

Usage notes
as a representation of lexical near-equivalence
for lexical choice

(*or* Making words of senses)

Chrysanne DiMarco
Department of Computer Science
University of Waterloo
Waterloo, Ontario
Canada N2L 3G1
519-888-4443; fax 519-885-1208
cdimarco@logos.uwaterloo.ca

and

Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, Ontario
Canada M5S 1A4
416-978-8747; fax 416-978-1455
gh@cs.toronto.edu

April 1993

Submitted to the 9th Annual Conference of the University of Waterloo Centre for the
New OED and Text Research, September 1993

Abstract

The task of choosing between lexical near-equivalents in text generation requires the kind of knowledge of fine differences between words that is typified by the usage notes of dictionaries. We discuss the structure and content of usage notes, give an example of the application of a usage note in lexical choice, and develop the concept of a formal, computational ‘usage note’.

1 Lexicons for lexical choice

1.1 Lexical choice and plesionymy

The problem of *lexical choice* in text generation is to determine the word that conveys most precisely the denotation and connotation that are to be expressed. The required meaning may be specified in either a language-independent knowledge representation or, in the case of non-interlingual machine translation, another natural language. (In the latter case, the problem is often known as *lexical transfer*.) Sometimes, the language will provide no suitable word, and a phrase will have to be constructed; at other times, it will offer many similar words, and the problem is one of discriminating between lexical near-equivalents. The goal of this work is to develop a representation for the lexical information that computational systems need in order to perform this discrimination.

Our specific interest is in representing the nuances that distinguish groups of *near-synonyms*, or *plesionyms* (Cruse 1986). Plesionyms are not fully inter-substitutable, as they may vary in their shades of denotation or connotation, or in the components of the meaning that they emphasize. (They may also vary in their grammatical or collocational constraints, but these won’t concern us in this paper). Our approach is to study, formalize, and extend the kinds of information about lexical differentiation that appears in dictionary usage notes or books of synonyms.

1.2 The conceptual domain and the lexical domain

Our goal is a representation for a lexicon in which semantic and stylistic distinctions can be made between synonyms and plesionyms, both within and across languages. The central idea is that coarse denotational differentiation occurs at the language-independent conceptual level, and connotational and fine denotational differentiation occurs at the language-dependent level, in the lexical entries themselves.

Our starting point is a familiar idea: a conventional KL-ONE-style taxonomic knowledge base that represents the conceptual knowledge of the system, and hence the basic conceptual distinctions made by words in all the languages under consideration. The relations used in the KB derive from standard semantic case theory, and sentences are represented as usual: as configurations of concepts and the relations that hold among them.

Lexical entries for words in the language or languages of interest are associated with the corresponding concepts in the KB; two synonyms (in the same or different languages) will map to the same concept. More precisely, lexical entries are associated with *configurations*

of a concept and various roles and fillers, in order to represent more finely grained semantic distinctions than those made by the concepts only: similar lexical items all map onto the same, fairly general, semantic predicate, and the associated roles and fillers represent the smaller denotational differences (for details, see Stede 1993 or DiMarco, Hirst, and Stede 1993).

However, the strength of a purely conceptual approach is also an inherent weakness: the differences between plesionyms must be represented as differences between concepts, and this is not always easy or natural. But there has to be some place at which we slip from concepts to words. Our proposal here is that it should be earlier rather than later. That is, the conceptual hierarchy is fairly coarse-grained—in effect, it records relatively language-independent concepts—and the fine tuning is then done in the lexical entries for each separate language.

1.3 An operational view of synonymy and near-synonymy

Given this representation, we can sidestep for our purposes the question that has long been debated by linguists and lexicographers as to exactly what degree of inter-substitutability should count as synonymy or near-synonymy (*cf.* Egan 1942, Sparck Jones 1986, Cruse 1986). For us, synonymy and near-synonymy arise at the point in the conceptual network at which a (language-independent) concept diverges into the set of (language-dependent) lexical entries for the words that, in one way or another, denote that concept. That is, the groups of words among which we need to discriminate are exactly the groups of words (in each language of interest) that correspond to each single concept in the taxonomic hierarchy of the knowledge base of the system. For convenience, we shall refer to each group as a set of plesionyms or near-synonyms, but we intend by these labels no theoretical import beyond that of this operational definition.

2 Plesionym discrimination in dictionaries

2.1 Information for discrimination from dictionary definitions

The task of discriminating among synonyms and near-synonyms requires extensive lexicographic information, but little of this is explicitly given in conventional dictionaries. Lexicographers have traditionally placed their emphasis on creating comprehensive lists of words of the language and distinguishing the separate senses of each word. What we need is almost the exact opposite: a comprehensive list of the senses of a language, distinguishing the separate words of each.¹ Nevertheless, at least some of the information that we require for distinguishing between near-synonyms is explicit, or nearly so, in the entries in a conventional dictionary. For example, the differences between *gawp*, *gaze*, and *stare*, all of which denote a kind of prolonged look, may be found from their entries in the *Oxford advanced learner's*

¹So-called 'reverse' or 'thematic' dictionaries (*e.g.*, Kahn 1990, Glazier 1992) try to do something like this; but of course, the 'senses' by which they are organized are really just alphabetized words. A thesaurus such as *Roget's* imposes a modest conceptual hierarchy on words, but does not list senses *per se*.

gawp stare impolitely or stupidly.

gaze look long and steadily (at sb/sth), usu in surprise or admiration.

stare look (at sb/sth) with the eyes wide open in a fixed gaze (in astonishment, wonder, fear, etc).

Figure 1: Definitions (abridged) of *gawp*, *gaze*, and *stare* from the *OALD*.

dictionary (*OALD*; fourth edition, 1989) (see Figure 1). Specifically, we find that to gaze is to look long and steadily; to stare is to do this with the eyes wide open; and gawping has the additional requirement that the act be impolite or stupid.

2.2 Usage notes

Where dictionaries *do* explicitly discriminate between near-synonyms is in their occasional usage notes and lists of synonyms.² An example of a usage note from the *OALD* is shown in Figure 2. But lexicographers do not seem to hold such notes in high esteem, regarding them perhaps more as a marketing gimmick than as true lexicographical scholarship, for they occur mainly in dictionaries aimed at the language-learner and college-student market, while the *Oxford English Dictionary* has none at all! Moreover, the number of notes is relatively small compared to the size of the dictionaries; a choice is made of a relatively small number of cases where discrimination is thought to be potentially troublesome for humans.³ In addition, there are on the market a number of books whose sole purpose is the discrimination of near-synonyms; that is, they contain usage note-style discrimination of many more groups of near-synonyms than dictionaries do. However, many of these books seem to be intended as much for entertainment as information (*e.g.*, *When is a pig a hog?*, Randall 1991), and even the largest can make no attempt to cover all the plesionym groups of the language.⁴

It is interesting to note that, even within a single dictionary, information in a usage note can be quite different from, and indeed sometimes contradictory to, that in the definitions. Compare Figure 2 with Figure 1. Unlike the definitions, the usage note says that gawping requires an open mouth, but doesn't mention that it is a subclass of staring; the definition of *stare* says that staring involves gazing, though the usage note implies that the two are distinct; and for *gaze*, the definition emphasizes surprise or admiration, while the usage note

²Some usage notes, of course, cover other aspects of language, such as grammar and pronunciation, that do not concern us here. And the *American heritage dictionary of the English language* (third edition, 1992) uses the term *synonym paragraph*, reserving *usage note* for remarks mostly concerning grammar and 'correctness'.

³In the *OALD*, with 27,000 entries, there are just 200 usage notes; about half are discriminations of near-synonyms, each distinguishing, typically, four or five words. The *Longman dictionary of contemporary English* (second edition, 1987) has 400 usage notes; the *American heritage* has 900 synonym paragraphs (not all of which include discriminations), as well as occasional information about nuance in its usage notes.

⁴*Modern guide to synonyms and related words* (Hayakawa 1968) claims to offer "more than 1000 essays [comparing] 6000 words"; *Webster's new dictionary of synonyms* (Gove 1984) makes no claim as to size, but by our estimate is about twice as big again.

look. **1 Look (at)** means to direct one's eyes towards a particular object: *Just look at this beautiful present.* ◦ *I looked in the cupboard but I couldn't find a clean shirt.* **2 Gaze (at)** means to keep one's eyes turned in a particular direction for a long time. We can gaze at something without looking at it if our eyes are not focussed: *He spent hours gazing into the distance.* ◦ *She sat gazing unhappily out of the window.* **3 Stare (at)** suggests a long, deliberate, fixed look. Staring is more intense than gazing, and the eyes are often wide open. It can be impolite to stare at somebody: *I don't like being stared at.* ◦ *She stared at me in astonishment.* **4 Peer (at)** means to look very closely and suggests that it is difficult to see well: *We peered through the fog at the house numbers.* ◦ *He peered at me through thick glasses.* **5 Gawp (at)** means to look at someone or something in a foolish way with the mouth open: *What are you gawping at?* ◦ *He just sits there gawping at the television all day!*

Figure 2: Usage note for *look* from the *OALD*.

highlights the possibility of unfocused gazing.

3 A study of usage notes

Both the structure and content of the usage notes of dictionaries and books of synonyms provide a great deal of useful information about lexical discrimination that we wish to adapt for computational use.

3.1 The structure of usage notes

While the style and length of usage-note entries varies somewhat (for example, the discriminations of Hayakawa (1968) are longer and more like essays), the following structure is characteristic:

- A statement of the meaning that is central or common to the set of words being discriminated. This is sometimes omitted in the *OALD* and *LDOCE*, or is implicit in the choice of the headword under which the note appears; the *OALD* often speaks of the *most neutral* word, the *most general*, the *most usual*, the one with the *widest use*.
- A description of the factors that distinguish each word in the set (*cf.* Egan 1942):
 1. Implications: Denotational differences between the meanings of words.
 2. Connotations: Nuances that 'colour' a word's meaning.
 3. Applications: Restrictions on a word's use.
- Examples of the use of each word in the set.

(An individual entry might also list analogous words, contrasted words, and antonyms; these do not concern us here.)

3.2 The content of usage notes

The descriptions of distinguishing factors follow a style or ‘language’ particular to the notes. The elements of the language include the denotative and connotative dimensions and features that we describe below, an infinite (but constrained) class of emphases, and a set of ‘operators’ such as *most general*, *most usual*, *mostly used*, *not normally used*, *neutral word*, *strong*, *emphasizes*, *suggests*, and *usually associated with*.

In a study described more fully by DiMarco, Hirst, and Stede (1993), we classified the factors that were given in the *OALD* and *LDOCE* to explain the differences between the words covered by each note. We observed that there were certain dimensions that were used quite frequently as denotative or connotative differentiae. Altogether, we noted 26 such dimensions for denotation and 12 for connotation (including a few that we added from the discussion of Vinay and Darbelnet (1958)). (We don’t, of course, claim this set to be complete or definitive.) Some of the dimensions are simple binary choices; others are continuous. Some examples are listed in Figure 3. Each line of the figure shows a dimension of differentiation (named, in most cases, for its endpoints), followed by example sentences in which two plesionyms or synonyms vary along that dimension. We have tried to show ‘pure’ examples, but often, of course, pairs of words will vary in several features simultaneously.

It should be noted that these lexical features for differentiation are not intended to be any kind of primitives for decompositional semantics. We are not using them to represent whole meanings, but rather to represent *differences* between meanings.

4 An example of lexical choice with the help of a usage note

Let us consider the process that would be followed in employing a usage note to discriminate between lexical near-equivalents. For example, the lexical-choice problem might occur within a machine translation system; for this kind of task, each lexical choice should be the one that best preserves the intent of the original source text.

From the following examples, we observe that, even in everyday writing, the same source-language word can be translated by different target-language words, according to the particular translation situation. In the following pairs of sentences, the word *cuts* has the same denotation in each case, but the translator has chosen *coupes* in one place, emphasizing the size of the cuts, and *réductions* in the other, emphasizing the result of cutting:

- (1) (a) The Hon. Don Mazankowski’s first budget featured huge cuts in the transfer payments for post-secondary education.
- (b) Le premier budget de l’honorable Don Mazankowski prévoit des coupes sombres dans les paiements de transfert au titre de l’enseignement postsecondaire.⁵

⁵ *CAUT/APCU bulletin*, 39(4), April 1992, p. 1.

DENOTATIONAL DIMENSIONS	CONNOTATIVE DIMENSIONS
Intentional/accidental: She { <i>stared at</i> <i>glimpsed</i> } him through the window.	Formal/informal: He was { <i>inebriated</i> <i>drunk</i> }.
Continuous/intermittent: Wine { <i>seeped</i> <i>dripped</i> } from the barrel.	Abstract/concrete: The { <i>error</i> <i>blunder</i> } cost him dearly.
Immediate/iterative: She { <i>struck</i> <i>beat</i> } the drum.	Pejorative/favorable: That suit makes you look { <i>skinny</i> <i>slim</i> }.
Sudden/gradual: The boy { <i>shot</i> <i>edged</i> } across the road.	Forceful/weak: The building was completely { <i>destroyed</i> <i>ruined</i> } by the bomb.
Terminative/non-terminative: Elle { <i>fripa</i> <i>chiffona</i> } la chemise.	Emphasis: I { <i>arranged</i> <i>organized</i> } a meeting of the committee.
Immediate/iterative: She { <i>crumpled up</i> <i>crumpled</i> } the note.	He { <i>cried</i> <i>wept</i> } in pain.
Emotional/non-emotional: Their { <i>relationship</i> <i>acquaintance</i> } has lasted for many years.	They had been { <i>enemies</i> <i>foes</i> } for many years.
Degree: We often have { <i>mist</i> <i>fog</i> } along the coast.	

Figure 3: Examples of features that dictionary usage notes adduce in word differentiation.

- (2) (a) To date the cumulative loss (up to 1992–93) in transfers will be \$4.8 billion. The Finance Minister did not announce these cuts in his speech.
 (b) Á ce jour et jusqu’en 1992–1993, la perte cumulative s’élèvera à 4.8 milliards de dollars. Le ministre des Finances n’a pas annoncé ces réductions dans son discours du budget.⁶

As a particular scenario, we will suppose that we need to choose from among English *looking* words that will translate the French *regarder* in a newspaper article dealing with a case of possible sexual harassment. Thus, the following sentence might occur in this context:

- (3) Elle pensait qu’il la regardait fixement.

According to Atkins *et al* (1978), Bailly (1970), and Bénac (1956), *regarder*, by itself, can suggest deliberate looking, or, combined with modifiers, can suggest various nuances related to time and intention. So, choosing whether *regarder* is translated by simple *looking*, close *peering*, or intense *staring* will convey subtly different meanings.

We will employ the usage note given in Figure 2 above to help us choose between *look*, *gaze*, *stare*, *peer*, and *gawp*. Our starting point is *look*, the central, or unmarked, word in this set; it is the canonical translation of *regarder*. The characterization of *look* as a directing of one’s

⁶ *ibid.*

eyes towards a particular object provides an initial *discrimination context* that, by default, will be a suitably neutral choice for the given situation. We can think of a discrimination context as a description, at different levels of abstraction, of the lexical-choice situation. In this description, we might store information about:

- Sentential context of the word being translated. We will look at the characteristics of other words in the sentence or surrounding text to see whether they suggest nuances that can help us to discriminate among the set of near-synonyms. We would look to the usage notes for these neighbouring words to obtain such semantic and stylistic information.
- Global context. We will also need to look at the properties of the events and entities described in the text. Such information would be obtained from the knowledge base that, as we described in Section 1.2, forms the basis of our lexical representation.

As we continue to evaluate possible lexical choices, we will update the discrimination context, to build up a description of the most appropriate choice for the given situation.

After *look*, the next possibility in the usage note is *gaze*. We note that *gaze* has an additional nuance denoting an extended period of looking. We would examine the current sentential context to see whether we can determine that the action of looking was prolonged. Cue words or phrases (*e.g.*, *fixement*, *longuement*, *attentivement*) might occur in the sentential context, or we might infer, directly or indirectly, the length of the looking from the global context (*e.g.*, there might be a closely related event that implies this conclusion, such as the looker remaining in the same physical location for some time). If we can determine, as in this particular case, that the looking *was* prolonged, then we would update the discrimination context by adding this new evidence.

In contrast to the neutrality of *look* and the simple nuance of *gaze*, *stare* can imply several different kinds of denotative and connotative aspects. The action of looking must be “long”; again, this could be inferred from the sentential context, as in the example, or the global context. The looking may also be “deliberate” or “fixed”. These denotative distinctions can be more difficult to verify than simple length of time. We might check the sentential context for cue words, but there might be no overt assertion of the deliberateness of the action. We might then try to reason from information in the global context—for example, can we infer from the surrounding text that the writer intends to imply that the looker made a conscious decision to move to a position close to the person? We know that staring is “more intense” than looking or gazing, and sometimes even impolite. What information from the sentential or global contexts might suggest that impoliteness is denoted?

Continuing in the usage note, we know also that peering can suggest difficulty of seeing well. If the looker is myopic or known to wear thick glasses, then we could surmise that he *peered* at the person.

Thus, some of the information about the appropriateness of a word may be available directly from the sentential context, but other information must be inferred from the global context. As we work through the cases in the usage note, making reasoned assumptions and inferences, we gradually build up evidence. Eventually, we will produce a ‘portrait’ of the

word most suitable for the given situation. In the example given, we would choose *stared* as the translation of *regardait*, given the information that the looking was done fixedly, but, if we had additional knowledge that the looker was extremely near-sighted, then *peered* would be the more appropriate choice.

But it is not always the case that we will want to choose the most specific word; sometimes, the unmarked word will be the most appropriate choice. One reason for this is that we may be dealing with uncertain and inexact information. It may be difficult to determine, with strong certainty, that a condition implying a certain lexical nuance does exist. The evidence we build up during the consideration of the various lexical choices may not be conclusive, and so the unmarked word may be the best fit for the given situation.

From this example of employing a usage note for lexical discrimination, it is apparent that the process of choosing among lexical near-equivalents can be characterized by the following specifications:

- Various kinds of information about the lexical-choice situation are needed, including sentential and global contexts.
- A discrimination context is built up as the various lexical choices are considered.
- Assumptions and inferences using the various levels of context provide evidence to update the discrimination context.
- These inferences may be based on incomplete and inexact information.

5 The formal, computational usage note

How could this kind of lexical choice be automated for use in computational text-generation and machine translation systems as described in Section 1 above? The requirements fall into three classes: data, representation, and process.

The data requirement is that for each language of interest, for each concept in our hierarchy that has more than one word attached, except in the rare case of absolute inter-substitutability there must be some kind of ‘computational’ or ‘formal’ usage note that discriminates among the words. This information is essentially lexicographic, and could be based, at least in part, on the usage notes that already exist in dictionaries and books of synonym discrimination. Section 6 below discusses the lexicographic implications of this.

The representation and process requirements are that this data be included in the lexical entries in a form that the system is able to use. Literal natural-language usage notes are obviously not suitable; nevertheless, the representation must be able to express essentially the same information.

Our study of denotative and connotative lexical discrimination features in Section 3.2 above is a start in determining exactly what needs to be included. Given a set of features such as these, a very natural representation is a *discrimination tree* or *network* (much like choosers in the Penman project (1989)), or, perhaps, more generally, a set of *pattern-action rules*. But in practice, this is too simplistic. First, such a method assumes that we always

wish to use the most specific word available, and we saw in Section 4 above that this is not always so. Second, as we also showed, lexical choice often involves a trade-off among a set of imperfect choices after all have been evaluated; this is not easily modelled in such systems. We are continuing, therefore, to study possible representations and their associated decision processes.

6 Implications for lexicography

If our model of lexical choice is to be practical, we need formal usage notes for most concepts in our knowledge base. However, the usage notes of conventional dictionaries and dictionaries of synonyms are little more than a conceptual starting point for this.

First, no set of usage-note data is complete; even Gove 1984 omits many words. Creating the union of the various data sets would be a major lexicographic task, and there would still, of course, be gaps. Second, the notes in their present natural language form are, in general, not amenable to automatic conversion to any kind of computationally usable representation. Third, the set of discriminations that is required is completely dependent on the particular knowledge base. A lexical discrimination centred upon a single concept cannot be easily adapted to centre upon a separate, related concept for a knowledge base that carves the world up in a slightly different way. We can expect to see, in the not-too-distant future, the development of some large, general, and widely used taxonomic hierarchies for AI and natural language systems; it is only with the coming of such ontologies that the effort of constructing accompanying lexical discriminations would be worthwhile (and it is in anticipation of that day that we undertake this research).

Nevertheless, it is clear that representing lexical discriminations for lexical choice is an important problem, and that dictionary usage notes as they presently exist give both a large kernel of data upon which to build, and insight into the nature of lexical discrimination itself. We believe that with the development of large taxonomic hierarchies and systems for lexical choice, the construction of large-coverage sets of lexical discriminations for text-generation systems will become an important area of research in lexicography.

Acknowledgements

For discussion, encouragement, inspiration, clever ideas, or good advice, we are indebted to Manfred Stede, Frank Tompa, Eduard Hovy, and Anne Marie Miraglia. Our research is supported by grants from the Natural Sciences and Engineering Research Council of Canada.

References

American heritage dictionary of the English language, third edition. Houghton Mifflin, 1992.
Atkins, Beryl T. *et al* (1978). *Collins-Robert French–English English–French dictionary*. London: Collins and Paris: Société du Nouveau Littré.

- Bailly, René (1970). *Dictionnaire des synonymes de la langue française*. Paris: Librairie Larousse.
- Bénac, Henri (1956). *Dictionnaire des synonymes*. Paris: Librairie Hachette.
- Cruse, D.A. (1986). *Lexical semantics*. Cambridge University Press.
- DiMarco, Chrysanne; Hirst, Graeme; and Stede, Manfred (1993). “The semantic and stylistic differentiation of synonyms and near-synonyms.” *AAAI Spring Symposium on Building Lexicons for Machine Translation*, Stanford, March 1993, 114–121.
- Egan, Rose F. (1942). “Survey of the history of English synonymy” and “*Synonym: Analysis and definition*”. Reprinted in Gove 1984, pp. 5a–25a.
- Glazier, Stephen (1992). *Random House word menu*. New York: Random House.
- Gove, Philip, editor (1984). *Webster’s new dictionary of synonyms*. Springfield, Mass.: Merriam-Webster.
- Hayakawa, S.I., editor (1968). *Funk & Wagnall’s modern guide to synonyms and related words*. New York: Funk & Wagnall. Reprinted as *Use the right word: Reader’s Digest modern guide to synonyms and related words*, Reader’s Digest Association, no date. Reprinted as *Choose the right word: A modern guide to synonyms*, New York: Harper & Row (Perennial Library), 1987. Edition for British English published as *Cassell’s modern guide to synonyms* (revised by Paul J. Fletcher), London: Cassell, 1971.
- Kahn, John Ellison, editor (1990). *Illustrated reverse dictionary*. Pleasantville, NY: Reader’s Digest Association. [Based on his *Reader’s Digest reverse dictionary*, London: Reader’s Digest Association; Americanized and expanded by Walter D. Glanze.]
- Longman dictionary of contemporary English*, second edition. Harlow: Longman Group, 1987.
- Oxford advanced learner’s dictionary of current English*, fourth edition. Oxford University Press, 1989.
- The Oxford English dictionary*, second edition. Oxford University Press, 1992.
- Penman project (1989). The Penman documentation. Information Sciences Institute, University of Southern California.
- Randall, Bernice (1991). *When is a pig a hog? A guide to confoundingly related English words*. New York: Prentice-Hall.
- Sparck Jones, Karen (1986). *Synonymy and semantic classification*. Edinburgh: Edinburgh University Press (Edinburgh information technology series 1).
- Stede, Manfred (1993). “Lexical choice criteria in language generation.” *Proceedings, Sixth conference of the European Chapter of the Association for Computational Linguistics*, Utrecht, April 1993, ???–???
- Vinay, J.-P. and Darbelnet, J. (1958). *Stylistique comparée du français et de l’anglais*. Montreal: Beauchemin.