

Against Idiosyncrasy in Ontology Development

Barry Smith¹

*Department of Philosophy and National Center for Ontological Research,
University at Buffalo, USA and Institute for Formal Ontology and Medical Information
Science, Saarbrücken, Germany*

Abstract. The world of ontology development is full of mysteries. Recently, ISO Standard 15926 (“Lifecycle Integration of Process Plant Data Including Oil and Gas Production Facilities”), a data model initially designed to support the integration and handover of large engineering artefacts, has been proposed by its principal custodian for general use as an upper level ontology. As we shall discover, ISO 15926 is, when examined in light of this proposal, marked by a series of quite astonishing defects, which may however provide general lessons for the developers of ontologies in the future.

Keywords. Upper-level ontology, data models for the oil and gas industry, ISO.

Forthcoming in B. Bennett and C. Fellbaum (Eds.), *Formal Ontology and Information Systems*, (FOIS 2006), Baltimore November 9-11 2006, in press.

What Happens When Data Models and Ontologies are Confused

Ontologies are, in one respect at least, comparable to telephone networks: they are designed to support exchange of information. The value of an ontology therefore depends, at least in part, on the quality of the network for shared communication which it provides, and on the number of users who agree to adopt this common network. This means that it depends also on the existence of a straightforward learning path for new users, and of clear and easily accessible documentation.

Before proposing an ontology for a given domain, accordingly, its custodians have a duty to maximize the likelihood that it will provide for the needs of maximally large numbers of potential users. This duty is all the more palpable where the ontology in question is advanced as an *upper-level ontology*, which is to say: an ontology that is designed for general adoption, as in the case of ISO Standard 15926 (“Lifecycle Integration of Process Plant Data Including Oil and Gas Production Facilities”), which is now being advanced as an upper-level ontological framework for ‘integrating diverse information systems’ and ‘integrating [and] analyzing mid-level ontologies’ without restriction.¹

¹ Corresponding author: Barry Smith, Department of Philosophy, 130 Park Hall, University at Buffalo, NY 14260, USA; E-mail: phismith@buffalo.edu.

I do not address here the question whether ISO 15926 is able to meet the specific data management needs of the community for which it was built. When examined in light of its potential use as an upper-level ontology, however, it is no less clear that ISO 15926 is marked by a series of defects, of a type which are, sadly, all too familiar in the ontology domain. Many of these defects flow from the terminological confusions which arise when the authors of an ontology do not take account of the fact that expressions such as ‘instance’, ‘entity’, ‘object’, ‘represent’, etc., are used in different ways by different (database, programmer, general user) communities. Others flow from the employment of philosophical and logical tools and theories which, although perhaps of some interest in their own right, are so counterintuitive from the perspective of the general users of ontologies as to constitute serious obstacles to learnability and accessibility. Yet others flow from simple use-mention confusions, where entities in reality are confused with their names or ‘representations’. Because these and related defects are still so common in ontology development work, I have used ISO 15926 as a source of examples of the characteristic ways in which ontology developers can go wrong. My goal is thus not one of mere criticism; rather, it is to draw out certain general principles which a good ontology should satisfy if it is to even reach the starting gate to be considered for adoption in the future, paying special attention to the role of ontologies in supporting the exchange of information. I have, surely, misunderstood many things in my attempts to come to grips with what I still see as the dark mysteries of ISO 15926; but even this provides evidence in favor of our first general principle:

1. **The principle of intelligibility:** an ontology that is advocated for general use should be understandable to those familiar with ontology development work who are willing to invest a reasonable amount of effort in mastering its documentation.

The major part of ISO 15926 is copyrighted by the International Organization for Standardization, from where it can be purchased, at a not inconsiderable sum, as a pdf file.² This brings us to a further general lesson, which we can formulate as follows:

2. **The principle of openness:** An ontology should be open and available to be used by all potential users without any constraint, other than (1) its origin must be acknowledged and (2) it should not to be altered and subsequently redistributed except under a new name.³ In addition the ontology should be (3) explained in ways which make its content intelligible to human beings, and (4) implemented in ways which make this content accessible to computers.

This principle implies not only that an ontology, if it is recommended for general use, should be in the public domain, but that the ontology should be marked by openness also in the wider sense that its features should be explained in clear, simple English, extended, where necessary, with technical terms.

In a domain like ontology, as is already clear for independent reasons, adoption by ISO does not guarantee that an artifact satisfies all the requirements which might reasonably be placed on an international standard.⁴ Indeed the attempt to enforce adoption of an ontology by taking the route of ISO standardization may bring costs: it makes it harder to correct errors; it often involves the making of less than ideal compromises, turning on the fact that adoption by ISO requires compatibility with prior ISO standards, many of which are – particularly in the informatics area – low in quality.

4. **The principle of re-using available resources:** if an ontology deals in a systematic way with entities or operators which are dealt with perfectly well already in some recognized resource used also by other ontology developers, then it should utilize this recognized resource.

We see another contravention of principle 4. in the treatment of terms like 'class_of_relationship_with_related_end_1', which is defined as follows (you will need to read this twice):

DEFINITION: A <class_of_relationship_with_related_end_1> is a <class_of_relationship> where a particular <thing> is related in the <class_of_relationship>, rather than the members of a <class>. The related <thing> plays the <role_and_domain> indicated by the class_of_end_1 attribute.

There is a perfectly good theory of relations, ranges, domains, ordered pairs, and of the transitivity, symmetry, etc. of relations, which is part of standard set theory. But because this resource was apparently ignored by the developers of ISO 15926, the result is gobbledygook, which no one (or at least: no one outside the oil and gas industry data-modeling community) would ordinarily feel the need to use, and definitions which no ordinary person would be in a position to understand.

Matters are made worse by the fact that some of the definitions are associated with terms with well-established meanings. That old terminological habits die hard is, unfortunately, a lesson still all too seldom taken account of in ontology development. It implies:

5. **The principle of terminological moderation:** Stay as close as possible to the terms already used by your intended audience and to their already established meanings. Use only terms for which either (1) there is a reasonable expectation that intended users of the ontology will have a need for them, or (2) such terms are required to fill gaps in the ontology in order to create a complete hierarchy.
6. **The principle of intelligible definitions:**⁷ Use definitions which are both (1) humanly intelligible (to avoid error in human use and maintenance) and (2) formally specifiable (as far as possible in such a way as to support one or other standard type of software).

A Rose is a Rose

The publicly accessible portions of the ISO 15926 documentation⁸ consist on the one hand of a list of terms together with definitions, and on the other hand of a set of diagrams. Neither the terms, nor the definitions, nor the diagrams are marked by a high degree of intelligibility.

Consider the sample term 'class_of_cause_of_beginning_of_class_of_individual', for which we are provided with the following:

DEFINITION: A <class_of_cause_of_beginning_of_class_of_individual> is a <class_of_relationship> that indicates that a member of a <class_of_activity> causes the beginning of a member of a <class_of_individual>.

Note the characteristic confusion of use and mention here. It is not, as the definition implies, the *class* which 'indicates', but rather (as common sense would suggest) the

corresponding *term*. (This problem is only made worse by the fact that it is not clear from its documentation whether ISO 15926 makes a distinction between a term and its referent.)

The term ‘class_of_cause_of_beginning_of_class_of_individual’ itself indicates further that we are to focus here on the *causes of beginnings of classes*. Yet the definition (in its strange, roundabout way) seems to be about the causes of the beginnings of *individuals* (it is about the members of the class of individuals). This is fortunate, because under the entry for ‘class’⁸ we are told that classes do not have beginnings, so that there could not literally *be* a ‘class_of_cause_of_beginning_of_class_of_individual’. By ‘beginning_of_class_of_individual’, therefore, ISO 15926 in fact means: *beginning of individual*. Its authors were accordingly not adhering to:

7. **The principle of terminological coherence:** for any expression ‘E’ in an ontology, ‘E’ means E.

From this it follows immediately that each expression in an ontology should have the same meaning on every occasion of use.

The requirement of univocity⁹ would normally, and for good reason, be regarded as a trivial constraint on the sensible use of language. Departures therefrom lead to a variety of familiar types of confusion and contribute much to the fact that (as will become all too painfully clear in what follows) the documentation of ISO 15296 will be unintelligible to almost all conceivable users of an upper-level ontology.¹⁰

One implication of the principle of terminological coherence is that an ontology should construct its complex terms in such a way that their constituent parts preserve their ordinary meanings. This principle is violated almost everywhere in the ISO 15926 documentation; thus for example the expression ‘individual’ is very often used (in order to save space?) to mean, not: *individual*, but rather: *possible individual*. The term ‘class_of_individual’ is defined as ‘a class whose members are instances of <possible_individual>’. The term ‘possible individual’ itself is defined, oddly, as meaning ‘thing that exists in space and time’.

Respect Compositionality

The most conspicuous puzzle raised by the treatment of many ‘class_of_X’ terms in ISO 15926 (as also of the many ‘class_of_class_of_X’ terms) turns on the very fact that these terms are included at all. For if classes or sets are needed, and if one needs to iterate the ‘class of’ (or ‘set of’) operator, then one will surely do this by means of some general facility, rather than by giving names in *ad hoc* fashion to just those 81 ‘class of’ or ‘class of class of’ terms one thinks one needs. (This is another application of the principle of re-use of available resources.)

In addition to ‘class_of_cause_of_beginning_of_class_of_individual’, ISO 15926 includes many other ‘class_of X’ entries for which the underlying ‘X’ term is itself, for whatever reason, missing from the ontology:

| | |
|------------------------------|-------------------------------|
| class_of_composite_material | class_of_molecule |
| class_of_compound | class_of_number |
| class_of_dimension_for_shape | class_of_organism |
| class_of_feature | class_of_organization |
| class_of_feature_whole_part | class_of_particulate_material |

| | |
|--|--|
| class_of_functional_object | class_of_person |
| class_of_inanimate_physical_object | class_of_property_space |
| class_of_indirect_connection | class_of_relationship_with_related_end_1 |
| class_of_individual | class_of_relationship_with_related_end_2 |
| class_of_information_object | class_of_relationship_with_signature |
| class_of_information_presentation | class_of_representation_translation |
| class_of_information_representation | class_of_scale_conversion |
| class_of_isomorphic_functional_mapping | class_of_sub_atomic_particle |

Thus while we have ‘class_of_organism’ and ‘class_of_person’ in the ontology, we do not also have ‘organism’ and ‘person’. Why not? Are there no *persons* in the world of the ISO 15926 ontology (which was developed by the oil and gas industries, we will remember, ‘to support the integration and handover of large engineering artefacts’)? More importantly still, is it appropriate to leave out ‘person’ and ‘organism’ in an *upper-level* ontology, when ‘stream’ and ‘representation of Gregorian date and UTC time’ are included?

These problems arise because the developers of ISO 15926 were not adhering to:

8. **The principle of compositional term construction:** if an ontology uses in a systematic way terms of the form ‘ $a \uparrow b$ ’ (where ‘ \uparrow ’ stands in for some term-binding operator like ‘of’ or ‘with’) then it should include also the corresponding a and b terms (or it should link to treatments of the latter in some other standard ontology).

The arguments for this principle are, I hope, clear. Not only does it contribute to intelligibility (users will more readily understand what ‘ a of b ’ or ‘ a with b ’ means if they are first of all provided with elucidations of the meanings of ‘ a ’ and ‘ b ’); it helps also to ensure completeness of the ontology (and in a way that also simplifies the business of error checking) – as contrasted with the mystifying randomness in term selection by which the ISO 15926 ontology is currently marked.

Exploit Recursion

In addition to the ‘class of’ terms in the ontology, we are also provided with an odd list of ‘class of class of’ terms:

| | |
|----------------------------------|---|
| class_of_class_of_composition | class_of_class_of_relationship |
| class_of_class_of_definition | class_of_class_of_relationship_with_signature |
| class_of_class_of_description | class_of_class_of_representation |
| class_of_class_of_identification | class_of_class_of_representation_translation |
| class_of_class_of_individual | class_of_class_of_responsibility_for_ |
| class_of_class_of_information_ | representation |
| representation | class_of_class_of_usage_of_representation |

Again, I could find no rationale for including just these items in the list rather than others. It is however worth noting that two of them, namely ‘class_of_class_of_composition’ and ‘class_of_class_of_representation’ have no corresponding ‘class of’ term in the ontology, though the first of these contains a reference to such a term in its definition:

DEFINITION: A <class_of_class_of_composition> is a <class_of_class_of_rela-

tionship> whose members are instances of <class_of_composition>. It indicates that a member of a member of the class_of_class_of_part is a part of a member of an instance of the class_of_class_of_whole,

which yields a nice gallimaufry of use-mention confusions in the provided

EXAMPLE: Toxicity description is a class_of_class_of_part of a material data sheet, where the description “has carcinogenic components” is a class_of_part on the Mogas Material Safety Data Sheet, and copy #5 of the Mogas Material Safety Data Sheet has “has carcinogenic components” as a part.

From this we learn that a description is a class (what, then, are the *members* of a description?); the rest of the example text departs too far from grammatical English to make sense.

Don't Confuse Types and Instances

It is a widespread problem with almost all contemporary work on ontologies and terminologies that inadequate attention is paid to the distinction between types (kinds, universals) and instances (individuals, particulars). Thus for example we find in the ANSI standard for controlled vocabularies¹¹ that the same relation of part to whole is asserted to obtain both between what are called ‘general concepts’, for example *brain* and *central nervous system*, and between what are called ‘specific instances’, for example *Toronto* and *Ontario*, thereby entrenching as part of an international standard what is in fact a well-documented confusion.^{12,13}

In the same confused vein, ISO 15926 defines ‘class of information object’ to mean: ‘a <class_of_arranged_individual> whose members are members of zero or more <class_of_information_representation> and of zero or more <class_of_information_presentation>’, informing us that ‘[u]sually, it is a physical_object (like a paper document) that is classified as a <class_of_information_object> ... Newspaper is a <class_of_information_object>.’

Why do we have ‘a paper document ... is classified as a <class_of_information_object>’ rather than the seemingly more sensible: a paper document is classified as an information_object’?

9. **The principle of types and instances:** An ontology should clearly mark whether given expressions are referring to types (universals, kinds, generals) or to instances (particulars, tokens, individuals).

What is meant by ‘Newspaper is a <class_of_information_object>’ is of course something like: *newspaper is_a information object*, or in other words: *the type newspaper is_a_subtype_of the type information_object* – something which can be said also, and more directly, and using English grammar, as follows: *a newspaper is an information object*.

Don't Confuse Mass Nouns and Count Nouns

ISO 15926 defines ‘class_of_compound’ to mean: ‘a <class_of_arranged_individual> whose members consist of arrangements of molecules of the same or different types,

bound together by intermolecular forces'. We are told that '[t]his includes both mixtures and alloys ... Water, sulphuric acid, sand, limestone, and steel can be represented by instances of <class_of_compound>.' What we are not told is whether it is some given *portion* of water or rather the corresponding substance-*type* which is an instance of this class.

If the former, should not the ontology, given its purpose, provide (or better: refer to, or link to) a serviceable theory of portions and masses of stuff (and indeed a link to some ontology of liquids¹⁴)? Instead ISO 15926 has developed its own theory of portions and masses, which are called 'batches', and which satisfy axioms like:

A Batch is a type of Material.
All Batches are Materials.
All Equipments are Materials.
Each Material must be either a Batch or an Equipment – but not both.

Avoid Circularity

Like all good top-level ontologies, our "Integration of life-cycle data for oil and gas production facilities" ontology contains its own tiny, hand-crafted ontology of mathematics, constructed out of terms such as 'class_of_number', which is defined as meaning: 'a <class_of_class> whose members are members of <arithmetic_number>'; and 'integer_number' for which we are provided with the helpful:

DEFINITION: An <integer_number> is an <arithmetic_number> that is an integer number.

The latter reminds us also of:

10. **The principle of non-circularity:** a good ontology should recognize the distinction between defined and primitive terms; it should avoid circular definitions; and, *a fortiorissimo*, it should avoid HL7-style nonsense-definitions of the forms: '*an a is the b of an a*', or: '*an a is an a which is b*'.¹⁰

Leaving aside certain very special contexts,¹⁵ circular definitions provide benefits neither to human beings nor to machines. They arise because ontology developers, who have not realized that not all terms in an ontology can be defined, are seeking a spurious completeness.

Don't Use Plural Nouns with Singular Verbs

ISO 15926 comes also with its own home-built geometry, as for example in:

DEFINITION: A <class_of_dimension_for_shape> is a <class_of_class_of_relationship> that indicates that members of the class_of_shape have a dimension that is a member of the class_of_dimension.

We are told in elucidation that 'Specifying that members of the "class of circle" have members of "class of diameter" is an instance of <class_of_dimension_for_shape>.' This is (I think) a roundabout way of saying: circles have diameters.

Note that ‘class of circle’ and ‘class of diameter’ are themselves not included in the ontology, and neither is the term ‘dimension_for_shape’. There is however a term ‘dimension_of_shape’, defined as ‘a <class_of_class_of_relationship> that indicates that members of the <shape_dimension> are dimensions of the <shape> members’. Life is made even harder by the fact that the example text provided for the above definition – ‘The sets of 10m lines that are diameters of 10m circles is an example of <dimension_of_shape>’ – conforms only loosely to the rules of English grammar. In particular, it reflects a departure from:

11. **The principle of singular nouns:** the terms of an ontology should be formulated in the singular, and the ontology’s documentation should pay careful attention to the distinction between singular and plural nouns and to the requirement of noun-verb agreement.

Combine Terms Coherently

The chaotic switching around in the use of ‘of’ and ‘for’ in the geometric corner of ISO 15926 reminds us of another general lesson:¹⁶

12. **The principle of coherence in the use of generic term-building operators:** If an ontology uses in a systematic way terms of the form ‘*a* † *b*’ (where ‘†’, again, stands in for ‘with’, ‘without’, ‘of’, etc.), then it should specify clearly the syntax of ‘†’, provide a statement of what expressions of the form ‘*a* † *b*’ mean in terms of the meanings of ‘*a*’ and ‘*b*’, and use each such form in the same way throughout.

An analogous principle applies of course also to unary operators such as ‘class of’. In several places ISO 15926 has pairs of terms ‘X’ and ‘class of X’, which are such that the definition of the former stands in no obvious relation to the definition the latter (in a way which would create serious obstacles, were the ontology ever to reach the point where it was required to support automatic reasoning). Thus for example we have, in addition to the pair ‘dimension_of_shape’ and ‘class_of_dimension_for_shape’, also the pair ‘shape_dimension’ and ‘class_of_shape_dimension’. The last two terms are defined, in seeming independence of each other, as follows:

DEFINITION: A <shape_dimension> is a <class_of_class_of_individual> that is a set of <individual_dimension> that define an aspect of a shape.

DEFINITION: A <class_of_shape_dimension> is a <class_of_class> that is a dimension of a <class_of_shape>.

Similarly we have the two terms ‘responsibility_for_representation’ and ‘class_of_responsibility_for_representation’, which are defined as follows:

DEFINITION: A <responsibility_for_representation> is a <relationship> that indicates that the controller <possible_individual> administers the controlled <representation_of_thing>.

DEFINITION: A <class_of_responsibility_for_representation> is a <class_of_relationship> whose members indicate that a <possible_individual> (usually an organization) deems that members of the pattern can be used as representations of the represented thing.

In each such case, in a properly constructed ontology, the ‘class of’ term would be introduced, not by means of its own special definition, but rather in the obvious recursive way, bringing (again) obvious benefits of formal coherence along the way.

Check Your Work for Errors

Note the use in the above definitions of problematic expressions such as ‘indicate’, ‘deems’, ‘usually’, and so forth, a pattern which is illustrated also in ISO 15926’s own miniature theory of mereology, which contains definitions like:

DEFINITION: A <feature_whole_part> is an <arrangement_of_individual> that indicates that the part is a non-separable, contiguous part of the whole.

The general lesson here is:

13. **The principle of non-subjective definitions:** when formulating definitions avoid the use of phrases like ‘which may ...’, ‘that indicates ...’, ‘... characterize ...’, ‘an aspect of ...’ which invite subjective interpretation.

For another example of the problem which this principle is designed to prevent, consider the ISO 15926 term ‘class_of_relationship_with_signature’, which is defined as: ‘a <class_of_relationship> that may have a <role_and_domain> specified for each end’. Is a <class_of_relationship> which does *not* have a <role_and_domain> specified for each ‘end’ also a class_of_relationship_with_signature? In every case? Only in some cases?

ISO 15926 comes also with its own miniature theory of physics:

DEFINITION: A <class_of_sub_atomic_particle> is a <class_of_arranged_individual> whose members are constituent particles of atoms.

EXAMPLE: Proton, electron, meson, neutron, positron, muon, quark, and neutrino can be represented by instances of <class_of_sub_atomic_particle>

whereby:

DEFINITION: An <arranged_individual> is a <possible_individual> that has parts that play distinct roles with respect to the whole. The qualities of an <arranged_individual> are distinct from the qualities of its parts.

What are the *parts* of a neutrino? What distinct *roles* do they play? What roles do quarks play in the integration and handover of large engineering artefacts?

DEFINITION: A <class_of_feature_whole_part> is a <class_of_arrangement_of_individual> whose members are instances of <feature_whole_part>.

EXAMPLE Thermowells have stems, and tables have tops are examples of <class_of_feature_whole_part>.

The two just-mentioned definitions tell us that the entities which serve as wholes in instances (?) of <feature_whole_part> should be non-separable; yet the *examples* include tables (wholes) and tops (parts), where surely many tops *are* separable. So what does ‘non-separable’ mean? And how does its use here relate to its use in the definition of ‘composite material’, where we are told that fibreglass and carbon fibre consist of

‘separable compounds’?

Don’t Confuse Definitions with Comments

14. **The principle of non-redundant definitions:** do not include clauses in definitions which contribute nothing to the application of the definition.

This principle is violated for example in:

DEFINITION: An <event> is a <possible_individual> with zero extent in time. An <event> is the temporal boundary of one or more <possible_individual>s, although there may be no knowledge of these <possible_individual>s.

DEFINITION: A <possible_individual> is: a <thing> that exists in space and time. This includes:

- things where any of the space-time dimensions are vanishingly small,
- those that are either all space for any time, or all time and any space,
- the entirety of all space-time,
- things that actually exist, or have existed,
- things that are fictional or conjectured and possibly exist in the past, present or future,
- temporal parts (states) of other individuals,
- things that have a specific position, but zero extent in one or more dimensions, such as points, lines, and surfaces.

In this context existence is based upon being imaginable within some consistent logic, including actual, hypothetical, planned, expected, or required individuals.

Question: are things which look like small flies from a distance actual or possible individuals?

DEFINITION: An <actual_individual> is a <possible_individual> that is a part of the space-time continuum that we inhabit. It exists in the present, past, or future of our universe, as opposed to some imagined universe.

Question: what is the difference between ‘being part of the space-time continuum that we inhabit’ (= being actual) and ‘existing in space and time’ (= being possible)? Why are fictional things included in the list of entities which exist in space and time? Is this because ‘space’ and ‘time’ themselves refer to possible space and possible time? If so, then are actual individuals themselves more properly to be conceived as entities which exist in possible, or in actual, space-time? Note how the confusions here stem from contravention of the principle of term coherence. If ‘A’ does not mean: A, but rather: *possible A*, then ‘*possible A*’ itself means something like: *possible possible A*, and so on, *ad exasperandum*.

Conclusion: ISO 15926 Is Not An Ontology

We can come closer to an understanding of ISO 15926 if we consider its treatment of qualities, such as length or temperature or color; or of roles, such as the status of

someone in an organisation. ISO 15926 does not recognize entities of these sorts. It deals with the color or length of an entity X, rather, by talking about X's relationships to strings or number-representations. My suspicion is that something similar applies to all the entries in ISO 15926. If this is so, of course, then we do not have here anything which could properly be described as an ontology. Rather, we have the equivalent of a coding scheme, rather like the Standard Algebraic Notation for Chess. The latter is, to be sure (unlike ISO 15926), elegant and efficient. But it is not an ontology of chess.

Acknowledgements

This work was funded in part by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant 1 U 54 HG004028. Thanks are due also to the Alexander von Humboldt and Volkswagen Foundations, to Bill Andersen, and, last but not least, to Matthew West.

References

- [1] M. West, "ISO 15926 – Integration of Lifecycle Data", Presentation at Upper Ontology Summit, National Institute of Standards and Technology (NIST), Gaithersberg, MD, March 15, 2006, http://ontolog.cim3.net/file/work/UpperOntologySummit/UO-Summit-Meeting_20050315/UOS--west_20060315.ppt (last visited April 23, 2006).
- [2] *ISO Store: Browse, Search and Purchase ISO Standards*: <http://www.iso.org/iso/en/> (last visited April 23, 2006).
- [3] <http://obofoundry.org/>. Last accessed July 1, 2006.
- [4] B. Smith, W. Ceusters and R. Temmerman, "Wüsteria", Medical Informatics Europe (MIE 2005), Geneva, *Stud Health Technol Inform.* 2005;116:647–652.
- [5] P. Aczel. *Non-well-founded sets*, CSLI Lecture Notes 14, Stanford: Center for the Study of Language and Information, 1988.
- [6] Common Logic Standard, <http://cl.tamu.edu> (last visited April 23, 2006).
- [7] J. Köhler, K. Munn, A. Rüegg, A. Skusa, B. Smith, "Quality Control for Terms and Definitions in Ontologies and Taxonomies", *BMC Bioinformatic*, s2006, 7:212.
- [8] ISO/FDIS 15926-2 - Lifecycle integration of process plant data including oil and gas production facilities: Data model: EXPRESS and EXPRESS-G listing: http://www.tc184-sc4.org/wg3ndocs/wg3n1328/lifecycle_integration_schema.html (last visited April 23, 2006).
- [9] B. Smith, J. Köhler and A. Kumar "On the application of formal principles to life science data: A case study in the Gene Ontology", *Data Integration in the Life Sciences* (DILS 2004), 79-94.
- [10] B. Smith and W. Ceusters, "HL7 RIM: An Incoherent Standard", *Medical Informatics Europe 2006*, in press.
- [11] ANSI/NISO Z39.19-2005, "Guidelines for the Construction, Format, and Management of Monolingual Controlled Vocabularies", <http://www.niso.org/standards/resources/Z39-19-2005.pdf>.
- [12] B. Smith, et al., "Relations in Biomedical Ontologies", *Genome Biology* (2005), 6 (5), R46.
- [13] M. Donnelly, T. Bittner and C. Rosse, "A Formal Theory for Spatial Representation and Reasoning in Biomedical Ontologies", *Artificial Intelligence in Medicine*, Vol. 36, Nr. 1, 2006, 1–27.
- [14] P. Hayes, "Naive Physics 1: Ontology for Liquids", in: J. Hobbs and R. Moore. (eds.) *Formal Theories of the Commonsense World*, Norwood, N.J.: Ablex Pubs, 1985, 71–107.
- [15] A. Gupta and N. Belnap, *The Revision Theory of Truth*, Cambridge MA: MIT Press, 1993.
- [16] C. J. Mungall, "Obol: Integrating Language and Meaning in Bio-Ontologies", *Comparative and Functional Genomics*, 2004, 5(7):509-520.