Mapping AGROVOC and the Chinese Agricultural Thesaurus: Definitions, tools, procedures

A. C. LIANG †, M. SINI†

†*Library and Documentation Systems Division, Food and Agriculture Organization of the United Nations (FAO), Rome, Italy,*
acliang@alum.mit.edu, margherita.sini@fao.org

Abstract

This paper describes the procedures for a concept-based mapping of two agricultural thesauri, the multilingual AGROVOC, created and maintained by the Food and Agricultural Organization, and the bilingual Chinese Agricultural Thesaurus, created and maintained by the Chinese Academy of Agricultural Science. Conducted under the auspices of FAO's Agricultural Ontology Service, the mapping project aims to extend AGROVOC with an additional set of perspectives on the agricultural domains, enrich its domain and language coverage, and to make use of AGROVOC as a common data model for data exchange among a wide range of multilingual repositories within agriculture.

Key words: Mapping, Thesaurus, AGROVOC, CAT, SKOS.

1. Introduction

The objective of this project is to link two terminology sources related to agriculture, the multilingual AGROVOC and the bilingual Chinese Agricultural Thesaurus (CAT) by mapping corresponding concepts from one terminology to the other. The immediate motivation for this work is to provide AGROVOC users access to the vast repositories of Chinese agricultural knowledge, hitherto inaccessible to non-Chinese speakers, by enabling them to use the vocabulary with which they are familiar (AGROVOC) in the languages with which they are familiar (any of the languages in which AGROVOC occurs) as an entry point to CAT-indexed resources. But there are also several other reasons for undertaking such a mapping. First, mapping allows for the extension of AGROVOC with an additional set of perspectives on the agricultural domains (e.g., herbs used in Chinese medicine). Second, in terms of domain and language coverage AGROVOC becomes enriched through the mapping to CAT since it may not go into as much depth in some domains as CAT does (e.g., AGROVOC has a single concept for mathematics whereas CAT has nearly 200 subconcepts). Third, the process of mapping can help improve the content and structure of both the terminologies by revealing areas that neither covers in depth and structures that are in need of revision. Fourth, in terms of the Food and Agricultural Organization's (FAO) Agricultural Ontology Service (AOS) initiative, mapping from AGROVOC to CAT begins to make use of AGROVOC as the common data model for data exchange among a wide range of multilingual repositories within agriculture. In addition, this project will demonstrate the

extent to which incorporating other terminologies within the AGROVOC Concept Server via mapping is a feasible exercise.

In this paper, concept refers to any mental representation or thought that can be expressed through a word or a phrase. Each thesaurus consists of concepts, each of which has a term code, a unique alphanumeric string, assigned to it. Each term code, or concept, is associated with terms, or language specific lexical realizations. Mapping consists of identifying concepts in a target vocabulary that are semantically equivalent to the concepts constituting a source vocabulary; mappings require references to the original vocabularies to be completely useable. A merge, in contrast, fully integrates the two vocabularies into a new structure. The mapping presented in this project is concept-based.

2. The problem of heterogeneity in vocabularies

In the previous section, we enumerated some of the reasons for mapping AGROVOC and CAT. There are however a considerable number of obstacles to surmount. The factors that hinder the ability to integrate any set of vocabularies have been noted by Doerr (2001), Zeng and Chan (2004), McCulloch, Shiri, Nicholson (2005), and others. In addition to sources of heterogeneity such as level of granularity, differences in degree of pre- and post-coordination, the degree of "shattering" where concepts from a single domain may occur spread out and in isolation from each other in the context of a general terminology system (Koch et al. 2001), and the use of different terms for the same concept, there are also issues specific to the integration of multilingual thesauri.

Zeng and Chan, citing Hudon (1997), note sources of heterogeneity arising from multilingual thesauri that are of particular relevance for this project. For example, the more distant the languages and cultures, the more disparate are the conceptual structures. AGROVOC's perspective on all of the agricultural subdomains it covers is decidedly from a Western, scientific point of view. CAT also adopts such a perspective for many of its subdomains; however it also includes notions of Chinese cosmology, such as earth, water, fire, metal, and wood, which refer to categories that are used to understand phenomena and organize essential activities rather than referring to the actual material themselves. Another source of heterogeneity is translations into the target language, which may also introduce problems for integration. Although the project does not involve translation per se, every single concept in AGROVOC was translated to Chinese (Chang and Lu 2002). Here it is not clear to what extent literal translations were created that don't make sense in the target language.

Zeng and Chan also describe methods for associating/integrating vocabularies, including mapping by computational and intellectual methods, which is pertinent to the approach we take, described below.

3. The structure of AGROVOC and CAT

Initially, AGROVOC was used for indexing information materials produced within the international cooperative information systems AGRIS and CARIS, and for data retrieval from those systems. Nowadays, AGROVOC is used more widely to index agricultural information in repositories all over the world. FAO published the first edition of AGROVOC in 1982, and then issued the second, third and fourth editions in 1988, 1995, and 1999, respectively. FAO issued the Web edition in 2000 which now has 12 languages (Arabic, Chinese, English, French, Spanish, Czech, Portuguese, Thai, Japanese, Slovak, Lao, Hindi). AGROVOC has a

total number of 16769 descriptors and 10968 non-descriptors in the area of agricultural science. The Chinese translations were provided by the same thesaurus experts who manage CAT (Chang and Lu 2002).

CAT was developed as a knowledge management tool for the agricultural forestry and biological fields. It is the second largest multi-disciplinary thesaurus in China. It has been approved by the authority of the Agricultural Ministry of China as a criterion for agricultural document retrieval systems and for archiving administration and scientific research resources. CAT has been extended by the Taiwan Agricultural Science Information Center. CAT contains 64638 terms in Chinese, including 51614 descriptors and 13024 non-descriptors. It has BT/NT, UF/USE and RT relations. Most descriptors have English translations. Some biological taxonomic names have only Latin translations. Only 200 descriptors and nearly all non-descriptors have no translation. There are 2332 top terms and terms are organized in 40 categories and sub-categories (indicated by codes 01, 19, 50, etc., e.g. crops) up to a maximum depth of three.

3.1 AGROVOC and CAT as multilingual thesauri

We have referred to AGROVOC as a multilingual thesaurus and CAT as a bilingual one. Doerr (2001) distinguishes several types of multilingual thesauri:

> (1) "translated thesauri," where concepts are lexicalized in one or more languages so that they can be interpreted by speakers of those languages; a distinction is made between the indexing term, or concept, and the translations, which are not used for indexing;
> (2) "correlated thesauri," where corresponding concepts from different thesauri are mapped to each other; and where mappings are the means by which queries using concepts from one thesaurus may be translated to concepts from another thesaurus;
> (3) "interlingua", where concepts may be independent of any individual thesaurus mapping to it and which allows for exchange of concepts of participating thesauri.

CAT falls under the category of the translated thesaurus, where the concepts are expressed in Chinese and the English (or Latin) translations simply added. AGROVOC originally fell into this category. However, the common practice regarding the use of AGROVOC for indexing records at AGRIS Centers is that translations themselves have been applied as index terms, i.e., instead of using the term code to represent the concept, the lexicalization corresponding to the language of the host institution is used[1]. In addition, as new, non-Western European languages, in particular, were added, and it became clear that there were semantic problems in trying to fit terms into a Western European (English)-biased template, we have changed the overall data model of AGROVOC to allow for more language-specific expression of concepts (e.g., Thai AGROVOC has numerous terms for mango species for which there are no correspondences in the other languages). As a result, the structure of AGROVOC has started to take on characteristics of both correlated thesauri and an interlingua. Although not associated with distinct databases, new language versions of AGROVOC are like mapped thesauri in that they are not obliged to provide translations for all currently existing AGROVOC concepts if no lexicalization exists, and they may introduce new concepts for which there is no corresponding AGROVOC concept. Viewed as a single entity, AGROVOC

[1] This practice is changing. The XML files representing the AGRIS metadata refer to the URI of the concepts associated with the indexing terms.

has some of the characteristics of an interlingual thesaurus in that it consists of concepts that may be shared among language groups independently of whether or not a given language group has lexicalizations for it (cf. use of Esperanto as an interlingua for machine translation in Hutchins 1986, Chapter 16).

However, what makes this particular project unique and also problematic is the fact that the thesauri being mapped are multilingual. Previous mapping work has only dealt with monolingual thesauri or terminologies that are in the same language (e.g., Vizine-Goetz 2004 on the OCLC terminology services) or monolingual thesauri or terminologies in different languages (e.g., Infolab 2000 on MACS, Koch, Neuroth, and Day 2001 on Renardus). In contrast, both CAT and AGROVOC are not only multilingual, but they both contain Chinese and English terms. Thus, the question arises, if two concepts in two independently developed thesauri both have the same Chinese and English lexicalizations, can they be considered equivalent concepts (setting aside for the time being, the issue of relations to surrounding concepts)? Or to put it differently, to what extent do both English and Chinese lexicalizations need to be considered in deciding whether or not the two concepts are equivalent?

We decided to approach this problem by considering the following: AGROVOC, although a multilingual thesaurus, began its development with terminologists most familiar with English/Western European languages, the addition of non-European languages being added (over a decade) later. The Chinese Agricultural Thesaurus, in turn, was developed by Chinese terminologists for Chinese users. When analyzing the semantic equivalences or near-equivalences across the two thesauri, the Chinese terms in the CAT are given primary consideration over any existing English translation as lexicalizing a given concept, and the English terms in AGROVOC are likewise considered the primary lexicalization for a given concept. Thus, the analysis of the mapping relations is approached as a multilingual thesaurus mapping, from AGROVOC's English lexicalizations to CAT's Chinese lexicalizations. However, one more factor must be taken into consideration. The Chinese lexicalizations in AGROVOC were added recently by the CAAS group, the same organization doing the mapping. These translations were done at least in part through cross-referencing with the Chinese Agricultural Thesaurus. Thus, most AGROVOC terms are already "linked" in some sense to CAT terms. Whether they can be considered conceptual equivalents remains to be seen (through intellectual review), but the possibility that AGROVOC contains the beginnings of a mapping to the CAT will be taken into account below when we specify the mapping procedure.

4. The notion of equivalence

The mapping is carried out with AGROVOC designated as the source terminology and CAT as the target terminology. This means that the mapping relationship between AGROVOC and the CAT concepts is always one-to-one (or more), for those AGROVOC concepts that have CAT equivalents.

The notion of equivalence is treated in Doerr (2001), Koch et al. (2001), and in the Simple Knowledge Organization System (SKOS) (Matthews and Miles 2004, based on Doerr) literature in terms of set theory. That is, two concepts are deemed to be equivalent if they are associated with, or classify the same set of objects. In principle, this approach makes sense given that concepts used for the purposes of indexing are indeed serving a classifying

function.[2] And it provides a framework for talking about concepts and their relations. However, for the practical purposes of establishing a mapping between two concepts in two different terminology systems, it is extremely difficult to ascertain to what extent such concepts classify the same set of objects (or documents). For example, it is impossible to determine whether the degree of overlap (Doerr 2001, Koch et al. 2001) between a concept in AGROVOC that overlaps with one in CAT is major or minor (i.e., more or less than 50%) because neither is associated with the same database. Even considering that the set theory paradigm is meant as a set of explicit criteria for (potentially) quantitatively assessing equivalence and other kinds of relations between concepts, for our purposes, it is too abstract to be of practical use to address the immediate needs of the terminologists who will do the mapping. In the absence of a set of objects with which a concept can be identified, the actual analysis that is carried out in this project to determine whether a concept in one vocabulary is equivalent to one in another is done so on the basis of "comprehension of the term itself," (Doerr 2001:6) in light of its relations to other concepts in the thesaurus in which it occurs, and not in terms of the mathematical notion of set.

Therefore, as a starting point to the identification of equivalence, we take an approach similar to Vizine-Goetz (2004) by which we initially identify, programmatically, a list of term matches that serve as candidate equivalent pairs. By term match, we mean, for any given pair of concepts, the occurrence of any of the following in the thesauri:

- the same Chinese and English terms;
    - AGROVOC Fertilizers / 肥料 -- CAT 肥料 / Fertilizer
- the same Chinese terms but different English terms (in case of synonymy);
- the same English terms but different Chinese terms (in case of synonymy);
    - AGROVOC Tanning / 鞣制 – CAT 染网 / Tanning

These candidates then undergo intellectual review to validate (or reject) the matches as conceptual equivalents. Review involves the evaluation of concepts based on their hierarchical and associative relations to other concepts within the thesauri. Rather than applying (exclusively) the notions of set theory per se, concepts may be (also) considered in terms of their semantic features. It is more likely that concept pairs with matching Chinese and English terms, as well as those with matching Chinese terms will be concept equivalents for the reasons mentioned above.

5. Relations

We examined mapping relations in Doerr (2001), Koch (2001), and SKOS (Miles and Brickley 2004). We retained three relations that all three sources posit: exact, broader and narrower, and dispensed with minor and major overlap for the reason mentioned above. We

[2] Although the set theory approach to concept meaning is the dominant paradigm in discussions on knowledge organization systems, it is not the only way to analyze concepts. Concepts may also be associated with semantic features, and mappings based on equivalences between feature sets. For example, a medical vocabulary used to index laboratory results, the Logical Observation Identifier Codes and Names (LOINC), consists of concepts made up of six semantic features (Huff et al. 1998), each of which contains a value. The combination of six feature-value pairs yields a controlled vocabulary item that may have synonyms and other term variants associated with it. Any laboratory result can be defined in terms of these features and mapped to this vocabulary.

also retained the three Boolean operators, AND, OR, and NOT to combine target concepts to which source concepts can be mapped using the three aforementioned relations.

5.1 Exact match

When a concept in AGROVOC corresponds exactly to one or more concepts in CAT, the relation is an exact match. Unlike the other two relations (cf. below), this one is always applied when a correspondence can be identified. Figure 1 exemplifies an exact equivalence.

AGROVOC **Social classes**社会阶级 CAT阶级 **CLASS**

Figure 1. Exact matched concepts

That relation is supported by the occurrence of the following respective child concepts, shown in Figure 2.

AGROVOC **Ruling classes**统治阶CAT 统治阶级 **RULING CLASS**

Figure 2. Exact matched concepts children of concepts in Figure 1.

5.2 Broad/Narrow matches

A broad equivalence is one where the source concept is more specific in some sense than the target concept. A narrow equivalence is one where the source term is more general in some sense than the target term or expression. Following Doerr (2001), to the extent possible, we carry out mappings that are at least complete, and ideally optimal. A complete mapping is one where a source concept, having no exact equivalence in the target, is matched to at least one target concept that is semantically broader and at least one target concept that is semantically narrower. An optimal mapping is one where the aforementioned broader target concept is the nearest broader concept to the source concept, and where the aforementioned narrower target concept is the nearest narrower concept to the source concept.

Figure 3 shows that there are several possible concepts in the target vocabulary for establishing complete and optimal matches.

AGROVOC Physician	activity
	NT human activity
	NT1 professional
	NT2 neurologist
	NT3 neurooncologist

Figure 3. Optimal mapping.

Although AGROVOC Physician can be broad matched with several of the fictional target concepts listed, namely, activity, human activity, and professional, the last one, professional is mapped to since it is the narrowest of all the broad matches. Conversely, the source concept can be narrow matched with both neurologist and neurooncologist. The actual mapping is made to neurologist since it is the broadest of the narrow matches.

5.3 Boolean operators

These allow the expression of complex target concepts to which equivalences may be made.

The AND operator is a binary operator expressing conjunction. ‘AND’ is used to identify a concept formed from the intersection of two or more concepts. The simplest use of AND is an equivalence corresponding to the conjunction of two named concepts. Formally, the relationship is exemplified as follows.

Woman exactMatch Female AND Adult human

The OR operator is a binary operator expressing disjunction. ‘OR’ represents the union of the semantics of two or more concepts.

For example, In Chinese there are different concepts corresponding to the grain and the plant. The term ‘大米’ is used as rice as a grain which can be eaten; whereas ‘稻’ is used for the plant ‘Oryza sativa’. English does not make such a distinction. Thus, the mapping is

Rice exactMatch ‘rice as grain’ 大米 OR ‘Oryza sativa’ 稻

The NOT operator is a unary operator expressing complementarity or negation. ‘NOT’ can be used to create a mapping target from which one or more elements of meaning are excluded.

Mainland China exactMatch China NOT (Hong Kong OR Macau OR Taiwan)

6. OWL Web Ontology Language as the vocabulary encoding

The OWL Web Ontology Language (OWL) is used to represent both the vocabulary and the mappings. This is because OWL contains the apparatus to express the relations that are required for mapping equivalences across vocabularies. It is also supported by numerous tools that can manage the mapping work. In addition, for reasons which will not be delved into here, the FAO authors have decided, independently of this project, that AGROVOC is to be represented in OWL. Some of the decisions we have made regarding the conversions, described below, are due to our priorities concerning OWL rather than the mapping per se.

The following elements from OWL are used:

<class>:	concept
<equivalentTo>	exactMatch
<subClassOf>	broad/narrow match

Further, the following conversions are made to each thesaurus.

- All BT/NT relations are converted to super/subclass relations.
- RT is retained and converted to an OWL object property. For this project, RT is simply a placeholder and will not be relevant to the mapping other than to provide semantic associations that may be of use in analyzing the meaning of a concept.
- UF/USE are eliminated. All terms, whether descriptor or non descriptor, are converted to OWL classes. Non-descriptors have the same parent(s,), i.e., are converted into siblings of the corresponding descriptors. This is motivated by the fact that there appear to be at least as many non-descriptors that are conceptually distinct from the selected term as those that are synonymous with it. Thus, these terms will

require further analysis to identify those that are concepts and those that are synonyms.

All other information (e.g., scope notes, etc.) is not included in the source files for the mapping work.

7. Preparation

Both AGROVOC and CAT have been converted to OWL. AGROVOC will be divided into files, each corresponding to an AGRIS category in the AGRIS/CARIS classification scheme.

Candidate equivalences have been extracted. There are over 4000 matching Chinese-English pairs, nearly 4000 matching Chinese-different English pairs, and about 2000 matching English-different Chinese pairs.

8. Tools

The Semantic Standards group within the Library and Documentation Systems Division at FAO, in collaboration with the Chinese Academy of Agricultural Sciences (CAAS) analyzed several possibilities for executing the mapping: Protégé[3], Vocabulary Integration Environment (VINE)[4], and a custom solution (MS Excel/MS Access).

Protégé was the most promising of the three possibilities. It is a free Java-based tool that can be downloaded and installed. It has an easy and user-friendly way for representing concepts and their lexicalizations.

We anticipated that the merging work would be done with the Prompt merging tool, a Protégé plug-in designed to manipulate branches and concepts of both the thesauri within the same environment as well as to propose suggestions for which concepts and relations to merge. Prompt turned out to be extremely unstable with several of the versions of Protégé that we tried.

Then, given that the work would be done collaboratively, we anticipated that the ideal architecture would be a distributed client-server environment where terminologists would access and do the mappings via connections to a single server using a Protégé client. So even if the Prompt tool didn't work at least we could exploit the functionalities of Protégé. However, numerous tests, particularly those concerning ease of set up and performance, were done on the Protégé server and were found to be extremely problematic. For example, it was not possible to successfully load the full AGROVOC OWL file without generating errors.

In addition to these technical problems, according to the CAAS team, the network connection at CAAS is not always optimal, and working on a central client-server system wasn't envisioned to be compatible with the work habits of the CAAS team.

VINE was also tested by loading the two OWL files and executing a mapping using the definition of new relationships as allowed by the tool. But it was considered inadequate as it lacks many of the relationships and level of sophistication needed to produce a complete

[3] http://protege.stanford.edu/
[4] http://marinemetadata.org/examples/mmihostedwork/ontologieswork/vine/

mapping. Additionally, the initial feedback (mid-2005) among users for whom the tool was originally developed (participants of the Marine Metadata Interoperability project) indicated performance problems within the Windows environment.

The third alternative was to do the whole mapping by hand using an Excel spreadsheet formatted to allow for (easy) input. The mapping would be executed by specifying in separate columns the ID of the term from the source concept and the ID of the term to be mapped to in another column. The only advantages of this alternative were that no internet connection was needed and that the time to prepare the spreadsheet was less than that needed to customize an existing tool such as Protégé. However, the mapping process is extremely cumbersome, and the output would have to be converted into a useable format. It was also hard to identify the representation of a concept of both the source and the target, as no clear graphical representation was possible.

Finally, we returned to Protégé, and decided to do the mapping simply by using the functionalities available without benefit of special plug-ins. Both the source and the target terminiologies were imported in the tool, and specific relationships were created between the so created concepts. Each domain specific AGROVOC file would be loaded, and the entire CAT would be imported. This would allow each terminologist to work independently on their segment of AGROVOC on a constant version of CAT. All mappings would already be in the OWL format.

For this work, the following versions are being used:

AGROVOC (1st November 2005)
CAT (1994)
Protégé 3.2b, Build 278, Windows version
Java J2SE 1.4.2.05

9. Procedure

All candidate exact equivalences will be reviewed first. Then the top terms will be mapped. Terminologists will then proceed down the hierarchy to do the rest of the mappings. It is intended that everything in AGROVOC will, if possible, eventually be mapped.

Any error that is encountered, any problematic concept that cannot be mapped, will be noted by the terminologists.

It is likely that the mapping procedure will need to be revised as the work proceeds and problems are encountered that were not forseen.

10. Conclusion

We have described a method for mapping two multilingual thesauri. As mentioned, this is an ongoing work. Most of the methodology has been identified as well as the mapping rules. What is lacking is an evaluation procedure. Short evaluation cycles need to be established in order to determine the quality of work and to be able to adjust the mapping procedures before considerable time and effort has been invested. In the longer term, the management of changes in the source and target vocabularies, and how that affects the mapping will have to be addressed. Tests of the usefulness of the mapping will also need to be implemented.

Acknowledgement

The authors are grateful to Mr. Chang Chun of the Chinese Academy of Agricultural Sciences for his help in providing mapping examples.

References

Chinese Academy of Agricultural Sciences. 1994. Agricultural Thesaurus. Chinese Agricultural Press.

Chang Chun, Lu Wenlin. 2002. The translation of agricultural multilingual thesaurus. Proceedings of the Third Asian Conference for Information Technology in Agriculture. Edited by Mei Fangquan. China Agricultural Scientech Press, pp. 526-528.

Doerr, Martin. 2001. Semantic Problems of Thesaurus Mapping. *Journal of Digital Information*, Volume 1 Issue 8, Article No. 52, March 26. http://jodi.tamu.edu/Articles/v01/i08/Doerr/

Food and Agricultural Organization of the United Nations. Multilingual Agricultural Thesaurus (AGROVOC). http://www.fao.org/agrovoc/ (Accessed November 11, 2005).

Hudon, Michele. (1997) Multilingual thesaurus construction: integrating the views of different cultures in one gateway to knowledge concepts. *Knowledge Organization* 24(2): 84-91.

Huff, Stanley M. M., Rocha, R. A., McDonald, C. J., De Moor, G. J., Fiers, T., Bidgood, W. D., Jr., Forrey, A. W., Francis, W. G., Tracy, W. R., Leavelle, D., Stalling, F., Griffin, B., Maloney, P., Leland, D., Charles, L., Hutchins, K., & Baenziger, J. 1998. Development of the Logical Observation Identifier Names and Codes (LOINC) vocabulary. *Journal of the American Medical Informatics Association* 5(3):276-292.

Hutchins, John. 1986. Machine translation: past, present, future. (Ellis Horwood Series in Computers and their Applications.) Chichester: Ellis Horwood. 382pp.

Infolab. 2000. MACS: Multilingual Access to Subjects. infolab.uvt.nl/pub/hoppenbrouwersj-2001-23.pdf.

Koch, Traugott. 2001. The Renardus Broker: A meta-subject gateway. http://www.lub.lu.se/tk/renardus/tokyoren.html.

Koch, Traugott, Neuroth, Heike and Day, Michael (2001). Renardus: Cross-browsing European subject gateways via a common classification system (DDC). 8pp. In: "Subject Retrieval in a Networked Environment". Papers presented at an IFLA Satellite Meeting sponsored by the IFLA Section on Classification and Indexing and the IFLA Section on Information Technology,14-16 August 2001, Dublin, OH, USA. http://www.ukoln.ac.uk/metadata/renardus/papers/ifla-satellite/ifla-satellite.pdf.

McCulloch, Emma and Shiri, Ali and Nicholson, Dennis. 2005. Challenges and issues in terminology mapping: a digital library perspective. Electronic Library 23(6) pp.671-677.

Miles, A. and Brickley, D. 2004. SKOS Mapping Vocabulary Specification. http://www.w3.org/2004/02/skos/mapping/spec/ (Accessed May 16, 2005)

Miles, A., and B. Matthews. 2004. Inter-Thesaurus Mapping. http://www.w3c.rl.ac.uk/SWAD/deliverables/8.4.html (Accessed April 16, 2005)

Vizine-Goetz, Diane, Carol Hickey, Andrew Houghton, and Roger Thompson. 2004. "Vocabulary Mapping for Terminology Services." *Journal of Digital Information*,

Volume 4, Issue 4, Article no. 272, March 11.
http://jodi.tamu.edu/Articles/v04/i04/Vizine-Goetz/

Wang Amy Y., J. Barrett, T. Bentley, D. Markwell, C. Price, K. A. Spackman, and Michael Q. Stearns. 2001. Mapping Between SNOMED RT and Clinical Terms Version 3: A Key Component of the SNOMED CT Development Process. *Proceedings of the 2001 AMIA Fall Symposium*. November 3, Washington, D.C., pp. 741-745.

Zeng M. and Chan L. 2004. Trends and issues in establishing interoperability among knowledge organization systems. Journal of American Society for Information Science and Technology, 55(5): 377 - 395.