Word Association Testing and Thesaurus Construction

Louise F. Spiteri

School of Library and Information Studies

Dalhousie University

Halifax , Nova Scotia

Canada

Louise.Spiteri@dal.ca

Abstract

This paper examines the suitability of word association tests to generate user-derived descriptors, descriptor hierarchies, and categories of inter-term relationships. The typical assumption underlying these word association tests is that the response terms function either as synonyms or antonyms, an assumption that restricts unnecessarily the potential value of such tests. Rather than assuming how people inter-relate two terms, it may be more useful to ask participants to explain why they think these two terms are related. In this study, thirty library and information science practitioners were asked to provide as many response words as they could for fifteen stimulus terms and to describe how the response and stimulus terms were inter-related. The word association test was successful in generating a set of user-derived descriptors . Participants identified twenty types of inter-term relationships, the most commonly-cited of which are type, part, synonym, activity, and tool. That the participants identified a total of twenty types of relationships suggests also that word association tests can serve as a valuable tool in examining the different ways users group terms and the types of inter-term relationships that end users most commonly associate with any given concept and its response terms.

Introduction

Word association testing is a technique developed by Carl Jung to explore complexes in the personal unconscious. Jung came to recognize the existence of groups of thoughts, feelings, memories, and perceptions, organized around a central theme, that he termed psychological complexes. This discovery was related to his research into word association, a technique whereby words presented to patients elicit other word responses that reflect related concepts in the patients'

psyche, thus providing clues to their unique psychological make-up (Schultz and Schultz, 2000).

Word association testing has been used extensively in psychology to assess the personality of the test subjects (Galton, 1880; Kent and Rosanoff, 1910; Russell, 1970). Projective techniques, of which word association is a type, typically present respondents with an ambiguous stimulus and ask them to disambiguate this stimulus. The underlying principle behind most projective techniques is that respondents project aspects of their own personalities in the process of disambiguating test stimuli. The interpreter of the projective technique can thus examine answers to these stimuli for insights into the respondents' personality dispositions. In a typical word association test, subjects are asked to respond to a stimulus word with the first word that comes to mind. These associative responses have been explained by the principle of learning by contiguity: "objects once experienced together tend to become associated in the imagination, so that when any one of them is thought of, the others are likely to be thought of also, in the same order of sequence or coexistence as before" (Wettler and Rapp, 1996).

Word association tests present a potentially useful tool in the construction of information retrieval (IR) thesauri, especially for involving end users in the process. The design of thesauri normally employs a deductive approach: broad categories of terms are selected and then sub-divided into narrower sets based upon the application of a series of pre-ordained inter-term relationships. In a previous paper (Spiteri, 2002), the author proposed a theoretical framework by which word association tests could be used to generate user-derived descriptors and term hierarchies for IR thesauri. The focus of this paper is to explore the results of a pilot study, based upon this theoretical framework, that examines the extent to which word association tests can be used to

- Generate user-derived descriptors, i.e., terms that are most commonly associated with a given concept by the majority of respondents. End-users are provided with a list of domain-specific stimulus terms and are then asked to provide response terms;
- Generate user-derived descriptor hierarchies, i.e., the most commonly-associated attributes, properties, characteristics, parts, etc., of a given concept as identified by the majority of respondents. End-users are asked not only to provide response terms but also to specify how they think these terms are related to the stimulus terms; and
- Generate user-derived categories of inter-term relationships, i.e., the most commonly-associated types of relationships identified by the majority of respondents.

Rationale

Word association tests have been used in the construction of such lexical tools as ontologies, taxonomies, and thesauri to elicit the most typical terms that people associate with a given stimulus term in order to understand how end users categorize vocabulary around a central concept (Spiteri, 2002). The assumption

underlying a number of the uses of word association tests is that the response terms function either as synonyms or antonyms; the interpretation of these relationships is made by the researchers rather than by the participants (Deese, 1965; Nielsen, 1997; Miller et al., 1993). To a limited extent, Word association tests have been used to ask participants to provide attributes and activities associated with the stimulus terms (Battig and Montagu, 1959; Smith and Mark, 1999; Tversky and Hemenway, 1983). Once again, however, the researchers, rather than the participants, categorized how the response terms were related specifically to the stimulus terms.

IR thesauri, however, contain more than mere listings of antonyms and synonyms; they contain terms that are bound in a variety of hierarchical and associative relationships (e.g., whole-part, an object and the tools used to produce it, etc.). Given this, the assumption that response terms are necessarily synonyms or antonyms of stimulus terms restricts unnecessarily the potential value of word association tests. When presented with the word *dogs*, for example, many people respond with the word *cats*. A cat is clearly not a synonym for a dog, neither is it an antonym, yet in the minds of many people, these two terms are closely connected. Rather than assume how people inter-relate these two terms, it may be more useful to ask the participants to explain why they think these two terms are related (e.g., they are both types of domestic animal).

Word association tests often restrict participants to providing only one response term per stimulus term, which could also be too restrictive. Is *cat*, for example, the only term that people associate commonly with *dogs*? Since IR thesauri act as tools to assist in indexing and searching, it would be useful to use word association tests to elicit as large a set as possible of inter-related terms, a set that reflects the variety of ways in which end users approach a given concept.

IR thesauri rely typically upon the use of symbols such as USE/UF, BT, NT, and RT to demonstrate inter-term relationships. The exact nature of the inter-term relationship expressed by any one of these symbols is not, however, necessarily obvious. For example, is the BT/NT relationship based upon a whole-part, instance, or a genus-species division? ISO and NISO guidelines suggest that the symbols BTG/NTG, BTP/NTP, and BTI/NTI be used to distinguish respectively amongst the genus-species, whole-part, and instance hierarchical relationships, but, for the most part, the more generic BT/NT symbols are used (ISO, 1986; NISO, 1993). The equivalence relationship can include synonyms, quasi-synonyms, and even antonyms; the use of USE/UF indicates only that some type of equivalence relationship exists but not the exact nature of this relationship. The situation becomes ever murkier with the associative relationship, where the generic RT is used to express up to eleven different types of inter-term relationships (ISO, 1986; NISO, 1993).

Word association testing could thus also be used to generate sets of relationship labels (or facet indicators) based upon the terminology participants use to describe how their response terms are related to the respective stimulus terms. Some ontologies, for example, specify the exact nature of inter-term relationships through the use of labels such as "IS A," "IS A TOOL OF," "IS A DOMAIN OF," and so forth ([Theory-Frame Ontology](#), 1997; [OpenCyc Selected Vocabulary and Upper Ontology](#), 2002). The lexical database [WordNet](#), for example, displays terms according to

such stated relationships as synonyms, hypernyms/hyponyms (is a kind of), and holonyms/meronyms (is a part of). The term *dog*, for example, has the following hierarchy:

Hypernyms (dog is a kind of): => canine, canid

Hyponyms (... is a kind of dog): => dalmatian, coach dog, carriage dog

Holonyms (dog is part of ...): => Member of: canis, genus canis

Meronyms (... is a part of dog): => Has part: paw

By using end user generated relationship labels, IR thesauri could follow the model set by such lexical tools to design hierarchies that display more clearly and intuitively the nature of inter-term relationships.

Methodology

Since most thesauri are domain specific, it is essential that the stimulus terms chosen for the word association test be drawn from the domain at hand. For this pilot project, the subject domain of Library and Information Studies (LIS) was chosen although this methodology could be applied to a variety of domains as needed. A test bed of stimulus words for LIS was drawn from the following sources:

- [Open directory project](#)
- [ASIS Thesaurus](#)
- [Government of Canada Core Subject Thesaurus](#)
- [ERIC Thesaurus](#)
- [Legislative Indexing Vocabulary](#)

Stimulus terms were chosen if they were common to at least two-thirds of the sources consulted. In this way, some degree of term familiarity amongst the participants could be anticipated. The total number of stimulus terms chosen was fifteen. Participants were drawn from the library-practitioner population in Atlantic Canada. Calls for participation were communicated via the listservs of the Atlantic Provinces Library Association (APLA) and the Nova Scotia Library Association (NSLA). The total number of participants was thirty. For each stimulus term, the participants were asked to take no more than two minutes (self-administered) to write down as many response terms as they thought were related to the stimulus term. Participants were also asked to explain in written form how they thought each of their response terms related to the respective stimulus term. The stimulus terms used were:

Authority Control
Cataloguing
Censorship
Digital Libraries
Information Literacy
Information Retrieval
Intellectual Freedom

Intellectual Property
Information Services
Librarians
Library Science
Reference Materials
Special Collections |
Technical Services
Thesauri

User-derived response terms

For each stimulus term, all the response terms provided by each participant were noted. These terms were divided into two categories: (a) terms that occurred uniquely (i.e., that were cited by only one participant) and (b) terms that were cited by two or more participants. It should be noted that the singular and plural forms and variant spellings of the same response term were considered to constitute one term (e.g., librarian/librarians, cataloguing/cataloging). The average number of response terms assigned by the participants per stimulus term was calculated. Stimulus terms were ranked in order of (a) the total number of unique response terms assigned to them and (b) the average number of unique response terms assigned to them per participant. A list of the most commonly-occurring stimulus term/response term word pairs was generated. Since one of the foci of word association tests is to examine consensus in the way participants react to a stimulus term, a response term had to be cited by at least fifty percent of the participants to make it a candidate for a word pair.

Inter-term relationships

For each stimulus term, a list of participant-defined inter-term relationships was derived. The inter-term relationships were divided into two categories: (a) those that occurred uniquely (i.e., were cited by only one participant) and (b) those that were cited by two or more participants. The matching of inter-term relationships was rather more complicated than the matching of response terms since, in the latter case, the possible overlaps in the types of relationships expressed needed to be determined. In other words, if one participant says that Term A is a *type of* Term B and another participant says that Term A is a *form of* Term B, is this, in fact, the same type of relationship? The relationship labels cited by the participants were thus examined independently by the principal researcher and a research assistant. The two evaluators determined independently which of these labels constituted unique types of relationships and which constituted overlapping types of relationships and then compared their results. From this exercise, a single list of user-derived types of relationships was established, and the frequency with which these types were cited by the participants was noted.

Findings: Incidence of response terms

Figure 1 shows the stimulus terms ranked in order of the total number of unique response terms assigned by the participants, with an average of seventy unique response terms per stimulus term.

Stimulus term	Number of unique response terms
Digital libraries	97
Cataloguing	93
Censorship	88
Librarians	84
Information literacy	77
Information services	74
Authority control	72
Library science	71
Thesauri	66
Intellectual freedom	59
Reference materials	59
Special collections	59
Intellectual property	55
Information retrieval	53
Technical services	48

Figure 1: Stimulus terms ranked in order of total number of unique response terms

Figure 2 shows the stimulus terms ranked in order of the average number of response terms assigned by each participant, with an average of 4.1 response terms per stimulus term.

Stimulus term	Average no. of response terms per participant
Cataloguing	5.3
Reference materials	5.0
Information retrieval	4.6
Censorship	4.5
Information services	4.3
Digital libraries	4.1
Intellectual freedom	4.1
Thesauri	4.0
Information literacy	3.9
Authority control	3.7
Library science	3.7
Intellectual property	3.6
Librarians	3.6
Special collections	3.6
Technical services	3.0

Figure 2: Stimulus terms ranked in order of average number of response terms per participant

Figure 3 shows the stimulus term/response term pairings that were cited by at least fifty percent of the participants; the complete list of word pairs is found in Appendix 1.

Stimulus Term	Response Term	Frequency
Intellectual property	Copyright	82%
Information services	Reference services	79%
Technical services	Cataloguing	74%
Reference materials	Encyclopedias	71%
Intellectual freedom	Censorship	58%
Technical services	Acquisitions	56%
Reference materials	Dictionaries	51%

Figure 3: Word pairs cited by $\geq 50\%$ of the participants

The large number of response terms (Figure 1) compared to the average number of response terms per participant (Figure 2) suggests that there is not always a high degree of overlap in response terms; in fact, each stimulus term contains response terms that are mentioned only once. On the other hand, that, on average, participants cited 4.1 response terms per stimulus term means that restricting responses to only one term can, in fact, place a limit on the full potential of word association tests. The word pair *Reference materials/Encyclopedias* is a case in point: 71% of the participants cited *Encyclopedias* as a response term to *Reference materials*, yet *Encyclopedias* was not always the first term cited by the participants, as is the pattern with all the word pairs that appear in Figure 3. Figure 3 also indicates that a stimulus term may frequently be associated with more than one response term as is the case with *Reference materials/Encyclopedias*, *Reference materials/Dictionaries*, *Technical services/Cataloguing*, and *Technical services/Acquisitions*.

Another factor to be noted is that in Figure 3 a number of the word pairs do not constitute incidences of synonyms or antonyms; in fact, perhaps only *Information services/Reference services* could be considered as synonyms. The only seemingly-obvious antonyms are *Intellectual freedom/Censorship*, which may serve to support the suggestion that restricting word association tests to the derivation of only synonyms and antonyms is too restrictive and fails to make full use of the potential of these tests.

Findings: Incidence of inter-term relationships

The two evaluators agreed that the following labels constituted the same type of relationship, to which they assigned the label that had been cited the most frequently by the participants:

- Type of/Form of = Type of
- Participant/Member/Advocate = Participant
- Component of/Part of = Part of
- Goal of/Aim of/Purpose = Purpose
- Action/Activity = Activity
- Equivalent term/Synonym = Synonym
- Place/Location = Location

Figure 4 shows the stimulus terms ranked in order of the total number of unique relationships assigned to their response terms by the participants, with an average of 11.7 relationships per stimulus term.

Stimulus term	Total no. of unique types of <u>inter-term relationships</u>
Authority Control	16
Cataloguing	16
Censorship	14
Intellectual Freedom	14
Information Retrieval	14
Digital Libraries	13
Information Literacy	13
Information Services	13
Librarians	13
Library Science	13
Intellectual Property	11
Reference Materials	11
Thesauri	11
Technical Services	08
Special collections	08

**Figure 4: Stimulus terms ranked in order of total
number of unique inter-term relationships**

Figure 5 shows the participant-defined inter-term relationships ranked in order of frequency, with a total number of twenty unique types of relationships.

Inter-term relationship	Total no. of occurrences
Type	216
Part	209
Synonym	166
Activity	144
Tool	87
RT	75
Attribute	65
Product	59
Participant	58
NT	45
Is	31
BT	30
Antonym	26
Location	22
Purpose	17
Format	11
Skill	05
Use	07
Requirement	03
Source	03

Figure 5: Inter-term relationships ranked in order of frequency of occurrence

Although the incidence of synonymous relationships is high, which is in keeping with the more traditional uses of word association testing, Figure 5 indicates that the *Type* and *Part* relationships are cited the most frequently by the participants, thus suggesting that word association tests may not necessarily produce only synonyms and antonyms. That the participants identified a total of twenty types of relationships also suggests that word association tests can serve as a valuable tool in examining the different ways in which users group terms and the types of inter-term relationships that end users most commonly associate with any given concept and its response terms. More significant, perhaps, is the importance of asking participants to explain how their response terms are related to the stimulus terms. True synonyms, e.g., *elevators/lifts*, and antonyms, e.g., *life/death*, may be relatively easy to identify by the researcher; but without the explanations provided by the participants, it would be difficult for the researcher to interpret the inter-term relationships of most of the word pairs found in Figure 3.

The application of this word association test resulted in a total of 192 incidences of equivalence relationship (synonyms and antonyms), 531 incidences of hierarchical relationship (part of, type of, BT, NT, Is), and 556 incidences of associative relationships (all remaining relationships). The total number of inter-term relationships is 1279: 15% equivalent, 42% hierarchical, and 43% associative. As can be seen, the equivalence relationship constitutes the minority of inter-term relationships identified by the participants, which is not in keeping with typical assumptions about the results of word association tests. The hierarchical and

associative relationships constitute almost identical proportions of the participants' relationships. What is also clear is that participants do distinguish amongst different types of hierarchical relationships, which suggests that they go beyond the simple BT/NT distinctions one finds in most thesauri.

Conclusions

The word association test applied in this study was successful in generating a set of user-derived descriptors. Although the response terms provided by the participants varied quite significantly at times, areas of consensus did emerge where at least fifty percent of the participants provided the same terms in response to a given stimulus term. Participants provided an average of 4.1 response terms per stimulus term, which suggests that the normal restriction of one response term per stimulus term can serve to limit the contribution of word association testing to the development of a collection of descriptors for a thesaurus.

The findings suggest that word association tests could be used to generate user-derived term hierarchies. Results indicate that synonyms/antonyms (i.e., the equivalence relationship, according to ISO and NISO) is not, in fact, the only type of inter-term relationship reflected in the response terms although this has often been the underlying assumption of previous applications of word association tests. Participants, in fact, provided in practically equal measure instances of equivalence, hierarchical, and associative relationships. The importance of asking the participants to explain how their response terms are related to the stimulus terms cannot be overlooked. Without this explanation, any interpretation of the relationship between, say, *librarians* and *information professionals* would reflect the mental model of the researcher rather than that of the participants. This application of word association therefore lends itself to the use of inductive reasoning in the construction of thesauri. Rather than start with a general concept and assume an existing relationship between or among terms associated with that concept, which is the typical procedure used in the construction of many thesauri, word association allows thesaurus designers to study the patterns of inter-term relationships that may emerge. Word association tests could also be used to test existing term hierarchies: do the end-users (or even the thesaurus designers themselves) inter-relate terms in the same way as these hierarchies do?

If word association tests are to be used as aids to thesaurus construction, it would also be very useful to examine the degree of consensus amongst the types of relationship proposed between word pairs. All participants cited *information professionals*, for example, as a synonym of *librarians*, *copyright* was always cited as an antonym of *intellectual freedom*, and *dictionaries* as a type of *reference materials*.

Save for the occasional use of the generic BT, NT, and RT labels, the participants had no difficulty making clear distinctions between how different response terms were related to the same stimulus term; in other words, they did not say that a response term was merely broader or narrower than a stimulus term or that it was simply related to the stimulus term. It may therefore be helpful if thesauri could show an equal degree of clarity in the way they display inter-term relationships. The relationship labels suggested by the participants could be used as follows:

Librarians

Synonym(s)	Information Professionals
Types	Academic Librarians Public Librarians Reference Librarians Special Librarians
Activity	Acquisitions Cataloguing Collections Development Information Services

The typical thesaurus display for the hierarchy above would be:

Librarians

UF	Information Professionals
NT	Academic Librarians Public Librarians Reference Librarians Special Librarians
RT	Acquisitions Cataloguing Collections Development Information Services

The labels used would vary among the displays since not all descriptors may have *parts* or *tools*, but then again, not all descriptors have BTs, NTs, or RTs. The thesaurus could be designed to allow users to sort displays according to type of relationship, e.g., all *activities* associated with the term *librarian*.

Reaching true consensus in the design of thesaurus displays is a near-impossible task due to the potential variety within the population served. The admittedly limited application of the word association test in this study has, however, provided a degree of consensus among the participants, which suggests that there would be merit in conducting further studies with larger numbers of participants and with more varied populations. This study has not attempted to measure the potential impact that the social and ethnographic composition of the participants could have on the latter's selection of response terms and of their determination of inter-term relationships. Further studies in this area could provide interesting and valuable insight into the degree to which term hierarchies may be affected by cultural, educational, and social factors.

APPENDIX 1

Incidence of overlap in unique response terms per stimulus term

Stimulus Term	Response Term	Frequency
Authority control	Standardization	21%
Cataloguing	AACR2	28%
	<u>Cataloging</u>	21%
	Organization	21%
	Classification	21%
	MARC	21%
Censorship	Banned books	43%
	Book burning	29%
	Intellectual freedom	21%
Digital libraries	Digital collections	30%
	Virtual libraries	30%
Intellectual property	Copyright	82%
	Patents	43%
	Trademarks	36%
Librarians	Information professionals	21%
	Information specialists	21%
Library science	Information science	45%
	Library schools	24%
	Library studies	24%
Reference materials	<u>Encyclopedias</u>	71%
	Dictionaries	54%
	Books	32%
	Directories	32%
	Almanacs	21%
	Atlases	21%
	Reference collections	21%
Special collections	Rare books	46%
	Archives	38%
	Manuscripts	24%
Technical services	Cataloguing	74%
	Acquisitions	56%
Thesauri	Controlled vocabulary	32%
	Synonyms	32%
	Subject headings	25%
	Descriptors	21%

References

- Battig, W. F., & Montague, W.E. (1969). Category norms for verbal items in 56 categories: A replication and extension of the Connecticut Category Norms. *Journal of Experimental Psychology Monograph* 80(3), Part 2, 1-46.
- Deese, J. (1965). *The structure of associations in language and thought*. Baltimore , MD : The Johns Hopkins Press.
- Galton, F. (1880). Psychometric experiments. *Brain*, 2, 49-162.
- ISO. (1986). *Documentation-guidelines for the establishment and development of monolingual thesauri*. ISO 2788:1986. International Organization for Standardization.
- Kent, G. H., & Rosanoff, A. J. (1910). A study of association in insanity. *American Journal of Insanity*, 67, 37-96, 317-390.
- Miller, G., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. (1993). *Introduction to WordNet: An on-line lexical database*. Retrieved April 27, 2004 , from <http://www.cogsci.princeton.edu/~wn/obtain/5papers.pdf>
- Nielsen, M. L. (1997). The word association test in the methodology of thesaurus construction. *Proceedings of the 8 th ASIS SIG/CR Classification Research Workshop*, 43-58. Washington , DC : American Society for Information Science.
- NISO. (1994). *Guidelines for the construction, format, and management of monolingual thesauri*. ANSI/NISO Z39.19-1993 . Bethesda , MD : National Information Standards Organization.
- OpenCyc selected vocabulary and upper ontology . 2002. Retrieved April 27, 2004 , from <http://www.cyc.com/cycdoc/vocab/vocab-toc.html>
- Russell, W.A. (1970). The complete German language norms for responses to 100 words from the Kent-Rosanoff Word Association Test. In L. Postman & G. Keppel (Eds.), *Norms of word association*, 53-94. New York : Academic Press.
- Schultz, D. P., & Schultz, S. E. (2000). *The history of modern psychology*. Seventh edition. Harcourt College Publishers.
- Smith, B., & Mark, D. (1999). Ontology with human subjects testing: An empirical investigation of geographic categories. *American Journal of Economics and Sociology*, 58(2), 245-272.
- Spiteri, L. F. (2002). Word association testing and thesaurus construction: Defining inter-term relationships. In L. Howarth, C. Cronin, & A. Slawek (Eds.), *Advancing knowledge: Expanding horizons for information science. Proceedings of the 30 th Annual Conference of the Canadian Association for Information Science, 30 May-01 June 2002* . Toronto , ON : Faculty of Information Studies, University of Toronto .

Theory-Frame Ontology . (1997). Retrieved April 27, 2004 , from
<http://www.ksl.stanford.edu/people/brauch/demo/frame-ontology>

Tversky, B., & Hemenway, K. (1983). Categories of environmental scenes. *Cognitive Psychology* , 15(1), 121-149.

Wettler, M., & Rapp, R. (1996). *Computation of word associations based on the co-occurrence of words in large corporations*. Retrieved April 27, 2004 , from
<http://www.fask.uni-mainz.de/user/rapp/papers/wvlc93/latex2html/wvlc93.html>