Frank, A. U. "The Rationality of Epistemology and the Rationality of Ontology." In *Rationality and Irrationality, Proceedings of the 23rd International Ludwig Wittgenstein Symposium, Kirchberg Am Wechsel, August 2000*, edited by Barry Smith and Berit Brogaard, 110-20. Vienna: Hölder-Pichler-Tempsky, 2001.

Andrew U. Frank
Dept. of Geoinformation
Technical University Vienna
frank@geoinfo.tuwien.ac.at

1 Introduction

Philosophers have proposed many different ontologies. Despite hundreds of years of effort, it has been impossible to reconcile the differences between them and to establish a single, widely accepted ontology. For practical purposes a consistent and comprehensive ontology is necessary: information systems which manage adequate descriptions of the world must be constructed on the basis of some ontology, even if this ontology is never explicitly described. This was not clear in the early years of information systems and many practical problems were discovered which could later be traced back to inappropriate ontological assumptions. The connection between information systems and ontology was at the foundation of the CYC project (Lenat, Guha *et al.* 1990) and has since gained substantial acceptance among theoretical and practical thinkers in information systems (Guarino 1998; Sowa 1998). The construction of re-usable ontologies (Frank 1997) has become an interesting, rapidly growing business and 'ontologist' is an acceptable job description in forward-looking IT companies.

The design of Geographic Information Systems, which cover information about objects and properties in the world with respect to their location (Longley, Goodchild *et al.* 1999) involves ontologies too. Indeed, such systems are ontologically more demanding than ordinary administrative information systems. They span a much larger diversity of kinds of things: from the description of the elevation of the surface of the earth with a regular grid of points to the description of the natural land cover (woods, fields, etc.) and morphology (mountains, valleys, etc.). They also include man-made features like roads and buildings as well as artificial boundaries between a range of different sorts of political and administrative units (Smith 1995), etc. There is no ready-made single ontology to cover all of these most diverse aspects of reality. Therefore we propose here the construction of an ontology consisting of several coordinated tiers.

An ontology constructed from tiers can integrate different ontological approaches in a unified system. In particular, it can merge a plenum, continuous space ontology with Aristotle's 'natural kind' ontology of objects. We can also integrate the ontology of 'social reality' described by Searle (1995). It seems possible also to overcome some of the differences between competing proposals, differences which we can understand as motivated by the examples the authors have in mind. From our practical experience, we have learned that a single ontology, which applies to all situations and the most diverse kinds of phenomena in the world or in our imagination, is not achievable. Therefore we propose here an orderly integration of otherwise contradictory proposals.

I am not interested here in terminological discussions, and I use terms like 'ontology' in a generic way; Guarino (1997) has shown the many different uses of the term by different authors and I do not want to add to this list. My approach is empirical and stresses our daily experience in interacting with the world as a source of knowledge to build ontologies. The goal is a computational model of an ontology, which can be used for the construction of information systems.

The remainder of this paper first gives an overview of the tiers and then discusses each of them in turn. It sketches how a computational model of ontology could be built and draws some conclusions about its usefulness.

2 The Five Tiers of the Ontology

An ontology for an information system is necessarily based on a realist position. Therefore tier 0 of the ontology assumes that there exists a physical reality, which may best be imagined as a four-dimensional continuous field of attribute values. This could be looked on as the ontology proper, where the next tiers are perhaps more similar to what some authors would assign to the realm of epistemology. Tier 1 covers the point-wise observation of this reality by cognitive agents. Tier 2 discusses how agents form objects from point-wise observations; this is somewhat similar to Aristotle's metaphysics. Tier 3 embraces social reality in the sense of Searle (1995) and other similar socially constructed elements (Berger and Luckmann 1996). Tier 4, finally, deals with the ideas cognitive agents have about the world.

Tier O: human -independent reality
Tier E1: observation of physical world
Tier E2: objects with properties
Tier E3: social reality
Tier E4: subjective knowledge

Fig. 1. The five tiers of ontology

The discussion here excludes the effects of learning on the ontology; it describes what is true when we consider a short period (days, weeks) and excludes the changes which are possible through extended experiences in an environment.

3 Physical Reality Seen as an Ontology of a Four-Dimensional Field

The physical laws which describe the behavior of the macroscopic world can be expressed as differential equations, which describe the interaction of a number of properties in space – the whole seen as forming a continuum. For each point in space and time a number of properties can be observed: color, the forces acting at that point, the material and its properties (like mass, melting temperature at the point, and so on). Movement of objects can be described as changes in these properties; even the movement of solid objects can be described, the cohesive forces in the body maintaining its shape. The description of reality via differential equations (e.g., the description of forces in a plate under a load) is widely used in mechanical and civil engineering, geology, etc. This view is also quite natural for most 'global systems' studies (Mounsey and Tomlinson 1988).

A field can be observed at every point in space and time for different properties:

$$f(x, y, z, t) = a.$$

Abstracting from the temporal effects, a snapshot of the world can be described by the formula which Goodchild called 'geographic reality' (Goodchild 1992).

$$f(x,y,z) = a$$

The processes occurring in this physical reality have spatial and temporal extensions: some are purely local and happen very fast; others are very slow and affect very large regions. The processes of objects moving on the tabletop are fast (m/sec) and the spatial extent is small (m); movement of persons in cities is again fast (m/sec) and the movements of the buildings very slow (mm/annum); geological processes are very slow (mm/annum) and affect large areas (1000 km^2). One can thus associate different processes with different frequencies in space and time (Fraser 1981). Each science has a certain scope: it is concerned with processes in a specific spectrum of space and time which interact strongly; other processes, not included in this scope, appear then to be either so slow or so fast that they can be considered constant.

Space and time form together a four-dimensional space in which other properties are organized. Giving space and time a special treatment results in simpler formulations of the physical laws that are of particular interest to humans. For example, the mechanics of solid bodies, e.g., the movement of objects on a tabletop, is explainable by Newtonian mechanical laws, which relate phenomena which are easily observable for humans in a simple form ($s = v\ t$, etc.). Other sciences, for example, astrophysics, prefer other coordinate systems in which mass is included.

However, the assumption that the formula $a = f(x,y,z,t)$ describes a regular function in the sense of a function which yields only one single value is equivalent to the assumption that there is only one single space-time world and excludes 'parallel universes' as parts of reality.

4 Observation of Physical Reality

Agents can - with their senses or with technical instruments - observe the physical reality at the current time, the 'now'. Results of observations are measurement values on some measurement scale (Stevens 1946), which may be quantitative or qualitative.

Observation with a technical measurement system such as remote sensing comes very close to an objective, human-independent observation of reality. A subset of the phenomena in reality is objectively observed. Many technical systems allow the synchronous observations of an extent of space at the same time, for example, remote sensing of geographic space from satellite. A regular grid is used and the properties observed are energy reflected in some bands of wavelength (typically the visible spectrum plus some part of infrared).

Observation through sampling of many points is effected also by our eyes, but it is also used by robots, where TV cameras which sample the field in a regular grid are used to construct 'vision' systems to guide the robot's actions in manipulating objects or guiding the robot's movements through buildings (Kuipers 1998).

Observations of reality are always marked by imprecision – the knowledge we acquire is never perfect. The technical effects of our measurement systems allow us at best measurements up to $^{10-13}$, which is, incidentally, much worse than the theoretical limits imposed by the Heisenberg uncertainty principle.

5 Objects with Properties

Our cognitive system is so effective because, from the array of sensed values, it forms individuals, which are usually called objects, and it reasons about them. Thinking of tables and books and people is much more effective than seeing the world as consisting of data values for sets of cells, regularly subdivided across a grid (i.e., three-dimensional cells, often called voxels). It is economical to store properties of objects and not deal with individual raster cells. As John McCarthy and Patrick Hayes have pointed out:

> ...suppose a pair of Martians observe the situation in a room. One Martian analyzes it as a collection of interacting people as we do, but the second Martian groups all the heads together into one subautomaton and all the bodies into another. ...How is the first Martian to convince the second that his representation is to be preferred? ...he would argue that the interaction between the head and the body of the same person is closer than the interaction between the different heads. ...when the meeting is over, the heads will stop interacting with each other but will continue to interact with their respective bodies. (McCarthy and Hayes 1969, p. 33)

Our experience in interacting with the world has taught us that the most appropriate subdivision of continuous reality is that into individuals. The latter are most often continuous in space and endure in time. Instead of reasoning with arrays of connected cells, as is done, for example, in computer simulations of strain analysis or oil spill movements, we select the more economical and more direct mode of reasoning with individuals: The array on the tabletop is divided into objects at the boundaries where cohesion between cells is low; a spoon consists of all the material which moves with the object when I pick it up and move it to a different location. This is obviously more effective than individual efforts to reason about the content of each cell. In an ever changing world, objects are typically formed in such a way that many of their properties remain invariant over time, which further simplifies reasoning. Animals and most plants form individuals in a natural way.

The cognitive system operates very quickly in identifying objects with respect to typical interactions. We see things as chairs or cups if they are presented in situations where sitting or drinking are of potential interest. Under other circumstances, the same physical objects may be seen as a box and a vase. The detection of 'affordances' of objects is immediate and not a product of conscious reasoning. The identification of affordances implies a breakup of the world into objects: the objects are what we can interact with (Gibson 1979).

Cognitive science has demonstrated that small infants as early as three months have a tendency to group what they observe in terms of objects and to reason in terms of objects. It has been shown that animals do the same. Most of the efforts of our cognitive system to structure the world into objects are unconscious and so it is not possible for us to scrutinize them. There are a number of well-known effects where the same image is interpreted in different ways, for example, the well known Necker-cubes which can be seen as cube or a corner, but not both at once. But such examples are rare. The default process assigns objects univocally.

Efforts to explain the categorization of phenomena in terms of common nouns based on a fixed set of properties were initiated by Aristotle. These occasionally lead to contradictions. Dogs are often specified as 'can bark', 'have four legs', etc., but from such a set of attributes it does not follow that my neighbor's dog, which lost a leg in an accident, is no longer a dog. Modern linguistics and psychology assume generally that prototype effects make some exemplars better examples for a class than others. A robin is a better example for a bird than a penguin or an ostrich (Rosch 1973; Rosch 1978).

Linguistic analysis suggests that the ways objects are structured are closely related to operations one can perform with them, and empirical data support this (Jackendoff 1983; Fellbaum 1998).

Humans have a limited set of interactions with the environment – five senses to perceive it and operations like walking, picking up, etc. - and these operations are common to all humans. Therefore the object structure - at least at the level of direct interaction - is common to all humans and it provides the foundation on which to build the semantics of common terms (Lakoff 1988). In general, the way individual objects and object types are formed varies with the context, but is not arbitrary. This commonality in the basic experiences of all humans gives sufficient grounding for the semantics of everyday words.

6 Social Ontology

Human beings are social animals; language allows us to communicate and to achieve high levels of social organization and division of labor. These social institutions are stable, evolve slowly and are not strongly observer dependent. Conventionally fixed names for objects, but also much more complex arrangements which are partially modeled according to biological properties, for example, the kin system (Lévi-Strauss 1967), or property rights derived from physical possession, can be refined and elaborated to the complex legal system of today's society.

6.1 Names

The types of proper and common names used in our various natural languages are clearly the result of a social process: proper names are words used for individuals, which identify objects in ways which are different from predicates to select individuals based on unique sets of properties. Such socially agreed identifiers seem to be a property of the individual, because they exist outside of the observing agent. Pointing out that 'chien', 'Hund' and 'cane' are equally good words to describe what in English is called a dog, should make it clear that none of these names is more natural than any other. Examples for proper names and similar identifiers reach from names for persons and cities to license plates for cars; there are also short-lived names created, like 'my for', during a single dinner.

6.2 Institutions

Social systems construct rules for their internal organization (Berger and Luckmann 1996), for example, laws, rules of conduct and manners, ethics, etc. Such rules are not only procedural ("thou shalt not kill"), but often create new conceptual objects (e.g., marriage in contradistinction to cohabitation without social status), adult person (as a legal definition and not a biological criterion), and so on. Institutions are extremely important in our daily life and appear to us as real; who would deny the reality of companies, such as the Microsoft Corporation.

Much of what administration and therefore administrative databases deal with are facts of law - the classification of reality in terms of the categories of the law. The ontology of these objects is defined by the legal system and is only loosely related to the ontology of physical objects; for example, legal parcels behave in some ways similar to liquids: one can merge them, but it is not possible to recreate the exact same parcels again (without the agreement of the mortgage holders) (Medak 1999; Medak in press).

7 Ontology of Cognitive Agents

Cognitive agents - persons and organizations - have incomplete and partial knowledge of reality, but they use this knowledge to deduce other facts and they make decisions based on such deductions. Agents are aware of the limitations of the knowledge of other agents; social games, social interaction and business are to a very large degree based on the reciprocal limitations of knowledge. Game theory explores rules for behavior under conditions of incomplete knowledge (von Neumann and Morgenstern 1944; Davis 1983; Baird, Gertner *et al.* 1994).

The knowledge possessed by a person or an organization increases over time, but the knowledge lags necessarily behind the changes on the side of reality. Decisions are made based on this not quite up-to-date knowledge. Fairness dictates that the actions of agents are judged not with respect to perfect knowledge but rather with respect to the incomplete knowledge the agent had or should have had if he had shown due diligence. Sometimes the law protects persons who have no knowledge of certain facts. The popular saying is "Hindsight is 20/2" or "afterwards, everybody is wiser". A fundamental aspect of modern administration is the concept of an audit: administrative acts must be open to inspection so that it can be established whether they were performed according to the rules and regulations. Audits must be based on the knowledge available to the agent, not on the facts discovered later. For audits it must therefore be possible to reconstruct the knowledge which an agent, for example, in public administration, had at a certain time. This leads to the bi-temporal perspectives usually differentiated in a database: the time a fact becomes true in the world and the time the agent acquires knowledge of this fact (Snodgrass 1992).

8 Computational Model of a Tiered Ontology

The design of the tiered ontology is oriented towards the construction of a computational model. The demonstration of misunderstandings and terminological difficulties in various texts on ontology but also the observation of problems in practice with differences in the interpretation of terms have led us to investigate computational models which reduce our reliance on natural language terminology. Algebras define terms up to an isomorphism without regress to other, previously defined terms, – which is exactly what is necessary to define the behavior of objects in reality or their simulated behavior in an information system. Between reality and information system we should have as far as possible an isomorphism. The two realms – reality and information – are connected by the experience of the agent interacting with the world based on his knowledge.

Certain parts of the ontology have been translated into computational models in a multi-agent setting (Weiss 1999). Multi-agent systems, the way we use them, are systems in which we simulate agents, including their bodies and perceptual and cognitive systems, in a simulated reality. We have completed one such simulation in which one agent explores a simplified city and then draws a map, which is later used by another agent to navigate (Figure 2). We have also completed a simulation for social reality (Bittner in progress) wherein the meanings of terms like 'ownership' and 'land' are defined. Agents then follow the rules of real estate law in dealing with the simulation. It seems possible to construct a computational model of the complete five tiers of the ontology in this framework.

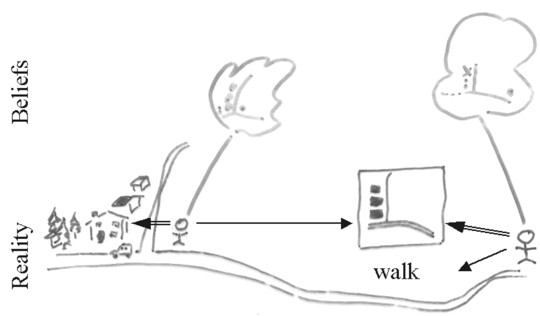

Fig.2. An agent producing a map and another agent using a map for navigation

We found it extremely useful to have a way of formally checking that the descriptions are complete, i.e., that all parts which are used to define a concept are in turn defined somewhere else in terms of a very simple set of primitives. Checking that the types of inputs and outputs correspond – something which can be done automatically – gives additional confidence that the model is logically consistent (Milner 1978; Jones 1994). Running the computational model allows us finally to test whether the model reflects correctly the intended behavior. We found the public domain functional language Haskell (Peterson, Hammond *et al.* 1996) extremely useful for this purpose.

9 Conclusion

In today's world of networked information systems, the clarification of the ontological bases used to collect and manage data becomes ever more important. Questions of interoperability (Goodchild, Egenhofer *et al.* 1998) are very often essentially ontological questions. In this environment practical ontologies - ontologies, which work - become necessary. They can help us to understand how to integrate data from different sources, and possibly in a single system. This topic will be further explored in the REVIGIS project (REVIGIS 2000).

We have sketched here a program of a tiered ontology, where different approaches are used on each tier. We follow an empirical approach, and integrate different ways of forming an ontology to achieve a practically useful solution. Our experiments so far suggest that computational models for ontology are possible. This would be a substantial step towards practically useful ontologies for information systems.

This text is a brief version of a much more detailed description which will be published in a book from the EU project Chorochronos, where it provides the ontology for the design of spatio-temporal databases (Sellis and Koubarakis to appear).

10 Acknowledgments

The work by Barry Smith and Nicola Guarino on ontologies has influenced my thinking in many details. I thank Werner Kuhn, Max Egenhofer, Annette von Wolff, Martin Raubal and Christine Rottenbacher for the time they have taken to discuss these ideas with me. Roswitha Markwart has improved the text of this contribution.

Support from the Chorochronos and REVIGIS project (both European Commission) and a project on formal ontology for land registration systems (financed by the Austrian Science Fund) are gratefully acknowledged.

References

Baird, D. G., R. H. Gertner, et al. (1994). Game Theory and the Law. Cambridge, Mass., Harvard University Press.

Berger, P. L. and T. Luckmann (1996). The Social Construction of Reality. New York, Doubleday.

Bittner, S. (in progress). An agent-based model of reality in a cadastre. Ph.D. thesis, Department of Geoinformation. Vienna, Technical University Vienna.

Davis, M. D. (1983). Game Theory. Minneola, NY, Dover Publications.

Fellbaum, C., Ed. (1998). WordNet: An Electronic Lexical Database. Language, Speech, and Communication. Cambridge, Mass., The MIT Press.

Frank, A. U. (1997). Spatial Ontology: A Geographical Information Point of View. Spatial and Temporal Reasoning O. Stock. Dordrecht, Kluwer: 135-153.

Fraser, J. T., Ed. (1981). The Voices of Time. Amherst, The University of Massachusetts Press.

Gibson, J. (1979). The ecological approach to visual perception. Hillsdale, NJ, Erlbaum.

Goodchild, M. F. (1992). "Geographical Data Modeling." Computers and Geosciences **18**(4): 401-408.

Goodchild, M. F., M. Egenhofer, et al., Eds. (1998). Interoperating Geographic Information Systems (Proceedings of Interop '97, Santa Barbara, CA). Norwell, MA, Kluwer.

Guarino, N. (1997). "Understanding, building, and using ontologies: A commentary to "Using Explicit Ontologies in KBS Development", by van Heijst, Schreiber, and Wielinga." International Journal of Human and Computer Studies **46**: 293-310.

Guarino, N. (1998). Formal Ontology and Information Systems. Formal Ontology in Information Systems (Proceedings of FOIS98, Trento, Italy, 6-8 June, 1998). N. Guarino. Amsterdam, IOS Press: 3-15.

Jackendoff, R. (1983). Semantics and Cognition. Cambridge, Mass., MIT Press.

Jones, M. P. (1994). Qualified Types: Theory and Practice, Cambridge University Press.

Kuipers, B. (1998). A hierarchy of qualitative representations for space. Spatial Cognition - An Interdisciplinary Approach to Representing and Processing Spatial Knowledge. C. Freksa, C. Habel and K. F. Wender. Berlin Heidelberg, Springer-Verlag. **1404**: 337-350.

Lakoff, G. (1988). Cognitive Semantics. Meaning and Mental Representations. U. Eco, M. Santambrogio and P. Violi. Bloomington, Indiana University Press: 119-154.

Lenat, D. G., R. V. Guha, et al. (1990). "CYC: Toward programs with common sense." Communications of the ACM **33**(8): 30-49.

Lévi-Strauss, C. (1967). Structural Anthropology, Basic Books.

Longley, P., M. Goodchild, et al., Eds. (1999). Geographical Information Systems - Volume 1: Principles and Technical Issues; Volume 2: Management Issues and Applications. New York, John Wiley & Sons.

McCarthy, J. and P. J. Hayes (1969). Some Philosophical Problems from the Standpoint of Artificial Intelligence. Machine Intelligence 4. B. Meltzer and D. Michie. Edinburgh, Edinburgh University Press: 463-502.

Medak, D. (1999). Lifestyles - A Paradigm for the Description of Spatiotemporal Databases. Ph.D. thesis, Department of Geoinformation. Vienna, Technical University Vienna.

Medak, D. (in press). Lifestyles. Life and Motion of Socio-Economic Units. A. U. Frank, J. Raper and J.-P. Cheylan. London, Taylor & Francis.

Milner, R. (1978). "A Theory of Type Polymorphism in Programming." Journal of Computer and System Sciences **17**: 348-375.

Mounsey, H. and R. F. Tomlinson, Eds. (1988). Building Databases for Global Science - Proceedings of the IGU Global Database Planning Project, Tylney Hall, Hampshire, UK, 9-13 May 1988. London, Taylor & Francis.

Neumann von, J. and O. Morgenstern (1944). Theory of Games and Economic Behavior. Princeton, NJ, Princeton University Press.

Peterson, J., K. Hammond, et al. (1996). Report on the functional programming language Haskell, Version 1.3. http://haskell.cs.yale.edu/haskell-report/haskell-report.html - Research Report YALEU/DCS/RR-1106, Yale University.

REVIGIS (2000). The REVIGIS Project Web Page. URL: http://www.cmi.univ-mrs.fr/REVIGIS/.

Rosch, E. (1973). "Natural categories." Cognitive Psychology **4**: 328-350.

Rosch, E. (1978). Principles of Categorization. Cognition and Categorization. E. Rosch and B. B. Lloyd. Hillsdale, NJ, Erlbaum.

Searle, J. R. (1995). The Construction of Social Reality. New York, The Free Press.
Sellis, T. and M. Koubarakis, Eds. (to appear). Spatio-Temporal Databases. Berlin Heidelberg, Springer-Verlag.
Smith, B. (1995). On drawing lines on a map. Spatial Information Theory - A Theoretical Basis for GIS (Int. Conference COSIT'95). A. U. Frank and W. Kuhn. Berlin Heidelberg, Springer-Verlag. **988** : 475-484.
Snodgrass, R. T. (1992). Temporal Databases. Theories and Methods of Spatio-Temporal Reasoning in Geographic Space (Int. Conference GIS - From Space to Territory, Pisa, Italy). A.U. Frank, I. Campari and U. Formentini. Berlin, Springer-Verlag. **639**: 22-64.
Sowa, J. F. (1998). Knowledge Representation: Logical, Philosophical , and Computational Foundations. Boston, PWS Publishing.
Stevens, S. S. (1946). "On the theory of scales of measurement." Science **103**(2684): 677-680.
Weiss, G. (1999). Multi-Agent Systems: A Modern Approach to Distributed Artificial Intelligence. Cambridge, Mass., The MIT Press.