Language and Cognitive Processes, Vol. 1, No. 3, pp. 171–185 (1986)

Dictionaries in the mind*

GEORGE A. MILLER
Psychology Department, Princeton University, Princeton, NJ 08544, USA

Received June 1985; accepted in revised form September 1985

Abstract—How lexical knowledge is acquired, and how it is organized in memory for rapid retrieval during language use, are central questions for congitive psychologists. Research into these questions has revealed interesting differences and similarities between the subjective dictionaries in our heads and the objective dictionaries on our shelves. The differences might be reduced in the future by publishing dictionaries as computer programs.

INTRODUCTION

When you learn a language, you must learn its pronunciation, vocabulary, grammar, and usage, each of which is characterized in a separate body of linguistic theory. All of these theories—phonology, syntax, semantics, pragmatics—describe things that a beginner must learn and that a skillful speaker must know.

Now, what people know and how they learn it are basic questions for cognitive psychology; when those questions concern what people learn and know about languages, we call it psycholinguistics.

In principle, a psycholinguist should be interested in all these kinds of learning. In fact, however, lexical learning is usually dismissed as uninteresting, as little more than a tedious exercise in memorizing a long list of conventional associations between sounds and meanings. That evaluation has been a mistake: there are many fascinating lexical puzzles worth studying.

What does it mean to know the vocabulary of a language? A naïve answer might run as follows:

(1) *Lexical knowledge*: To know the vocabulary of a language is to know the sounds and meanings of all important words in the language.

There are many difficulties with (1)—the difficulties of determining what a word is, what a meaning is, what makes a word important—but the principal concern of the present discussion is the deeper problem of what it means to 'know' any lexical element.

The operational test for knowing a word eventually comes down to satisfying one or both of the following two criteria:

(2) *The receptive criterion*: You know a word if you can respond appropriately when you hear it used.

(3) *The productive criterion*: You know a word if you can use it naturally and appropriately in sentences that express your own thoughts.

*This paper was first presented in a revised form as the author's Lewis Clark Vanuxem Public Lecture at Princeton University, 29 March 1984. (A shortened version was read at the *35th Annual Georgetown University Round Table in Languages and Linguistics*, Washington, D.C., 16 March 1984.)

For now, let us avoid specifying what 'appropriately' entails in (2) and (3). For a beginning, it is enough to sketch a domain that is worthy of study; it is not necessary set down formal axioms from which the domain could be developed.

THE RECEPTIVE CRITERION

To be able to satisfy the receptive criterion for knowing a word is to be able to recognize the word and retrieve its meaning from lexical memory. For some words, proof that you know a word can be given by pointing to an instance. More often, however, proof involves giving (or at least recognizing) the word's definition. Consequently, retrieving a word's meaning from memory is commonly likened to looking it up in a dictionary, an analogy that implicitly presupposes:

(4) *The basic psycholexical analogy*: A language user's lexical memory is organized into independent lexical entries the way a printed dictionary is.

Of course, with a printed dictionary you must use alphabetical retrieval, whereas you can get access to your own lexical knowledge by way of either sound or meaning. Except for that, however, a printed dictionary is commonly assumed to be an explicit theory of people's lexical knowledge. In other words, all competent speakers of English are assumed to have an English dictionary in their hands.

When you consider (4) more carefully, however, you begin to realize how improbable it is. Lexicographers do not think about psychology—why should anyone assume that they create psychological theories? The way a lexicographer goes about writing a dictionary is not at all the way a psychologist would go about constructing a psychological theory. And it is certainly not the way a child goes about learning a vocabulary.

What, then, is the relation between the objective dictionaries on our shelves and the subjective dictionaries in our heads? The analogy (4) contains the beginning of an answer, but it will have to be formulated more abstractly before it can be accepted.

A printed dictionary is simply a lexical database. The language in which to discuss databases is the language of data retrieval systems—a technical language in terms of which we might formulate a general theory of lexical data management. Then we could reformulate the basic psycholexical analogy along the following lines:

(5) *The data processing hypothesis*: A language user's lexical memory and a printed dictionary are both specific instantiations of some more general theory of the storage and processing of lexical information.

Can this formulation be tested? Since the nature of lexical memory is still uncertain, let's begin with printed dictionaries, whose nature is reasonably familiar. Given an abstract characterization of a printed dictionary, we could then consider whether it also fits, or could be adapted to, the little that is known about lexical memory.

Characterizing dictionaries abstractly is a matter for lexicology, and so is part of the general theory that underlies the practical art of lexicography.

LEXICOLOGY

In its essence, a printed dictionary is simply a mapping of senses onto words, and a mapping can be conveniently represented as a matrix. I have called it a *vocabulary matrix*.

Imagine a huge matrix with all the words in a language along the top of the matrix, and all the different senses that those words can express down the side. If a particular sense can be expressed by a word, then the cell in that row and column contains an entry; otherwise it contains nothing. The entry itself can provide syntactic information, or examples of usage, or even a picture—whatever the lexicographer deems important enough to include. Table 1 contains a fragment of a vocabulary matrix by way of illustration.

Table 1.
Illustrative fragment of a vocabulary matrix. Columns represent modal verbs; rows represent modal senses; 'E' in a cell means that the word in that column can be used to express the senses in that row

Senses	Words can	may	must	should	will
be able to	E	·	·	·	·
be permitted to	E	E	·	·	·
be possible	E	E	·	·	·
be obliged to	·	·	E	·	·
certain to be	·	·	E	·	E
be necessary	·	·	E	E	·
expected to be	·	·	E	E	E

Several comments can be made about the vocabulary matrix. First, it should be apparent that any printed dictionary can be represented as a vocabulary matrix: simply add a separate column to the matrix for every word, and add a separate row to the matrix for every sense of every word that is given in the printed dictionary. A lexical matrix can be viewed as an impractical way of printing a dictionary on a single, very large sheet of paper.

Second, entering such a matrix consists of searching down some column or across some row. So a vocabulary matrix can be entered either with a word or with a sense. Thus, one difference between printed dictionaries, which can only be entered with a word, and subjective lexicons, which can be entered with either words or senses, disappears when dictionaries are represented in this more abstract form.

Third, if you enter the matrix with a sense and search along a row, you find all the words that can express that sense. When different words express the same sense, we say they are *synonymous*. On the other hand, if you enter the matrix with a word and look down that column, you find all the different senses that that word can express. When one word can express two or more senses, we say that it is ambiguous, or *polysemous*. Thus, two major complications of lexical knowledge, synonymy and polysemy, are seen as complementary aspects of a single abstract structure.

Finally, since the vocabulary matrix serves only to represent the mapping between the two domains, it is free to expand as new words, or new senses for familiar words, are added. Of course, the number of columns is relatively fixed by the size of the vocabulary, so a theorist's major degrees of freedom are in deciding what the senses are and how to represent them.

Most printed dictionaries have a great deal of *pseudopolysemy*—that is to say, the number of different senses that a word can express appears to be inflated, in

part because good lexicography is compulsive lexicography and the compulsive lexicographer tries to draw distinctions as fine as possible. On the other hand, synonymy tends to be underestimated, at least in part because the same sense is expressed slightly differently when related words are defined independently, thus inflating the number of apparently different defining phrases in the matrix.

If the vocabulary matrix can be accepted as an abstract characterization of printed dictionaries, then the next question is whether it can also provide an acceptable characterization of the subjective lexicon.

An obvious and immediate objection to such a proposal is that a vocabulary matrix cannot escape the world of words—in such a matrix, single words are related to the strings of words that define them. The lexical concepts that are associated with words in our mental lexicon do not have such a wordy feel to them. Indeed, it is often difficult to produce a passable definition for a common word whose meaning you know perfectly well. It is by no means clear what stands in for defining phrases in our mental dictionaries, but whatever it is, it somehow enables us to break out of the circle of words and to denote things, people, situations, or events in the world around us.

One's immediate impulse, therefore, is to deny that a vocabulary matrix could provide a satisfactory description of the subjective lexicon. Nevertheless, something like a vocabulary matrix might suffice if only the relation between defining phrases and lexical concepts could be satisfactorily resolved. Before rejecting it out of hand, therefore, we should consider the differences between objective and subjective dictionaries in more detail. Let us begin with a discussion of the growth of vocabulary in children.

THE GROWTH OF VOCABULARY

Recent studies of cognitive development have not totally ignored vocabulary, but it is fair to say that the principal interest in vocabulary has been as a convenient window on conceptual development. Interest in the words themselves has been secondary.

It was not always so. During the 1930s and 40s, no doubt stimulated by the observation that vocabulary is the most dependable part of any intelligence test, psychologists devoted considerable attention to the growth of vocabulary in children. It was found that the first word appears between 12 and 18 months of age, then new words begin to appear more and more rapidly until by the age of three a parent can no longer tell whether the child knows a particular word. Thereafter, vocabulary grows so rapidly that dictionary sampling techniques are required to keep track of it.

Several careful, cross-sectional studies of vocabulary growth were conducted. Consider some data that were collected by Mildred Templin and published in 1957. According to Templin, a six-year-old child of average intelligence knows 13000 words and the average eight-year-old child knows 28300 words.

In order to appreciate these numbers, it helps to convert them into words learned per day. Between six and eight years of age an average child learns 21 words per day. Children of superior intelligence probably learn words at twice that rate.

This estimate seems high, but other studies of vocabulary growth give even higher rates of learning. Apparently, a fantastically broad and rapid learning

process goes on during childhood and psychologists know very little about it. At least two things are clear, however:

(6) *Hypothesis of lexical sensitivity*: Young children are alert to new words and very good at acquiring vocabulary.

And, since nobody teaches children 21 words every day:

(7) *Hypothesis of contextual learning*: Young children are able to educe meanings of words from the contexts in which they hear them used.

In 1974 Carey and Bartlett invented a way to observe this rapid learning in more detail (Bartlett, 1977; Carey, 1978; Carey and Bartlett, 1978; Miller, 1977). They used color names. First, they established that a group of three-year-old children did not know the color olive; most called it green, some called it brown. Carey and Bartlett decided to teach the children a nonsense name for olive, a name that they could not have heard anywhere but in our nursery school. So they painted one tray olive and another tray blue and asked each child casually. 'Hand me the chromium tray. Not the blue one, the chromium one.' The child would pause, perhaps point to the olive tray and ask, 'This one?' A week later, with no further guidance, the children were again asked to name the colors. This time when olive was presented they paused. They couldn't remember the nonsense name for it, but now they knew it wasn't green or brown. A single exposure was enough to begin a reorganization of their color lexicon.

Carey thought she saw a two-step learning process:

(8) *A two-step learning process*: First, young children notice a new word and immediately assign it to a semantic category. Second, they discover and learn the distinctions among words assigned to the same category.

After hearing 'chromium' just once, for example, they had assigned it to the semantic field of color names. Children are very good at keeping these fields separate, even when they don't really know what the words mean (Dockrell, 1983). If asked their age, for example, they almost never say 'round' or 'lunch'. They answer with a number, a word from the appropriate category. Prior to age six or seven children prefer thematic relations (*table-eat*) to categorical relations (*table-chair*), but they will use categories when they think they are learning the meaning of a new word (Markman and Hutchinson, 1984).

The second step is to work out distinctions among words within a category; that step may take a long time. When you calculate that an average child learns 21 words a day it does not mean that the learning process is complete in one day. A child will be alert to and learning about many words at the same time, although it is not known how many words they are sensitive to at any given time.

One thing is reasonably sure, however. Children could not memorize 21 arbitrary, unrelated facts every day. Words can be learned so rapidly because they are *not* unrelated, because they form conceptually related patterns. A vocabulary is a coherent, integrated *system* of concepts. In other words, a feature that tends to be ignored in objective dictionaries—the organization of words into semantic fields—is what makes rapid vocabulary learning possible. If we hope to understand this process, therefore, we must have some reasonably definite characterization of lexical organization.

LEXICAL ORGANIZATION

Most cognitive psychologists assume that lexical memory is organized in at least

two ways, phonologically and semantically. A variety of common observations—that it is easy to think of words starting with the same sound or of words that rhyme, that confusions in memory and slips of the tongue often occur between words that sound alike, and so on—all support the assumption of a phonological organization. But an equal array of common observations—that people associate words with similar meanings, that similar meanings are used to define new words, that relations of entailment between words are easily recognized, and so on—all support the assumption of a conceptual organization, an organization into semantic fields. Children exploit both kinds of organization in their effort to learn words, but the present interest is only in the conceptual organization of the subjective lexicon.

Psychologists have studied conceptual relations among words in many different ways. Probably the oldest and best known approach is the word association experiment: 'Give me the first word you think of when I say . . .'. Although this technique was developed in order to study differences among people, the most obvious fact about the results is how well people agree—how few different words they think of in response to the stimulus word. Moreover, associations given to semantically related words can be analyzed to reveal underlying conceptual groupings (Deese, 1965).

Printed dictionaries do recognize specialized lexical fields, although the recognition is buried under the alphabetical ordering and very difficult to use. For example, technical senses of a word will be marked 'anatomy', 'jurisprudence', 'botany', or whatever. If you gathered together all those senses with a given marking, you would have a reasonably good glossary of technical terms in that specialty. Unfortunately, however, with a printed dictionary there is no easy way to retrieve all and only those senses bearing a particular marking. And English language dictionaries do not use markers in this way to distinguish among nontechnical semantic fields. What they *do* use, of course, are synonyms in the definitions—if you trace out the network of cross references provided by synonyms and antonyms, you can usually get some idea of the relevant semantic field. This is tedious work with a printed dictionary, whereas conceptual relations come immediately to mind in the subjective lexicon. That is to say, the subjective lexicon seems to be organized like a collection of dictionaries on many different subjects.

Attempts have been made to explain semantic fields as a natural reflection of the semantic decomposability of lexical concepts:

(9) *Origin of semantic fields*: A semantic field is a set of words whose semantic decompositions all share at least one nuclear concept.

The semantic field of motion verbs, for example, can be defined as all those verbs whose semantic decomposition includes the primitive (or atomic) concept of change-of-location-over-time (Miller, 1972; Miller and Johnson-Laird, 1976). In addition to motion, however, some motion verbs indicate the method of moving—*walk* versus *run*, for example—some require knowledge of the speaker's location—*come* versus *go*—many indicate the direction of motion—*rise*, *fall*, *approach*—or the medium through which the motion occurs—*swim* versus *fly*. Still other motion verbs express something about causation—*raise* versus *rise*, for example.

About a dozen different semantic components can be teased out of the field of motion verbs, so that each motion verb looks like a particular package of concepts; the semantic field consists of many related but slightly different packages. To the extent that this kind of semantic decomposition is valid, it helps to explain semantic

fields in terms of shared semantic components.

This explanation can be (and has been) misunderstood. To claim that semantic decomposition is possible is not to claim that it is necessary. In particular, the following hypothesis is highly dubious:

(10) *On-line synthesis of lexical meaning*: A word that can be semantically decomposed cannot be understood in ordinary discourse until its meaning has been synthesized by combining its elementary component concepts.

When one thinks of a motion verb, one does not necessarily synthesize its meaning by compounding primitive concepts, any more than one synthesizes the concept of 20 out of the prime numbers that it can be factored into (Miller, 1978a).

It should be pointed out that not every shared concept can serve as the nucleus of a semantic field. Although the set of verbs that share a concept of movement do seem to form a semantic field, the set of verbs that share a concept of causation do not. For example, *move*, *kill*, *give*, and *tell* are all causative verbs, but they clearly do not form a semantic field. Miller and Johnson-Laird (1976) have speculated that there are some basic concepts, like motion or possession, that can provide a nucleus around which a semantic field can develop, whereas other concepts, like cause or instrument, are used to elaborate concepts within many different semantic fields.

But to return to the central argument: one important difference between objective and subjective dictionaries has now been identified:

(11) *The multiple field hypothesis*: The dictionary in your head is organized into many sub-dictionaries for different semantic fields, with minimal polysemy within any single field.

The alphabetical ordering of a printed dictionary, on the other hand, combines all semantic fields and so calls attention to the polysemy across fields. In actual use, words are never as ambiguous as printed dictionaries would lead us to expect, because the topic of a discourse usually determines the semantic field that is appropriate.

LEARNING WORDS FROM CONTEXT

Learning to connect words with concepts that they can be used to express seems so simple—like one of those automatic associative processes that once made behaviorism sound feasible. Perhaps that is why cognitive psychologists have expressed little interest in it. Whatever your theoretical preferences, however, the fact is that words and concepts do become associated, rapidly and in large numbers. If we are not to account for this learning in terms of operant conditioning, we must offer plausible cognitive alternatives.

An appreciation of semantic fields can provide part of a cognitive theory of vocabulary growth. Another part must be provided by a better understanding of how words are learned from context, by an account of the knowledge-based processing of information provided by the contexts in which the word occurs.

In 1950 Werner and Kaplan conducted an experiment designed to reveal how well children can guess the meaning of a new word after hearing it used in a succession of sentences. After hearing 'A corplum may be used for support', 'Corplums may be used to close off an open space', 'A corplum may be long or short, thick or thin, strong or weak', and so on, the children offered their opinion of what a corplum was. Children eight years old seemed to have difficulty disentan-

gling the new word from the context in which it occurred, but 13-year-old children could usually solve the problems.

The question raised by Werner and Kaplan seems to have been ignored until 1981 when two workers in Amsterdam, van Daalen-Kaptjens and Elshout-Mohr, reopened it. What they contributed was an insightful analysis of protocols that they collected from college students who were asked to think out loud while they solved such problems.

The process of transforming several contexts of use into an acceptable definition was called 'decontextualization' by van Daalen-Kaptjens and Elshout-Mohr. They claimed that at least two steps are required. First, the context is reformulated into a sentence about the unknown word. For example, 'The painter used a corplum to stir his paints', which is a sentence about a painter, might be reformulated as 'A corplum can be used to stir paints', which is a sentence about corplums. Second, this reformulated information is transformed into an aspect of the meaning: 'Corplums can be implements.'

Subjects who could perform both steps generally succeeded in formulating an acceptable definition; subjects who performed only the first step were not as successful. They speculated that their successful subjects used the context plus their general knowledge to select a schematic model, then used successive contexts to narrow down the exact properties of the model.

The preferred way to write a definition is to give a superordinate name and to follow it with a relative clause that differentiates this particular instance of the class from all other instances. This strategy is not always available, but when it is, it seems to lead to clear, simple definitions. The general format is:

(12) *Preferred definitional format*: 'An X is a Y that . . .', where the phrases following 'that' provide the distinguishing particulars.

For example, 'A *noise* is a sound that . . . is loud, confused, indistinct, or disagreeable.' Or 'A *porter* is a person who . . . is employed to carry travelers' luggage.' And so on. In order to formulate such a definition, of course, one must know what the superordinate category is. If encouraged to give definitions, children under five years of age seldom mention a superordinate (Anglin, 1977), even though there is reason to believe that such categorical relations are known to them. Between five and ten years, however, the conventional definitional format begins to emerge. The older children's increased use of superordinates is presumably a consequence of adopting this format, not of some change in their underlying knowledge (Watson, 1982).

Although college students have far more world knowledge to draw on than do young children, it seems that both the Carey and Bartlett three-year-olds and the Dutch college students were trying to impose this conventional format (12) on the meanings of new words via the same, two-step path: first identify a general category, later work out the distinguishing particulars. No doubt there are exceptions to this rule, but a substantial fraction of word learning probably does follow some such pattern.

A more psychometric approach has been adopted by Sternberg and Powell (1983), who found that high school students' scores on a learning-from-context task correlate about 0.6 with IQ, vocabulary, and reading comprehension scores. They suggest that context is valuable insofar as it provides information about certain general aspects of the target word's meaning. Indeed, the aspects that they list

resemble the kind of meaning components that Miller and Johnson-Laird used to characterize motion verbs, and that others have developed for other kinds of semantic decomposition. Apparently, Sternberg and Powell share the intuition:

(13) *The contextual learning assumption*: An ability to perform conceptual decompositions is valuable for learning word meanings from context.

If this assumption is correct, it suggests how semantic fields and learning-from-context might work together in some eventual cognitive theory of vocabulary growth.

These remarks should suffice to indicate some of the differences between objective and subjective lexicons, differences in organization, use, and method of construction. In view of these differences, it is necessary to reject, or at least to revise considerably, the common assumption (4) that lexical retrieval from a printed dictionary provides a good model for lexical retrieval from personal memory.

POLYSEMY AND THE PRODUCTIVE CRITERION

The productive criterion for knowing a word is less easily applied than is the receptive criterion, because the investigator must wait for the speaker to introduce the word—which may entail a very long wait, indeed. However, school teachers have the authority to reduce this wait by assigning pupils the task of using particular words in sentences. Indeed, this argument is frequently used as part of the students' training in 'dictionary skills'.

Most of the troubles that children have with dictionaries stem from the polysemy of their lexical entries. As dictionaries are written, many different senses are distinguished and the reader is expected to be able to tell which one is appropriate in any given context. Most intelligent adults are able to solve that problem, but it easy to find evidence that children have great difficulty with it. For example, Deese (1967) reported on one teacher of seventh grade English who gave her pupils an assignment to look up certain words in the *Webster's Collegiate Dictionary* and to use them in a sentence. Here is a slightly simplified summary of the definitions that the children found when they looked up the word 'chaste':

> *chaste*: 1: innocent of unlawful sexual intercourse. 2: celibate. 3: pure in thought and act, modest. 4: severely simple in design or execution, austere.

With that lexical entry in mind, consider Deese's account of what some children did with it: 'Here are some sentences written by these youngsters after they had looked up the word *chaste*. You will have to admit that they are all consistent with at least one of the senses supplied by the dictionary.

'1. "The amoeba is a chaste animal." Evidently, the youngster who wrote this sentence is following that part of the entry that says chaste means simple in design . . .

'2. "The plates were still chaste after much use." Here the notion of being unstained seems to be critical.

'3. "The milk was chaste." Evidently, the sense of pure is meant here.' (Deese, 1967.)

Such examples demonstrate that dictionaries are reference books, not teaching instruments. If you already know something about the word, a dictionary will

remind you of it. But if you are totally innocent, you cannot trust the dictionary as your sole source of lexical information.

Miller and Gildea (1985) are following up on Deese's project and trying to extend it. They have found that misunderstandings are disturbingly common, and they are beginning to build up an inventory of the kinds of mistakes that children make when they use a dictionary. For example, what leads a child to think that two boxers *engaged* in fighting are betrothed, that to stir a cake is to *stimulate* it, or that a family *erodes* a lot because it eats out?

Along the same line, Jorgensen (1984) has analyzed what you do when you look up an unfamiliar word in the dictionary. It is a challenging task. The continuity of your reading is interrupted, of course, but the context in which the word occurred must be kept in mind. Once the word is found in the dictionary, you may need to choose among alternative lexical entries on the basis of part of speech; then, within the right entry, you may have to choose among several alternative senses on the basis of context. That is to say, the context of the original passage must be compared with a succession of contexts suggested in the dictionary until a best guess can be made as to the intended sense.

In order to study this process in a simplified form, Jorgensen developed a questionnaire based on 40 words (20 frequent words and 20 relatively rare words). For each word, phrases defining two different senses, with a sentence illustrating each sense, were taken from a children's dictionary. The words were then replaced in the illustrative sentences by nonsense syllables, and subjects were asked which definition of the nonsense syllable was most appropriate in the context provided by the sentence.

For example, one item presented the sentence: 'The snow is hattay with the windows.' Then two senses for 'hattay' were given: '1: level, flat, smooth; 2: at the same level.' Subjects were asked to choose the better of the two senses of the word as it was used in the illustrative sentence.

The questionnaire was given to 20 Princeton undergraduates and to 20 fourth grade children. College students identified the correct sense on 83% of the items, fourth graders on 63%: chance, of course, was 50%. The fact that adult performance was not perfect suggests that the illustrative sentences found in dictionaries leave something to be desired. That the children made even lower scores confirms the difficulty they have with such a context-matching task.

After going through the questionnaire, the subjects then went through it a second time, but this time trying to guess what the real words were for which nonsense syllables had been substituted. (For the example just given, the original word was 'even'.) College students guessed right 59% of the time for high-frequency words, but only 24% of the time for low-frequency words; the comparable figures for children were 18% and 5%. Moreover, the ones they guessed correctly were not always the same ones that they had gotten right in the context matching task—so their success could not be attributed to seeing through the nonsense substitutions.

The consistent and reliable difference between high- and low-frequency words was especially interesting. Since Jorgensen had substituted nonsense syllables, she thought they would all be treated like very, very low-frequency words. But subjects did better on items that had been based on high-frequency words. One possible explanation is that dictionaries do not define high-frequency words the same way

they define low-frequency words. Another possibility is that different senses of frequently used words are more distinctively different than are those of words more rarely used.

THE SENSE–FREQUENCY RELATION

It has long been known that lexical entries for frequently used words tend to have more senses than do entries for low-frequency words (Zipf, 1945):

(14) *The sense–frequency relation*: On the average, the most frequently used words are the most polysemous.

This relation is sometimes presented as an example of the innate contrariness of human beings—the words that people use most often are the most ambiguous words in the language. When you remember how dictionaries are written, however, an alternative explanation suggests itself.

The basic method for compiling monolingual dictionaries has changed little in 100 years. Lexicographers comb through the best writing to collect instances of the use of every word. These citations are usually written on cards and filed alphabetically. When they have collected enough citations they start writing definitions. They take all the instances of a given word and sort them into piles representing different senses; then for each pile they try to write a phrase that captures the sense expressed by the word in every instance in that pile.

Jorgensen noted the following point about this procedure: when a lexicographer picks up the collection of citations for a given word, the piles of cards for frequently used words will be bigger than the piles for the infrequent words. For a common word like *line* there may be several hundred cards to be sorted into different senses, whereas for an infrequent word like *coherent*, say, there may only be a dozen or so. It occurred to us that the most commonly used words might be the most polysemous simply because the lexicographer has more citations available for the common words.

Apparently, nobody had investigated the effect of number of instances on the number of senses that a lexicographer will discriminate. So Jorgensen conducted such an experiment. She retrieved all the instances of certain words in the Brown University corpus of one million running words of text (Francis and Kucera, 1982), and systematically controlled the number of instances that people were asked to categorize. Some subjects worked with 20 citations, some with 100, some with 200.

The number of senses distinguished by a lexicographer might increase as a function of the number of instances available for two different reasons: (1) one reason would be the sampling problem—a small number of citations drawn at random would probably not include all the different uses of a word. In order to estimate the sampling effect, Jorgensen ran a Monte Carlo simulation. (2) A second reason for the sense–frequency relation would be the categorization effect already mentioned—the more citations a person is given, the more senses he or she will try to distinguish. It was this second effect that Jorgensen wanted to test.

The results were clear. There is an expected sampling effect, but there is no other evidence that the size of the pile affects the number of distinctions that a lexicographer will draw.

As a by-product of this study, however, Jorgensen observed that:

(15) *Analysis into senses*: Amateur lexicographers distinguish fewer different senses of frequently occurring words than do professional lexicographers.

Jorgensen's amateurs agreed with one another rather well on what the 'important' senses were, thus supporting the suspicion that professional lexicographers make a virtue of multiplying meanings beyond necessity.

THE QUALITY OF A DEFINITION

A problem that recurs repeatedly in attempts to experiment with lexical materials is the problem of evaluating the quality of a definition. Do children have difficulty using dictionaries because the definitions are poor? Do lexicographers write better quality definitions for senses that are represented by a greater number of instances of use? If you wanted to improve the quality of dictionary definitions, how would you evaluate your work?

One answer to questions about definitional quality is the subjective rating scale. Sternberg and Powell (1983) found that judges could agree in rating the quality of definitions written by high school students. But that method hardly seems sensitive enough to evaluate the work of professional lexicographers.

Jorgensen (1984) achieved a somewhat clearer criterion of quality by the following method. A nonsense syllable was substituted for the word whose sense had been defined, and subjects were requested to write sentences using the nonsense syllable in the intended sense. For example, they might be given something like: '*kleb*: a kleb is a person who is in charge, a leader' and a subject might write: 'John was the kleb of the basketball team.' Next, the original word was substituted back into the sentences that these subjects wrote. For our imaginary example, this would yield: 'John was the head of the basketball team.' The results were then submitted to another group of subjects to evaluate for normality and typicality. This example would not be rated as completely typical, since basketball teams are not typically said to have heads, although they could have.

The assumptions underlying this procedure is that a good definition is one that enables a person to write an acceptable sentence using the word in its intended sense. Otherwise said:

(16) *The productive criterion of definitional quality*: A good definition is one that provides the information required to satisfy the productive criterion for knowing the word.

By using this criterion, Jorgensen found that her amateurs could write definitions for words denoting concrete objects that were every bit as good as the definitions found in published dictionaries; for abstract concepts, however, the professionals did better. Her results also showed that definitional quality improves as the number of citations available increases, but the relation is strongest for the more frequent words. That is to say, for a good phrasing of some sense of a commonly used word, more citations are required than are required for a good definition of some relatively rare word.

Note that Jorgensen's productive criterion imposes a much more modest requirement than does the traditional assumption that a good definition should be synonymous with the word it defines, that is to say, that the word and its definition should be intersubstitutble *salva vertitate*:

(17) *The synonymity criterion of definitional quality*: A good definition is one that is substitutable for the word it defines in all contexts without altering truth values.

An operational test of synonymity might require people to pair words with their

definitions—a task that should be easy if the words and their definitions were synonymous. However, experience with such tasks has shown that they become difficult when the words involved become at all similar in meaning. Many students of the synonymy relation have expressed serious doubts about the practical realization of such equivalences.

Still other criteria could be proposed, of course. For example, we might require that the definition be useful for helping someone—a child, perhaps—learn the meaning of an unfamilar word:

(18) *The pedagogical criterion of definitional quality*: A good definition is one that provides the information required to add the word to one's vocabulary.

A pedagogical criterion would impose different requirements on definitions intended for learners of different sophistication. For a given learner, however, a definition satisfying the pedagogical criterion (18) should also satisfy the productive criterion (16), although the reverse would not necessarily hold.

It is difficult to defend any strong opinion about which criterion is best, since goodness in a definition must depend in part of the use you plan to make of it: definitions good for one purpose might not be optimal for some other.

AUTOMATED DICTIONARIES

In 1978 Miller proposed that it might be useful to put a dictionary into a computer (Miller, 1978b). The initial thought was that an automated dictionary might be useful in helping children build larger vocabularies—particularly children from socio-economic backgrounds where the kind of academic language used in the schools is not normally spoken. A lexical database might provide a basis for various kinds of computer games that might get children interested in words. There was no guarantee that computer words games would really accelerate vocabulary growth, but they could hardly be worse than doing nothing.

The first question, of course, was whether or not it is technically feasible to put a dictionary into a computer. The answer is Yes (Fox *et al.*, 1980). Since memory is what you need in order to store a dictionary in a computer, and since the cost of memory continues to decline, it makes good technological sense.

If you had such a lexical database, you would just type in the word you want, and the dictionary entry for it would appear instantly (more or less) on your video screen. But such a system would not represent a clever use of computer—the machine would be little more than a fast page-turner. Anyone who goes to the trouble of putting a dictionary into a computer should make better use of the computer than that.

How to use the computer intelligently set the problem. The second thought was to organize the lexical entries like a thesaurus. For example, in addition to giving the dictionary entry for the word, the system might also give the thesaurus entry: a menu of semantically related words. Then you could search the database conceptually as well as alphabetically.

To a psychologist, of course, a conceptually organized dictionary sounds much more interesting, since that is the way people search through the dictionaries in their minds. But a conceptual organization does not resolve all of the difficulties.

Many of the problems with existing dictionaries can be traced, directly or indirectly, to the importance of alphabetization. It is not merely that a user has to search for words alphabetically. The whole project of writing a dictionary is

organized alphabetically—that is how the enormous database is organized. For example, in writing the great *Oxford English Dictionary*, different lexicographers took responsibility for different parts of the alphabet (Murray, 1977). This division of labor had some odd consequences. For example, if you look up the modal verbs, you find wonderful essays on *can* and *could* in one volume, then another group of wonderful essays on *may*, *might*, and *must* in another volume, and finally a third set of wonderful essays on *shall*, *should*, *will* and *would*. Because modal verbs are scattered through the alphabet, the essays were written by different persons with different theories of modal expressions; it is impossible to combine them.

Another use for a computer might be to organize the lexical database while a dictionary is being compiled and written—to use the machine as kind of lexicographer's assistant. When words are defined one at a time, in alphabetical order, semantically related words are seldom considered together, but with properly designed assistance, a lexicographer could work with whole sets of semantically related words. Words are organized into semantic fields; within each field polysemy is minimal, but meanings are so interrelated that it is difficult to understand one word without seeing it in relation to all the other, related words. Would it not increase the value of a dictionary if those relations between meanings were considered explicitly as the dictionary was written?

Dictionaries published for computer use—perhaps as part of the word processing systems of the future—will surely come. Perhaps we will be smart enough to take full advantage of this new form of publication, perhaps we will not be satisfied with installing 19th century dictionaries into 21st century machines. Lexicology is as challenging a problem for computer technology as it is for psycholinguistics.

If these remarks leave the impression that far more questions have been raised than have been answered, then they will have accomplished their purpose, which was to sketch out an area for research, an area where psychologists, linguists, and computer scientists could profitably collaborate. If all the answers were known, nothing would be left to do. The answers are not known, but there are interesting questions here for everyone.

Acknowledgements

This paper has been written with the help of grants to Princeton University from the Alfred P. Sloan Foundation and the Spencer Foundation.

REFERENCES

Anglin, J. (1977). *Word, Object and Conceptual Development*. Norton, New York.

Barlett, E. J. (1977). The acquisition of the meaning of color terms: A study of lexical development. In: *Proceedings of the Stirling Conference on the Psychology of Language*, P. Smith and R. Campbell (Eds). Plenum, New York.

Carey, S. (1978). The child as word learner. In: *Linguistic Theory and Psychological Reality*, M. Halle, J. Bresnan and G. A. Miller (Eds). MIT Press, Cambridge, Mass.

Carey, S. and Bartlett, E.J. (1978). Acquiring a single new word. *Papers and Reports on Child Language Development*. Department of Linguistics, Stanford University, No. 15, pp. 17–29.

Daalen-Kaptjens, van, M. M. and Elshout-Mohr, M. (1981). The acquisition of word meaning as a cognitive learning process. *Journal of Verbal Learning and Verbal Behaviour 20*, 386–399.

Deese, J. (1965). *The Structure of Associations in Language and Thought*. Johns Hopkins University Press, Baltimore.

Deese, J. (1967). Meaning and change of meaning. *American Psychologist 22*, 641–651.

Dockrell, J. (1983). The child's acquisition of unfamiliar words: An experimental study. Ph.D. thesis,

University of Stirling. Cited by J. McShane and J. Dockrell, Lexical and grammatical development, in: *Language Production. Volume 2. Development, Writing and Other Language Processes*, B. Butterworth (Ed.). Academic Press, New York.

Fox, M. S., Bebel, D. J. and Parker, A. C. (1980). The automated dictionary. *Computer 13*, 35–48.

Francis, W. N. and Kucera, H. (1982). *Frequency Analysis of English Usage: Lexicon and Grammar*. Houghton Mifflin, Boston.

Jorgensen, J. (1984). Polysemy and psycholexicology. Ph.D. thesis, Princeton University.

Markman, E. M. and Hutchinson, J. E. (1984). Children's sensitivity to constraints on word meaning: Taxonomic versus thematic relations. *Cognitive Psychology 6*, 1–27.

Miller, G. A. (1972). English verbs of motion: A case study in semantics and lexical memory. In: *Coding Processes in Human Memory*, A. W. Melton and E. Martin (Eds). Winston, Washington, D.C.

Miller, G.A. (1977). *Spontaneous Apprentices: Children and Language*. Seabury Press, New York.

Miller, G.A. (1978a). Semantic relations among words. In: *Linguistic Theory and Psychological Reality*, M. Halle, J. Bresnan and G.A. Miller (Eds). MIT Press, Cambridge, Mass.

Miller, G.A. (1978b). Automated dictionaries. Paper prepared for *NIE Conference on Testing, Teaching, and Learning*, 17–26, August 1978.

Miller, G.A. and Gildea, P.M. (1985). How to misread a dictionary. *AILA Bulletin* (in press).

Miller, G.A. and Johnson-Laird, P.N. (1976). *Language and Perception*. Harvard University Press, Cambridge, Mass.

Murray, K.M.E. (1977). *Caught in the Web of Word: James Murray and the Oxford English Dictionary*. Yale University Press, New Haven, Conn.

Sternberg, R.J. and Powell, J.S. (1983). Comprehending verbal comprehension. *American Psychologist 38*, 878–893.

Templin, M.C. (1957). *Certain Language Skills in Children: Their Development and Interrelationships*. University of Minnesota Press, Minneapolis.

Watson, R.P.M. (1982). From meaning to definition: The development of word meaning in the school-aged child. Ph.D. thesis, University of Toronto.

Werner, H. and Kaplan, E. (1950). Development of word meaning through verbal context: An experimental study. *Journal of Psychology 29*, 251–257.

Zipf, G.K. (1945). The meaning-frequency relationship of words. *Journal of General Psychology 33*, 251–256.