

- Trotter, M. (1992). *Life writing: Exploring the practice of autoethnography in anthropology*. Unpublished master's thesis, University of Illinois, Urbana-Champaign.
- True confessions: The age of the literary memoir [Special issue]. (1996, May 12). *New York Times Magazine*.
- Turner, V., & Bruner, E. (Eds.). (1986). *The anthropology of experience*. Urbana: University of Illinois Press.
- Tyler, S. (1986). Post-modern ethnography: From document of the occult to occult document. In J. Clifford and G. E. Marcus (Eds.), *Writing culture: The poetics and politics of ethnography* (pp. 122-140). Berkeley: University of California Press.
- Van Maanen, J. (1988). *Tales of the field: On writing ethnography*. Chicago: University of Chicago Press.
- Van Maanen, J. (1995). An end to innocence: The ethnography of ethnography. In J. Van Maanen (Ed.), *Representation in ethnography* (pp. 1-35). Thousand Oaks, CA: Sage.
- Van Maanen, M. (1990). *Researching lived experience: Human science for an action sensitive pedagogy*. Albany: State University of New York Press.
- Wittgenstein, L. (1953). *Philosophical investigations* (G. Anscombe, Trans.). New York: Macmillan.
- Zola, I. K. (1982). *Missing pieces: A chronicle of living with a disability*. Philadelphia: Temple University Press.
- Zussman, R. (1996). Autobiographical occasions. *Contemporary Sociology*, 25, 143-148.

7

Data Management and Analysis Methods

Gery W. Ryan and H. Russell Bernard

◆ Texts Are Us

This chapter is about methods for managing and analyzing qualitative data. By *qualitative data* we mean text: newspapers, movies, sitcoms, e-mail traffic, folktales, life histories. We also mean narratives—narratives about getting divorced, about being sick, about surviving hand-to-hand combat, about selling sex, about trying to quit smoking. In fact, most of the archaeologically recoverable information about human thought and human behavior is text, the “good stuff” of social science.

Scholars in content analysis began using computers in the 1950s to do statistical analysis of texts (Pool, 1959), but recent advances in technology are changing the economics of the social sciences. Optical scanning today makes light work of converting written texts to machine-readable form. Within a few years, voice-recognition software will make light work of transcribing open-ended interviews. These technologies are blind to epistemological differences. Interpretivists and positivists alike are using these technologies for the analysis of texts, and will do so more and more.

Like Tesch (1990), we distinguish between the *linguistic tradition*, which treats text as an object of analysis itself, and the *sociological tradition*, which treats text as a window into human experience (see Figure 7.1). The linguistic tradition includes narrative analysis, conversation (or




Figure 7.1. Typology of Qualitative Analysis Techniques

discourse) analysis, performance analysis, and formal linguistic analysis. Methods for analyses in this tradition are covered elsewhere in this *Handbook*. We focus here on methods used in the sociological tradition, which we take to include work across the social sciences.

There are two kinds of written texts in the sociological tradition: (a) words or phrases generated by techniques for systematic elicitation and (b) free-flowing texts, such as narratives, discourse, and responses to open-ended interview questions. In the next section, we describe some methods for collecting and analyzing words or phrases. Techniques for data collection include free lists, pile sorts, frame elicitations, and triad tests. Techniques for the analysis of these kinds of data include componential analysis, taxonomies, and mental maps.

We then turn to the analysis of free-flowing texts. We look first at methods that use raw text as their input—methods such as key-words-in-context, word counts, semantic network analysis, and cognitive maps. We then describe methods that require the reduction of text to codes. These include grounded theory, schema analysis, classical content analysis, content dictionaries, analytic induction, and ethnographic decision models. Each of these methods of analysis has advantages and disadvantages. Some are appropriate for exploring data, others for making comparisons, and others for building and testing models. Nothing does it all.

◆ Collecting and Analyzing Words or Phrases

Techniques for Systematic Elicitation

Researchers use techniques for systematic elicitation to identify lists of items that belong in a cultural domain and to assess the relationships among these items (for detailed reviews of these methods, see Bernard, 1994; Borgatti, 1998; Weller, 1998; Weller & Romney, 1988). Cultural domains comprise lists of words in a language that somehow “belong together.” Some domains (such as animals, illnesses, things to eat) are very large and inclusive, whereas others (animals you can keep at home, illnesses that children get, brands of beer) are relatively small. Some lists (such as the list of terms for members of a family or the names of all the Major League Baseball teams) are agreed on by all native speakers of a language; others (such as the list of carpenters’ tools) represent highly specialized knowledge, and still others (like the list of great left-handed baseball

pitchers of the 20th century) are matters of heated debate. Below we review some of the most common systematic elicitation techniques and discuss how researchers analyze the data they generate.

Free Lists

Free lists are particularly useful for identifying the items in a cultural domain. To elicit domains, researchers might ask, "What kinds of illnesses do you know?" Some short, open-ended questions on surveys can be considered free lists, as can some responses generated from in-depth ethnographic interviews and focus groups. Investigators interpret the frequency of mention and the order in which items are mentioned in the lists as indicators of items' salience (for measures of salience, see Robbins & Nolan, 1997; Smith, 1993; Smith & Borgatti, 1998). The co-occurrence of items across lists and the proximity with which items appear in lists may be used as measures of similarity among items (Borgatti, 1998; Henley, 1969; for a clear example, see Fleisher & Harrington, 1998).

Paired Comparisons, Pile Sorts, Triad Tests

Researchers use paired comparisons, pile sorts, and triads tests to explore the *relationships* among items. Here are two questions we might ask someone in a paired comparison test about a list of fruits: (a) "On a scale of 1 to 5, how similar are lemons and watermelons with regard to sweetness?" (b) "Which is sweeter, watermelons or lemons?" The first question produces a set of fruit-by-fruit matrices, one for each respondent, the entries of which are scale values on the similarity of sweetness among all pairs of fruits. The second question produces, for each respondent, a perfect rank ordering of the set of fruits.

In a pile sort, the researcher asks each respondent to sort a set of cards or objects into piles. Item similarity is the number of times each pair of items is placed in the same pile (for examples, see Boster, 1994; Roos, 1998). In a triad test, the researcher presents sets of three items and asks each respondent either to "choose the two most similar items" or to "pick the item that is the most different." The similarity among pairs of items is the number of times people choose to keep pairs of items together (for some good examples, see Albert, 1991; Harman, 1998).

Frame Substitution

In the frame substitution task (D'Andrade, 1995; D'Andrade, Quinn, Nerlove, & Romney, 1972; Frake, 1964; Metzger & Williams, 1966), the researcher asks the respondent to link each item in a list of items with a list of attributes. D'Andrade et al. (1972) gave people a list of 30 illness terms and asked them to fill in the blanks in frames such as "You can catch _____ from other people," "You can have _____ and never know it," and "Most people get _____ at one time or other" (p. 12; for other examples of frame substitution, see Furbee & Benfer, 1983; Young, 1978).

Techniques for Analyzing Data About Cultural Domains

Researchers use these kinds of data to build several kinds of models about how people think. *Componential analysis* produces formal models of the elements in a cultural domain, and *taxonomies* display hierarchical associations among the elements in a domain. *Mental maps* are best for displaying fuzzy constructs and dimensions. We treat these in turn.

Componential Analysis

As we have outlined elsewhere, componential analysis (or feature analysis) is a formal, qualitative technique for studying the content of meaning (Bernard, 1994; Bernard & Ryan, 1998). Developed by linguists to identify the features and rules that distinguish one sound from another (Jakobson & Halle, 1956), the technique was elaborated by anthropologists in the 1950s and 1960s (Conklin, 1955; D'Andrade, 1995; Frake, 1962; Goodenough, 1956; Rushforth, 1982; Wallace, 1962). (For a particularly good description of how to apply the method, see Spradley, 1979, pp. 173-184.)

Componential analysis is based on the principle of distinctive features. Any two items (sounds, kinship terms, names of plants, names of animals, and so on) can be distinguished by some minimal set ($2n$) of binary features—that is, features that either occur or do not occur. It takes two features to distinguish four items ($2^2 = 4$, in other words), three features to distinguish eight items ($2^3 = 8$), and so on. The trick is to identify the smallest set of features that best describes the domain of interest. Table 7.1 shows that just three features are needed to describe kinds of horses.

TABLE 7.1 A Componential Analysis of Six Kinds of Horses

Name	Female	Neuter	Adult
mare	+	-	+
Stallion	-	-	+
Gelding	-	+	+
Foal	-	+	-
Filly	+	-	-
Colt	-	-	-

SOURCE: Adapted from D'Andrade (1995).

Componential analysis produces models based on logical relationships among features. The models do not account for variations in the meanings of terms across individuals. For example, when we tried to do a componential analysis on the terms for cattle (*bull, cow, heifer, calf, steer, and ox*), we found that native speakers of English in the United States (even farmers) disagreed about the differences between *cow* and *heifer*, and between *steer* and *ox*. When the relationships among items are less well defined, taxonomies or mental models may be useful. Nor is there any intimation that componential analyses reflect how "people really think."

Taxonomies

Folk taxonomies are meant to capture the hierarchical structure in sets of terms and are commonly displayed as branching tree diagrams. Figure 7.1 presents a taxonomy of our own understanding of qualitative analysis techniques. Figure 7.2 depicts a taxonomy we have adapted from Pamela Erickson's (1997) study of the perceptions among clinicians and adolescents of methods of contraception. Researchers can elicit folk taxonomies directly by using successive pile sorts (Boster, 1994; Perchonock & Werner, 1969). This involves asking people to continually subdivide the piles of a free pile sort until each item is in its own individual pile. Taxonomic models can also be created with cluster analysis on the similarity data from paired comparisons, pile sorts, and triad tests. *Hierarchical cluster analysis* (Johnson, 1967) builds a taxonomic tree where each item appears in only one group.




Figure 7.2. Clinicians' Taxonomy of Contraceptive Methods
SOURCE: Based on Erickson (1997).

Interinformant variation is common in folk taxonomies. That is, different people may use different words to refer to the same category of things. Some of Erickson's (1997) clinician informants referred to the "highly effective" group of methods as "safe," "more reliable," and "sure bets." Category labels need not be simple words, but may be complex phrases; for example, see the category in Figure 7.2 comprising contraceptive methods in which you "have to pay attention to timing." Sometimes, people have no labels at all for particular categories—at least none that they can dredge up easily—and categories, even when named, may be fuzzy and may overlap with other categories. *Overlapping cluster analysis* (Hartigan, 1975) identifies groups of items where a single item may appear in multiple groups.

Mental Maps

Mental maps are visual displays of the similarities among items, whether or not those items are organized hierarchically. One popular method for making these maps is by collecting data about the cognitive similarity or dissimilarity among a set of objects and then applying multi-dimensional scaling, or MDS, to the similarities (Kruskal & Wish, 1978).

Cognitive maps are meant to be directly analogous to physical maps. Consider a table of distances between all pairs of cities on a map. Objects (cities) that are very dissimilar have high mileage between them and are placed far apart on the map; objects that are less dissimilar have low mileage between them and are placed closer together. Pile sorts, triad tests, and paired comparison tests are measures of cognitive distance. For example, Ryan (1995) asked 11 literate Kom speakers in Cameroon to perform successive pile sorts on Kom illness terms. Figure 7.3 presents an MDS plot of the collective mental map of these terms. The five major illness categories, circled, were identified by hierarchical cluster analysis of the same matrix used to produce the MDS plot.¹

Data from frame substitution tasks can be displayed with correspondence analysis (Weller & Romney, 1990).² Correspondence analysis scales both the rows and the columns into the same space. For example, Kirchler (1992) analyzed 562 obituaries of managers who had died in 1974, 1980, and 1986. He identified 31 descriptive categories from adjectives used in the obituaries and then used correspondence analysis to display how these categories were associated with men and women managers over time. Figure 7.4 shows that male managers who died in 1974 and 1980 were seen




Figure 7.3. Mental Map of Kom Illness Terms

by their surviving friends and family as *active, intelligent, outstanding, conscientious, and experienced experts*. Although the managers who died in 1986 were still respected, they were more likely to be described as *entrepreneurs, opinion leaders, and decision makers*. Perceptions of female managers also changed, but they did not become more like their male counterparts. In 1974 and 1980, female managers were remembered for being nice people. They were described as *kind, likable, and adorable*. By 1986, women were remembered for their *courage and commitment*. Kirchler interpreted these data to mean that gender stereotypes changed in the early 1980s. By 1986, both male and female managers were perceived as working for success, but men impressed their colleagues through their knowledge and expertise, whereas women impressed their colleagues with motivation and engagement.




Figure 7.4. Correspondence Analysis of the Frequencies of 31 Disruptive Obituary Categories by Gender and Year of Publication

SOURCE: Erich Kirchler, "Adorable Woman, Expert Man: Changing Gender Images of Women and Men in Management," *European Journal of Social Psychology*, 22 (1992), p. 371. Copyright 1992 by John Wiley & Sons Limited. Reproduced by permission of John Wiley & Sons Limited.

◆ Methods for Analyzing Free-Flowing Text

Although taxonomies, MDS maps, and the like are useful for analyzing short phrases or words, most qualitative data come in the form of free-flowing texts. There are two major types of analysis. In one, the text is segmented into its most basic meaningful components: words. In the other, meanings are found in large blocks of text.

Analyzing Words

Techniques for word analysis include key-words-in-context, word counts, structural analysis, and cognitive maps. We review each below.

Key-Words-in-Context

Researchers create key-words-in-context (KWIC) lists by finding all the places in a text where a particular word or phrase appears and printing it out in the context of some number of words (say, 30) before and after it. This produces a *concordance*. Well-known concordances have been done on sacred texts, such as the Old and New Testaments (Darton, 1976; Hatch & Redpath, 1954) and the Koran (Kassis, 1983), and on famous works of literature from Euripides (Allen & Italie, 1954) to Homer (Prendergast, 1971), to Beowulf (Bessinger, 1969), to Dylan Thomas (Farrington & Farrington, 1980). (On the use of concordances in modern literary studies, see Burton, 1981a, 1981b, 1982; McKinnon, 1993.)

Word Counts

Word counts are useful for discovering patterns of ideas in any body of text, from field notes to responses to open-ended questions. Students of mass media have used word counts to trace the ebb and flow of support for political figures over time (Danielson & Lasorsa, 1997; Pool, 1952). Differences in the use of words common to the writings of James Madison and Alexander Hamilton led Mosteller and Wallace (1964) to conclude that Madison and not Hamilton had written 12 of the *Federalist Papers*. (For other examples of authorship studies, see Martindale & McKenzie, 1995; Yule 1944/1968.)

Word analysis (like constant comparison, memoing, and other techniques) can help researchers to discover themes in texts. Ryan and Weisner (1996) instructed fathers and mothers of adolescents in Los Angeles: "Describe your children. In your own words, just tell us about them." Ryan and Weisner identified all the unique words in the answers they got to that grand-tour question and noted the number of times each word was used by mothers and by fathers. Mothers, for example, were more likely to use words like *friends*, *creative*, *time*, and *honest*; fathers were more likely to use words like *school*, *good*, *lack*, *student*, *enjoys*, *independent*, and *extremely*. This suggests that mothers, on first mention, express concern

over interpersonal issues, whereas fathers appear to prioritize achievement-oriented and individualistic issues. This kind of analysis considers neither the contexts in which the words occur nor whether the words are used negatively or positively, but distillations like these can help researchers to identify important constructs and can provide data for systematic comparisons across groups.

Structural Analysis and Semantic Networks

Network, or structural, analysis examines the properties that emerge from relations among things. As early as 1959, Charles Osgood created word co-occurrence matrices and applied factor analysis and dimensional plotting to describe the relations among words. Today, semantic network analysis is a growing field (Barnett & Danowski, 1992; Danowski, 1982, 1993). For example, Nolan and Ryan (1999) asked 59 undergraduates (30 women and 29 men) to describe their "most memorable horror film." The researchers identified the 45 most common adjectives, verbs, and nouns used across the descriptions of the films. They produced a 45(word)-by-59(person) matrix, the cells of which indicated whether each student had used each key word in his or her description. Finally, Nolan and Ryan created a 59(person)-by-59(person) similarity matrix of *people* based on the co-occurrence of the words in their descriptions.

Figure 7.5 shows the MDS of Nolan and Ryan's data. Although there is some overlap, it is pretty clear that the men and women in their study used different sets of words to describe horror films. Men were more likely to use words such as *teenager*, *disturbing*, *violence*, *rural*, *dark*, *country*, and *hillbilly*, whereas women were more likely to use words such as *boy*, *little*, *devil*, *young*, *horror*, *father*, and *evil*. Nolan and Ryan interpreted these results to mean that the men had a fear of rural people and places, whereas the women were more afraid of betrayed intimacy and spiritual possession. (For other examples of the use of word-by-word matrices, see Jang & Barnett, 1994; Schnegg & Bernard, 1996.) This example makes abundantly clear the value of turning qualitative data into quantitative data: Doing so can produce information that engenders deeper interpretations of the meanings in the original corpus of qualitative data. Just as in any mass of numbers, it is hard to see patterns in words unless one first does some kind of data reduction. More about this below.

As in word analysis, one appeal of semantic network analysis is that the data processing is done by computer. The only investigator bias intro-




Figure 7.5. Multidimensional Scaling of Informants Based on Words Used in Descriptions of Horror Films

duced in the process is the decision to include words that occur at least 10 times or 5 times or whatever. (For discussion of computer programs that produce word-by-text and word-by-word co-occurrence matrices, see Borgatti, 1992; Doerfel & Barnett, 1996.) There is, however, no guarantee that the output of any word co-occurrence matrix will be meaningful, and it is notoriously easy to read patterns (and thus meanings) into any set of items.

Cognitive Maps

Cognitive map analysis combines the intuition of human coders with the quantitative methods of network analysis. Carley's work with this technique is instructive. Carley argues that if cognitive models or schemata exist, they are expressed in the texts of people's speech and can be represented as networks of concepts (see Carley & Palmquist, 1992,

p. 602), an approach also suggested by D'Andrade (1991). To the extent that cognitive models are widely shared, Carley asserts, even a very small set of texts will contain the information required for describing the models, especially for narrowly defined arenas of life.

In one study, Carley (1993) asked students some questions about the work of scientists. Here are two examples she collected:

Student A: I found that scientists engage in research in order to make discoveries and generate new ideas. Such research by scientists is hard work and often involves collaboration with other scientists which leads to discoveries which make the scientists famous. Such collaboration may be informal, such as when they share new ideas over lunch, or formal, such as when they are coauthors of a paper.

Student B: It was hard work to research famous scientists engaged in collaboration and I made many informal discoveries. My research showed that scientists engaged in collaboration with other scientists are coauthors of at least one paper containing their new ideas. Some scientists make formal discoveries and have new ideas. (p. 89)

Carley compared the students' texts by analyzing 11 concepts: *I*, *scientists*, *research*, *hard work*, *collaboration*, *discoveries*, *new ideas*, *formal*, *informal*, *coauthors*, *paper*. She coded the concepts for their strength, sign (positive or negative), and direction (whether one concept is logically prior to others), not just for their existence. She found that although students used the same concepts in their texts, the concepts clearly had different meanings. To display the differences in understandings, Carley advocates the use of maps that show the relations between and among concepts. Figure 7.6 shows Carley's maps of two of the texts.

Carley's approach is promising because it combines the automation of word counts with the sensitivity of human intuition and interpretation. As Carley recognizes, however, a lot depends on who does the coding. Different coders will produce different maps by making different coding choices. In the end, native-language competence is one of the fundamental methodological requirements for analysis (see also Carley, 1997; Carley & Kaufer, 1993; Carley & Palmquist, 1992; Palmquist, Carley, & Dale, 1997).

Key-words-in-context, word counts, structural analysis, and cognitive maps all reduce text to the fundamental meanings of specific words. These reductions make it easy for researchers to identify general patterns and




Figure 7.6. Coded Maps of Two Students' Texts

SOURCE: Kathleen Carley, "Coding Choices for Textual Analysis: A Comparison of Content Analysis and Map Analysis," in P. Marsden (Ed.), *Sociological Methodology* (Oxford: Blackwell, 1993), p. 104. Copyright 1993 by the American Sociological Association. Reproduced by permission of the American Sociological Association.

make comparisons across texts. With the exception of KWIC, however, these techniques remove words from the contexts in which they occur. Subtle nuances are likely to be lost—which brings us to the analysis of whole texts.

Analyzing Chunks of Text: Coding

Coding is the heart and soul of *whole-text analysis*. Coding forces the researcher to make judgments about the meanings of contiguous blocks of text. The fundamental tasks associated with coding are sampling, identifying themes, building codebooks, marking texts, constructing models (relationships among codes), and testing these models against empirical data. We outline each task below. We then describe some of the major coding traditions: grounded theory, schema analysis, classic content analysis, content dictionaries, analytic induction, and ethnographic decision trees. We want to emphasize that no particular tradition, whether humanistic or positivistic, has a monopoly on text analysis.

Sampling

Investigators must first identify a *corpus of texts*, and then select the *units of analysis within the texts*. Selection can be either random or purposive, but the choice is not a matter of cleaving to one epistemological tradition or another. Waitzkin and Britt (1993) did a thoroughgoing interpretive analysis of encounters between patients and doctors by selecting 50 texts at random from 336 audiotaped encounters. Trost (1986) used classical content analysis to test how the relationships between teenagers and their families might be affected by five different dichotomous variables. He intentionally selected five cases from each of the 32 possible combinations of the five variables and conducted $32 \times 5 = 160$ interviews.

Samples may also be based on extreme or deviant cases, cases that illustrate maximum variety on variables, cases that are somehow typical of a phenomenon, or cases that confirm or disconfirm a hypothesis. (For reviews of nonrandom sampling strategies, see Patton, 1990, pp. 169-186; Sandelowski, 1995b.) A single case may be sufficient to display something of substantive importance, but Morse (1994) suggests using at least six participants in studies where one is trying to understand the essence of experience. Morse also suggests 30-50 interviews for ethnographies and grounded theory studies. Finding themes and building theory may require fewer cases than comparing across groups and testing hypotheses or models.

Once the researcher has established a sample of texts, the next step is to identify the basic units of analysis. The units may be entire texts (books,

interviews, responses to an open-ended question on a survey), grammatical segments (words, word senses, sentences, themes, paragraphs), formatting units (rows, columns, or pages), or simply chunks of text that reflect a single theme—what Krippendorff (1980, p. 62) calls *thematic units*. In general, where the objective is to compare across texts (as in the case of classical content analysis), the units of analysis need to be non-overlapping. (For discussion of additional kinds of units of analysis, see Krippendorff, 1980, pp. 57-64; Tesch, 1990.)

Finding Themes

Themes are abstract (and often fuzzy) constructs that investigators identify before, during, and after data collection. Literature reviews are rich sources for themes, as are investigators' own experiences with subject matter. More often than not, however, researchers induce themes from the text itself.

There is more than one way to induce themes. Grounded theorists suggest a careful, line-by-line reading of the text while looking for processes, actions, assumptions, and consequences. Schema analysts suggest looking for metaphors, for repetitions of words, and for shifts in content (Agar & Hobbs, 1985). Content analysts have used KWIC to identify different meanings. Spradley (1979, pp. 199-201) suggests looking for evidence of social conflict, cultural contradictions, informal methods of social control, things that people do in managing impersonal social relationships, methods by which people acquire and maintain achieved and ascribed status, and information about how people solve problems. Each of these arenas is likely to yield major themes in cultures. Barkin, Ryan, and Gelberg (1999) had multiple coders independently sort informants' statements into thematic piles. They then used multidimensional scaling and cluster analysis on the pile-sort data to identify subthemes shared across coders. (For another example, see Patterson, Bettini, & Nussbaum, 1993.)

Willms et al. (1990) and Miles and Huberman (1994) suggest that researchers start with some general themes derived from reading the literature and add more themes and subthemes as they go. Shelley (1992) followed this advice in her study of how social networks affect people with end-stage kidney disease. She used the *Outline of Cultural Materials* (Murdock, 1971) as the basis of her coding scheme and then added additional themes based on a close reading of the text. Bulmer (1979) lists 10 different sources of themes, including literature reviews, professional

definitions, local commonsense constructs, and researchers' values and prior experiences. He also notes that investigators' general theoretical orientations, the richness of the existing literature, and the characteristics of the phenomena being studied influence the themes researchers are likely to find.

No matter how the researcher actually *does* inductive coding, by the time he or she has identified the themes and refined them to the point where they can be applied to an entire corpus of texts, a lot of interpretive analysis has already been done. Miles and Huberman (1994) say simply, "Coding is analysis" (p. 56).

Building Codebooks

Codebooks are simply organized lists of codes (often in hierarchies). How a researcher can develop a codebook is covered in detail by Dey (1993, pp. 95-151), Crabtree and Miller (1992), and Miles and Huberman (1994, pp. 55-72). MacQueen, McLellan, Kay, and Milstein (1998) suggest that a good codebook should include a detailed description of each code, inclusion and exclusion criteria, and exemplars of real text for each theme. If a theme is particularly abstract, we suggest that the researcher also provide examples of the theme's boundaries and even some cases that are closely related but *not* included within the theme. Coding is supposed to be data *reduction*, not proliferation (Miles, 1979, pp. 593-594). The codes themselves are mnemonic devices used to identify or mark the specific themes in a text. They can be either words or numbers—whatever the researcher finds easiest to remember and to apply.

Qualitative researchers working as a team need to agree up front on what to include in their codebook. Morse (1994) suggests beginning the process with a group meeting. MacQueen et al. (1998) suggest that a single team member should be designated "Keeper of the Codebook"—we strongly agree.

Good codebooks are developed and refined as the research goes along. Kurasaki (1997) interviewed 20 *sansei*—third-generation Japanese Americans—and used a grounded theory approach to do her analysis of ethnic identity. She started with seven major themes. As the analysis progressed, she split the major themes into subthemes. Eventually, she combined two of the major themes and wound up with six major themes and a total of 18 subthemes. (Richards & Richards, 1991, discuss the theoretical principles related to hierarchical coding structures that emerge out of the data.

Araujo, 1995, uses an example from his own research on the traditional British manufacturing industry to describe the process of designing and refining hierarchical codes.)

The development and refinement of coding categories have long been central tasks in classical content analysis (see Berelson, 1952, pp. 147-168; Holsti, 1969, pp. 95-126) and are particularly important in the construction of concept dictionaries (Deese, 1969; Stone, Dunphy, Smith, & Ogilvie, 1966, pp. 134-168). Krippendorff (1980, pp. 71-84) and Carey, Morgan, and Oxtoby (1996) note that much of codebook refinement comes during the training of coders to mark the text and in the act of checking for intercoder agreement. Disagreement among multiple coders shows when the codebook is ambiguous and confusing. The first run also allows the researcher to identify good examples to include in the codebook.

Marking Texts

The act of coding involves the assigning of codes to contiguous units of text. Coding serves two distinct purposes in qualitative analysis. First, codes act as *tags* to mark off text in a corpus for later retrieval or indexing. Tags are not associated with any fixed units of text; they can mark simple phrases or extend across multiple pages. Second, codes act as *values* assigned to fixed units (see Bernard, 1991, 1994; Seidel & Kelle, 1995). Here, codes are nominal, ordinal, or ratio scale values that are applied to fixed, nonoverlapping units of analysis. The nonoverlapping units can be texts (such as paragraphs, pages, documents), episodes, cases, or persons. *Codes as tags* are associated with grounded theory and schema analysis (reviewed below). *Codes as values* are associated with classic content analysis and content dictionaries. The two types of codes are not mutually exclusive, but the use of one gloss—*code*—for both concepts can be misleading.

Analyzing Chunks of Texts: Building Conceptual Models

Once the researcher identifies a set of things (themes, concepts, beliefs, behaviors), the next step is to identify how these things are linked to each other in a theoretical model (Miles & Huberman, 1994, pp. 134-137). Models are sets of abstract constructs and the relationships among them

(Bulmer, 1979). Grounded theory, schema analysis, ethnographic decision modeling, and analytic induction all include model-building phases.

Once a model starts to take shape, the researcher looks for negative cases—cases that don't fit the model. Negative cases either disconfirm parts of a model or suggest new connections that need to be made. In either instance, negative cases need to be accommodated. Negative case analysis is discussed in detail by Becker, Geer, Hughes, and Strauss (1961, pp. 37-45), Strauss and Corbin (1990, pp. 108-109), Lincoln and Guba (1985, pp. 309-313), Dey (1993, pp. 226-233), Miles and Huberman (1994, p. 271), and Becker (1998), and is used by schema analysts (Quinn, 1997), ethnographic decision modelers (Gladwin, 1989), and scholars who use analytic induction (Bloor, 1976; Cressey, 1953/1971; Lindesmith, 1947/1968).

In ethnographic decision modeling and in classical content analysis, models are built on one set of data and tested on another. In their original formulation, Glaser and Strauss (1967) emphasized that building grounded theory models is a step in the research process and that models need to be validated. Grounded theorists and schema analysts today are more likely to validate their models by seeking confirmation from expert informants than by analyzing a second set of data. For example, Kearney, Murphy, and Rosenbaum (1994) checked the validity of their model of crack mothers' experiences by presenting it to knowledgeable respondents who were familiar with the research.

Regardless of the kind of reliability and validity checks, models are simplifications of reality. They can be made more or less complicated and may capture all or only a portion of the variance in a given set of data. It is up to the investigator and his or her peers to decide how much a particular model is supposed to describe.

Below we review some of the most common methods researchers use to analyze blocks of texts. These include grounded theory, schema analysis, classical content analysis, content dictionaries, analytic induction, and ethnographic decision tree analysis.

Grounded Theory

Grounded theorists want to understand people's experiences in as rigorous and detailed a manner as possible. They want to identify categories and concepts that emerge from text and link these concepts into

substantive and formal theories. The original formulation of the method (Glaser & Strauss, 1967) is still useful, but later works are easier to read and more practical (Charmaz, 1990; Lincoln & Guba, 1985; Lonkila, 1995; Strauss, 1987). Strauss and Corbin (1990), Dey (1993), and Becker (1998) provide especially useful guidance. (For some recent examples of grounded theory research, see Hunt & Ropo, 1995; Irurita, 1996; Kearney et al., 1994; Kearney, Murphy, Irwin, & Rosenbaum, 1995; Sohler, 1993; Strauss & Corbin, 1997; Wilson & Hutchinson, 1996; Wright, 1997.)

Grounded theory is an iterative process by which the analyst becomes more and more "grounded" in the data and develops increasingly richer concepts and models of how the phenomenon being studied really works. To do this, the grounded theorist collects verbatim transcripts of interviews and reads through a small sample of text (usually line by line). Sandelowski (1995a) observes that analysis of texts begins with proof-reading the material and simply underlining key phrases "because they make some as yet inchoate sense" (p. 373). In a process called "open coding," the investigator identifies potential themes by pulling together real examples from the text (Agar, 1996; Bernard, 1994; Bogdan & Biklen, 1992; Lincoln & Guba, 1985; Lofland & Lofland, 1995; Strauss & Corbin, 1990; Taylor & Bogdan, 1984). Identifying the categories and terms used by informants themselves is called "*in vivo* coding" (Strauss & Corbin, 1990). As grounded theorists develop their concepts and categories, they often decide they need to gather more data from informants.

As coding categories emerge, the investigator links them together in theoretical models. One technique is to compare and contrast themes and concepts. When, why, and under what conditions do these themes occur in the text? Glaser and Strauss (1967, pp. 101-116) refer to this as the "constant comparison method," and it is similar to the contrast questions Spradley (1979, pp. 160-172) suggests researchers ask informants. (For other good descriptions of the comparison method, see Glaser, 1978, pp. 56-72; Strauss & Corbin, 1990, pp. 84-95.)

Another useful tool for building theoretical models is the conditional matrix described by Strauss and Corbin (1990, pp. 158-175). The conditional matrix is a set of concentric circles, each level corresponding to a different unit of influence. At the center are actions and interactions; the outer rings represent international and national concerns, and the inner rings represent individual and small group influences on action. The

matrix is designed to help investigators to be more sensitive to conditions, actions/interactions, and consequences of a phenomenon and to order these conditions and consequences into theories.

Memoing is one of the principal techniques for recording relationships among themes. Strauss and Corbin (1990, pp. 18, 73-74, 109-129, 197-219) discuss three kinds of memos: code notes, theory notes, and operational notes. Code notes describe the concepts that are being discovered in "the discovery of grounded theory." In theory notes, the researcher tries to summarize his or her ideas about what is going on in the text. Operational notes are about practical matters.

Once a model starts to take shape, the researcher uses negative case analysis to identify problems and make appropriate revisions. The end results of grounded theory are often displayed through the presentation of segments of text—verbatim quotes from informants—as exemplars of concepts and theories. These illustrations may be prototypical examples of central tendencies or they may represent exceptions to the norm. Grounded theory researchers also display their theoretical results in maps of the major categories and the relationships among them (Kearney et al., 1995; Miles & Huberman, 1994, pp. 134-137). These "concept maps" are similar to the personal semantic networks described by Leinhardt (1987, 1989), Strauss (1992), and D'Andrade (1991) (see below).

Schema Analysis

Schema analysis combines elements of the linguistic and sociological traditions. It is based on the idea that people must use cognitive simplifications to help make sense of the complex information to which they are constantly exposed (Casson, 1983, p. 430). Schank and Abelson (1977) postulate that schemata—or scripts, as they call them—enable culturally skilled people to fill in details of a story or event. It is, says Wodak (1992, p. 525), our schemata that lead us to interpret Mona Lisa's smile as evidence of her perplexity or her desperation.

From a methodological view, schema analysis is similar to grounded theory. Both begin with a careful reading of verbatim texts and seek to discover and link themes into theoretical models. In a series of articles, Quinn (1982, 1987, 1992, 1996, 1997) has analyzed hundreds of hours of interviews to discover concepts underlying American marriage and to show how these concepts are tied together. Quinn's (1997) method is to

"exploit clues in ordinary discourse for what they tell us about shared cognition—to glean what people must have in mind in order to say the things they do" (p. 140). She begins by looking at patterns of speech and the repetition of key words and phrases, paying particular attention to informants' use of metaphors and the commonalities in their reasoning about marriage. Quinn found that the hundreds of metaphors in her corpus of texts fit into just eight linked classes, which she calls lastingness, sharedness, compatibility, mutual benefit, difficulty, effort, success (or failure), and risk of failure.

Metaphors and proverbs are not the only linguistic features used to infer meaning from text. D'Andrade (1991) notes that "perhaps the simplest and most direct indication of schematic organization in naturalistic discourse is the repetition of associative linkages" (p. 294). He observes that "indeed, anyone who has listened to long stretches of talk—whether generated by a friend, spouse, workmate, informant, or patient—knows how frequently people circle through the same network of ideas" (p. 287).

In a study of blue-collar workers in Rhode Island, Claudia Strauss (1992) refers to these ideas as "personal semantic networks." She describes such a network from one of her informants. On rereading her intensive interviews with one of the workers, Strauss found that her informant repeatedly referred to ideas associated with greed, money, businessmen, siblings, and "being different." She displays the relationships among these ideas by writing the concepts on a page of paper and connecting them with lines and explanations.

Price (1987) observes that when people tell stories, they assume that their listeners share with them many assumptions about how the world works, and so they leave out information that "everyone knows." Thus she looks for what is *not* said in order to identify underlying cultural assumptions (p. 314).

For more examples of the search for cultural schemata in texts, see Holland's (1985) study of the reasoning that Americans apply to interpersonal problems, Kempton's (1987) study of ordinary Americans' theories of home heat control, Claudia Strauss's (1997) study of what chemical plant workers and their neighbors think about the free enterprise system, and Agar and Hobbs's (1985) analysis of how an informant became a burglar. We next turn to the two other methods used across the social sciences for analyzing text: classical content analysis and content dictionaries.

Displaying Concepts and Models

400

Visual displays are an important part of qualitative analysis. Selecting key quotes as exemplars, building matrices or forms, and laying theories out in the form of flowcharts or maps are all potent ways to communicate ideas visually to others. Models are typically displayed using boxes and arrows, with the boxes containing themes and the arrows representing the relationships among them. Lines can be unidirectional or bidirectional. For example, taxonomies are models in which the lines represent the super- and subordinate relationships among items. Relationships can include causality, association, choices, and time, to name a few.

A widely used method for describing themes is the presentation of direct quotes from respondents—quotes that lead the reader to understand quickly what it may have taken the researcher months or years to figure out. The researcher chooses segments of text—verbatim quotes from respondents—as exemplars of concepts, of theories, and of negative cases. Ryan (1999) has used multiple coders to identify typical quotes. He asks 10 coders to mark the same corpus of text for three themes. Ryan argues that the text marked by all the coders represents the central tendency or typical examples of the abstract constructs, whereas text marked by only some of the coders represents less typical examples and is more typical of the “edges” of the construct.

Tables can be used to organize and display raw text or can be used to summarize qualitative data along multiple dimensions (rows and columns). The cells can be filled with verbatim quotes (Bernard & Ashton-Voyoucalos, 1976; Leinhardt & Smith, 1985, p. 254; Miles & Huberman, 1994, p. 130), summary statements (Yoder, 1995), or symbols (Fjellman & Gladwin, 1985; Van Maanen, Miller, & Johnson, 1982). (For a range of presentation formats, see Bernard, 1994; Miles & Huberman, 1994; Werner & Schoepfle, 1987.)

Classical Content Analysis

Whereas grounded theory is concerned with the discovery of data-induced hypotheses, classical content analysis comprises techniques for reducing texts to a unit-by-variable matrix and analyzing that matrix quantitatively to test hypotheses. The researcher can produce a matrix by applying a set of codes to a set of qualitative data (including written texts as well as audio and video media). Unlike grounded theory or schema

analysis, content analysis assumes that the codes of interest have already been discovered and described.

Once the researcher has selected a sample of texts, the next step in classical content analysis is to code each unit for each of the themes or variables in the codebook. This produces a unit-by-variable matrix that can be analyzed using a variety of statistical techniques. For example, Cowan and O'Brien (1990) tested whether males or females are more likely to be survivors in slasher films. Conventional wisdom about such films suggests that victims are mostly women and slashers are mostly men. Cowan and O'Brien selected a corpus of 56 slasher films and identified 474 victims. They coded each victim for gender and survival. They found that slashers are mostly men, but it turned out that victims are equally likely to be male or female. Women who survive are less likely to be shown engaging in sexual behavior and are less likely to be physically attractive than their nonsurviving counterparts. Male victims are cynical, egotistical, and dictatorial. Cowan and O'Brien conclude that, in slasher films, sexually pure women survive and “unmitigated masculinity” leads to death (p. 195).

The coding of texts is usually assigned to multiple coders so that the researcher can see whether the constructs being investigated are shared and whether multiple coders can reliably apply the same codes. Typically, investigators first calculate the percentage of agreement among coders for each variable or theme. They then apply a correction formula to take account of the fact that some fraction of agreement will always occur by chance. The amount of that fraction depends on the number of coders and the precision of measurement for each code. If two people code a theme present or absent, they could agree, *ceteris paribus*, on any answer 25% of the time by chance. If a theme, such as wealth, is measured ordinally (low, medium, high), then the likelihood of chance agreement changes accordingly. Cohen's (1960) kappa, or K , is a popular measure for taking these chances into account. When K is zero, agreement is what might be expected by chance. When K is negative, the observed level of agreement is less than one would expect by chance. How much intercoder agreement is enough? The standards are still ad hoc, but Krippendorff (1980, pp. 147-148) advocates agreement of at least .70 and notes that some scholars (e.g., Brouwer, Clark, Gerbner, & Krippendorff, 1969) use a cutoff of .80. Fleiss (1971) and Light (1971) expand kappa to handle multiple coders. For other measures of intercoder agreement, see Krippendorff (1980, pp. 147-154) and Craig (1981).

Reliability "concerns the extent to which an experiment, test, or any measuring procedure yields the same results on repeated trials" (Carmines & Zeller, 1979, p. 11). A high level of intercoder agreement is evidence that a theme has some external validity and is not just a figment of the investigator's imagination (Mitchell, 1979). Not surprisingly, investigators have suggested many ways to assess validity (for reviews of key issues, see Campbell, 1957; Campbell & Stanley, 1963; Cook & Campbell, 1979; Denzin, 1997; Fielding & Fielding, 1986; Guba, 1981; Guba & Lincoln, 1982; Hammersley, 1992; Kirk & Miller, 1986; Lincoln & Guba, 1985). Bernard (1994) argues that, ultimately, the validity of a concept depends on the utility of the device that measures it and the collective judgment of the scientific community that a construct and its measure are valid. "In the end," he says, "we are left to deal with the effects of our judgments, which is just as it should be. Valid measurement makes valid data, but validity itself depends on the collective opinion of researchers" (p. 43). *Generalizability* refers to the degree to which the findings are applicable to other populations or samples. It draws on the degree to which the original data were representative of a larger population.

For reviews of work in content analysis, see Pool (1959); Gerbner, Holsti, Krippendorff, Paisley, and Stone (1969); Holsti (1969); Krippendorff (1980); Weber (1990); and Roberts (1997). Examples of classical content analysis can be found in media studies (Hirschman, 1987; Kolbe & Albanese, 1996; Spiggle, 1986), political rhetoric (Kaid, Tedesco, & McKinnon, 1996), folklore (Johnson & Price-Williams, 1997), business relations (Spears, Mowen, & Chakraborty, 1996), health care delivery (Potts, Runyan, Zerger, & Marchetti, 1996; Sleath, Svarstad, & Roter, 1997), and law (Imrich, Mullin, & Linz, 1995). Classical content analysis is also the fundamental means by which anthropologists test cross-cultural hypotheses (Bradley, Moore, Burton, & White, 1990; Ember & Ember, 1992; White & Burton, 1988). For early, but fundamental, criticisms of the approach, see Kracauer (1953) and George (1959).

Content Dictionaries

Computer-based, general-purpose content analysis dictionaries allow investigators to automate the coding of texts. To build such dictionaries, researchers assign words, by hand, to one or more categories (there are typically 50-60 categories in computerized content analysis dictionaries)

according to a set of rules. The rules are part of a computer program that parses new texts, assigning words to categories.

Work on content dictionaries began in the 1960s with the General Inquirer and continues to this day (Kelly & Stone, 1975; Stone et al., 1966; Zuell, Weber, & Mohler, 1989). The General Inquirer is a computer program that uses a dictionary (the *Harvard Psychosocial Dictionary*) to parse and assign text to coded categories. Over time, the dictionary has been updated. The latest version (*Harvard IV*) contains more than 10,000 words and can distinguish among multiple meanings of words (Rosenberg, Schnurr, & Oxman, 1990, p. 303). Because such dictionaries do not contain all the words in the English language, investigators can assign unrecognized words to categories as they see fit, a process of further modifying the "codebook."

How effective are computer-based dictionaries? An early version of the General Inquirer was tested on 66 suicide notes—33 written by men who had actually taken their own lives and 33 written by men who were asked to produce simulated suicide notes. The program parsed the texts and picked the actual suicide notes 91% of the time (Ogilvie, Stone, & Schneidman, 1966). Content dictionaries do not need to be very big to be useful. Colby (1966) created a simple dictionary to distinguish between Navaho and Zuni responses to thematic apperception tests. For additional examples of special-purpose dictionaries in content analysis, see Fan and Shaffer (1990), Furbee (1996), Holsti (1966), Jehn and Werner (1993), Laffal (1990, 1995), McTavish and Pirro (1990), and Schnurr, Rosenberg, Oxman, and Tucker (1986).

Content dictionaries are attractive because they are entirely reliable and automated, but, as Shapiro (1997) argues, this may be offset by a decrease in validity. For the time being, only humans can parse certain subtleties of meaning reflected in context (Viney, 1983), but computer-based dictionaries are getting better all the time. For example, texts are now scored by computer for the Gottschalk-Gleser psychological scales (measuring various forms of anxiety and hostility) with greater than .80 reliability (Gottschalk & Bechtel, 1993).

Analytic Induction and Boolean Tests

Analytic induction is a formal, nonquantitative method for building up causal explanations of phenomena from a close examination of cases. It was proposed as an alternative to statistical analysis by Znaniecki (1934,

pp. 249-331), modified by Lindesmith (1947/1968) and Cressey (1953/1971), and is discussed by Denzin (1978), Bulmer (1979), Manning (1982), and Becker (1998), among others. (For critiques of the approach, see Robison, 1951.) The method is a formal kind of negative case analysis.

The technique can be described in a series of steps: First, define a phenomenon that requires explanation and propose an explanation. Next, examine a case to see if the explanation fits. If it does, then examine another case. An explanation is accepted until a new case falsifies it. When a case is found that doesn't fit, then, under the rules of analytic induction, the alternatives are to change the explanation (so that you can include the new case) or redefine the phenomenon (so that you exclude the nuisance case). Ideally, the process continues until a universal explanation for all known cases of a phenomenon is attained. Explaining cases by declaring them all unique is a tempting but illegitimate option. Classic examples of analytic induction include Lindesmith's (1947/1968) study of drug addicts, Cressey's (1953/1971) study of embezzlers, and McCleary's (1978) study of how parole officers decide when one of their charges is in violation of parole. For a particularly clear example of the technique, see Bloor's (1976, 1978) analysis of how doctors decide whether or not to remove children's tonsils.

Ragin (1987, 1994) formalized the logic of analytic induction, using a Boolean approach, and Romme (1995) applies the approach to textual data. Boolean algebra involves just two states (true and false, present and absent), but even with such simple inputs, things can get very complicated, very quickly. With just three dichotomous causal conditions (A and not A, B and not B, and C and not C) and one outcome variable (D and not D), there are 16 possible cases: A, B, C, D; A, not B, C, D; A, B, not C, D; and so on. Boolean analysis involves setting up what is known as a truth table, or a matrix of the actual versus the possible outcomes. (For more on truth tables and how they are related to negative case analysis, see Becker, 1998, pp. 146-214.)

Schweizer (1991, 1996) applied this method in his analysis of conflict and social status among residents of Chen Village, China. (For a discussion of Schweizer's data collection and analysis methods, see Bernard & Ryan, 1998.) All the data about the actors in this political drama were extracted from a historical narrative about Chen Village. Like classic content analysis and cognitive mapping, analytic induction requires that human coders

read and code text and then produce an event-by-variable matrix. The object of the analysis, however, is not to show the relationships among all codes, but to find the minimal set of logical relationships among the concepts that accounts for a single dependent variable. With more than three variables, the analysis becomes much more difficult. Computer programs such as QCA (Drass, 1980) and ANTHROPAC (Borgatti, 1992) test all possible multivariate hypotheses and find the optimal solution. (QCA is reviewed in Weitzman & Miles, 1995.)

Ethnographic Decision Models

Ethnographic decision models (EDMs) are qualitative, causal analyses that predict behavioral choices under specific circumstances. An EDM, often referred to as a decision tree or flowchart, comprises a series of nested *if-then* statements that link criteria (and combinations of criteria) to the behavior of interest (Figure 7.7). EDMs have been used to explain how fishermen decide where to fish (Gatewood, 1983), what prices people decide to place on their products (Gladwin, 1971; Quinn, 1978), and which treatments people choose for an illness (Mathews & Hill, 1990; Ryan & Martínez, 1996; Young, 1980).

EDMs combine many of the techniques employed in grounded theory and classic content analysis. Gladwin (1989) lays out the fundamental steps for building an ethnographic decision tree model. (For other clear descriptions of the steps, see Hill, 1998; Ryan & Martínez, 1996.)

EDMs require exploratory data collection, preliminary model building, and model testing. First, researchers identify the decisions they want to explore and the alternatives that are available. Typically, EDMs are done on simple yes/no types of behaviors. They can be used, however, to predict multiple behaviors (Mathews & Hill, 1990; Young, 1980) as well as the order of multiple behaviors (Ryan & Martínez, 1996).

Next, the researchers conduct open-ended interviews to discover the criteria people use to select among alternatives. The researchers first ask people to recall the most recent example of an actual—not a hypothetical—behavior and to recall *why* they did or did not do the behavior. Here is an example from a study we've done recently: "Think about the last time you had a can of something to drink in your hand—soda, juice, water, beer, whatever. Did you recycle the can? Why [Why not]?" This kind of question generates a list of *decision criteria*. To




Figure 7.7. Decision Model of Constraints on the Use of Pills or Liquid Medications for Mothers Treating Children with Diarrhea in Rural Mexico

SOURCE: Based on data in Ryan and Martínez (1996).

understand how these criteria might be linked, EDM researchers ask people to compare the latest decision with other similar decisions made in the past. Some researchers have used vignettes to elicit the relationships among criteria (e.g., Weller, Ruebush, & Klein, 1997; Young, 1980).

With a list of decision criteria in hand, the researchers' next step is to systematically collect data, preferably from a new group of people, about how each criterion applies or does not apply to a recent example of the behavior. "Was a recycling bin handy?" and "Do you normally recycle cans at home?" are 2 of the 30 questions we've asked people in our study of recycling behavior. The data from this stage are used to build a preliminary model of the decision process for the behavior under scrutiny. Cases that do not fit the model are examined closely and the model is modified. Researchers tweak, or tune, the model until they achieve a satisfactory level of postdictive accuracy—understood to be at least 80% among EDM researchers. Parsimonious models are favored over more complicated ones. (For automated ways of building and pruning decision trees, see Mingers, 1989a, 1989b.)

The process doesn't end there—the same data are used in building a preliminary model and in testing its postdictive accuracy. When EDM researchers feel confident in their model, they test it on an independent sample to see if it predicts as well as it postdicts. Typically, EDMs predict more than 80% of whatever behavior is being modeled, far above what we expect by chance. (For more detailed arguments on how to calculate accuracy in EDMs, see Ryan & Martínez, 1996; Weller et al., 1997.)

Because of the intensive labor involved, EDMs have been necessarily restricted to relatively simple decisions in relatively small and homogeneous populations. Recently, however, we found we could effectively test, on a nationally representative sample, our ethnographically derived decision models for whether or not to recycle cans and whether or not to ask for paper or plastic bags at the grocery store (Bernard, Ryan, & Borgatti, 1999).

EDMs can be displayed as decision trees (e.g., Gladwin, 1989), as decision tables (Mathews & Hill, 1990; Young, 1980), or as sets of rules in the form of *if-then* statements (Ryan & Martínez, 1996). Like componential analysis, folk taxonomies, and schema analysis, EDMs represent an aggregate decision process and do not necessarily represent what is going on inside people's heads (Garro, 1998).

◆ Breaking Down the Boundaries

Text analysis as a research strategy permeates the social sciences, and the range of methods for conducting text analysis is inspiring. Investigators examine words, sentences, paragraphs, pages, documents, ideas, meanings, paralinguistic features, and even what is missing from the text. They interpret, mark, retrieve, and count. By turns, they apply interpretive analysis and numerical analysis. They use text analysis for exploratory and confirmatory purposes. Researchers identify themes, describe them, and compare them across cases and groups. Finally, they combine themes into conceptual models and theories to explain and predict social phenomena.

Figure 7.1 depicts a broad range of analysis techniques found across the social sciences. To conform our presentation with the literature on qualitative methods, we have organized these techniques according to the goals of the investigators and the kinds of texts to which the techniques are typically applied.

In this chapter, we focus on the sociological tradition that uses text as a "window into experience" rather than the linguistic tradition that describes how texts are developed and structured. Texts such as conversations, performances, and narratives are analyzed by investigators from both the sociological and linguistic traditions. Although the agendas of the investigators may differ, we see no reason why many of the sociological techniques we describe could not be useful in the linguistic tradition and vice versa.

We also distinguish between those analyses associated with systematically elicited data and those associated with free-flowing texts. We argue, however, that these data-analytic pairings are ones of convention rather than necessity. Investigators want to (a) identify the range and salience of key items and concepts, (b) discover the relationships among these items and concepts, and (c) build and test models linking these concepts together. They use free-listing tasks, KWIC, word counts, and the exploratory phases of grounded theory, schema analysis, and EDM to discover potentially useful themes and concepts.

Researchers use pile sorts, paired comparisons, triads tests, frame substitution tasks, semantic networks, cognitive maps, content analysis and content dictionaries, and the modeling phases of grounded theory, schema analysis, and EDM to discover how abstract concepts are related to each other. They display the relationships as models or frameworks. These

frameworks include formal models that rely on Boolean logic (componential analysis and analytic induction), hierarchical models (taxonomies and ethnographic decision models), probabilistic models (classic content analysis and content dictionaries), and more abstract models such as those produced by grounded theory and schema analysis. Below we describe two important examples of studies in which researchers combined methods to understand their data more fully.

Jehn and Doucet (1996, 1997) used word counts, classical content analysis, and mental mapping to examine conflicts among Chinese and U.S. business associates. They asked 76 U.S. managers who had worked in Sino-American joint ventures to describe recent interpersonal conflicts with business partners. Each person described a situation with a same-culture manager and a different-cultural manager. The researchers made sure that each manager interviewed included information about his or her relationship to the other person, who was involved, what the conflict was about, what caused the conflict, and how the conflict was resolved.

After collecting the narratives, Jehn and Doucet asked their informants to help identify the emic themes in the narratives. First, they generated separate lists of words from the intercultural and intracultural conflict narratives. They asked three expatriate managers to act as judges and to identify all the words that were related to conflict. They settled on a list of 542 conflict words from the intercultural list and 242 conflict words from the intracultural list. Jehn and Doucet then asked the three judges to sort the words into piles or categories. The experts identified 15 subcategories for the intercultural data (things like *conflict*, *expectations*, *rules*, *power*, and *volatile*) and 15 categories for the intracultural data (things like *conflict*, *needs*, *standards*, *power*, *contentious*, and *lose*). Taking into consideration the total number of words in each corpus, conflict words were used more in intracultural interviews and resolution terms were more likely to be used in intercultural interviews.

Jehn and Doucet also used traditional content analysis on their data. They had two coders read the 152 conflict scenarios (76 intracultural and 76 intercultural) and evaluate (on a 5-point scale) each on 27 different themes they had identified from the literature. This produced two 76×27 scenario-by-theme profile matrices—one for the intracultural conflicts and one for the intercultural conflicts. The first three factors from the intercultural matrix reflect (a) interpersonal animosity and hostility, (b) aggravation, and (c) the volatile nature of the conflict. The first two

factors from the intracultural matrix reflect (a) hatred and animosity with a volatile nature and (b) conflicts conducted calmly with little verbal intensity.

Finally, Jehn and Doucet identified the 30 intracultural and the 30 intercultural scenarios that they felt were the clearest and pithiest. They recruited 50 *more* expatriate managers to assess the similarities (on a 5-point scale) of 60-120 randomly selected pairs of scenarios. When combined across informants, the managers' judgments produced two aggregate, scenario-by-scenario, similarity matrices—one for the intracultural conflicts and one for the intercultural conflicts. Multidimensional scaling of the intercultural similarity data identified four dimensions: (a) open versus resistant to change, (b) situational causes versus individual traits, (c) high- versus low-resolution potential based on trust, and (d) high- versus low-resolution potential based on patience. Scaling of the intracultural similarity data identified four different dimensions: (a) high versus low cooperation, (b) high versus low confrontation, (c) problem solving versus accepting, and (d) resolved versus ongoing.

The work of Jehn and Doucet is impressive because the analysis of the data from these tasks produced different sets of themes. All three emically induced theme sets have some intuitive appeal, and all three yield analytic results that are useful. The researchers could have also used the techniques of grounded theory or schema analysis to discover even more themes.

Jehn and Doucet are not the only researchers ever to combine different analytic techniques. In a series of articles on young adult "occasional" drug users, Agar (1979, 1980, 1983) used grounded theory methods to build models of behavior. He then used classical content analysis to test his hypotheses. Agar conducted and transcribed three interviews with each of his three informants. In his 1979 article, Agar describes his initial, intuitive analysis. He pulled all the statements that pertained to informants' interactions or assessments of other people. He then looked at the statements and sorted them into piles based on their content. He named each pile as a theme and assessed how the themes interacted. He found that he had three piles. The first contained statements in which the informant was expressing negative feelings toward a person in a dominant social position. The second was made up of statements emphasizing the other's knowledge or awareness. The statements in the third small cluster emphasized the importance of change or openness to new experiences.

From this intuitive analysis, Agar felt that his informants were telling him that those in authority were only interested in displaying their authority unless they had knowledge or awareness; that knowledge or awareness comes through openness to new experience; and that most in authority are closed to new experience or change.

To test his intuitive understanding of the data, Agar (1983) used all the statements from a single informant and coded the statements for their role type (kin, friend/acquaintance, educational, occupational, or other), power (dominant, symmetrical, subordinate, or undetermined), and affect (positive, negative, ambivalent, or absent). Agar was particularly interested in whether negative sentiments were expressed toward those in dominant social roles. For one informant, Agar found that out of 40 statements coded as dominant, 32 were coded negative and 8 were coded positive. For the 36 statements coded as symmetrical, 20 were coded positive and 16 negative, lending support to his original theory.

Next, Agar looked closely at the deviant cases—the 8 statements where the informant expressed positive affect toward a person in a dominant role. These counterexamples suggested that the positive affect was expressed toward a dominant social other when the social other possessed, or was communicating to the informant, knowledge that the informant valued.

Finally, Agar (1980) developed a more systematic questionnaire to test his hypothesis further. He selected 12 statements, 4 from each of the control, knowledge, and change themes identified earlier. He matched these statements with eight roles from the informant's transcript (father, mother, employer, teacher, friend, wife, coworker, and teammate). Agar then returned to his informant and asked if the resulting statements were true, false, or irrelevant. (In no case did the informant report "irrelevant.") Agar then compared the informant's responses to his original hypotheses. He found that on balance his hypotheses were correct, but discrepancies between his expectations and his results suggested areas for further research.

These examples show that investigators can apply one technique to different kinds of data and they can apply multiple techniques to the same data set. Text analysis is used by avowed positivists and interpretivists alike. As we have argued elsewhere (Bernard, 1993; Bernard & Ryan, 1998), methods are simply tools that belong to everyone.

◆ What's Next?

We do not want to minimize the profound *intellectual* differences in the epistemological positions of positivists and interpretivists. We think, however, that when researchers can move easily and cheaply between qualitative and quantitative data collection and analysis, the distinctions between the two epistemological positions will become of *less practical* importance. That is, as researchers recognize the full array of tools at their disposal, and as these tools become easier to use, the pragmatics of research will lessen the distinction between qualitative and quantitative data and analysis.

The process is under way—and is moving fast—with the development of increasingly useful software tools for qualitative data analysis. Useful tools create markets, and market needs create increasingly useful tools. Qualitative data analysis packages (ATLAS/ti, NUD•IST, Code-A-Text, the Ethnograph, AnSWR, and others) have improved dramatically over the past few years (Fischer, 1994; Kelle, 1995; Weitzman & Miles, 1995). These products, and others, make it easier and easier for researchers to identify themes, build codebooks, mark text, create memos, and develop theoretical models. Based loosely on a grounded theory type of approach to qualitative analysis, many program suites have recently folded in techniques from classical content analysis. Several programs, for example, allow researchers to export data to matrices that they can then analyze using other programs.

Investigators, however, remain constrained by program-defined units of analysis—usually marked blocks of text or informants. Researchers need the flexibility to create matrices on demand, whether they be word-by-theme or word-by-informant matrices for word analysis and sentence-by-code or paragraph-by-code matrices for content analysis. A series of word analysis functions would greatly enhance the automated coding features found in programs that are geared to the interests of scholars in the grounded theory school. Investigators should be able to code a section of text using grounded theory, then identify the key words associated with each theme. They should be able to use key words to search for additional occurrences of the theme in large corpuses of text.

When programs make it easy to use multiple coders and to identify intercoder agreements and disagreements systematically, researchers will be better able to describe themes and to train assistants. Adding a variety of measures for calculating intercoder agreement, which only some pro-

grams do, would also be helpful. Some programs offer researchers the option of recording the marking behavior of multiple coders, yet offer no direct way to measure intercoder agreement.

The evolution of text analysis software is just beginning. Some 15 years ago, spell checkers, thesauruses, and scalable fonts were all sold separately. Today, these functions are integrated into all full-featured word-processing packages. Just 10 years ago, graphics programs were sold separately from programs that do statistical analysis. Today, graphics functions are integrated into all full-featured packages for statistical analysis. As programmers of text analysis software compete for market share, packages will become more inclusive, incorporating methods from both sides of the epistemological divide. It can't happen too soon.

◆ Notes

1. MDS displays are highly evocative. They beg to be interpreted. In fact, they *must* be interpreted. Why are some illnesses at the top of Figure 7.3 and some at the bottom? We think the illnesses at the top are more of the chronic variety, whereas those at the bottom are more acute. We also think that the illnesses on the left are less serious than those on the right. We can test ideas like these by asking key informants to help us understand the arrangement of the illnesses in the MDS plot. (For more examples of mental maps, see Albert, 1991; D'Andrade et al., 1972; Erickson, 1997.) (There is a formal method, called property fitting analysis, or PROFIT, for testing ideas about the distribution of items in an MDS map. This method is based on linear regression. See Kruskal & Wish, 1978.)

2. Alternatively, profile matrices (the usual thing-by-variable attribute matrix ubiquitous in the social sciences) can be converted to similarity matrices (thing-by-thing matrices in which the cells contain measures of similarity among pairs of things) and then analyzed with MDS (for step-by-step instructions, see Borgatti, 1999).

◆ References

- Agar, M. (1979). Themes revisited: Some problems in cognitive anthropology. *Discourse Processes*, 2, 11-31.
- Agar, M. (1980). Getting better quality stuff: Methodological competition in an interdisciplinary niche. *Urban Life*, 9, 34-50.
- Agar, M. (1983). Microcomputers as field tools. *Computers and the Humanities*, 17, 19-26.
- Agar, M. (1996). *Speaking of ethnography* (2nd ed.). Thousand Oaks, CA: Sage.