would seem a good very idea for formal ontologists to take account of what is known about the nature of human concepts and cognition, so that the ontologies that they construct can be as useful and comfortable to human users as possible.

Fauconnier and Turner [17] have studied **blending**, or **conceptual integration**, claiming it is a basic human cognitive operation, invisible and effortless, but fundamental and pervasive, appearing in the construction and understanding of metaphors, as well as many other cognitive phenomena, including grammar and reasoning. Simple examples of blends are two word phrases like "houseboat," "roadkill," "jazz piano," "computer virus" and "classical composer." Blending theory says that concepts come in clusters, called **conceptual spaces**, consisting of elements and relation instances among them [18]; note that this abstraction necessarily omits the qualitative, experiential aspects of what is represented. **Conceptual mappings** are partial functions from the item and relation instances of one space to those of another.

The simplest blends[3] have the form of Figure 1, where I_1 and I_2 are called the **input spaces**, B the **blend space**, and G the **generic space**; the latter contains conceptual structure that is shared by the two input spaces[4]. A **blendoid** of I_1, I_2 **over** G consists of a space B together with conceptual mappings $I_1 \to B$, $I_2 \to B$, and $G \to B$. There may be many such blendoids, but relatively few are likely to be interesting. Therefore additional principles are needed for identifying the most interesting possibilities, so that we can define a **blend** to be a blendoid that is *optimal* with respect to these principles. Fauconnier and Turner suggest a number of "optimality principles" for this purpose (see Chapter 16 of [17]), but they are too vague to be fully formalized, and seem mainly applicable to "common sense" blends, but not to generating creative poetic blends [19]. Whereas the CMT view of metaphor maps aspects of one domain to another, where the target domain concerns what the metaphor is "about," blending theory views metaphors as "cross-space mappings" that arise from blending conceptual spaces. For example, understanding "my love is a rose" involves blending spaces for "my love" and "rose," where identifying "love" and "rose" in the blend creates a correspondence between items in the input spaces, which is part of the "virtual" cross-space map.

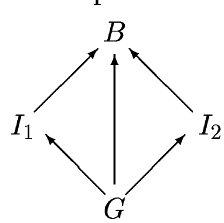

Figure 1: A Blend Diagram

A mathematical definition of blending is given in [20], based on a modification of the category theoretic notion of "pushout" [21] that takes advantage of an ordering relation on morphisms, with respect to their quality [18]. This notion of blending does not always give a

[2] The source UP is grounded in our experience of gravity, and the schema itself is grounded in everyday experiences, such as that when there is more beer in a glass, or more peanuts in a pile, the level goes up, and that this is a state we often prefer; therefore the image schema MORE IS UP, discussed in [15], is even more basic.

[3] This diagram is "upside down" from that used by Fauconnier and Turner, in that our arrows go up, with the generic G on the bottom, and the blend B on the top. This is due to a pervasive and natural duality between theories and models, in the sense that these terms are used in mathematical logic. Our convention is also consistent with the way that such diagrams are usually drawn in mathematics, as well as with the image schema MORE IS UP (since B is "more"). Also, Fauconnier and Turner do not include the map $G \to B$.

[4] However, [18] uses the term "**base space**", because it is more descriptive of how this space is used in applications to user interface design.

unique result. For example, four different blends of conceptual spaces for "house" and "boat" are houseboat, boathouse, amphibious RV, and boat for moving houses; there are also 44 other, less obvious blends [22].

Before introducing algebraic semiotics and structural blending, it is good to be clear about their philosophical orientation. The reason for taking special care with this is that, in Western culture, mathematical formalisms are often given a status beyond what they deserve. For example, Euclid wrote, "The laws of nature are but the mathematical thoughts of God." Similarly, the "situations" in the situation semantics of Barwise and Perry, which resemble conceptual spaces (but are more sophisticated – perhaps *too* sophisticated), are considered to be actually existing, real entities [23], even though they may include what are normally considered judgements[5]. The classical semiotics of Charles Sanders Peirce [24] also tends towards a Platonist view of signs. The viewpoint of this paper is that all formalisms are constructed in the course of some task, such as scientific study or engineering design, for the heuristic purpose of facilitating consideration of certain issues in that task. Under this view, all theories are situated social entities, mathematical theories no less than others; of course, this does not mean that they are not useful.

Algebraic semiotics, originally developed as a foundation for user interface design, attempts to overcome limitations of classical semiotics and blending theory, by addressing dynamic signs, social issues such as arise in collaboration, and the systematic mapping of signs in one system to signs in another. Details omitted here can be found in [18, 25]. A **semiotic system** or **semiotic theory** consists of: a **signature**, which gives names for sorts[6], subsorts, and operations; some **axioms**; a **level ordering** on sorts having a maximum element called the **top** sort; a **priority ordering** on the constructors at each level, where constructors are operations that build new signs from given parts; and a priority ordering on axioms. Sorts classify the parts of signs, among which data sorts provide values for attributes of signs (such as color and size). Axioms are constraints on the possible signs of a system. Levels express the whole-part hierarchy of complex signs, whereas priorities express the relative importance of constructors and their arguments; social issues play an important role in determining these orderings. This approach has a rich mathematical foundation, e.g., [20], since a signature plus equational axioms is an algebraic theory, on which there is a large literature. Conceptual spaces correspond to the very special case of semiotic theories where there is only one sort, there are no operations except those representing atomic elements and relations, and axioms only assert that a relation holds of certain constants.

Representations are uniform mappings of signs in a source space to signs in a target space. Since we formalize sign systems as algebraic theories with additional structure, we should formalize **semiotic morphisms** as mappings of theories that preserve the additional structure; however, these mappings must be partial, because in general, not all of the sorts, constructors, etc. are preserved in real examples. For example, the semiotic morphism from the rose space to the blend space for the metaphor "My love is a rose" (most likely) omits fertilizer and insects, while (possibly) preserving at least one of perfume and thorns. In addition to the structure of algebraic theories, semiotic morphisms should also (partially) preserve the priorities and levels of the source space. The extent to which a morphism preserves the various features of semiotic theories is an important determinant of its quality [20]. The simple form of blend in Figure 1 applies just as well to semiotic spaces and semiotic morphisms, in which case B called a **structural blend**; blending also extends to multiple spaces and morphisms. In the UCSD Meaning and Computation Lab, Fox Harrell and I have used a blending algorithm to generate novel metaphors for use in poems; see [22] for details. It seems likely

[5]The "types" of situation theory are even further removed from concrete reality.

[6]The word "sort" is used to avoid the ambiguities of the heavily overloaded word "type."

that this theory and its technology can be applied to computer ontologies, although the criteria used to optimize the quality of blends will likely differ from those respectively suggested for prosaic and for creative poetic metaphors, in [17] and in [22]. It also seems likely that structural blending [26] will be more useful than conceptual blending, because it can be used to build structured situations, not unlike those of [23], except of course for its very different underlying philosophial perspective.

4 Ethnomethodology

Traditional social science methods stand outside the situation being studied, applying methods different from those by which group members make sense of their world. In part, this reflects a misunderstanding of research in the hard sciences, since quantum measurements necessarily disturb the system measured, and modern philosophy of science claims that all measurements are "theory laden" [27]. Ethnomethodology argues that social scientists should use the same sense-making methods as group members [28], and denies that analysts have a unique access to objectivity. For example, if you study Balinese music by transcribing onto Western music paper, using the modern Western 12 tone equal tempered scale, you may conclude that Balinese micro-tonal scales are flawed and "primitive." But in fact, Balinese musicians are highly accomplished; they have their own methods for teaching their music, and their own musical theory, according to which their scales, rhythms, and structures are correct; they do not orient to the twelfth root of two.

Ethnomethodology [28] and its outgrowth of conversation analysis [29] consider that social order is accomplished by members in their moment by moment interactions. For example, although the word "seminar" suggests a pre-existing category, it is in fact constructed by members' use of a room with a certain arrangement of chairs, in their orientation towards someone understood to be the speaker, in their allotment of a very long turn to the speaker, etc. The idea of **member's categories** is to find the categories that members themselves use to order their social world, rather than to impose an analyst's order on it. The fundamental idea is that the social world is already orderly, and this order is an on-going creation of the participants. Further, we don't know in advance what the relevant categories are, so we should not come to the data with a pre-given coding scheme. It is implicit in the notion of members' categories as organizing activity that analysts do not reconstruct intentions or mental processes, except in so far as these are evident to those involved in the activity. Analysts cannot simply construct subjects' mental models or intentions. Rather, it is necessary to demonstrate what participants are doing that allows other participants to infer their intentions. Thus, the activity of the analyst in postulating intentions is not different from that of the participants, and proceeds on the same evidence. An extended application to the photocopy industry is given in [30], and an excellent ethnomethodology of mathematics is given by Livingston [31].

5 Sociology of Science

A brilliant book [32] by Bowker and Star on classification systems demonstrates the intensely political and ethical nature of classification, as well as its malleability, evolution, and local interpretation. Examples considered include racial classification in South African apartheid, the International Classification of Diseases, and the Nursing Intervention Classification. Such systems hardly resemble the neat equivalence relations of pure mathematics and computer science. On the contrary, they are inherently ambiguous, and typically have anomalous cases (e.g., classified as "Other" or "N/A"); they are highly political, and they embody values; they require ongoing work to apply and to maintain, work which is often invisible, e.g., done by “backroom” committees. Some phenomena are highlighted and others are ignored, some people suffer and others exalt, e.g,, when boundaries shift and property tax rates change. In one infamous case, a South African jazz musician was reclassified five times, each with serious personal consequences. Some countries delayed recognizing the severity of AIDS for political reasons. Here is part of the summary from [32]:

> We have seen throughout this book that people (and the information systems that they build) routinely conflate formal and informal, prototypical and Aristotelian aspects of classification. There is no such thing as an unambiguous, uniform classification system. (Indeed, the deeper one goes into the spaces of classification expertise – for example, librarianship or botanical systematics – the more perfervid one finds the debates between rival classificatory schools.)

My research on data integration for ecologists has seen taxonomists arguing at length over what appear to outsiders as very small points.

6 Conclusions

Previous sections have surveyed results from a variety of contemporary fields that tend to support skepticism about the computer-based ontologies for the Semantic Web, database integration, etc. Although we do not question the applicability of this technology to many specific problems, we do suggest that potential problems should be carefully considered. In particular, we hope to promote a greater awareness of the situated, embodied, embedded, enacted nature of human concepts, and of some techniques that have been developed to deal with this, ranging from the extreme abstraction of category theory to the extreme concreteness of ethnomethodology. This final section attempts a more philosophical perspective, beyond assessing what is possible and what is not.

Traditional approaches to categorization tend to decontextualize experience. Attaching a formal label to a real entity necessarily omits an enormous amount of relevant information. For example, no score, nor even a spectral analysis, can capture all the nuances of an actual musical performance, which will include particular actions by particular musicians, musical instruments, listeners, and rooms. One approach to avoiding such problems is to reify the notion of context. But phenomenologists and ethnomethodologists emphasize that context is dynamically emergent from activity, rather than fixed, definable in advance, formally representable, or separable from activity. This implies that it is better to speak of *situated actions* [30] or *occasions of action*, rather than of *contextualized representations*, as for example do many researchers in ubiquituous computing, because neither situations nor their contexts are specifiable, representable, stable, or separable from their actual uses; however, even this is better than trying to find *decontextualized* representations, which can never truly exist, even though social activity often creates the appearance that they do. When such apparent stability occurs, it is of course greatly to the advantage of ontologists, system designers, sociologists, and others, including ordinary members of society, but it should not be taken for granted.

Paul Dourish [33] gives an insightful discussion of context in connection with current trends towards ubiquitous or “context aware” computing; the problem addressed by this field is how to use powerful new sensor technologies to make computational systems more responsive to their users’ physical and social settings, as those users move through and modify these settings. This has turned out to be unexpectedly difficult, and Dourish claims this is essentially for reasons like those described in the previous paragraph.

Roland Barthes [34] combined and extended the structuralist semiotic theories of Saussure [35], creating a powerful language for cultural and media studies, which as it evolved through various stages, has been called semiotics, semiology, structuralism, and finally post-structuralism at the hands of Derrida and others who introduced intertextuality, deconstruction, and other controversial concepts. Although Barthes was a literary theorist, Derrida views himself as a philosopher working, among other things, to update Heidegger, a philosopher who indeed had similar, perhaps even more radical, ideas, in carrying out his project to reveal the history of Being without reference to a self or ego, either personal or transcendental.

Heidegger's *The Onto-theo-logical Constitution of Metaphysics* [36] criticizes "ontotheology," which might be briefly described as taking Being as the universal ground of all beings, and thus as above all beings; Heidegger also claimed that ontotheology is the essence of modern technology. He believed that the West lost its heritage, reducing the original Greek notion of being as "unconcealment" (*aletheia*) to mere existence in the ontotheological sense, and in the process losing the mysterious immediacy and power of being, to the alienation of both experience and object by technological reduction. This goes beyond the doubts raised by the cognitive linguists, ethnomethodologists, sociologists of science, etc. about the possibility of static, precise, complete, eternally valid categorizations of what is, even in some limited domain. It asks us to consider if we are living the right way, if we find our lives meaningful, and if not, what is the cause. It even suggests that there may be ways to be happier, more fulfilled, and more balanced, by questioning the presuppositions of technology, and attending to our actual experience.

References

[1] Joseph Goguen, Guilian Wang, Young-Kwang Nam, and Kai Lin. Abstract schema morphisms and schema mapping generation. Technical report, Dept. Computer Science and Engineering, UCSD, 2004. Submitted for publication.

[2] Young-Kwang Nam, Joseph Goguen, and Guilian Wang. A metadata integration assistant generator for heterogeneous distributed databases. In Robert Meersman and Zahir Tari, editors, *Proc. Intl. Conf. on Ontologies, DataBases, and Applications of Semantics for Large Scale Information Systems*, volume 2519 of *Lecture Notes in Computer Science*, pages 1332–1344. Springer, 2002.

[3] Guilian Wang, Joseph Goguen, Young-Kwang Nam, and Kai Lin. Critical points for interactive schema matching. In Jeffrey Xu Yu, Xuemin Lin, Hongjun Lu, and YanChun Zhang, editors, *Advanced Web Technologies and Applications*, pages 654–664. Springer, 2004.

[4] Joseph Goguen and Rod Burstall. Institutions: Abstract model theory for specification and programming. *Journal of the Association for Computing Machinery*, 39(1):95–146, January 1992.

[5] Joseph Goguen and Grigore Roşu. Composition of hidden information modules over inclusive institutions. In *From Object-Orientation to Formal Methods: Essays in Honor of Johan-Ole Dahl*. Springer, to appear 2003.

[6] Răzvan Diaconescu. Grothendieck institutions. *Applied Categorical Structures*, 10:383–402, 2002.

[7] Suad Alagic and Philip Bernstein. A model theory for generic model management. In Giorgio Ghelli and Gösta Grahne, editors, *Proc. Database Programming Languages 2001*, pages 228–246. Springer, 2002.

[8] Joseph Goguen and Răzvan Diaconescu. Towards an algebraic semantics for the object paradigm. In Hartmut Ehrig and Fernando Orejas, editors, *Proceedings, Tenth Workshop on Abstract Data Types*, pages 1–29. Springer, 1994. Lecture Notes in Computer Science, Volume 785.

[9] Răzvan Diaconescu. Interpolation in Grothendieck institutions. *Theoretical Computer Science*, 311:439–461, 2004.

[10] Joseph Goguen. Data, schema and ontology integration. In *Proceedings, Workshop on Combinination of Logics*, pages 21–31. Center for Logic and Computation, Instituto Superior Tecnico, Lisbon, Portugal, 2004.

[11] Stevan Harnad. The symbol grounding problem. *Physica D*, 42:335–346, 1990.

[12] Hubert Dreyfus. *What Computers Still Can't Do*. MIT, 1992.

[13] Francisco Varela, Evan Thompson, and Eleanor Rosch. *The Embodied Mind*. MIT, 1991.

[14] Eleanor Rosch. On the internal structure of perceptual and semantic categories. In T.M. Moore, editor, *Cognitive Development and the Acquisition of Language*. Academic, 1973.

[15] George Lakoff. *Women, Fire and Other Dangerous Things: What categories reveal about the mind*. Chicago, 1987.

[16] George Lakoff and Mark Johnson. *Metaphors We Live By*. Chicago, 1980.

[17] Gilles Fauconnier and Mark Turner. *The Way We Think*. Basic, 2002.

[18] Joseph Goguen. An introduction to algebraic semiotics, with applications to user interface design. In Chrystopher Nehaniv, editor, *Computation for Metaphors, Analogy and Agents*, pages 242–291. Springer, 1999. Lecture Notes in Artificial Intelligence, Volume 1562.

[19] Joseph Goguen and Fox Harrell. Style as a choice of blending principles. In Shlomo Argamon, Shlomo Dubnov, and Julie Jupp, editors, *Style and Meaning in Language, Art Music and Design*, pages 49–56. AAAI Press, 2004.

[20] Joseph Goguen and Grant Malcolm. *Algebraic Semantics of Imperative Programs*. MIT, 1996.

[21] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 1971.

[22] Joseph Goguen and Fox Harrell. Foundations for active multimedia narrative: Semiotic spaces and structural blending, 2004. To appear in *Interaction Studies: Social Behaviour and Communication in Biological and Artificial Systems*.

[23] Jon Barwise and John Perry. *Situations and Attitudes*. MIT (Bradford), 1983.

[24] Charles Saunders Peirce. *Collected Papers*. Harvard, 1965. In 6 volumes; see especially Volume 2: Elements of Logic.

[25] Joseph Goguen and Fox Harrell. Information visualization and semiotic morphisms. In Grant Malcolm, editor, *Multidisciplinary Studies of Visual Representations and Interpretations*, pages 93–106. Elsevier, 2004. Proceedings of a workshop held in Liverpool, UK.

[26] Joseph Goguen. Semiotic morphisms, representations, and blending for interface design. In *Proceedings, AMAST Workshop on Algebraic Methods in Language Processing*, pages 1–15. AMAST Press, 2003. Conference held in Verona, Italy, 25–27 August, 2003.

[27] Imre Lakatos. *The Methodology of Scientific Research Programmes*. Cambridge, 1978.

[28] Harold Garfinkel. *Studies in Ethnomethodology*. Prentice-Hall, 1967.

[29] Harvey Sacks. *Lectures on Conversation*. Blackwell, 1992. Edited by Gail Jefferson.

[30] Lucy Suchman. *Plans and Situated Actions: The Problem of Human-machine Communication*. Cambridge, 1987.

[31] Eric Livingston. *The Ethnomethodology of Mathematics*. Routledge & Kegan Paul, 1987.

[32] Geoffrey Bowker and Susan Leigh Star. *Sorting Things Out*. MIT, 1999.

[33] Paul Dourish. What we talk about when we talk about context. *Personal and Ubiquitous Computing*, 2003. Online edition.

[34] Roland Barthes. *Elements of Semiology*. Hill and Wang, 1968. Trans. Annette Lavers and Colin Smith.

[35] Ferdinand de Saussure. *Course in General Linguistics*. Duckworth, 1976. Translated by Roy Harris.

[36] Martin Heidegger. *Identity and Difference*. Chicago, 1969. Translated and introduced by Joan Stambaugh.