A Proposal for a Gene Functions Wiki

Robert Hoehndorf[1,2,3], Kay Prüfer[2], Michael Backhaus[1,2], Heinrich Herre[1,3],
Janet Kelso[2], Frank Loebe[1,3], and Johann Visagie[2]

[1] Research Group Ontologies in Medicine, Institute for Medical Informatics, Statistics and
Epidemiology, University of Leipzig
[2] Department of Evolutionary Genetics, Max-Planck-Institute for Evolutionary Anthropology
[3] Institute for Informatics, University of Leipzig

Abstract. Large knowledge bases integrating different domains can provide a
foundation for new applications in biology such as data mining or automated
reasoning. The traditional approach to the construction of such knowledge bases
is manual and therefore extremely time consuming. The ubiquity of the internet
now makes large-scale community collaboration for the construction of knowl-
edge bases, such as the successful online encyclopedia "Wikipedia", possible.

We propose an extension of this model to the collaborative annotation of mole-
cular data. We argue that a semantic wiki provides the functionality required for
this project since this can capitalize on the existing representations in biological
ontologies. We discuss the use of a different relationship model than the one pro-
vided by RDF and OWL to represent the semantic data. We argue that this leads
to a more intuitive and correct way to enter semantic content in the wiki. Further-
more, we show how formal ontologies could be used to increase the usability of
the software through type-checking and automatic reasoning.

1 Background

Recent technology developments have lead to the availability of genome sequence and
annotation data for a wide variety of species. More than twelve mammalian and thou-
sands of non-mammalian genomes have been sequenced and are publicly available.
Large volumes of biological data including sequences, structures, functions, pathways
and networks are now available. One of the major challenges in the field of bioinfor-
matics is to store and represent this data in a way which enables researchers to analyze
data integrated from diverse domains[1–3].

Understanding the relationship between phenotype (identifiable traits) and genotype
(heritable information), and the influence of environmental factors on both, remains a
major research area. The interaction between various factors resulting in the phenotype
is highly complex and requires a detailed understanding of multiple areas of biology.
Among these, gene function is considered central.

One approach to solve the problem of representing and structuring data about genes
and gene products is the Gene Ontology[4]. In its current release, the Gene Ontology
has more than 19000 concepts. Each concept names either a molecular function, or a
biological process to which a gene product is associated, or the location of the gene
product within the cell. The concepts are linked by two relations, part-of and is-a. Both

relations satisfy the transitivity property (the statement, that if some gene is associated with a category it is also associated with all the linked super-categories)[1]. The Gene Ontology is maintained by a team of curators and relies heavily on input from the community for its content and correctness. Linking gene products to the GO is performed independently for various databases and model organisms. While the GO provides information about genes that have certain functions, or which are known to be in some way associated with a function, process or cellular component, it does not provide a review of the research and discussions leading to this hypothesis – information which is crucial in the analysis of the function of a gene product and in discussing it in context with previous assumptions.

While the current biological databases and ontologies provide information about many features of a gene, the discussion of hypotheses regarding the biological function is missing. Wang [5] raised the question whether this problem could be solved by a wiki similar to wikipedia[6]. In his opinion the striking advantage of a wiki is the implicit community involvement. He concludes that "[A] wiki on gene function, which utilizes the collective brain power of biologists around the world, would be an invaluable tool for biological sciences."

2   The Application of Wikis and the Need for Structure

Wang's idea to use a wiki for the collaborative work on a comprehensive description of the function of genes is inspired by the success story of the online encyclopedia Wikipedia. A recent comparison of the traditional to the Wikipedia way of gathering information shows that this approach is competitive [7]. While these results are encouraging and show that the idea has potential, we see one major obstacle to this approach: the users of Wikipedia use the encyclopedia in a more or less traditional way to find information about one keyword (and relevant further articles by following links) while large scale analyzes in biology require the extraction of information regarding many items at the same time in a format suitable for computation. We propose the use of an extended semantic wiki as suitable for addressing these issues and providing data in biologically relevant formats. In this semantic wiki, instances or concepts (which are instances of some meta-category "concept") are treated as wiki pages and relationships between them are treated as hyperlinks between wiki pages.

Furthermore, since ontologies play a major role in the description of data about genes, the structure of existing ontologies must be compatible with the wiki. Therefore, a semantic wiki will allow for the description of information about transcripts and for the collaborative development of a biomedical knowledge base which is used for describing transcripts and other biological entities. In this paper, we will call this wiki the "gene function wiki".

2.1   Semantic Wikis

In order to enable the users of the gene function wiki to extract the contents in a machine readable format, the syntax of the standard wiki page has to be enhanced. Many wikis

---

[1] In the terminology of the Gene Ontology this is called the "True-Path-Rule".

which use a formal model to represent content have been developed in recent years. Most of these semantic wikis use Semantic Web technology (OWL and RDF) as their underlying representation formalism. For our purposes, semantic wikis can be divided into two main categories: ontology editors with features to support collaboration or transactions[8, 9], and wikis which are extended to allow for the semantic annotation of links or attributes[10–13]. The first often support more expressive constructs, such as OWL-DL, and are intended for users who have experience in the creation and use of ontologies and knowledge bases and are interested in the collaborative features, while the latter support RDF and rarely more expressive formalisms, and tend to be focused on users with a main interest in adding semantic context to the edited text.

The gene function wiki is intended for use by biologists with limited knowledge of formal logic, the Semantic Web or ontologies. It is therefore vital to keep the front-end intuitive, while representing the complexity of interactions between genes. Intuitive and commonsense ways for knowledge acquisition are of major importance if the application is to be widely adopted.

2.2   Requirements

We summarize here the requirements for building a gene function wiki. First, it must be possible to use and represent the structure of ontologies in the wiki in a way which can be queried rapidly. For example, a query for all the genes which are involved in apoptosis in neural crest cells relies heavily on the structure and semantics of the relations in the Gene Ontology[4] and the Celltype Ontology[14], and requires a structured representation of the information in our wiki describing gene functions.

A different type of formal information is necessary to answer questions about the exact way in which a gene is taking part in a biological process, i.e., whether the gene product just supports the function, or if it is an integral part of the chemical reaction. We need to be able to distinguish between a gene product which participates in a process from one which results from it. Ontologies such as the Ontology of Functions[15, 16] require n-ary relations, so queries for the exact role a gene product plays in a relation are relevant. We may want to restrict these queries even further, for example by requesting a specific author or evidence. All this requires a way to add structured information to the wiki and well-defined semantics for the relations and attributes used. Furthermore, access to the relations used in biomedical ontologies and their semantics are needed to represent and query the description of gene functions.

Because a major part of the knowledge which is developed is represented as text, the use of a simple frontend is essential to the application. We will therefore extend a semantic wiki to be adapted specifically to the problem at hand, since none of the existing prototypes satisfies our needs.

In addition to the use of existing biomedical ontologies in the description of the functions of genes and gene products, new concepts which are not yet part of any biomedical knowledge base may be required to describe the functionality of a gene, and it must be possible to add them to the wiki and interrelate them with the existing biomedical ontologies. Therefore, another application for the gene function wiki will be the use as a collaborative ontology and knowledge base curation system[17]. For this task, it is

crucial to provide intuitive ways to enter semantic content in the wiki, and to implement automatic checks for logical errors.

Finally, it must be possible to describe concepts as well as individuals in the gene function wiki. Concepts are used to categorize findings in an experiment, which is an individual. Also, annotating some experiment requires the description of individuals, while the conclusions are often abstractions and generalizations, therefore concepts.

3   Representation Language

Most semantic wikis allow only for the representation of binary relations, due to the restriction of the RDF format to binary relations. However, many relations in biomedicine, such as the annotation relation, can take more arguments. What is needed is an intuitive way to model n-ary relations in a semantic wiki. This can be done by keeping the original understanding of the semantic relations, as a typed link to another page in the wiki, but adding argument slots to the relation, which may be filled by arguments of an appropriate type (further discussed later on).

3.1   OWL and RDF

OWL and RDF are specifications for a metadata model maintained by the World Wide Web Consortium (W3C). RDF allows one to make statements about things in the form of subject – predicate – object triples, where the subject is the resource which is described, the relation represents a specific aspect of this resource, and the object is the value of the relationship.

OWL is an extension of RDF, and can be used to share and publish ontologies. OWL comes in three flavors, OWL-Lite, OWL-DL and OWL-Full. The expressive power of OWL-DL is equivalent to the description logic $SHOIN(D)$, and the expressive power of OWL-Lite to $SHIF(D)$ [18].

It is possible to reify statements made in RDF, and treat them as a new resource. This can be used to introduce n-ary relations. It is further possible to introduce a concept in OWL that takes n attributes as an n-ary relation in OWL. This can be used to export our data model, which is described later, to OWL.

3.2   Relations and Roles

K. Devlin [19] describes a model for relationships which is close to the everyday use of relations in, for example, natural language expressions, and which meets our requirements. Relations in [19] are specified by means of a *name* and named *argument roles*, which are slots in which objects of a specified *type* can be placed. It is possible to omit arguments in a use of the relation. However, a minimality condition is defined for each relation, defining which argument slots must be filled in order for a relation to be meaningful. For example, the relation *partOf* could be described as $< partOf|part, whole, context >$ or the relation *eats* as $< eats|eater, eatenObject, means, location, time >$ where *means* denotes the means used to eat (such as a knife). The statement "John eats an apple now" would be represented as $\ll eats, eater \rightsquigarrow John, eatenObject \rightsquigarrow apple,$

time $\rightsquigarrow t_{now}$ ≫. Note that the *means* and *location* argument roles are unfilled. A minimality condition here could state that either the *eater* or the *objectEaten* must be filled. These argument roles will also restrict the type of object which may fill the role, as we will discuss in section 4.

RDF can be integrated in this view: for some RDF triple *subject*, *predicate*, *object*, we can define the relation < *relation*, *subject*, *object* >. The relationship model we use is similar to UML associations[20] and topic maps[21, 22]. Our representation of the structure of relations is also close to the account on roles given in [23, 24]. We will show this similarity in the discussion of our data-model in section 5.

4   Ontological Type System

Part of the strategy which lead to the success of wikis is that they leave their users a maximum of liberty; there is no structure in a wiki except for the one provided by the users of the wiki. We, however, want to use a wiki for the creation of a structured knowledge base, in a domain in which rich representations of structures exist in the form of ontologies. In this section, we address the question how to add such a structure without limiting the ability a user has to edit information in the gene function wiki. Instead we wish to provide an easy reference to the structure which is available in the wiki itself, and to structures which have been developed outside the wiki and will be used to annotate content of the wiki, such as the Gene Ontology.

But first, let us collect some examples of what kind of structure we talk about. Biological processes in the Gene Ontology, for example, are related using two relations, *part-of* and *is-a*. Additionally, conceptualizations and formalizations of the most general entities in biology are developed[25, 26]. Some of these conceptualizations are new and still need to gain wide acceptance in the biological community[15], but others such as the need for the concepts of "function", "process" and "localization" as included in the Gene Ontology are accepted throughout the scientific biological community. At least the terminology and structure that the Gene Ontology provides must be usable for the description of the gene products.

We provide a structural layer of the wiki in the form of a biomedical core ontology. This core ontology gives natural language and formal definitions of the most general biomedical concepts, such as *biological process*, *biological function*, or *organism*. Additionally, the core ontology defines relations between these concepts, for example a relation *Realizes* between processes, functions, and objects. Furthermore, it defines the upper categories of all the biomedical domain-ontologies which are used in conjunction with the wiki.

Since a core ontology is a rigorous yet abstract formalization of the entities and relations of a domain, all of the (semantic) information in the wiki can be embedded in the core ontology. Making a general set of relations available leads to less redundancy in the definition of new relations. For example, the relation *is-a* could also be named *subclass-of*, *specialization-of*, or *subsumed-by*. By providing one relation in the core ontology, all these names will be derivates of this one relation, which has been formally defined.

Furthermore, the core ontology can be used as a type system for the relations and concepts in the gene function wiki as described in [27]. For example, a relation *has-function* could be specified as relating only biological functions, biological objects, and a context. With concepts in the gene function wiki that have types, two things can be done whenever some concepts are related using the *has-function* relation: (1) verifying whether the arguments of this relation are of the type which is specified by the *has-function* relation, and (2) automatically classifying the arguments of the relation as arguments of the appropriate type.

As an example of what can be done, assume that a *hasCellFunction* relation was defined in the following way, where the arguments are specified as (*role*, *type*) tuples, and *cell*, *bioFunction*, and *situation* are concepts defined in the core ontology:

$$< hasCellFunction, (bearer, cell), (function, bioFunction), (context, situation) >$$

A concept (wiki page) *A* which occurs in the *bearer* role of this relation is automatically classified as a cell, and inherits all properties of cells that are defined in the core ontology (such as the potential to be part of some tissue) and the ontologies which are embedded in the core ontology (such as the existence of a part which is a membrane as defined in the Gene Ontology). If the same concept *A* occurs in a different relation filling a role which is typed as *organism*, *A* is reclassified as *cell ⊓ organism*, a mono-cellular organism (e.g., a bacteria).

Also, if cells in the core ontology were defined as a superclass of things which have as part a nucleus, and the concept *A* in the wiki had a nucleus as part, then *A* can be classified as cell.

This type system can also be used for information retrieval: using the core ontology, it is possible to query for all the processes to which a gene is associated.

We use the core ontology GFO-Bio[26] which is based on the top-level ontology General Formal Ontology[28] (GFO). This is particular useful because the GFO includes a well-developed analysis of ontological categories such as universals, concepts and symbols. Symbols play an important role in the description of genes, DNA or RNA. Although the GFO is a top-level ontology which is formalized in first order logic, an OWL-DL version is available for conceptual modeling purposes. GFO-Bio, which is based on the GFO, is available in OWL-DL as well.

5   Data-Model

In addition to storage for the text in the wiki articles, we need an additional place to store the semantic data, similar to other semantic wikis[10].

A UML diagram of the data model is shown in figure 1. What can be seen is the distinction between concepts on the left hand side, and their instances – individuals – on the other. Relations can be defined as concepts by a relation name and a number of role–type pairs. The type of a role can be a disjunctive type, which is the reason for the $n : m$ relationship between roles and types. Roles are bound to one relation and cannot be reused.

On the other side are instances. Instances of relations – called relators in the General Formal Ontology (GFO) [28] – can be split into instances of roles, called
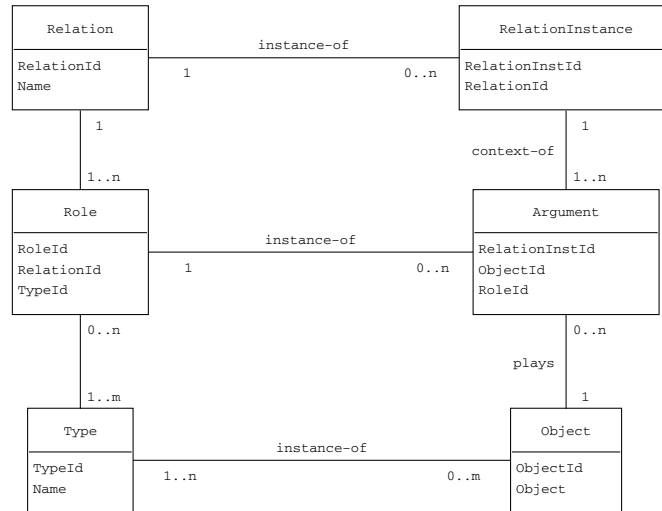
Fig. 1. UML Diagram of the data model

qua-individuals in [29], which are played by objects. In the terminology of [24, 23], roles – here represented in the argument table – are dependent on a context (the instance of the relation, or the relator) and a player (the object). These constraints are represented as restrictions on the cardinality of the relation arguments have to objects and relators.

6   Discussion and Future Research

There are still many problems which require a solution. Although there are some biomedical core ontologies already available, none of them has ever been applied as discussed in our proposal. It is to be expected that in addition to the concepts of the biomedical domain, concepts which are not part of this domain must be added. Types such as numbers, coordinates, strings, author names, et cetera are not a part of the current core ontologies for biomedicine, but will be required for application in the wiki.

Another question is how far the idea we discussed here could be generalized to other domains. It is tempting to use the same architecture for other domains such as chemistry, physics or even social sciences, by adding different core ontologies for other domains, and embed them into the same top-level ontology, GFO. However, it remains open, whether interesting parts of other domains can be formalized with the language we provide, or if richer formalisms are required. For example, the core ontology of mathematics is set theory. But formalizing interesting theorems requires a rich – and undecidable – language. In biology, useful information can be formalized in languages that can never express logical inconsistencies, e.g. in the Gene Ontology[4], and is therefore particularily suited for collaborative knowledge aquisition. A similar problem arises if the type system was reduced to a top-level ontology such as the GFO[28], in order to modify the

core ontology itself. The language in which the core ontology is formalized is at least description logic, and despite the fact that the language to formalize this knowledge would become highly complex, the problem of how to treat inconsistencies arises and ways to resolve these must be found.

Another issue to be solved before a gene function wiki could become successful is that of trust in the information in such a wiki. The advantage of the ontologies provided by a consortium is that domain and ontology experts collaborated in the creation of these ontologies which results in a high quality for them. A way to solve this in a wiki is to use a confidence or reputation system for the users of the wiki. However, the representation of this confidence in the knowledge model is problematic: it may result in fuzzy truth values for parts of the ontology[30], which would have to be defuzzicated for most applications. We also consider a web-of-trust approach[31] where users can trust particular users, e.g., ontology experts, and obtain a view on the knowledge in the wiki based on their choice of whom to trust. However, it remains open what the neutral perspective on the wiki[32] would be, as this would divide the wiki in a number of personal knowledge bases, different for each user depending on whom she trusts.

Finally, the integration of external databases and ontologies may require modifications on automated reasoners. Many ontologies which exist at present are formalized using a trivial knowledge representation language, often based on directed acyclic graphs. Reasoning on these structures is highly efficient, which is necessary as these ontologies tend to have a large number of concepts. Ontologies formalized in description logic, on the other hand, require more sophisticated reasoners which are much less efficient, while they usually have much fewer concepts. For efficiency, it would be beneficial to employ a hybrid reasoner which uses the most efficient reasoning algorithm for each part of the knowledge base. E.g., it performs a graph search on a directed acyclic graph and reuses the results from this query when performing queries which require the core ontology (which is formalized in description logic), while preserving the semantics and definitions which are given by the integration of the domain ontologies in the core ontology.

7   Conclusion

Let us revisit what we have discussed so far. First, we argued that a wiki can be enhanced by semantic relations, and that this addition is necessary for our application in order to search for genes and other biological entities, automatically check consistency, classify and group genes together, and for the integration with other knowledge bases. Second, we introduced our data model for storing the semantic content of the wiki, using n-ary relations.

Finally, a type system which is based on a formal core ontology for bio-medicine is beneficial. Because the information in the gene function wiki is highly structured due to the representation of semantic relations, it is necessary to provide the most general building blocks of the semantic content in the wiki. We will use the types provided by the core ontology as type system for the assertion of semantic relations. We can use a top-level ontology as foundation for the biomedical core ontology in order to allow the content of the gene function wiki to be used in a wider scientific context – such as chemistry.

In summary we have discussed requirements and some theoretical aspects of the implementation of a gene function wiki which we believe may provide new insights for biologists, as well as the Semantic Web and the wiki community. A work-in-progress, prototypical implementation, which, however, may not be (fully) functional, can be found for evaluation purposes on `http://onto.eva.mpg.de/bowiki`.

An implementation of a gene function wiki has the potential to provide a powerful tool for the annotation of gene data in biology. Additionally, the integration of formal ontologies and wikis may lead to new applications for wikis and ontologies in areas where their use has been rather limited until now. Further, using the framework introduced here for curation and maintainance of biomedical ontologies will enable the possibility to use information and ontology extraction methods from computer linguistics in order to create prototypical ontologies, or to generally make ontology curation and annotation faster and cost-efficient[33].

With the rapid growth of biological knowledge increasingly sophisticated methods are needed in order to close the gaps in storing, processing and representing this knowledge. Using a wiki for these purposes is far more than just an appealing idea. The needs of researchers force us to consider novel approaches such as data organization using ontologies, and data mining combined with reasoning over the given facts.

References

1. Chicurel, M.: Bioinformatics: bringing it all together. Nature **419**(6908) (2002) 751–755
2. Birney, E., et al.: Ensembl 2006. Nucleic Acids Res **34**(Database issue) (2006) 556–561
3. Wheeler, D.L., et al.: Database resources of the national center for biotechnology information. Nucleic Acids Res **34**(Database issue) (2006) 173–180
4. The Gene Ontology Consortium: Gene ontology: tool for the unification of biology. Nat Genet **25**(1) (2000) 25–29
5. Wang, K.: Gene-function wiki would let biologists pool worldwide resources. Nature **439**(7076) (2006) 534
6. The Internet Community: Wikipedia, the free encyclopedia. `http://wikipedia.org` (2006)
7. Giles, J.: Internet encyclopaedias go head to head. Nature **438** (2005) 900–901
8. Auer, S.: Powl – a web based platform for collaborative semantic web development. In: Proc. of 1st Workshop Workshop Scripting for the Semantic Web (SFSW'05). (2005)
9. Fischer, J., Gantner, Z., Rendle, S., Stritt, M., Schmidt-Thieme, L.: Semantic wiki COW. `http://www.informatik.uni-freiburg.de/cgnm/software/cow/index_en.html` (2006)
10. Völkel, M., Krötzsch, M., Vrandecic, D., Haller, H., Studer, R.: Semantic wikipedia. In: Proceedings of the 15th international conference on World Wide Web. (2006) 585–594
11. Harris, D., Harris, N.: Kendra base. `http://base.kendra.org.uk/` (2006)
12. Dello, K., Tolksdorf, R., Paslaru, E.: Makna. `http://www.apps.ag-nbi.de/makna/` (2006)
13. Campanini, S.E., Castagna, P., Tazzoli, R.: Platypus wiki: a semantic wiki wiki web. In: Semantic Web Applications and Perspectives, Proceedings of 1st Italian Semantic Web Workshop. (2004)
14. Bard, J., Rhee, S.Y., Ashburner, M.: An ontology for cell types. Genome Biol **6**(2) (2005) R21

15. Burek, P., Hoehndorf, R., Loebe, F., Visagie, J., Herre, H., Kelso, J.: A top-level ontology of functions and its application in the Open Biomedical Ontologies. Bioinformatics **22**(14) (2006) e66–e73

16. Burek, P.: Ontology of Functions. PhD thesis, Institute of Informatics (IfI), University of Leipzig (2006) forthcoming.

17. Hoehndorf, R., Prüfer, K., Backhaus, M., Visagie, J., Kelso, J.: The design of a wiki-based curation system for the ontology of functions. In: The Joint BioLINK and 9th Bio-Ontologies Meeting. (2006)

18. Horrocks, I., Patel-Schneider, P.F.: Reducing OWL entailment to description logic satisfyability. In: Proc. ISWC 2003. Number 2870 in LNCS. Springer (2003) 17–29

19. Devlin, K.: Logic and Information. Cambridge University Press (1991)

20. Group, O.M.: UML 2.0 infrastructure specification. Document ptc/03-09-15 (2004)

21. Pepper, S., Moore, G.: XML topic maps (XTM) 1.0. `http://topicmaps.org/xtm/1.0/` (2001)

22. Garshol, L.M., Moore, G.: Topic maps – XML syntax. `http://www.isotopicmaps.org/sam/sam-xtm/` (2006)

23. Loebe, F.: An analysis of roles: Towards ontology-based modelling. Master's thesis, Institute of Informatics (IfI), University of Leipzig (2003)

24. Loebe, F.: Abstract vs. social roles: A refined top-level ontological analysis. In Boella, G., Odell, J., van der Torre, L., Verhagen, H., eds.: Proceedings of the 2005 AAAI Fall Symposium 'Roles, an Interdisciplinary Perspective: Ontologies, Languages, and Multiagent Systems', Nov 3-6, Arlington, Virginia. Number FS-05-08 in Fall Symposium Series Technical Reports, Menlo Park (California), AAAI Press (2005) 93–100

25. Rector, A., Stevens, R., Rogers, J.: Simple bio upper ontology. `http://www.cs.man.ac.uk/˜rector/ontologies/simple-top-bio/` (2006)

26. Loebe, F., Hoehndorf, R.: General Formal Ontology. `https://savannah.nongnu.org/projects/gfo/` (2006)

27. Vrandecic, D., Krötzsch, M.: Reusing ontological background knowledge in semantic wikis. In Völkel, M., Schaffert, S., Decker, S., eds.: Proceedings of the First Workshop on Semantic Wikis – From Wikis to Semantics. (2006)

28. Herre, H., Heller, B., Burek, P., Hoehndorf, R., Loebe, F., Michalek, H.: General Formal Ontology (GFO) – a foundational ontology integrating objects and processes. Onto-Med Report 8, University of Leipzig (2006)

29. Masolo, C., Guizzardi, G., Vieu, L., Bottazzi, E., Ferrario, R.: Relational roles and quaindividuals. In Boella, G., Odell, J., van der Torre, L., Verhagen, H., eds.: Proceedings of the 2005 AAAI Fall Symposium 'Roles, an Interdisciplinary Perspective: Ontologies, Languages, and Multiagent Systems', Nov 3-6, Arlington, Virginia. Number FS-05-08 in Fall Symposium Series Technical Reports, Menlo Park (California), AAAI Press (2005) 103–112

30. Straccia, U.: A fuzzy description logic. In: Proceedings of AAAI-98, 15th Conference of the American Association for Artificial Intelligence, Madison, US (1998)

31. Wikipedia: Web of trust — wikipedia, the free encyclopedia (2006) [Online; accessed 19-August-2006].

32. Wikipedia: Wikipedia:neutral point of view — wikipedia, the free encyclopedia (2006) [Online; accessed 19-August-2006].

33. Good, B.M., Tranfield, E., Tan, P., Shehata, M., Singhera, G., Gosselink, J., Okon, E., Wilkinson, M.: Fast, cheap and out of control: a zero curation model for ontology development. In: Proceedings of the PSB'06. (2006)