Ontotheology, Ontology, and Society

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

Abstract

This paper surveys some work currently little known in computer science suggesting limitations and/or improvements to the state of the art for computational ontology. Ethnomethodology emphasizes the negotiable, situated, embodied, emergent character of classification, as of all human activity. Cognitive linguistics and cognitive semantics study categorization, conceptual spaces, metaphor and blending, and reach similar conclusions. Sociology of science studies the intensely political and ethical aspects of classification systems, as well as their maleability, evolution, and local interpretation. Post-structuralism considers intertextuality, deconstruction, etc. Heidegger criticizes *ontotheology* as the alienated notion of "being" that is the essence of modern technology. These large bodies of work each suggest skepticism about extreme claims for ontologies, despite their undoubted applicability for certain appropriate problems. The paper also suggests possible approaches for circumventing various difficulties, some based on category theory. Careful consideration of skeptical arguments, hyperbolic claims, technical advances, and logical foundations can yield a balanced assessment of what may be possible and desirable, versus what seems impossible and undesirable, and may also lead to more humility, better ethics, better theory, and greater humanity.

1 Introduction

One of the most pressing problems for the technological application of ontology[1] is to understand its limitations. One reason for concern is that organizations, managers, and even experienced engineers, often expect too much, and know too little, 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, often with an underlying economic motivation.

Typical modern approaches to computational ontology follow principles originally due to Aristotle, arranging concepts in a taxonomy, i.e., a hierarchy organized by inclusion, where each concept is defined by necessary and sufficient conditions, preferably in the form of a conjunction of atomic unary predicates,

---

[1] There are two distinct but related senses of this word: its technical computer science sense, often associated with the Semantic Web (e.g., [4]), and its technical philosophical sense; when confusion is possible, we call the first *computational ontology*, and the second *philosophical ontology*. It is important that these two should not be confused.

which are assumed to be self-evident. It is the purpose of this paper to explore some limitations of this approach, by examining its philosophical foundations, and by comparing it with contemporary empirical results on human concepts and their role in social interaction.

Computational ontologies have genuine promise when restricted to appropriate, well-understood domains, such as B2B transactions in a car manufacturer's supply chain. The well known semantic web vision of Berners-Lee, Hendler and Lassila in [4] goes far beyond this, calling for automated agent negotiations about medical services, involving delicate tradeoffs among important human values. It is difficult to see how current technologies could achieve the agent performance envisioned in such scenarios; moreover, there are ethical grounds for questioning the extent to which humans should be encouraged to delegate responsibility to such agents, should they ever come to exist.

Though many engineers are no doubt skeptical, philosophy can make significant contributions to understanding the nature and limitations of computer ontologies. From about the mid twentieth century, philosophy has developed a deeping critical understanding of many methodological presuppositions that are largely unquestioned among engineers and scientists even today[2]. Other disciplines that can contribute to a more complete picture of the difficulties with computational ontologies include sociology of science, post-structuralism, ethnomethodology, and cognitive linguistics. Brief expositions are given of these, in hope of convincing some readers of their relevance, and moving them to learn more, perhaps motivated by their own encounters with difficulties like those that we discuss.

It is in the nature of this 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 uncritical acceptance of various forms of philosophical reductionism and realism.

Acknowledgements: I thank Jenny Wang and Young-Kwang Nam for collaboration on schema matching, and Kai Lin and Vitaliy Zavesov for help with implementation. This paper is based on work partially supported by the National Science Foundation under Grant No. ITR 0225676, the Science Environment for Ecological Knowledge (SEEK) project, which is building tools to support distributed research on long term ecological problems.

2 Philosophy

In philosophy, *reductionism* refers to the reduction of complex phenomena to simpler already understood phenomena. This can be a brilliantly effective research strategy, but it can also be a dangerous path to wishful thinking, which tries to avoid confronting the inherent difficulties of complex situations. The general reduction of concepts to context independent unary predicates is, in my

---

[2] However, there has also been considerable development in traditional (Cartesian, reductionist, etc.) philosophy.

opinion, an excellent example of such unwarranted reductionism, even though it can be successful in certain concrete instances. One difficulty with such approaches is that they decontextualize experience. Attaching a formal predicate to a real entity necessarily omits an enormous amount of potentially relevant information. For example, a written note “D” in a musical score cannot capture all the nuances of an actual performance, which will include contributions from particular musicians, musical instruments, rooms, and listeners.

One approach to avoiding such problems is to reify the notion of context, i.e., to treat contexts as real autonomous entities, whereas by definition, a *context* is the presupposed (and partially unknown) background of a given present manifestation. Some further discussion of problems with reifying context is given in Section 5.

The reification of context is a good example of *philosophical realism*, a doctrine which (following Plato) asserts the actual existence of ideal entities. Its most typical examples are from mathematics, ranging from the natural numbers up to transfinite sets, limit ordinals, etc. The combination of reductionism and realism encourages a comforting belief that all currently troubling difficulties with computational ontologies, such as the complex dependency on context exhibited by many concepts, can be predicted and controlled, or at least contained, by a more rigorous application of techniques that are already well developed and understood. One goal of this paper is to undermine such optimism, which has resulted in massive losses of time, energy, and money, through failure to cancel projects with undeliverable goals[3].

Philosophical reductionism and realism are manifestations of a very general, pervasive and enduring aspect of Western culture, which Martin Heidegger [32] called *ontotheology*: the (often unconscious) attitudes and presuppositions that are used to justify over optimistic reifications and reductions, particularly in technological projects. Heidegger belongs to the tradition called *phenomenology*, or continental philosophy, which takes actual human experience as its foundation, rather than some posited (but in practice unreachable) universal principles. In a typically precise but difficult way, Heidegger characterizes ontotheology as taking Being (as such) as the universal ground of all beings (in opposition to appreciating the particularities of actual experiences of beings as and when they actually arise). Considering the context of a given situation to be a self-existing eternal object, is a typical error inspired by ontotheological thinking. Another, which is highlighted by Heidegger [32], is to take the identity of beings as unproblematic.

In addition to cultural biases towards technological optimism, one can identify some more specific obstacles to reductionist knowledge representation. Karl Popper [44] coined the term “tacit knowledge” for knowledge that we have, but cannot express verbally (let alone symbolically), usual examples are how to tie laces on shoes, and how to ride a bicycle, though the phenomenon is far more common than these may suggest, including e.g., how to handle customer com-

---

[3] Several notorious examples can be found in software engineering, such as the late 1990s project to upgrade the US air traffic control system, costing billions of dollars [15].

plaints, how to wink, how to flirt, and how to shop. Among the reasons why tacit knowledge is so common are that actual human cognition is embodied, enacted, and socially and physically situated: that is, we have bodies, which act on and in the world, always in specific situations that have particular physical characteristics and social contexts: our conscious thoughts and actions are never abstract, disembodied, or without motivation; for example, even when we are doing mathematics, we deal with particular pens and papers (or computer terminals), and a particular problem, having a particular history.

Moreover, much of what we need to know, e.g., about riding a bicycle, is already present in the bicycle, and therefore does not need to be represented and stored; other knowledge relevant to a given task may be present in the structure of our bodies, in interactions with other people, in other artifacts, etc. It is therefore misleading to say that we "have" tacit knowledge, since it may not be in our minds at all, but rather located elsewhere. Terms like distributed knowledge, distributed cognition [34], ecological perception [16], being-in-the-world, and ready-to-hand [31] are more accurate.

All these relatively recent trends, and others discussed later in this paper, signal a departure from the rigid logic oriented stream of Western thought that perhaps reached its zenith in the logical positivism of the Vienna Circle, which famously declared that any statement that could not be reduced to formal logic over empirically grounded predicates was inherently meaningless. Similar ideas came to dominate the early development of cognitive science, in a movement often called *cognitivism*, a now receding 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 took a much more narrow view, which considered cognition as computation, and thus neglected body, emotion, and society, and caused the representation of knowledge to emerge as a central problem.

In one classic form of cognitivism, often called "good old fashioned AI" (GOFAI), knowledge is represented in symbolic logic, an approach of which the logical positivists would no doubt have approved. Its conspicuous failure, e.g., in the Japanese Fifth Generation project, has inspired a number of biologically motivated alternatives, such as neural nets and artificial life, which do not, however, abandon computational modeling, nor do they solve the *symbol grounding problem*, stated (but not satisfactorily solved) by Stevan Harnad [30]: how can the symbolic representations used in a computational model come to refer to the real world? While the information processing models of cognitivism might be adequate to allow computers to play 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 [9, 49], and inspire unwarranted confidence in reductionist visions of computational ontology, as well as impoverished views of what it means to be human.

From this point of view, formal philosophical ontology, as practiced by Barry Smith [41] and others, is a step backwards, embracing extreme forms of realism and reductionism, and attempting to fulfill the goals of the Vienna Circle by other means. For example, their complex axioms for "formal mereologies" render many decision problems unsolvable that are efficiently implemented in simpler formalisms (such as OWL [42]), and moreover, they are poorly adapted to typical applications, such as B2B transactions. Heidegger would surely have classified this work as ontotheology[4].

3 Data, Schema and Ontology Integration

One problem with Heidegger's notion of ontotheology is that it seems to leave little room for creative and constructive uses of technology. Is it really enough, as Heidegger seems to imply, for us to change our orientation to Being? How can that help us answer concrete questions, such as whether computer ontologies are likely to be helpful for particular applications, and how to improve their utility when they do help? A partial answer is our view that model building and system building should *not* ignore the particularities of the beings who are involved with them. We can accept that models (such as ontologies) are abstractions that necessarily omit many aspects of real situations, without accepting that they must subvert values that are important to users. This is perhaps the essence of movements such as usability (also known as user-centered design) and value-sensitive design [12]; my own value-driven design (previously called value-centered design) [20] takes this further by integrating values with the entire iterative process of design and development. My view is that values help to determine the relation between Being and beings, and in particular, the corresponding notions of identity.

This section addresses some specific issues in data integration, which has emerged as a major challenge in the early twenty-first century. The rise of inexpensive storage media, data warehousing, sensor nets, and the web, have made available vast amounts of data. But it can be very difficult to find what you want and then combine it properly to get what you really need. Difficulties include highly variable structure and quality of data and meta-data. Science labs and businesses often store data in spreadsheets, or even just formatted files, with little or no documentation of structure 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 automatic schema integration were feasible, then some interesting problems could be solved largely automatically, as in the semantic web vision. But these assumptions are far from true, and format is only part of the difficulty, which essentially involves social and value issues[5]: Contemporary science is highly competitive, so that scientists often do not want to share data until the research that uses it is complete. They prefer local storage and analysis, and they develop local data formats and conventions for interpretation; often these are specific to a particular project, to facilitate processing data for that project, even though the formats and conventions would not be useful for more general classes of

---

[4] It is amusing to notice that Barry Smith's homepage contains several anti-Heidegger links.

[5] The rest of this paragraph is based on the author's ethnographic observations during the SEEK project.

data. Little need is felt for meta-data, because the data are well understood by the local community during the lifetime of the project. Values here include control, privacy, and rapid processing of data, to ensure rapid publication of valid results, in order to ensure priority of discovery.

One proposed solution is to connect items in e-documents to concepts in ontologies. However, this 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 does the notion of identity. The latter problem manifests in exceptions (e.g., for the predicate "human," consider bionic appendages, cyborgs, androids, etc.).

Heterogeneity occurs at every level: different data sets often have different formats (e.g., schemas), there are often multiple ontologies for the same domain, and these ontologies[6] are often expressed in different languages (i.e., different logics). 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. Heterogeneity involves more than just different formats, ontologies and logics. For example, users must often use detailed contextual knowledge about data to "clean" it, correcting for transcription errors, faulty sensors, etc., and reconciling different accuracies, different rates and sources of collection, etc.; these are again issues of identity that involve values.

The diversity of data formats, ontologies, ontology languages, and ontology logics will not go away; it is a natural consequence of the diversity of beings and their values. 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 e-documents, through schemas and ontologies, up to ontology logics. Such issues can be addressed at each level using mappings or *morphisms*, to translate structures at that level [21, 23]. But of course, this cannot include all the details needed for all possible contingencies.

Our laboratory at UCSD has designed and built a tool called SCIA for defining schema mappings to integrate and transform data and DTD or XML meta-data [21, 40, 50]; it is now being extended to relational and other kinds of schema and ontologies. Fully automatic mapping generation is infeasible, but this tool tries to minimize total user effort by identifying the *critical points* where user input can yield the largest reduction of future matching effort [50]. SCIA also lets users define functions and conditions in mappings, whereas other tools require doing this manually in a separate phase [40].

The theory behind SCIA and our approach to ontology integration uses algebra and category theory [21, 23]; in particular, Grothendieck constructions can be used to integrate heterogeneous structures at every level. *Institutions* [24] use category theory to axiomatize the notion of logical system based the

---

[6] A computational ontology is just a theory over a logic, i.e., a set of sentences in that logic.

relation of satisfaction of a sentence by a model over a signature[7]; we call the satisfaction of S by M over a signature Σ a *triad*, in honor of Peirce. Applying the Grothendieck construction at this level allows integration over multiple logics. Institutions have also been used to give semantics for powerful module systems [29], multi-logic specification languages [6], databases [1, 21], ontologies [23], and to generalize many results in classical model theory, such as Craig interpolation [7].

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, and not to be confused with any concrete entities that it might be used to represent.

4 Cognitive Science

This section reviews some research from cognitive science that can help to evaluate the potential of computational ontology. We begin with work of Eleanor Rosch beginning in the late 1960s, on a systematic experimental study of categorization [45], which overturned then prevalent ideas about their propositional nature. In brief, basic-level categories (like "bird") are determined by similarity to prototypes, which then may be expanded radially by analogies. This research, brilliantly summarized in [36], became the foundation for the *conceptual metaphor theory* (abbreviated CMT) of George Lakoff and others, which has greatly deepened our understanding of metaphor [37]. One result of this research is that many metaphors come in families, called *image schemas*, that share a common pattern, based on how humans live in the world. 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[8] and are called *basic image schemas*; they tend to yield the most persuasive metaphors. It would seem a very good idea for computational 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 humans as possible.

Gilles Fauconnier and Mark Turner [10] have studied *blending*, or *conceptual integration*, claiming it is a basic human cognitive operation, invisible and seemingly effortless, but fundamental and pervasive, appearing in the construction and understanding of metaphors, as well as in many other cognitive phenomena, including grammar and reasoning. Simple examples of blends are two word

---

[7] This triadic relation can be seen as formalizing Peirce's triadic theory of meaning [43], including his notion of "interpretant" in that of signature.

[8] 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 [36], is even more basic.

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 [10]; note that this abstraction necessarily omits the qualitative, experiential aspects of what is represented. Whereas conceptual spaces are constructed on-the-fly for particular purposes, *conceptual domains*, though structurally similar, are large relatively stable configurations of related concepts and relations; conceptual spaces are constructed by selecting items from conceptual domains. *Conceptual mappings* are partial functions from the item and relation instances of one conceptual space to those of another.

The simplest blends[9] 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[10]. A *blendoid* [18] 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 in general, relatively few are interesting. Therefore additional principles are needed to identify 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 [10] suggest a number of *optimality principles* for this purpose (see Chapter 16 of [10]), but these are too vague to be fully formalized, and mainly apply to "common sense" blends, but not to creative poetic blends [26]. Whereas CMT views metaphors as maps from one domain to another (where the target domain concerns what the metaphor is "about"), blending theory views metaphors as "cross-space mappings" that emerge from blending conceptual spaces. For example, understanding "my love is a rose" involves blending spaces for "my love" and "rose," where the identification of "love" and "rose" in the blend space creates a correspondence between items in the input spaces, which is part of the emergent cross-space map from I_1 to I_2.

$$\begin{tikzcd}
 & B & \\
I_1 \arrow[ur] & & I_2 \arrow[ul] \\
 & G \arrow[uu] \arrow[ul] \arrow[ur] &
\end{tikzcd}$$

Figure 1: A Blend Diagram

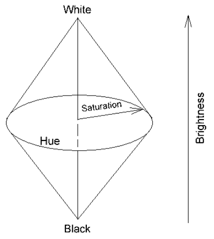

Figure 2: Human Color Space

A mathematical definition of conceptual blending is given in [28], and a general theory called "$\frac{3}{2}$-colimits" is also tentatively proposed; it enriches the

---

[9] 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. 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 $G \to B$.

[10] However, [18] uses the term *base space*, because it better describes of how this space is used in applications to user interface design.

category theoretic notion of colimit[11] with an ordering relation on morphisms; this relation can be used to reflect the quality of morphisms [18], and thus can represent certain values. Colimits apply beyond the simple situation of Figure 1, to diagrams with multiple spaces and morphisms. Unlike the usual category theoretic colimit, $\frac{3}{2}$-colimits are not always unique. For example, four different blends of conceptual spaces for "house" and "boat" are houseboat, boathouse, amphibious RV, and boat for moving houses; computations done with the Alloy blending algorithm developed by Fox Harrell and I show that there are also 44 other, less obvious blends, most of which are far from optimal [27].

Peter Gärdenfors has proposed an alternative to the conceptual spaces of Fauconnier, based on geometry rather than logic [13]. This raises the challenge of reconciling the two notions, which we temporarily call *G-spaces* and *F-spaces*, respectively. It can be done using triads (in the sense of Section 3) with F-spaces as sentences and G-spaces as models, though each must be enriched: F-spaces need to be given axioms and the ability to name structures such as the real numbers as in [22]; and G-spaces need to be provided with models of relations and functions as in [22]. Gärdenfors' most compelling example is the human color space, a 3D manifold shaped like a spindle (Figure 2). An F-space for colors would have names for some set of colors, and its corresponding G-space would have a sub-manifold[12] for each name [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" of 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 [3], even though they may include what are normally considered human judgements[13]. Modern interpretations of the classical semiotics of Charles Sanders Peirce [43] also tend 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 research 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 cannot be useful. This view is highly consistent with the pragmatism of Peirce.

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 a theory of representation, based on systematic mappings of signs in one system to signs in another [18, 25]. A *semiotic system* or *semiotic theory* consists of: a

---

[11] Colimits abstractly capture the notion of "putting together" objects to form larger objects, in a way that takes account of shared substructures; see [17] for an intuitive discussion of the mathematics, and [39, 11] for details.

[12] These should be convex, and perhaps fuzzy, as discussed in [22].

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

signature, which gives names for sorts[14], 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 reflect 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, since they are intended to reflect values. This approach has a rich mathematical foundation, since a signature plus equational axioms is an algebraic theory, on which there is a large literature, e.g., in algebraic semantics [28]. 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, axioms only assert that a relation holds of certain constants, and there are no levels or priorities.

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 scent or 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, thus giving a formalization of some of the optimality principles of [10], as further discussed in [27]. The simple form of blend in Figure 1 also applies nicely to semiotic spaces and semiotic morphisms, in which case it is called *structural blending*, since it can take account of structure, through constructor functions in the signature. Structural blending [19] is more useful than conceptual blending, because it can be used to build structured situations, not unlike those of [3], except for the very different underlying philosophical perspective. Structural blending is also subsumed under $\frac{3}{2}$-colimits [27].

The Alloy conceptual blending algorithm developed by Harrell and I has been used to generate novel metaphors in what we call *polypoems*, poems which are different each time they are read [27, 26]; structural blending is used to generate their syntax. We are currently extending this system to support interactive multimedia events. A key problem is how to implement more general optimality principles that better reflect artistic values; we believe geometric methods such as in nonlinear dynamical systems theory are applicable. It seems likely that the technology and theory we are developing for this will be useful for computational ontologies, by allowing more general and flexible optimality principles.

---

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

5 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 in addition, recent philosophy of science claims that all measurements are "theory laden" [35]. Ethnomethodology argues that social scientists should use the same sense-making methods as group members [14], and denies that analysts have a unique access to objectivity. For example, if you study Balinese music by transcribing it into Western musical notation, based on 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. Their music is sophisticated, complex, and beautiful.

Ethnomethodology [14] and its outgrowth conversation analysis [46] 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 categories. The basis for this is that the social world is already orderly, an on-going creation of the participants. Further, we cannot 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 should not reconstruct mental processes, e.g., intentions, except in so far as these are evident to those involved in the activity. Rather, analysts must determine what participants are doing that allows other participants to infer their intentions; this postulation of intentions does not differ from that of participants, and proceeds on the same evidence. An extended application to the photocopy industry is given by Lucy Suchman in [48], and an excellent ethnomethodology of mathematics is given by Eric Livingston [38].

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* [48] or *occasions of action*, rather than of *contextualized representations*, because neither situations nor their contexts are specifiable, representable, stable, or separable from their actual uses; however, even reified contexts are better than *decontextualized* representations, which never truly exist, even though social activity may create 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, and it cannot readily be exported to other contexts.

Paul Dourish [8] gives a very insightful discussion of context in connection with current trends towards ubiquitous "context aware" computing, which seeks 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 paragraphs.

6 Sociology of Science

A brilliant book [5] 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 in [5] include racial classification in South African apartheid, the International Classification of Diseases, and the Nursing Intervention Classification. Such systems hardly resemble the neat partial ordering relations of pure mathematics and computer science. On the contrary, they are inherently ambiguous, typically have anomalous cases (e.g., classified as "Other" or "N/A"), are highly political, and embody local 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, the race of a South African jazz musician was reclassified five times, each with serious personal consequences. Some countries have refused to accept the classification of AIDS as a serious communicable disease for political reasons. Here is a part of the summary from [5]:

> We have seen throughout this book that people (and the information systems that they build) routinely confuse 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 ethnographic research on ecologists confirms this, with observations of taxonomists arguing at enormous length over what appear to outsiders as very small points, but which in fact reflect deep underlying differences in values.

7 Post-Structuralism

Roland Barthes [2] extended the structuralist semiotic theories of Saussure [47], 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 Jacques Derrida and others who introduced intertextuality, deconstruction, and other controversial (and often misunderstood) concepts. Whereas structuralism focused on the analysis of single texts (understood in a very broad sense, to include possibly very large entities), *intertextuality* is the principle that no text can be understood except in the context of other prior and contextualizing texts. Deconstruction is a technique of literary analysis that seeks to undermine absolutes by demonstrating their dependence on context.

Although Barthes was a literary theorist, Derrida views himself as a philosopher working, among other things, to update Heidegger, who indeed had similar, even more radical, ideas, in carrying out his anti-ontotheological project of revealing the history of Being without reference to a self or ego, either personal or transcendental[15]. These ideas have been extremely influential in the humanities, and are generally consistent with the other areas described earlier; in particular, they are highly critical of reductionism and philosophical realism.

8 Conclusions

Previous sections have surveyed (with many simplifications) a variety of contemporary developments that support skepticism about the more ambitious visions for applications of computational ontologies to the Semantic Web, database integration, autonomous agents, etc. We do not question the potential of such technologies for many specific problems, but we do hope to promote a greater awareness of limitations that result from 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. Although it seems unlikely that any combination of such techniques can fulfill goals like those laid out for the Semantic Web in publications such as [4], no doubt new ideas will continue to emerge, such as tagsonomies, and it remains to be seen what is possible. An important particular conclusion is that we should carefully consider the extent to which agents can and should handle human values; the issues involved are both technical and ethical.

Heidegger claimed that ontotheology is the essence of modern technology [33]. He believed that the West has gradually lost its philosophical heritage, reducing the original Greek notion of being as "unconcealment" (*aletheia*) to mere existence, and in the process losing the immediacy and power of being, and alienating both experience and object through scientific and technological reduction. This position goes well beyond the doubts raised by cognitive linguists, ethnomethodologists, sociologists of science, deconstructionists, etc. about the possibility of static, precise, complete, eternally valid categorizations. It asks us to consider if we are living the right way, if we find our lives mean-

---

[15] Much of twentieth century philosophy can be seen as reacting to Heidegger's thought; for example, Habermas and Gadamer developed theories of communication in part to counteract Heidegger's tendencies towards pessimism and individualism.

ingful, and if not, what is the cause. It also suggests it may be possible to be happier, more fulfilled, and more balanced, by questioning the presuppositions of technology, and attending to the brilliance of our actual experience.

References

[1] Suad Alagic and Philip Bernstein. A model theory for generic schema management. In Giorgio Ghelli and Gosta Grahne, editors, *Proc. Database Programming Languages 2001*, pages 228–246. Springer, 2002. Lecture Notes in Computer Science, volume 2397.

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

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

[4] Tim Berners-Lee, James Hendler, and Ora Lassila. The semantic web. *Scientific American*, 284(5):34–43, May 2001.

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

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

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

[8] Paul Dourish. What we talk about when we talk about context. *Personal and Ubiquitous Computing*, 8(1):19–30, December 2003. Online edition.

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

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

[11] José Luiz Fiadeiro. *Categories for Software Engineering.* Springer, 2004.

[12] Batya Friedman, editor. *Human Values and the Design of Computer Technology.* Cambridge, 1997.

[13] Peter Gärdenfors. *Conceptual Spaces: The Geometry of Thought.* Bradford, 2000.

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

[15] W. Wyat Gibbs. Software's chronic crisis. *Scientific American*, pages 72–81, September 1994.

[16] James Gibson. *An Ecological Approach to Visual Perception.* Houghton Mifflin, 1979.

[17] Joseph Goguen. A categorical manifesto. *Mathematical Structures in Computer Science*, 1(1):49–67, March 1991.

[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. 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.

[20] Joseph Goguen. Groundlessness, compassion and ethics in management and design. In Richard Boland and Fred Callopy, editors, *Managing as Designing*, pages 129–136. Stanford, 2004.

[21] Joseph Goguen. Data, schema, ontology, and logic integration. *Logic Journal of the IGPL*, 13(6), 2006. Special issue on combining logics, edited by Walter Carnielli, Miguel Dionisio and Paulo Mateus.

[22] Joseph Goguen. Mathematical models of cognitive space and time. In Mitsu Okada and Daniel Andler, editors, *Reasoning and Cognition*. To appear, 2006.

[23] Joseph Goguen. Information integration in instutions. In Lawrence Moss, editor, *Memorial volume for Jon Barwise*. Indiana, to appear 2006.

[24] 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.

[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 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.

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

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

[29] Joseph Goguen and Grigore Roşu. Composing hidden information modules over inclusive institutions. In Olaf Owe, Soren Krogdahl, and Tom Lyche, editors, *From Object-Orientation to Formal Methods: Essays in Honor of Johan-Ole Dahl*, pages 96–123. Springer, 2004. Lecture Notes in Computer Science, Volume 2635.

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

[31] Martin Heidegger. *Being and Time*. Blackwell, 1962. Translated by John Macquarrie and Edward Robinson from *Sein und Zeit*, Niemeyer, 1927.

[32] Martin Heidegger. The onto-theo-logical constitution of metaphysics. In *Identity and Difference*, pages 42–74. Chicago, 1969. Translated and introduced by Joan Stambaugh.

[33] Martin Heidegger. The question concerning technology. In *Basic Writings*, pages 283–217. Harper and Row, 1977. Translated by David Krell; original from 1953.

[34] Edwin Hutchins. *Cognition in the Wild.* MIT, 1995.

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

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

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

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

[39] Saunders Mac Lane. *Categories for the Working Mathematician.* Springer, 1998. Second Edition.

[40] 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.

[41] Fabian Neuhaus, Pierre Grenon, and Barry Smith. A formal theory of substances, qualities, and universals. In Achille Varzi and Laure Vieu, editors, *Formal Ontology in Information Systems (FOIS'04)*, pages 49–59. IOS Press, 2004. Turin, Italy.

[42] Peter Patel-Schneider, Patrick Hayes, and Ian Horrocks. OWL web ontology language: Semantics and abstract syntax, 2002. `www.w3.org/TR/owl-semantics/`.

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

[44] Karl Popper. *Conjectures and Refutations.* Routledge, 2002. Classic reprint.

[45] 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.

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

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

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

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

[50] 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.