Appendix A

Basic Concepts Underlying the PARAFAC-CANDECOMP Three-Way Factor Analysis Model and Its Application to Longitudinal Data[1,2]

Richard A. Harshman and Sheri A. Berenbaum

In the following pages we present the basic concepts and assumptions underlying the PARAFAC-CANDECOMP factor analysis model and consider possible benefits and disadvantages of this approach to analyzing longitudinal data.

MULTIMODAL RELATIONSHIPS

Two-Way versus Three-Way Data

To understand three-way factor analysis, it is important first to understand the difference between two-way and three-way data. A

[1]The authors appreciate the opportunity to work with the IHD staff on possible applications of PARAFAC. In particular we would like to thank Dr. Dorothy Eichorn, Norma Haan, and Marjorie Honzik for their contribution and support throughout this study. During the preparation of this appendix, S.A.B. was supported by NIMH Grant # MH 14647.

[2]PARAFAC and CANDECOMP are two procedures for three-way analysis that were developed independently but are essentially equivalent. PARAFAC stands for *PARA*llel *FAC*tors factor analysis and is described in Harshman, 1970, 1972b, 1976. CAN-

ISBN 0-12-233680-1

multivariate data set usually consists of a rectangular array of observations. Such sets are organized according to two intersecting "modes" or "ways" of classification. For example, in a set of personality measurements, the data point x_{ij} might correspond to the score on the ith personality variable which was obtained by the jth individual and thus the two modes of classification would be "variables" and "persons." Frequently, factor analysis deals with a matrix of correlations between variables; in this case the data point r_{ij} would represent the correlation between variables i and j. Here, both "ways" correspond to the same basis of classification—"variables." (This case is considered two-way because it takes two intersecting classifications—the row variable and the column variable—to locate a single datum.)

It is becoming more common to collect three-way data. For example, a two-way array might be collected for each of several experimental conditions, producing a three-way set where "conditions" represents the third mode. Most relevant here is the fact that longitudinal multivariate data usually have a three-way form of organization. For example, in a set of longitudinal personality measurements, x_{ijk} might correspond to the score of the ith personality variable, as measured on the jth individual, on the kth occasion. Or, the data might consist of correlations, so that the data point r_{ijk} corresponds to the correlation between variables i and j, computed from their values on the kth occasion. (If the data are correlations, it is not necessary that the same individuals be measured on the different occasions; a cross-sectional approach is also possible.) (For useful discussion of these and related distinctions among types of data see Carroll & Arabie, 1980.)

Two-Way versus Three-Way Factor Analysis

Traditional forms of factor analysis can only be applied to a two-way array. For this reason, longitudinal researchers have had to factor their three-way data by indirect means, e.g., "collapsing" the three-way array into a two-way summary matrix, by averaging over one of the three modes. However, such a procedure entails considerable loss of information. In particular, it does not allow the analysis to incorporate any of the three-way structural relationships among the observations. To reduce loss of information, longitudinal analyses are often done by dividing the data set into a series of two-way "slices" and then perform-

---

DECOMP stands for *CAN*onical *DECOMP*osition and is described in Carroll and Chang, 1970. We sometimes refer to the basic model as the "PARAFAC-CANDECOMP model," but often call it simply the "PARAFAC model" for short.

ing a separate factor analysis on each slice. However, it is then necessary to compare the factors across slices in order to recover three-way information. Here again, any systematic three-way relationships are not available to the factor analysis procedure itself. Another limitation of such piecewise approaches is that there are several different ways to "slice" a three-way longitudinal array. The most obvious, perhaps, is to produce a two-way slice for each occasion, but one might instead want to make a slice for each person, so that each slice would be a two-way longitudinal matrix of variables by occasions. Alternatively, it is sometimes useful to consider a slice for every variable, so that each slice is a longitudinal matrix of persons by occasions. With each method, different aspects of the three-way structure of covariation are revealed (see, e.g., Cattell, 1952).

In the last 10–15 years several new types of factor and principal component analysis have been developed which can *directly* analyze three-way data sets (e.g., Bentler & Weeks, 1979; Carroll & Chang, 1970; Corballis, 1970; Harshman, 1970, 1972a, 1976; Jöreskog, 1971; Kroonenberg & de Leeuw, 1980; Sands & Young, 1980; Tucker, 1963, 1966). Yet, surprisingly, few attempts have been made to apply these methods to longitudinal data. The potential benefits of such application seem considerable: several of these three-way methods should be able to extract common factors that describe the covariation of variables *both* across persons and across time, without losing the distinctness of the two aspects of variation. Some of these models promise uniquely determined factors—potentially resolving the ambiguities of rotation that have hampered two-way factor analyses. Although the three-way methods vary in their assumptions, generality, and in the complexity of the resulting solutions, many of them should prove valuable for studying longitudinal multivariate patterns. We have concentrated on exploring applications of one of these techniques, PARAFAC.

THE PARAFAC-CANDECOMP MODEL

PARAFAC provides what is probably the simplest three-way generalization of traditional factor analysis, yet it is also one of the most powerful—when its particular assumptions are appropriate to the data. But these assumptions are somewhat restrictive. PARAFAC is not a completely general model appropriate for any type of three-way data. For example, some data may need to be transformed before they are suitable. Therefore, a clear understanding of the model is important not only for correct interpretation of PARAFAC solutions, but also for determining its correct application to various types of three-way data. We consider first a formulation of the model for analysis of raw data (or centered and rescaled data). Then we develop the equivalent model for analysis of covariances computed from the raw data.

Recall that the traditional *two-way factor model* for a raw data point can be written (for q factors) as follows:

$$x_{ij} = a_{i1}f_{j1} + a_{i2}f_{j2} + \cdots + a_{is}f_{js} + \cdots + a_{iq}f_{jq} + e_{ij} \tag{1}$$

Consider an interpretation of this equation in terms of personality factors. For such data, x_{ij} would represent the score obtained by the jth person on the ith personality variable or item. In the model, a_{i1} represents the loading of the ith variable on the first factor, that is, the amount that the ith variable is affected by (or measures) the first factor, while f_{j1} represents the loading of the jth person on the first factor, that is, the amount that the first factor is expressed in that individual—how much of that factor the individual has. (This person loading is traditionally known as the individual's "factor score" on that factor.) By multiplying these two loadings together, the term $a_{i1}f_{j1}$ represents the contribution of the first factor to the data point x_{ij}. (It is apparent, for example, that if either the jth person's loading or the ith variable's loading on that factor is very small, the factor will not contribute appreciably to the observed data point x_{ij}.) In like manner, the term $a_{i2}f_{j2}$ refers to the contribution of the second factor, $a_{is}f_{js}$ represents the contribution of factor s, and so on for the rest of the q factors. The term e_{ij} is an error term and represents the unique or unsystematic part of the data point that cannot be fit by the q common factors. (For further details on e_{ij}, see section on assumptions about error terms.)

The PARAFAC-CANDECOMP Model Applied to Score Matrices

To continue our example using personality item scores, suppose we are dealing with a *longitudinal* array of such scores, where a given set of items was administered to the same individuals on several different occasions. If we applied the PARAFAC model directly to such data, the underlying structure of a personality item score would be represented as follows:

$$x_{ijk} = a_{i1}f_{j1}w_{k1} + a_{i2}f_{j2}w_{k2} + \cdots + a_{is}f_{js}w_{ks} + \cdots + a_{iq}f_{jq}w_{kq} + e_{ijk} \tag{2}$$

In this equation, x_{ijk} represents the score on the ith personality item, as obtained by the jth person, on the kth occasion. On the right-hand side of the equation are $q + 1$ terms. The first q of these terms represents the contributions of the q common factors to the observed score; the last term is an error component. In this three-way model data point, the a and f coefficients have the same meaning as in the two-way model: they are variable loadings and person loadings. However, in each term there is an extra coefficient, the w_{ks} coefficient (e.g., w_{k1} for factor 1, w_{k2} for factor 2, etc.). These w_{ks} coefficients are occasion loadings, and are perfectly analogous to the a_{is} and f_{js} coefficients taken from the two-way model. For example, w_{k1} is the loading of the first factor on the kth occasion, that is, the amount that the first factor tends to be expressed on the kth occasion, the size or importance of the factor on that occasion. (In more precise terms, w_{k1} is proportional to the standard deviation of the contributions of the first factor on the kth occasion, provided these contributions have zero mean; otherwise, w_{k1} is proportional to the root-mean-square average of the contributions of factor 1 on the kth occasion.)

In terms of the *algebra* of the model, Eq. 2 represents perhaps the simplest and most mathematically straightforward three-way generalization of factor analysis possible. Conceptually, however, it involves more than might at first be apparent. PARAFAC was not developed simply because it seemed a natural extension of the two-way model. Rather, it is a formalization and extension of Cattell's (1944) insightful theoretical approach to solving the problem of factor rotation. This approach postulates that any "real" factor present on two occasions will maintain, across the two occasions, a simple proportionality of factor loadings or factor scores. In PARAFAC, the w_{ks} values are the proportionality coefficients which relate the loadings of a given factor across occasions. When this proportionality assumption is appropriate, the factor axes which reveal this proportionality are *uniquely* located in the factor space, the "real" factors are determined in a way not possible with two-way factor analysis. (This point is discussed in more detail in the section on uniqueness.)

Such proportionality is not, however, always a reasonable assumption. Consider, for example, the problem that arises when we try to apply the PARAFAC model in a simpleminded way to longitudinal personality data. The model implies an exceptional orderliness in the way that individuals change their personalities across time.

When interpreting the model, we can imagine that the w coefficient multiplies either the person loading (factor score) or the item loading—the results are mathematically equivalent. With the personality data example, it is perhaps more appropriate to consider the item loadings as fixed and the individuals' factor scores as changing across occasions. Under this interpretation, an individual's loading (factor score) at time k for factor s can be written $f_{js}w_{ks}$. This implies that on occasion k, everyone's loading on factor s goes up or down in proportion to the occasion loading w_{ks}. For example, if $w_{3s} = 1.0$ and $w_{4s} = 1.3$, so that factor s increases 30% from occasion 3 to occasion 4, then everyone's factor score is increased by exactly 30% between these two occasions. If factor s were Extroversion, then everyone would become 30% more extroverted during this period. Such an orderly pattern of change seems unlikely when one is dealing with individual personalities. It seems more reasonable to expect that the circumstances of each person's life will produce a change in any given personality factor that is specific to the individual. Each person may change by a different amount, and some may even change in the opposite direction from others.

The pattern of variation in which every data source shows the same percentage increase in the expression of a given factor across a particular time period may be more appropriate for the description of data drawn from points in some causally interconnected system. For example, if the data sources were points in an individual's brain, and the variables were measurements of different kinds of brain activity, it might be reasonable to expect all points of observation to show a coordinated proportional increase in the amount of a certain brain activity from one time to the next. For such reasons, the type of orderly variation implied by Eq. 2 has been called "system variation" (see Harshman, 1970, pp. 19–21 for further discussion). In contrast, when the sources of variation across time are thought to reside independently inside each data source (e.g., persons), the resulting variation is termed "object variation" (Harshman, 1970, pp. 22–25). In the object variation case, one might write the expression for person j's factor score for occasion k as follows: $f_{jsk} = (f_{js} + v_{jsk})$, where v_{jsk} represents the idiosyncratic shift or variation in the factor score for person j on occasion k.

Some longitudinal data sets may display a mixture of object variation and system variation. Even with personality data, there may be in the life span of most persons certain common experiences that cause a trend toward increases or decreases in particular factors at particular times (e.g., a factor of "responsibility" might generally tend to increase as individuals age). By applying the PARAFAC system variation model directly to the raw score or deviation score matrices for such data, one would obtain w_{ks} values that would represent the systematic part of the cross-time variation. Such an approach may be more appropriate with longitudinal growth data, for example, than with personality data. However, it is possible to factor analyze both the system variation and object variation parts of the temporal variation by treating the data in a slightly less direct fashion, that is, by taking covariances or deviation cross-products of the variables on each occasion and then analyzing the resulting three-way set of covariance or cross-product matrices according to the PARAFAC model. Both system and object variation will contribute to the covariances. Furthermore, though the distinction between system and objection variation disappears in covariance analysis, it can be recovered. If desired, the two types of variation can be reconstructed *after* the factor analysis, and studied separately or in combination. This is the approach that was taken in our work at IHD.

The PARAFAC-CANDECOMP Model Applied to Covariance Matrices

Before we discuss the PARAFAC model for covariances, let us first recall, for comparison, the traditional two-way factor model for cross-products or covariances:

$$c_{ij} = a_{i1}a_{j1} + a_{i2}a_{j2} + \cdots + a_{is}a_{js} + \cdots + a_{iq}a_{jq} + e_{ij} \qquad (3)$$

Here, a_{is} and a_{js} represent the loadings of personality variables i and j on factor s, as before. The observed covariance c_{ij} is simply the sum of the cross-products of the factor loadings for the two variables, plus error. (For simplicity, discussion here is restricted to the formulation for orthogonal factors.) We discuss covariances, rather than correlations, because the covariances properly reflect changes in variance across occasions: computation of correlations on each occasion would impose a standardization within each occasion that would disturb the desired proportional relationships between factors across different occasions.

The three-way generalization of this covariance model can be derived from Eq. 2 (Harshman, 1972a). It has the following simple form:

$$c_{ijk} = a_{i1}a_{j1}w^2_{k1} + a_{i2}a_{j2}w^2_{k2} + \cdots + a_{is}a_{js}w^2_{ks} + \cdots + a_{iq}a_{jq}w^2_{kq} + e_{ijk} \qquad (4)$$

Here c_{ijk} is the covariance between variables i and j, computed across persons, on occasion k. (Alternatively, it can be the sum rather than the average of deviation cross-products, the covariance "uncorrected for sample size," as in Chapter 5.) The a_{is} and a_{js} represent the loadings of variables i and j on factor s, as before. The w^2_{ks} coefficients are the occasion loadings, indicating the size or importance of factor s on occasion k. These occasion loadings correspond to those of Eq. 2, with one difference: whereas the raw-score formulation (Eq. 2) was in terms of w_{ks} coefficients that are proportional to the *standard deviation* of factor s on occasion k, the covariance or cross-product formulation of Eq. 4 involves w^2_{ks} coefficients that are proportional to the *variances* of the factors, that is, these loadings are exactly the squares of the w_{ks} loadings of Eq. 2. (The occasion loadings are squared in Eq. 4 because the cross-product of two variables contains a w_{ks} coefficient for each variable.)

Note that any reference to an individual's factor score has disappeared from Eq. 4. As a result, it can be shown (Harshman, 1972a) that fitting Eq. 4 to a three-way set of covariances does not entail the strong assumption of system variation that is entailed by Eq. 2. Although Eq. 4 allows the overall factor score variance to change from one occasion to the next, it does not require that the percentage change in factor score size be the same for each person; only the average change in factor size appears in the equation. Thus it does not matter whether individual shifts are coordinated or not. Because the weaker assumptions of Eq. 4 are consistent with data showing object variation, this model may provide an appropriate means of analyzing longitudinal personality score data.

It may seem that going from Eqs. 2 to 4 results in the loss of valuable information about the factor scores of individuals. In fact, however, this loss is only temporary. We show below that factor scores for individuals can be recovered, once the common factor loadings are obtained.

UNIQUENESS OF SOLUTION

Significance of Uniqueness

We now consider what is perhaps the most important property of the PARAFAC-CANDECOMP model: its ability to determine factors uniquely, without "rotation." The potential significance of such a capability is apparent to anyone familiar with the history of factor analysis and the controversy that has surrounded the issue of factor rotation. In the personality domain, for example, Cattell has advocated oblique simple structure; Eysenck and Guilford have preferred orthogonal solutions while differing on how to bring additional empirical information to bear on the rotation process. As a result, differences in rotational procedure have repeatedly led to debates about the correct description and interpretation of dimensions underlying the variables in a given domain (Eysenck, 1977, versus Guilford, 1977, provides a recent example of such a dispute). As Comrey (1967) has said, "The rotation process has been the target of much criticism, and continues to be the weakest link in the entire [factor analysis] process."

Because of the variety of rotations possible with any set of factors, some psychometricians suggest that there is no such thing as a "real" or "correct" set of factors for describing the patterns of covariation in a particular data set. They interpret the different possible rotations of a given set of factors as representing alternative, equally valid, ways of describing the same set of relationships.

Indeed, when factors are used only for convenient description of a given two-way data set and not for inference beyond that data set, this relativistic attitude may be justified. For some investigators, however, factor analysis provides one of the main sources of empirical information on which to build new explanatory constructs; such constructs are intended to have broad application beyond the limits of a particular set of data. For these researchers, the attempt to identify "real" or "explanatory" factors is simply another expression of the time-honored scientific endeavor of attempting to identify the underlying causes of observed regularities. Certainly, theorists such as Cattell and Eysenck consider their factors to have such causal implications (Cattell, 1952; Eysenck, 1970, Chapter 12), and thus, for them, the dispute over alternative sets of factors is both meaningful and important.

Differently rotated factors will often give rise to different scientific hypotheses about underlying causes in a given data domain and thus will lead to different predictions about the outcome of possible experiments in that domain. (For example, different factor theories of intellectual abilities give rise to different predictions about patterns of intellectual change due to drugs, brain damage, etc.) Clearly, in such cases those factors that lead to the most accurate predictions should be considered the most accurate descriptions of "reality" in that domain. Of course, any scientific construct is only an approximation to "reality," and the constructs derived from factors are probably rougher approximations than many. Yet there seems no reason why factor-derived constructs should not be granted the same empirical status as other scientific constructs, as long as they play the same sort of role in generating scientific hypotheses that are empirically testable.

One real objection to using factors as starting points for the development of explanatory constructs is that they are often dependent on relatively arbitrary decisions concerning rotation. The choice of a "simple structure" rotation, for example, is often fairly hard to defend on empirical grounds, and when this is so, any claim of explanatory validity for the particular factors that result from such rotation has little internal support. If it is claimed that such factors have greater empirical validity, this claim usually needs to be supported by a series of external verifications of their distinctive predictions (i.e., predictions that would be implied by those factors, but not by factors obtained through alternative rotations; see Harshman, 1970 or 1976, for further discussion). Certainly the extensive series of investigations and interlocking experiments conducted by researchers such as Eysenck, Cattell, and Guilford give their findings much greater weight than the results of any one or two factor analyses in isolation. Nonetheless, the differences in rotational philosophy of these investigators has apparently led them down different paths, and their conclusions, therefore, represent somewhat different perspectives on the nature of personality.

PARAFAC was developed to help overcome the problems of rotation by strengthening the factor model itself. PARAFAC incorporates into the three-way extension of the factor model an important principle first conceived by Cattell (1944), which takes advantage of the extra information about factors present in three-way data to obtain a unique set of factors without rotation.

Unlike two-way data, where factors produced by many alternative rotations will fit the data equally well (forcing one to go outside the data to test the implications of a given factor rotation), with the PARAFAC model for three-way data, differently rotated factors will not, in general, fit the data equally well. Thus one can empirically test different potential factor rotations *within the same data set* from which the factors are being extracted. (Of course, further confirmation from experiments that go beyond the data would also be important.) With longitudinal data, for example, different rotations of the possible factors underlying a set of variables will give rise to different predictions concerning patterns of change in covariances among the variables across time. If the data are adequate, there will be one "rotation" (i.e., one set of a, f, and w values) that fits the data *across time* better than any other. By seeking and finding this unique set of factors, PARAFAC provides an empirically grounded basis for selecting the best "rotation," that is, the best candidates for "real" or "explanatory" factors underlying a given domain.

The Basis for Uniqueness: The Principle of Proportional Profiles

In 1944 Cattell proposed what he called "the principle of parallel proportional profiles" as an alternative to the simple structure criterion for selecting a preferred rotation of factor axes (Cattell, 1944). His idea was simple, yet powerful: by comparing the factors extracted separately from two different but related data sets, it should be possible to discover the "real" orientation of axes in the two solutions by finding that orientation that brings their factor loadings (or factor scores) into parallel, proportional correspondence across solutions. As long as the two data sets are not equivalent, but instead possess the same common factors in different relative proportions, there is only *one* rotational position that will reveal this correspondence. It can be shown that such a correspondence of proportional (rather than identical) loading patterns or factor scores would be very unlikely to arise in two data sets by chance or as a mathematical artifact. Hence, its discovery in real data would indicate some common empirical influences acting to organize both data sets, but in different degrees in each. Thus, Cattell argued that factors determined by rotation to parallel proportional profiles would have stronger empirical meaning than those given by other rotations of the same axes (Cattell, 1944; Cattell & Cattell, 1955).

Because of mathematical and computational difficulties, the proportional profiles criterion was not successfully implemented as a rotation technique (Cattell & Cattell, 1955), and so was subsequently neglected. The mathematical difficulties have more recently been overcome, however, by noting that the principle of proportional profiles implies a particular three-way generalization of factor analysis, namely the PARAFAC models already discussed (Harshman, 1970, 1976). (Interestingly, the equivalent CANDECOMP model was developed without reference to Cattell's idea. It is based instead mainly on a rationale growing out of a consideration of individual differences in multidimensional scaling; see Carroll and Chang, 1970.)

There is a sense in which the discovery of substantial proportional changes in loadings across occasions can be taken to constitute confirmatory evidence both for the empirical "reality" of the factors defined by a particular rotation of axes and for the incorrectness of other rotational positions. To see this, imagine the following oversimplified case. A researcher is comparing covariances among a set of variables on occasion 1 with the corresponding covariances on occasion 2. If the investigator finds only small, random differences between the covariances on the two occasions, the comparison will provide no reason to choose a particular rotation of the factors. Appeal must be made instead to "simple structure" or other principles invoked to determine rotations in two-way factor analysis. If, on the other hand, changes in covariances across time showed certain systematic patterns, these patterns could strongly imply a particular rotation of the factor axes. Suppose, for example, that the covariances could be divided into three sets: one set that consists of covariances showing no real changes across the two occasions, only small random fluctuations; a second set of covariances, all of which increase by approximately 20% (with random fluctuations around this value); and a third set of covariances, all of which decrease by approximately 35%. Suppose further, that the three sets of covariances are, in fact, covariances among three distinct sets of variables. What could an investigator conclude?

First, one would be compelled to acknowledge that the pattern of changes is far too systematic and coordinated to have happened by chance. Second, to explain all the covariances that shifted by the same proportion, one would infer that the variables involved must share some common influence, that is, the variance of a common factor. Third, one would justifiably prefer a rotation of three underlying factors that assigned each factor to one of the three groups of variables that show different shifts in covariances across time. With such a rotation, the pattern of changes across time could be simply explained in terms of changes in the variances of the factors. Other rotations would not provide an explanation of the changes across time, and, in fact, would be inconsistent with them. It would not be possible to have the *same* factors on all three occasions unless they were rotated to be consistent with the patterns of shift across occasions. Normally the situation will be more complicated, because factors will probably overlap in influence and affect some of the same variables. Furthermore, a given factor will affect different variables to different degrees. Yet the mathematical idea is still applicable (though the resulting patterns are harder to describe verbally) and the conclusion is still the same: certain coordinated patterns of change in covariances across time (or across experimental conditions, or whatever else is represented by the third mode) can help to identify the underlying factors that are changing. Further, such coordinated patterns of change provide *confirmatory evidence* for the particular rotation of factors that they help to identify.

Reliability Check

But how does one distinguish patterns of covariance-change across time that are only random sampling fluctuations from those systematic enough to establish unambiguously a preferred rotation for the factors? One method is to compare factors obtained in random split halves of the data. The degree of similarity of factors found in two subsets of a given data set can provide an impression of the stability of the factors obtained and of the likelihood that they would be found by someone analyzing a new sample of similar data. If approximately the same set of factors and orientation of factors is determined in two halves of one's subject sample, the characteristic patterns are evidently reliable enough to be repeatable, even with half as many subjects in each sample. Features of the solution that do not replicate across split halves should be interpreted with caution. (For examples of the use of split-half validation of PARAFAC factors, see Gandour & Harshman, 1978; Meyer, 1980).

Precautions in addition to checking for the replicability of particular features of one's solution are advisable. The factor model represents a particular type of linear additive approximation to the patterns present in the data. Such a simple approximation is often quite useful and adequate for empirical investigations, but this will not always be the case. One must be aware of the limitations of the model one is fitting to the data and sensitive to any indications that a more complex model may be required. (For example, certain types of violations of the model's linearity assumptions can be recognized and interpreted by looking for nonlinear relationships between PARAFAC factors, an approach less likely to be feasible with two-way factor analysis; see Harshman, 1970, Chapter 7.)

Necessary Conditions for Uniqueness

As is apparent from the foregoing discussion, PARAFAC depends on distinctive patterns of variation of factors across the third mode (e.g., time) to separate the influences of the different factors and thus determine the best orientation of factor axes, that is, to identify the unique factors. If, however, the factors do not change in distinct ways, but instead certain factors always change their sizes together, and to the same degree, across the third mode, then PARAFAC will not be able to resolve the influence of these factors into unique components.

Consider the following algebraic analogy (borrowed from Harshman, Ladefoged, & Goldstein, 1977). In the simple equation $x + y = 20$, there is no unique solution because an infinite number of x, y pairs will satisfy the weak constraints imposed by this specified relationship. This example corresponds to the kind of nonuniqueness of loadings that occurs in two-way factor analysis. However, if we consider additional information about the unknown parameters, for example, if we consider several related equations in parallel, such as $x + y = 20$ and $2x + 3y = 55$, and if we require the same values of x and y to satisfy both equations, we obtain a unique solution. Similarly, parallel analysis of several data sets in terms of a common set of factors (unknown parameters) can provide a unique solution for the factors. But it is essential that the coefficients of x and y do not have the same ratio to one another in the two simultaneous equations. If the second equation were $2x + 2y = 40$ it would not provide any new constraints on the solution, and so would not provide a unique solution. The ratio of the factor sizes for two factors must differ across at least some of the occasions in order for the three-way data set to determine a unique solution for those factors. (If most of the factors vary in independent fashion across occasions, but two factors always change size together, then all factors except those two will be uniquely determined.) For more detailed discussion of uniqueness, and mathematical proofs, see Harshman, 1970, 1972b, 1976; Kruskal, 1976, 1977.

Before a factor analysis is performed, distinct patterns of variations for each factor might be *expected* but one could not be *certain* that they would be present. After performing the analysis, however, one can partially verify the presence or lack of distinct patterns of factor variation across time by examining the table of occasion loadings (the w_{ks} coefficients). If the patterns of variation of these loadings across time are distinct for each factor, then the interpretation of the solution can proceed with greater confidence. However, if several columns have very similar patterns of loadings, then those factors may not be uniquely determined. One must proceed with caution in interpreting those factors and certainly test the solution for uniqueness and reliability by methods described above and elsewhere (e.g., Harshman *et al.,* 1977).

IMPLEMENTING THE MODEL: COMPUTATIONAL CHARACTERISTICS OF PARAFAC-CONDECOMP

To properly evaluate any application of PARAFAC, one must appreciate not only the differences between the two-way and three-way factor *model,* but also the resulting basic differences in computational procedures. Even with two-way factor analysis, one cannot simply "input the data and output the results." It is important to test various alternatives, select the most appropriate options, and make informed decisions at various stages of the analysis. With PARAFAC, an intelligent understanding of the basic stages of the computational procedure is even more important. There is a specific series of steps for determining the correct number of factors and ensuring that the optimal solution has been obtained for any given number of factors. Understanding these steps and evaluating how well the results support the conclusions drawn by the investigator is an important part of evaluating any PARAFAC analysis.

The computational characteristics of PARAFAC differ from those of conventional factor analysis in two basic respects: (*a*) the factors must be extracted simultaneously rather than successively, and thus a separate analysis must be performed each time a solution with a different number of factors is to be examined; (*b*) each solution is obtained iteratively, rather than by a "closed form" direct computation; consequently checks must be made to ensure that the iterative procedure has reached an optimal solution and that the same optimal solution will be obtained consistently.

Finally, for longitudinal applications of PARAFAC involving analysis of covariance, we propose a modified method of computing factor scores. This could be considered a third computational difference between our approach and conventional two-way factor analysis. It is useful to consider all three of these differences more carefully.

Simultaneous Extraction of Factors: Testing Effects of Different Numbers of Factors

Conventional two-way factor analysis normally proceeds by extracting factors stepwise, that is, first extracting one factor, then extracting a second factor from the variance that is left over after removing the first factor, and in this way removing succeeding factors from a given data set. In this conventional procedure, the first factor will be the largest, explaining as much variance as is possible with one factor, with the next factor explaining as much residual variance as possible, and so on. Such an approach imposes certain arbitrary restrictions on the form of the factors. First, it causes the factors to be orthogonal to one another in all modes. Further, because it maximizes the variance extracted by each successive factor, it forces the early factors to represent a combination of influences whenever such a combination can explain more variance than can a factor that represents a single underlying pattern of influence. Such "raw" unrotated factors are not normally interpretable.

In two-way factor analysis, the stepwise approach causes no problems because a subsequent rotation is used to "untangle" the combinations of influences in the unrotated factors. However, for PARAFAC the incorrectly rotated (sequentially extracted) factors would actually capture less of the data variance than the same number of factors in the correct orientation. To take advantage of the important uniqueness of PARAFAC in determining the best factor orientations, it is necessary to extract all the desired factors *simultaneously*. This procedure provides the best estimate of each factor, given the presence of others, and removes the arbitrary restrictions imposed by sequential extraction. Because of the requirement of simultaneous extraction, one must perform a separate analysis at each dimensionality to be examined. Thus, for example, one must perform a four-dimensional analysis, and then a separate five-dimensional analysis, to compare the form of the factors in four- and five-dimensional solutions. Generally, it will not be the case that the factors in the five-dimensional solution are identical to those found in the four-dimensional solution plus one additional factor.

Iterative Computation: Checking for Convergence, Optimality, and Uniqueness

To perform a PARAFAC analysis at a given dimensionality, an iterative procedure is used. Unlike two-way factor analysis, there is no direct "closed form" procedure for obtaining the optimal factor loadings. Instead, each solution starts with a random set of values for the variable loadings, person loadings, and occasion loadings (the a, f, and w coefficients). PARAFAC then improves this random first guess incrementally in small steps, increasing the fit of the solution to the data at each step, until a final optimal solution is reached. Stepwise changes are comparatively large at first, but gradually become smaller as the solution approaches its final form. When these steps become vanishingly small, the process is said to have converged. Obviously it is unreasonable to try to interpret a solution that is far from convergence, because additional iterations may change the loadings enough to alter the interpretation. Therefore it is necessary to establish, for each solution, that it is close enough to convergence to provide accurate estimates of the final loadings.

Because PARFAC proceeds by successive iterations, different sets of initial random loadings will cause the program to pass through a different series of intermediate stages on the way to the final converged solution. It is important to determine, therefore, whether the final solution obtained is completely determined by the data or is instead partially a function of the random initial loadings. With data that have the necessary independent variations of the underlying factors across all three modes (as explained previously in the discussion of uniqueness), different starting points should converge to the same final solution, as long as the number of factors being extracted does not exceed the number underlying the data. If too many factors are extracted, the solution may be nonunique and differ with different starting points. (In certain circumstances nonuniqueness may also occur when too few factors are extracted, but this is less common. If, e.g., the "true" number of factors is 5, there can be a unique 5-dimensional solution but also sometimes two competing 4-dimensional solutions, each with 4 factors that resemble a different subset of the full 5 factors.) Local optima can also cause problems by "trapping" the iterative procedure in an invalid partial solution before it reaches the overall best fitting solution for a given dimensionality.

To check that an obtained solution is the desired global optimum for that dimensionality and that it is in fact unique and independent of starting position, several analyses are conducted from different random starting positions and the resulting solutions and fit values are compared. The obtained factors are also compared with factors extracted at lower and higher dimensionalities to gain further information on their stability. (For more detailed discussion of these, and related checking procedures, see Harshman *et al.,* 1977.)

Method of Computing Factor Scores

Given item loadings for q PARAFAC dimensions, factor scores can be computed by the same regression technique often used in two-way factor analysis. To obtain an individual's q factor scores for a given occasion, that person's column of all item *ratings* (his/her column of raw data for that occasion) is approximated by a weighted combination of the q columns of all item *loadings* (the PARAFAC dimensions obtained from the covariance analysis). The q multiple regression weights that yield the best-predicting combination of factor loadings become the factor score estimates. They represent the amounts of the q factors which, in combination, would produce a set of ratings most resembling the individual's obtained ratings for the given occasion. Thus, they represent estimates of the amount of each factor that the person has on that occasion.

As is well known, the regression procedure for estimating factor scores generates an intermediate set of values called the *factor score coefficients.* These coefficients provide a simple way of estimating any individual's factor scores from his data. For each factor, the coefficients describe a weighted linear combination of the individual's observed item ratings that gives the best estimate of that individual's factor score on that factor. If a person has high ratings on those items that count strongly toward the particular factor, he/she obtains a high score on that factor. With three-way data, our method of computing the factor score coefficients is the same as is used in two-way factor analysis, that is, we simply take the generalized inverse of the factor loading matrix. However, our method of applying these coefficients differs in two respects from the standard procedure for two-way factor analysis: (*a*) we apply the coefficients separately to the subject's data on each occasion, and thus compute several sets of factor scores for each individual; (*b*) we apply the coefficients directly to the *uncentered raw data* to recover information about factor means.

In many applications, and particularly in longitudinal analysis, it would be very useful to be able to compare factor score means across groups or occasions. Because there has been some controversy about the proper way to handle changes in factor means when analyzing longitudinal data (e.g., Bentler, 1973, discusses problems with various approaches), it may be useful to explain and justify the approach we advocate (and which is used in Chapter 5). We remove any information on changes in factor means before performing the PARAFAC analysis and then recover the information at the time of estimating factor scores. Recall that the first step in computing item covariances (or, in Chapter 5, sums of deviation cross-products) involves centering the data across persons (i.e., removing the mean from each column of each of the data matrices according to the formula $\overset{*}{x}_{ijk} = x_{ijk} - \bar{x}_{i.k}$). Naturally this removes from the resulting covariances any information concerning mean changes in the factor scores across occasions (or in different subsamples if these correspond to different levels of k). This is desirable because it avoids possible "contamination" of within-occasion covariances by the effects of any cross-occasion (or cross-subsample) variance of the factor means that is not simply proportional to within-occasion changes in factor variances (as might occur with certain sex or cohort effects).

It can be shown that the recovered item loadings are *not affected* by this centering (aside from avoidance of possible "contamination," as mentioned). However, centering across persons *does* affect the factor scores. If we applied the factor score coefficients to person-centered data, we would get person-centered factor scores, that is, each factor score would represent a person's deviation from the average person's score on that factor for that occasion and subsample (for occasion k we would obtain $\overset{*}{f}_{jsk} = f_{jsk} - \bar{f}_{.sk}$). To recover "raw" factors scores, that is, scores that include the mean component, we apply the factor score coefficients directly to the raw, uncentered data. We can then study changes in factor means across times or subsamples by comparing the means of these factor scores. (Because the loadings for items are not affected by centering, except as noted above, they provide an appropriate description of uncentered as well as centered data. Consequently, the factor score coefficients obtained from such loadings are also applicable to uncentered as well as centered data. Artificial additive constants in the data might introduce small additive constants into the factor scores, but as long as these constants do not change appreciably over time they should not affect the differences in factor means of groups or occasions.) We believe this approach will generally solve the problems noted by Bentler (1973) and others regarding the treatment of mean factor changes in longitudinal factor analysis.

SPECIAL ASSUMPTIONS UNDERLYING PARAFAC ANALYSIS OF LONGITUDINAL DATA

In addition to the usual assumptions of factor or component analysis (e.g., that scores are linearly decomposable into factors plus error), certain further assumptions about the behavior of factors across time are required for PARAFAC analysis of longitudinal data. It is important to consider the various special assumptions or limitations that may be implied by using particular forms of the PARAFAC model. One can then evaluate the advantages and disadvantages of PARAFAC compared to various other factor analytic approaches to longitudinal data.

Factor Loading Invariance

Only one set of variable or item loadings is obtained by PARAFAC analysis of the three-way data. These are the a_{is} values of Eqs. 2 or 4. Consequently, it is implicit in the PARAFAC model that the pattern of loadings of variables on a given factor remains unchanged from one occasion to the next. In other words, the factorial content of variables is assumed to remain constant across occasions of measurement. This provides a desirable base of common comparison across occasions, but is seldom exactly true with variables that span long time periods. For example, this assumption may not be strictly valid for the type of personality data that we have been taking as an example; the meaning of some items—their interpretation in terms of underlying traits—may differ when used to describe a 40-year-old, compared to their meaning when applied to an adolescent. For any longitudinal data set being considered for PARAFAC analysis, the researcher must decide how serious this problem is, that is, whether the probable changes in factorial content of the variables would be so great as to render useless the PARAFAC approximation in terms of constant factorial composition across time. Some flexibility here seems appropriate, however. Unless it is reasonable to assume that one's variables measure at least roughly the same thing across time periods, then not only PARAFAC but most other types of longitudinal analysis or comparison become almost impossible.

It is often useful to obtain empirical information on the stability of factorial content of items across occasions by doing separate two-way factor analyses in the individual occasions or three-way analyses on earlier versus later subsets of the data. For example, these methods were employed to check the appropriateness of the PARAFAC analysis of Q-sort items described in Chapter 5.

The Nature of Factor Score Changes

We have chosen an interpretation of PARAFAC in which the changes across time are attributed to changes in the person loadings or "factor scores" of individuals, rather than to changes in the variable or item loadings (i.e., the w_{ks} is taken to multiply the f_{js} rather than the a_{is}). Other interpretations are possible, but this one seems to us to be the most appropriate for the particular personality example we have been considering. Having adopted this convention, let us consider what limitations are imposed on the patterns of factor score change when the PARAFAC-CANDECOMP model is applied to different types of data.

We have already seen how the direct application of PARAFAC to raw-score matrices imposes the limitation that factor score changes be proportional across occasions, that is, $f_{jsk} = f_{js} w_{ks}$. This strictest model does not even allow for shifts of the factor means across time except those proportional to changes in factor standard deviations. If one suspects that other, additive baseline shifts may be present, the data can easily be made suitable for PARAFAC analysis by centering the data for each occasion across persons before performing the factor analysis. (This removes any baseline shifts in factor scores and restores the required proportionality of factor changes. If the data are otherwise appropriate for direct PARAFAC, the centering will not affect the form of the extracted *variable* loadings and will only remove the means from the extracted person loadings. Information on baseline shifts in the factor means can be recovered, if desired, by factor score estimation procedures described above.)

Because of the assumption of system-variation in the direct application of PARAFAC to raw data matrices, this kind of application should be restricted to data in which system-variation is likely to be present (e.g., longitudinal growth data, economic system data, physiological measurements, or data where the third mode represents various experimental conditions designed to alter systematically the relative influence of the underlying factors; for examples see Harshman *et al.*, 1977; Meyer, 1980). For such data, however, direct PARAFAC has advantages over the indirect approach involving covariances. The direct approach allows oblique solutions if they provide a better fit to the data, whereas indirect analysis via covariances always gives orthogonal factors—unless the more general model PARAFAC2 is used[3] (see the following for further discussion of this point). The direct approach also yields three distinct sets of loadings rather than two, removing the need for estimating "factor scores."

When system variation is not likely to be found, a more indirect application of PARAFAC is warranted. If PARAFAC is applied to covariance matrices computed from the data rather than to the score matrices themselves, there are no restrictive assumptions imposed concerning the patterns of changes in factor scores *across occasions.* The data can follow object or system variation, or any intermediate pattern (Harshman, 1972a). Furthermore, although the separate computation of covariances on each occasion necessarily removes any information concerning mean changes in the factor scores across occasions, this information can be recovered. In fact, it is possible to estimate the specific object-variation pattern of each person's changes in factor scores across time. (See discussion of factor scores in the preceding.) Nonetheless, one restrictive assumption about factor scores is implied by using PARAFAC to analyze covariances, that is, that the factor scores for the different factors are orthogonal across persons within each occasion.

Orthogonal versus Oblique Factors

In dealing with the "obliqueness" of factors and of correlations among factors, it is necessary to specify which mode of the three-way data set is being considered. For example, two factors can vary in a

[3]A generalization of the PARAFAC model allows oblique factors to be extracted from three-way sets of covariance matrices. This model, called PARAFAC2 (Harshman, 1972a, 1976), has only been implemented in an experimental program that is not fully perfected. It is also more expensive to use because, at this point, it requires substantially more computer time. To avoid additional complications that would have been introduced by considering the PARAFAC2 model as well as PARAFAC, only PARAFAC was used in these initial longitudinal investigations.

correlated manner across persons and yet vary independently across occasions. Traditionally, those doing factor analysis of personality data have been concerned with the independence or correlation of factors *across persons*. To them, "oblique" factors were ones that had correlated factor scores. With our more general three-way model, we might want to consider the orthogonality or obliqueness of factors in any of the three modes; in terms of the personality score example, we might consider the independence of person loadings, occasion loadings, or variable loadings.

When PARAFAC is applied directly to the three-way raw data (or centered data) matrix, the factor loadings that emerge may be either oblique or orthogonal in any or all of the three modes. The results in a particular case will be uniquely determined by the structure of the three-way data itself. But in the analysis of covariances, one of the three modes "disappears" from the data. In the personality score example, the person mode disappears because covariances among items are computed across persons. The remaining two modes (variables and occasions, in our example) are still free to display orthogonal or oblique patterns of factor loadings, but to use the simple form of Eq. 4 for analysis of covariances, it is necessary to assume that the factor loadings are orthogonal in the mode that "disappeared."

We are involved in a trade-off of restrictions. In terms of our personality score example, the analysis of covariances allows us to relax the restriction that the variation of factor scores *across time* follow the system-variation pattern but imposes a new restriction of orthogonality on the variation of factor scores *across persons* on any one given occasion (i.e., for n persons we assume that $\sum_{j=1}^{n} = (f_{jsk})(f_{jtk}) = 0$, if $s \neq t$). (This restriction can be relaxed; see Footnote 3, p. 455.)

Although the restriction to orthogonality across persons is probably not exactly appropriate for most data sets, it will often serve as a reasonable simplifying assumption, particularly for the purposes of an initial investigation. When, for the "true" factors, the factor scores would actually have been correlated, then the solution obtained by analysis of covariances will provide a best orthogonal approximation to those factors. (In practice, however, factor score estimates obtained for these "orthogonal" factors may still be slightly correlated.) Unless the "correct" factors are strongly correlated across persons, this orthogonality restriction usually produces only a modest shift in the loadings that does not substantially alter interpretation of the factors. For object-variation data such as the personality score example, the trade-off of restrictions involved in covariance analysis is a worthwhile one.

Assumptions about Error Terms: Factor versus Component Analysis

When we set forth the equations defining the PARAFAC model for raw data (Eq. 2) and covariance matrices (Eq. 4) there were no specifications of the properties of the error terms e_{ijk}. Two approaches to these terms are possible, one corresponding to a three-way generalization of common factor analysis and the other to a three-way generalization of principal component analysis. Requiring that the error terms be uncorrelated across tests and persons results in a three-way form of common-factor analysis. Alternatively, specifying that these errors are simply the (correlated) residuals from a least-squares fitting of the q factors produces what is technically a three-way form of principal components analysis. For raw-data factoring, only the component-like analysis is currently programmed. However, for analysis of covariances, program options exist to implement either approach.[4] If the diagonals of the covariance matrices are left unaltered in the process of analysis, then a principal-components-like solution is obtained. If, on the other hand, the program option to iterate on the diagonal is selected, so that the diagonal cells are continually re-estimated in the course of the analysis, a common-factor type of solution is obtained. For large covariance matrices such as were analyzed in Chapter 5, it makes little difference which approach is taken. Because the diagonals constitute a very small percentage of the total data, the changes in their values that result from iteration have very little effect on the values of the factor loadings. (However, when employing the common-factor model, we should speak of our subsequently obtained person loadings as "factor score estimates" rather than "factor scores.")

Linear Independence of Factor Variations across Time

There is an additional assumption which is not necessary for the factor model of Eq. 2 or Eq. 4 to be appropriate, but which *is* necessary if one is to interpret the PARAFAC factors as "explanatory" and meaningful without rotation. The patterns of change of factor "size" or variance across time must be distinct for each factor. As discussed in the section on uniqueness these distinct patterns of change across time are

[4]A portable FORTRAN computer program for PARAFAC may be leased from Scientific Software Associates, 48 Wilson Avenue, London, Ontario, Canada N6H 1X3. Or, write to Richard A. Harshman, Department of Psychology, University of Western Ontario, London, Ontario, Canada N6A 5C2.

the clues that allow PARAFAC to resolve the factors. Loosely speaking, factors can be correlated across time, but should not be perfectly correlated. (In fact, variations in a factor generally should not be perfectly predictable from any linear combination of the variations in other factors. For a more precise statement of the required conditions for uniqueness, see Harshman, 1972b; Kruskal, 1976, 1977.)

Because the potential for a unique solution is one of the most important advantages of a PARAFAC analysis, the user should (*a*) make an effort before the analysis to ensure that the required distinct patterns of changes in the factor variance across time will be present in the data collected, (*b*) verify these distinct changes after the analysis by examining the occasion loadings, and (*c*) check the reliability of the factor orientations by performing additional split-half or "jack-knifing" analyses.

SUMMARY AND CONCLUSION

New methods of factor analysis may have great value for analysis of longitudinal data (as well as many other kinds of "three-way" data). They allow the entire three-way longitudinal data set to be entered into the analysis, without the usual requirement of "collapsing" the data into a two-way array. Because a more complete set of information is available to the analysis, a more complete and accurate description of the data can emerge from the analysis. For example, factor loadings for *occasions* as well as for variables, and factor scores for persons, can all be obtained from the same analysis.

In this discussion we have focused on the PARAFAC-CANDECOMP model for three-way factor analysis. This approach offers a particularly important advantage over two-way factor analysis, specifically, unique solutions. The additional information present in the three-way data (i.e., the pattern of variation of factors across time, as well as across variables and individuals) is used to remove the rotational ambiguity present in two-way factor analysis. Provided that (*a*) the factors show distinct patterns of variation across time, (*b*) the data are appropriate for the PARAFAC model, and (*c*) the solution is stable (e.g., across split-halves of the data), then the PARAFAC-determined factor axes have an empirical justification considerably stronger than any available in conventional two-way factor analysis: any rotation of those factors would reduce their ability to simultaneously fit all of the data matrices across the successive time periods.

When using the PARAFAC-CANDECOMP procedure for three-way factor analysis, it is important to keep in mind certain general assumptions implied by the model, along with the special assumptions involved in its application to longitudinal data. These considerations can be summarized as follows: investigators must be able to assume that the factorial content of their variables is relatively constant across time (i.e., that approximately the same factors are present on the different occasions), and that the major systematic differences across occasions arise from changes in the relative importance of the factors from one occasion to the next. When these assumptions are appropriate, the PARAFAC-CANDECOMP method of three-way factor analysis can be a powerful way to uncover meaningful factors and study their changes across time.