Identifying Spatial Variation Patterns in Multivariate Manufacturing Processes: A Blind Separation Approach

Daniel W. Apley and Ho Young Lee

Department of Industrial Engineering
Texas A&M University
College Station, TX 77843-3131

Large sets of multivariate measurement data are now routinely available through automated in-process measurement in many manufacturing industries. These data typically contain valuable information regarding the nature of each major source of process variability. In this article we assume that each variation source causes a distinct spatial variation pattern in the measurement data. The model that we use to represent the variation patterns is of identical structure to one widely used in the so-called "blind source separation" problem that arises in many sensor-array signal processing applications. We argue that methods developed for blind source separation can be used to identify spatial variation patterns in manufacturing data. We also discuss basic blind source separation concepts and their applicability to diagnosing manufacturing variation.

KEY WORDS: Blind source separation; Factor rotation; Manufacturing variation reduction; Multivariate statistical process control; Principal components analysis.

1. INTRODUCTION

In-process, or in-line, measurement technology in manufacturing is now widespread in many industries. In automobile body assembly, for example, laser-optical measurement stations are commonly built into the assembly line at various stages. In each measurement station, well over 100 key dimensional characteristics distributed over the auto body or its subassemblies may be measured. Moreover, 100% of the auto bodies produced are measured. The availability of abundant in-process measurement data creates tremendous potential for monitoring, diagnosing, and eliminating major root causes of process variation.

In any manufacturing process, there are generally numerous independent variation sources that contribute to the overall level of variability in the measurement data. Each variation source may result in a distinct spatial variation pattern across any or all of the process characteristics measured. The presumption throughout this article is that if we can identify the precise nature of each spatial variation pattern, then we can use this information as a diagnostic tool to facilitate the ultimate goal of identifying and eliminating the root causes.

This was the objective considered by Apley and Shi (2001) (hereafter A&S), who used the following model to represent the variation patterns. Let $\mathbf{x} = [x_1, x_2, \ldots, x_n]'$ be an $n \times 1$ random vector that represents a set of n measured characteristics from the product or process. Let $\mathbf{x}_i$, $i = 1, 2, \ldots, N$, be a sample of N observations of $\mathbf{x}$. In auto body assembly, for example, $\mathbf{x}$ would represent the vector of all measured dimensional characteristics across a given auto body, and N would be the number of auto bodies in the sample. It is assumed that $\mathbf{x}$ obeys the model

$$\mathbf{x} = \mathbf{C}\mathbf{v} + \mathbf{w}, \tag{1}$$

where $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_p]$ is an $n \times p$ constant matrix with linearly independent columns. The vector $\mathbf{v} = [v_1, v_2, \ldots, v_p]'$ is a $p \times 1$ 0-mean random vector with independent components, each scaled (without loss of generality) to have unit variance. The vector $\mathbf{w} = [w_1, w_2, \ldots, w_n]'$ is an $n \times 1$ 0-mean random vector that is independent of $\mathbf{v}$.

The interpretation of the model is that there are p independent *variation sources* $\{v_i : i = 1, 2, \ldots, p\}$ that affect the measurement vector $\mathbf{x}$. Each source has a linear effect on $\mathbf{x}$ that is represented by the corresponding column of $\mathbf{C}$. Together, $\mathbf{c}_i v_i$ describes the effect of the ith source on $\mathbf{x}$. The variation *pattern vector*, $\mathbf{c}_i$, indicates the spatial nature of the variation caused by the ith source. Because the elements of $\mathbf{v}$ are scaled to have unit variance, $\mathbf{c}_i$ also indicates the magnitude or severity of the ith source. The random vector $\mathbf{w}$ represents the aggregated effects of measurement noise and any inherent variation not attributed to the sources. Unless otherwise noted, it is assumed throughout that the covariance matrix of $\mathbf{w}$ is $\boldsymbol{\Sigma}_{\mathbf{w}} = \sigma^2\mathbf{I}$, a scalar multiple of the identity matrix. In Section 8 we discuss how to apply the methods of this article in situations where this assumption would not be reasonable. All random variables are assumed to be 0 mean. If not, then the mean of $\mathbf{x}$ should be estimated and subtracted from the data.

The objective is to estimate each of the variation pattern vectors in $\mathbf{C}$, as well as the number of sources, p, based on a sample of observations of $\mathbf{x}$. The estimated pattern vectors could then be used to illustrate the nature of the spatial variation and provide insight into the root causes. If the physics of the process were well understood, and one were able to analytically model the variation pattern vectors off-line, then the on-line task would reduce to pattern recognition or classification. Apley and Shi (1998) and Ceglarek and Shi (1996) presented methods for accomplishing this. In many manufacturing processes there are far too many potential variation sources to model, however,

TECHNOMETRICS, AUGUST 2003, VOL. 45, NO. 3
DOI 10.1198/004017003000000041

although the number p of actual sources present at any given time may be reasonable. In addition, the physics of the process may be so complex that analytical modeling becomes impossible. Consequently, the more challenging and less restrictive objective addressed in this article—estimating $\mathbf{C}$ "blindly" from on-line data with no off-line modeling required—has practical importance.

Because the model (1) is similar to what is assumed in standard linear orthogonal factor analysis and principal components analysis (Jackson 1980, 1981; Johnson and Wichern 1998), one may consider using those methods to estimate $\mathbf{C}$. Principal components analysis (PCA) and factor analysis produce an estimate of $\mathbf{C}$ that is unique only up to postmultiplication by a $p \times p$ orthogonal rotation matrix. The objective of factor rotation (Johnson and Wichern 1998) is to find the rotation matrix that yields the clearest interpretation of the resulting estimate of $\mathbf{C}$. The most common factor rotation methods (e.g., the varimax method) involve the optimization of somewhat artificial, predefined interpretability criteria. Depending on the structure of the true $\mathbf{C}$, which is determined by the nature of the variation sources and the underlying physics of the process, these methods may or may not rotate the estimate of $\mathbf{C}$ into something closer to the true value. This is discussed in more detail in Section 6 (also see A&S).

The method proposed by A&S can be viewed as a form of factor rotation that, rather than using some predefined interpretability criterion, attempts to rotate the estimate of $\mathbf{C}$ so that it is as close as possible to the true $\mathbf{C}$. This will presumably result in the clearest interpretability, in the sense of leading to the clearest understanding of the true nature of the variation sources and their root causes. To accomplish this, the method of A&S assumes certain structural constraints on the true $\mathbf{C}$. Although their assumptions are less restrictive than the implicit assumptions involved in the varimax method (see Sec. 6), they limit the method's applicability to some extent. In addition, the method of A&S involves a level of user subjectivity that may prohibit its use by process operators and engineers who have limited statistical training. A more generic and black-box method (i.e., requiring less user input) of blindly estimating $\mathbf{C}$ would have broader applicability and could be used by a wider audience.

A class of signal processing methods, commonly termed *blind source separation* (Cardoso 1998; Haykin 2000), seems to provide a solution to this problem. Blind separation methods were originally developed for processing sensor array (e.g., radar, sonar, wireless communication) signals. Although these signal processing problems differ in many respects from the manufacturing variation diagnosis problem described earlier, they use a model with a structure nearly identical to (1). The main purpose of this article is to demonstrate that existing blind source separation methods provide a reasonably effective, generic, and black-box means of uniquely identifying manufacturing variation patterns in the aforementioned context.

The remainder of the article is organized as follows. Section 2 provides an example illustrating the suitability of (1) for representing spatial manufacturing variation patterns. Section 3 briefly introduces the blind source separation problem in sensor array signal processing. Section 4 provides an overview of blind source separation methods, and Section 5 provides an illustrative example from automotive crankshaft manufacturing. Whereas the method of A&S imposes assumptions on the structure of $\mathbf{C}$, blind separation methods impose assumptions on the distribution of the sources. Section 6 compares the various assumptions and discusses some of their implications in the context of diagnosing manufacturing variation. Section 7 investigates the effects of violating the blind separation assumptions. Finally, Section 8 discusses strategies for applying blind separation methods when the noise covariance matrix is not a scalar multiple of the identity matrix.

2. ILLUSTRATION OF THE MODEL

In this section, the model (1) is motivated with an example from auto body assembly. Apley and Shi (1998) and A&S have provided more detailed discussions of similar examples. Figure 1(a) shows the measurement layout on the rear quarter panel subassembly of an auto body. The quarter panel subassembly consists of a quarter panel, joined to a D-pillar reinforcement. The measurements are taken after the quarter panel subassembly is joined to the body side (although the body side is omitted from the figure). The y-direction is toward the front of the vehicle, and the z-direction is up. The measurement vector $\mathbf{x}$ consists of $n = 10$ y/z-plane dimensional coordinates of five features [numbered 1–5 in Fig. 1(a)] on the quarter panel subassembly. Although roughly 200 dimensional coordinates

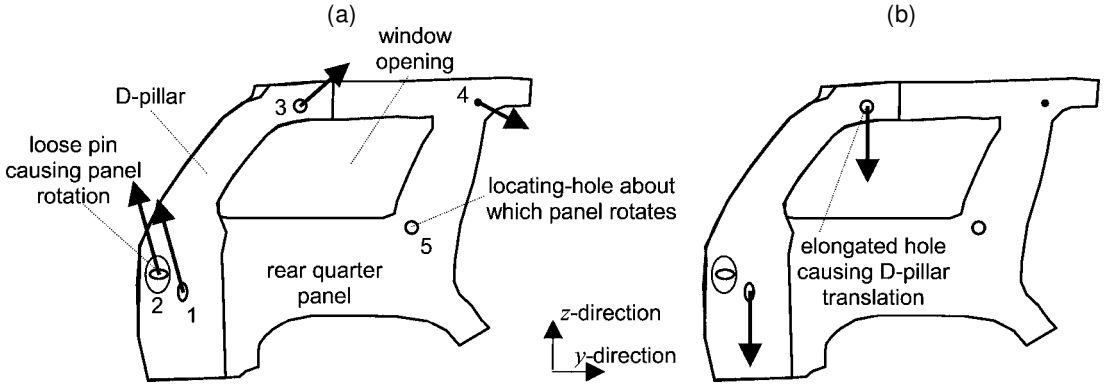

Figure 1. Two Linear Spatial Variation Patterns in Auto Body Assembly: (a) A Rotation of the Quarter Panel Subassembly About Locating Hole 5 and (b) a Translation of the D-Pillar in the z-Direction.

were measured on the entire auto body, this example considers only the quarter panel subassembly measurements for illustrative purposes.

Before the individual panels are welded together, they must be located accurately with respect to one another in fixtures and then clamped into place. The panels are positioned in the y/z plane using pins rigidly attached to the fixtures. Each pin mates with either a hole or a slot in the panel, so that the panel position is constrained but not overly constrained. For example, before the D-pillar is joined to the quarter panel, the quarter panel is located by pins that mate with one hole (feature 5) and one slot (feature 2) that are punched into the quarter panel after stamping. Likewise, the D-pillar is located by pins that mate with one hole (feature 3) and one slot (feature 1) that are punched into the D-pillar. When the complete quarter panel subassembly is subsequently joined to the body side, it is constrained by pins that mate with the same hole and slot (features 5 and 2) that were used to locate the quarter panel in the previous operation.

Through repeated use (1,000 panels per day may be placed into each fixture), the pins frequently become worn or loose. If this occurs, the panel will no longer be constrained to lie in its proper location when placed into the fixture. In one case study, the pin that mates with the feature 2 slot had become loose in the operation joining the quarter panel subassembly to the body side. When the quarter panel subassembly was placed into the fixture, it was free to rotate by a small amount in the y/z plane about the feature 5 hole. If an individual quarter panel subassembly rotated by (say) 1 degree in the clockwise direction, it retained the incorrect position when it was subsequently clamped into place and welded to the rest of the body side. For the next auto body, the quarter panel subassembly might rotate in the counterclockwise direction. From auto body to auto body, the loose pin caused a distinct spatial variation pattern in the measurement vector $\mathbf{x}$.

Referring to this as variation pattern 1, v_1 is the random variable (scaled to have unit variance) that represents the angle of rotation of a quarter panel subassembly. The vector $\mathbf{c}_1$ indicates that the pattern is a rotation of the quarter panel about feature 5, and the signs and magnitudes of the elements of $\mathbf{c}_1$ depend on the geometry of the panels and fixtures and the location of the measurements. The elements of $\mathbf{c}_1$ are plotted as arrows in Figure 1(a) at the locations of the features to which they correspond. The scaling is for visual convenience, and the y/z coordinates of each feature have been combined into a single arrow.

In the same case study, the feature 3 hole had been elongated in the z-direction due to improper stamping of the D-pillar. When each D-pillar was placed into the fixture before being welded to the quarter panel, the pin no longer completely constrained the D-pillar in the z-direction. This resulted in a second distinct variation pattern in which the D-pillar translated up on some autobodies and down on others. Referring to this as variation pattern 2, v_2 is the unit variance random variable proportional to the amount a D-pillar translates. The pattern vector $\mathbf{c}_2$ is shown in Figure 1(b). The elements of $\mathbf{c}_2$ are all 0, except for those that correspond to the z-direction coordinates of the D-pillar features.

This example illustrates not only the applicability of the linear model (1), but also how the results of estimating $\mathbf{C}$ might be used to gain insight into the root causes of the variation. The pattern vectors $\mathbf{c}_1$ and $\mathbf{c}_2$ illustrated in Figures 1(a) and 1(b) were, in fact, estimates produced by one of the methods described in Section 4. Based on this illustration, $\mathbf{c}_1$ and $\mathbf{c}_2$ were clearly interpreted as a rotation of the quarter panel subassembly about feature 5 and an up/down translation of the D-pillar. Operators and engineers familiar with the process tooling then identified a small set of potential root causes, and a brief investigation at the assembly line verified the actual root causes described earlier.

In the preceding example, it was implied that the variation patterns were linear. The rotation pattern depicted in Figure 1(a) is nonlinear, strictly speaking, because the exact relationship between the elements of $\mathbf{x}$ involves geometric functions such as sines and cosines. But the relationship is closely approximated as linear when small angles of rotation are involved. In practice, most manufacturing variation patterns would involve some degree of nonlinearity. As discussed by Apley and Shi (1998) and A&S, the model (1) may be viewed as the linearized approximation of a more exact nonlinear model.

3. THE BLIND SOURCE SEPARATION PROBLEM

Blind source separation is a term used to describe a number of related signal processing problems in which there is an array of spatially distributed sensors, each of which picks up signals from a number of distinct, signal-emitting sources (Cardoso 1998; Haykin 2000). Applications include radar and sonar signal processing, biomedical (e.g., electroencephalography, electrocardiography, fetal heartbeat) and geophysical signal monitoring, wireless communications, and speaker localization. We illustrate the situation with the classic blind separation example of speaker localization, sometimes referred to as the "cocktail-party problem." Suppose that there are a number of people (the sources) speaking simultaneously in a room, and that there are also a number of microphones (the sensors) spatially distributed throughout the room. Let p and n denote the number of speakers and microphones, as shown in Figure 2. Let $x_{i,t}$ denote the signal recorded by the ith microphone at time t, and let $v_{j,t}$ denote the speech signal emitted by the jth speaker at time t. Each microphone signal will generally be a mixture of source signals received from all of the speakers (typically assumed to

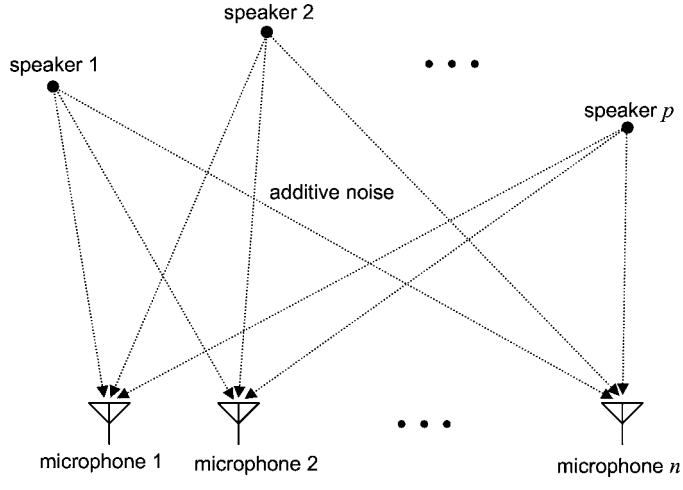

Figure 2. The Blind Separation Problem for Speaker Localization.

be a weighted linear combination) plus additive noise. Ignoring any time delays, we can write the model as

$$x_{i,t} = c_{i,1}v_{1,t} + c_{i,2}v_{2,t} + \ldots + c_{i,p}v_{p,t} + w_{i,t}, \qquad i = 1, 2, \ldots, n, \tag{2}$$

where each $c_{i,j}$ is a weighting coefficient that depends on the distance between the ith microphone and the jth speaker. The quantity $w_{i,t}$ is the noise affecting microphone i at time t. Combining (2) for $i = 1, 2, \ldots, n$, the form of the speaker localization model is identical to (1). Many other sensor-array signal processing problems yield a sensor/source model with this same linear structure. In radar signal processing, for example, the sources are p spatially distributed objects to be detected, and the sensors are an array of n spatially selective radar antennas. It is generally assumed that the p source signals are random and independent. The noise is sometimes assumed to be negligible in sensor-array signal processing, whereas in most manufacturing applications it would be nonnegligible.

The term "blind" in blind separation refers to the situation in which information on the sources must be determined solely from the data sample ($\mathbf{x}_t,\ t = 1, 2, \ldots, N$), with no prior knowledge of the relationship between the source and sensor signals other than the assumed structure of (1). To accomplish this, it is necessary to first estimate $\mathbf{C}$ "blindly" from the data. If the objective in the speaker localization problem were purely to identify and track the location of each speaker, then estimation of $\mathbf{C}$ would be the primary objective. Because the weighting coefficients contained in $\mathbf{C}$ depend on the distances between the speakers and the microphones, triangulation principles could be used to determine the location of each speaker. In radar, sonar, and other applications, the location of the sources can also be determined based on the weighting coefficients in $\mathbf{C}$ (Monzingo and Miller 1980). If, on the other hand, the objective in the speaker localization problem were to distinguish between the speech signals of each speaker, then estimation of $\mathbf{C}$ would be only an intermediate step. Once an estimate of $\mathbf{C}$ is obtained, straightforward linear regression could be used to estimate each of the p individual speech signals over the data sample. This is similar to the objective in typical wireless communications applications, where the original source signals bear some transmitted information.

In manufacturing variation diagnosis, it may be useful to estimate both $\mathbf{C}$ and the source signals. The columns of $\mathbf{C}$ provide information on the spatial nature of the variation patterns, and the estimated source signals provide information on the temporal nature of the variation over the data sample. Although the primary focus of this article is on estimating $\mathbf{C}$, the example in Section 5 illustrates how the source signals may be estimated and utilized.

4. BLIND SOURCE SEPARATION METHODS

Because the classic blind source separation model is identical to (1), many of the blind separation methods apply directly to manufacturing variation diagnosis. This section discusses two main classes of blind separation algorithms: second-order methods and fourth-order methods. Each class involves a different set of additional assumptions regarding the characteristics of the source distributions, discussed in Sections 4.2 and 4.3. Within each class there are many variants. Rather than provide a comprehensive survey of the different variants, we focus on a single method from each class that illustrates the main principles and that is relatively straightforward to implement. (Comprehensive discussions can be found in Cardoso 1998; Hyvarinen 1999; Hyvarinen and Oja 2000).

4.1 The Common Step: Principal Components Analysis

Most blind separation methods, including those discussed in this article, use a form of PCA as the first step. PCA, which involves the eigenvectors and eigenvalues of the covariance matrix of $\mathbf{x}$, is also the basis for most factor rotation methods. We describe the methods in the context that the true covariance matrix of $\mathbf{x}$, denoted by $\mathbf{\Sigma}_\mathbf{x}$, is known. In practice, one must work with an estimate of $\mathbf{\Sigma}_\mathbf{x}$ obtained from the data sample. To implement the methods discussed in this section, all distributional parameters would be replaced by their sample estimates.

From the model structure and assumptions, the covariance of $\mathbf{x}$ is

$$\mathbf{\Sigma}_\mathbf{x} = E[(\mathbf{Cv} + \mathbf{w})(\mathbf{Cv} + \mathbf{w})'] = \mathbf{CC}' + \sigma^2\mathbf{I}. \tag{3}$$

Let $\{\mathbf{z}_i : i = 1, 2, \ldots, n\}$ denote an orthonormal set of eigenvectors of $\mathbf{\Sigma}_\mathbf{x}$, and let $\{\lambda_i : i = 1, 2, \ldots, n\}$ denote the corresponding eigenvalues, arranged in descending order. It follows from (3) that $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_p > \sigma^2 = \lambda_{p+1} = \lambda_{p+2} = \cdots = \lambda_n$. Thus p is the number of dominant eigenvalues, and σ^2 is equal to any of the smallest $n - p$ eigenvalues.

A PCA decomposition of $\mathbf{\Sigma}_\mathbf{x}$ also yields

$$\begin{aligned}\mathbf{\Sigma}_x &= \sum_{i=1}^{n} \lambda_i \mathbf{z}_i \mathbf{z}_i' = \sum_{i=1}^{p} (\lambda_i - \sigma^2)\mathbf{z}_i\mathbf{z}_i' + \sigma^2 \sum_{i=1}^{n} \mathbf{z}_i\mathbf{z}_i' \\ &= \mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]\mathbf{Z}_p' + \sigma^2\mathbf{I},\end{aligned} \tag{4}$$

where $\mathbf{Z}_p = [\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_p]$, and $\mathbf{\Lambda}_p = \text{diag}\{\lambda_1, \lambda_2, \ldots, \lambda_p\}$. For the covariance structures in (3) and (4) to be consistent, $\mathbf{C}$ must be of the form $\mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{Q}$ for some $p \times p$ orthogonal matrix $\mathbf{Q}$.

Using PCA to find p, $\mathbf{Z}_p$, $\mathbf{\Lambda}_p$, and σ^2 is the first step of both the second-order and the fourth-order blind separation methods. The remainder of the problem reduces to finding $\mathbf{Q}$ and then using $\mathbf{C} = \mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{Q}$. Note that PCA provides no information regarding $\mathbf{Q}$, because any orthogonal matrix will result in a covariance structure (3) that is consistent with (4). When applying PCA to the sample covariance matrix, it may not be clear how many eigenvalues are dominant. A&S discussed various methods for estimating p. The average of the $n - p$ smallest eigenvalues is then used as an estimate of σ^2.

Rather than working directly with the measurements $\mathbf{x}$, blind separation methods usually work with a transformed version with spatially whitened (uncorrelated) components. The p-length whitened vector of measurements is defined as $\mathbf{y} = \mathbf{W}^{-1}\mathbf{x}$, where $\mathbf{W} = \mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}$, and $\mathbf{W}^{-1} = [\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{-1/2}\mathbf{Z}_p'$ is a left inverse of $\mathbf{W}$ ($\mathbf{W}^{-1}\mathbf{W} = \mathbf{I}$). From the relationship $\mathbf{C} = \mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{Q} = \mathbf{WQ}$, it follows that

$$\mathbf{y} = \mathbf{W}^{-1}\mathbf{x} = \mathbf{W}^{-1}[\mathbf{Cv} + \mathbf{w}] = \mathbf{Qv} + \mathbf{W}^{-1}\mathbf{w}. \tag{5}$$

Because $\mathbf{y}$ has diagonal covariance matrix $\mathbf{I} + \sigma^2[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{-1}$, $\mathbf{W}^{-1}$ is often referred to as the "whitening matrix."

4.2 Second-Order Methods

Second-order methods use only second-order statistics (covariance and autocovariance) of the data. These methods impose the additional assumption that of the p sources present, no pair has the exact same autocorrelation function. A necessary condition for this assumption to hold is that at least $p - 1$ of the p sources are temporally autocorrelated. The noise is assumed to be temporally uncorrelated.

Let $\mathbf{v}_t$ denote the source vector at time t. Define $\boldsymbol{\Sigma}_{\mathbf{v},\tau} = E[\mathbf{v}_t \mathbf{v}_{t+\tau}]$ to be the autocovariance matrix of $\mathbf{v}$ at lag $\tau \geq 0$. By the assumption of source independence, $\boldsymbol{\Sigma}_{\mathbf{v},\tau}$ is a diagonal matrix with diagonal elements $\rho_{1,\tau}, \rho_{2,\tau}, \ldots, \rho_{p,\tau}$, where $\rho_{i,\tau}$ is the autocorrelation function of v_i(the autocorrelation and autocovariance functions of the sources are equivalent, because they are scaled to have unit variance). From (5), the autocovariance matrix of $\mathbf{y}$ at lag $\tau \geq 1$ is $\boldsymbol{\Sigma}_{\mathbf{y},\tau} = E[(\mathbf{Q}\mathbf{v}_t + \mathbf{W}^{-1}\mathbf{w}_t)(\mathbf{Q}\mathbf{v}_{t+\tau} + \mathbf{W}^{-1}\mathbf{w}_{t+\tau})'] = \mathbf{Q}\boldsymbol{\Sigma}_{\mathbf{v},\tau}\mathbf{Q}'$. Because $\mathbf{Q}$ is an orthogonal matrix,

$$\mathbf{Q}'\boldsymbol{\Sigma}_{\mathbf{y},\tau}\mathbf{Q} = \boldsymbol{\Sigma}_{\mathbf{v},\tau} = \begin{bmatrix} \rho_{1,\tau} & & & \\ & \rho_{2,\tau} & & \\ & & \ddots & \\ & & & \rho_{p,\tau} \end{bmatrix} \tag{6}$$

is diagonal for all $\tau \geq 1$. Equation (6) forms the basis for second-order blind separation methods. Its significance is that the $p \times p$ orthogonal matrix $\mathbf{Q}$ that we seek is the matrix that jointly diagonalizes the entire set $\boldsymbol{\Sigma}_{\mathbf{y},\tau}$, $\tau \geq 1$. This joint diagonalizer is unique if the assumption holds that no pair of sources have the exact same autocorrelation function (see theorem 3 of Belouchrani, Abed-Meraim, Cardoso, and Moulines 1997).

When working with *sample* data, no single $\mathbf{Q}$ will jointly diagonalize the sample $\boldsymbol{\Sigma}_{\mathbf{y},\tau}$'s for all $\tau \geq 1$. Tong, Soon, Huang, and Liu (1990) proposed choosing $\mathbf{Q}$ to be the orthogonal matrix that diagonalizes $\boldsymbol{\Sigma}_{\mathbf{y},\tau}$ for a single specified τ. Belouchrani et al. (1997) improved the approach by choosing $\mathbf{Q}$ to be the orthogonal matrix that jointly approximately diagonalizes $\boldsymbol{\Sigma}_{\mathbf{y},\tau}$ for a set of τ's (e.g., $\tau = 1, 2, \ldots, 10$). Specifically, $\mathbf{Q}$ is chosen to minimize the sum of the squares of the off-diagonal elements of the set of matrices, $\mathbf{Q}'\boldsymbol{\Sigma}_{\mathbf{y},\tau}\mathbf{Q}$, for the specified set of τ's. Fortunately, there is a computationally efficient numerical method for accomplishing this. Belouchrani et al. (1997) have provided details of the algorithm, which is a generalization the Jacobi technique (Golub and Loan 1989) for exactly diagonalizing a single matrix.

The reason that the second-order method requires that the autocorrelation functions differ is somewhat apparent from the relationship $\boldsymbol{\Sigma}_{\mathbf{y},\tau} = \mathbf{Q}\boldsymbol{\Sigma}_{\mathbf{v},\tau}\mathbf{Q}'$. Consider the extreme case in which all p sources have the exact same autocorrelation function, $\rho_{i,\tau} = \rho_\tau$, $i = 1, 2, \ldots, p$. Then for each τ, $\boldsymbol{\Sigma}_{\mathbf{y},\tau} = \mathbf{Q}\boldsymbol{\Sigma}_{\mathbf{v},\tau}\mathbf{Q}' = \rho_\tau \mathbf{Q}\mathbf{Q}' = \rho_\tau \mathbf{I}$, a scalar multiple of the identity matrix. Thus *any* orthogonal matrix will diagonalize the entire set, and $\mathbf{Q}$ cannot be uniquely identified.

4.3 Fourth-Order Methods

As the name suggests, fourth-order methods use fourth-order statistics to uniquely estimate $\mathbf{Q}$ under their own specific set of additional assumptions. Whereas second-order methods impose assumptions on the source and noise autocorrelation, fourth-order methods impose the assumptions that no more than one of the p sources follows a Gaussian distribution, and that the noise is either negligible or follows a Gaussian distribution. Fourth-order methods can be derived as approximate maximum likelihood estimation (MLE) methods (Cardoso 1998). In addition to the foregoing assumptions, exact MLE methods typically assume that additional characteristics of the source distributions are known (e.g., that the sources follow uniform distributions). In this sense, the fourth-order methods involve a more relaxed set of assumptions and less a priori knowledge than exact MLE methods. There also is a computationally efficient algorithm for their implementation (Cardoso and Souloumiac 1993).

For an arbitrary 0-mean random vector $\mathbf{u} = [u_1, u_2, \ldots, u_p]'$, the fourth-order cumulant of its ith, jth, kth, and lth elements, $1 \leq i, j, k, l \leq p$, is defined as

$$C_{i,j,k,l}(\mathbf{u}) = E[u_i u_j u_k u_l] - E[u_i u_j]E[u_k u_l] - E[u_i u_k]E[u_j u_l] - E[u_i u_l]E[u_j u_k]. \tag{7}$$

Note that $C_{i,i,i,i}(\mathbf{u})$ is the kurtosis of u_i. Three important cumulant properties are as follows (Rosenblatt 1985; Stuart and Ord 1987): (a) If $\mathbf{u}$ is Gaussian, then all of its fourth-order cumulants are 0; (b) if $\mathbf{u}$ and $\mathbf{z}$ are independent and of equal dimension, then $C_{i,j,k,l}(\mathbf{u} + \mathbf{z}) = C_{i,j,k,l}(\mathbf{u}) + C_{i,j,k,l}(\mathbf{z})$; and (c) if the elements of $\mathbf{u}$ are independent, then all cross-cumulants of $\mathbf{u}$ are 0. A cross-cumulant is defined as $C_{i,j,k,l}(\mathbf{u})$ with $i, j, k, l \neq i, i, i, i$.

Let $\mathbf{U}$ be an arbitrary $p \times p$ orthogonal matrix, and consider the transformation $\mathbf{U}'\mathbf{y}$ of the whitened data. Because $\mathbf{w}$ is assumed to be Gaussian and independent of $\mathbf{v}$, properties (a) and (b) and eq. (5) imply that $C_{i,j,k,l}(\mathbf{U}'\mathbf{y}) = C_{i,j,k,l}(\mathbf{U}'\mathbf{Q}\mathbf{v})$. When $\mathbf{U}$ is the desired orthogonal matrix $\mathbf{Q}$, $\mathbf{U}'\mathbf{Q}\mathbf{v} = \mathbf{v}$ has independent components, and all cross-cumulants of $\mathbf{U}'\mathbf{y}$ are 0 by property (c). This fact forms the basis for fourth-order methods. The objective is to find the orthogonal matrix $\mathbf{U}$ that minimizes the cross-cumulants of $\mathbf{U}'\mathbf{y}$, and $\mathbf{Q}$ is then taken to be the minimizer. This can be viewed as finding the orthogonal transformation of (the already spatially uncorrelated) $\mathbf{y}$ whose components are as independent as possible, where the cross-cumulants provide the measure of independence. This bears a close relationship to PCA, in which the data are transformed to have uncorrelated, but not necessarily independent components. Because of this, Comon (1994) has referred to blind separation methods of this type as independent components analysis.

Comon (1994) has suggested taking $\mathbf{Q}$ to be the minimizer (over all $p \times p$ orthogonal matrices $\mathbf{U}$) of the sum of the squares of the entire set of cross-cumulants of $\mathbf{U}'\mathbf{y}$. Cardoso and Souloumiac (1993) proposed taking $\mathbf{Q}$ to be the minimizer of a similar criterion,

$$\sum_{\substack{1 \leq i,j,k,l \leq p \\ l \neq k}} C_{i,j,k,l}^2(\mathbf{U}'\mathbf{y}), \tag{8}$$

which involves only a subset of the cross-cumulants. Cardoso and Souloumiac (1993) showed that $\mathbf{U} = \mathbf{Q}$ is the unique minimizer of (8) if there is at most one Gaussian source (more precisely, if there is at most one source with zero kurtosis). The advantage of (8), known as the *joint approximate diagonalization of eigenmatrices* (JADE) criterion, is that there exists a computationally efficient method for finding its minimizer. Cardoso

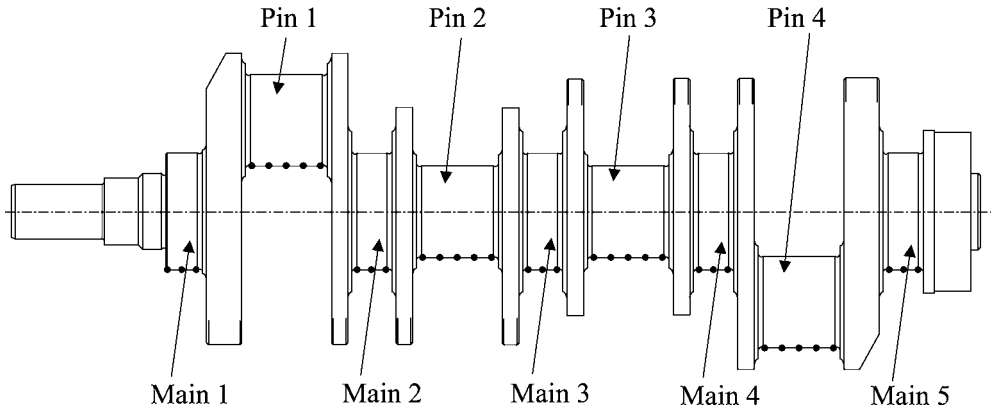

Figure 3. Geometry of the Crankshaft in the Example of Section 5.

and Souloumiac (1993) have shown that an equivalent expression for (8) is

$$\sum_{\substack{1\le i,j,k,l\le p\\ l\ne k}} C^2_{i,j,k,l}(\mathbf{U}'\mathbf{y}) = \sum_{\substack{1\le i,j,k,l\le p\\ l\ne k}} [\mathbf{U}'\mathbf{M}(i,j)\mathbf{U}]^2_{k,l}, \quad (9)$$

where $[\bullet]_{k,l}$ denotes the kth-row, lth-column element of a matrix, and each $p \times p$ cumulant matrix $\mathbf{M}(i,j)$ $(1 \le i,j \le p)$ is defined such that $[\mathbf{M}(i,j)]_{k,l} = C_{i,j,k,l}(\mathbf{y})$.

From (9), the JADE criterion is equivalent to finding the orthogonal matrix $\mathbf{U}$ that minimizes the sum of the squares of the off-diagonal elements of the set of transformed cumulant matrices $\{\mathbf{U}'\mathbf{M}(i,j)\mathbf{U} : 1 \le i,j \le p\}$. This gives rise to the "joint diagonalization" term in the JADE acronym. The "approximate" term in the acronym stems from the fact that with sample data, no orthogonal transformation will result in all sample cross-cumulants exactly equal to 0. The sample cumulant matrices can be only approximately diagonalized in the sense that (9) is minimized. The sample cumulants are defined in the obvious way by replacing the expectations of the quantities in (7) with their sample averages.

Because the second-order and fourth-order methods involve similar joint approximate diagonalizations, the same computationally efficient generalization of the Jacobi technique can be used for both cases. Cardoso and Souloumiac (1993) have provided details on the JADE algorithm, which is often used as a benchmark for evaluating other algorithms (Reed and Yao 1998; Wax and Sheinvald 1997). Matlab code is available on request from the authors.

5. AN ILLUSTRATIVE EXAMPLE

The automotive crankshaft manufacturing process comprises forging, rough cutting, finish cutting, drilling, grinding, and polishing. Figure 3 shows the geometry of a crankshaft produced in a manufacturing line in which an extensive amount of in-process measurement and inspection occur. Near the end of the line, for example, stylus traces around the circumference at a number of locations on the main bearings and pin bearings are taken automatically (for 100% of the crankshafts produced). The difference between the target diameter and the maximum, minimum, and average diameter at each location is then logged. (This example considers only the maximum diameter measurements.) The "$\bullet$" symbols in Figure 3 indicate the locations at which the diameter measurements are taken. The diameters are measured at three locations along each of the five main bearings (mains 1–5) and at five locations along each of the four pin bearings (pins 1–4). Thus the measurement vector $\mathbf{x}$ for each crankshaft consists of $n = 35$ diameter measurements.

Based on a sample of $N = 247$ crankshafts, it was estimated (using the methods discussed in A&S) that $p = 3$ major variation sources were present. Estimates of the three variation pattern vectors and the corresponding source signals using the fourth-order method are shown in Figures 4–6. The rotated version $\mathbf{Q}'\mathbf{y}_t$ of the whitened data was used as an estimate of $\mathbf{v}_t$. Each element of a pattern vector is represented as an arrow at the location of the corresponding diameter measurement. The length of the arrow is proportional to the magnitude of the element. (The same scaling was used in all three figures.) The sign of each element is represented by the direction of the arrow (pointing out of the crankshaft for a positive element and into the crankshaft for a negative element). Many elements of each pattern vector were negligibly small, in which case their arrows were omitted. Note that we could reverse the direction of all arrows without changing the meaning of the patterns. In other words, the ith pattern represents *variation* in $\mathbf{x}$ in both the positive $\mathbf{c}_i$ and the negative $\mathbf{c}_i$ directions. Whether the ith pattern causes a diameter to increase or to decrease on a particular crankshaft (say, crankshaft t) depends on whether $v_{i,t}$ is positive or negative. An arrow pointing out of the crankshaft, coupled with a positive source signal, represents an increase in diameter.

The three variation patterns were ordered in terms of decreasing severity, which is somewhat apparent by comparing the lengths of the arrows in Figures 4–6. We can quantify the severity of each source by noting that the *total variance* of $\mathbf{x}$ is

$$\sum_{i=1}^{n} E\left[x_i^2\right] = E[\mathbf{x}'\mathbf{x}] = \sum_{i=1}^{p} E[(\mathbf{c}_i v_i)'(\mathbf{c}_i v_i)] + E[\mathbf{w}'\mathbf{w}]$$

$$= \sum_{i=1}^{p} \mathbf{c}_i'\mathbf{c}_i + n\sigma^2.$$

Thus the contribution of the ith source to the total variance is $\mathbf{c}_i'\mathbf{c}_i$. The total variance for this example was 113, and the

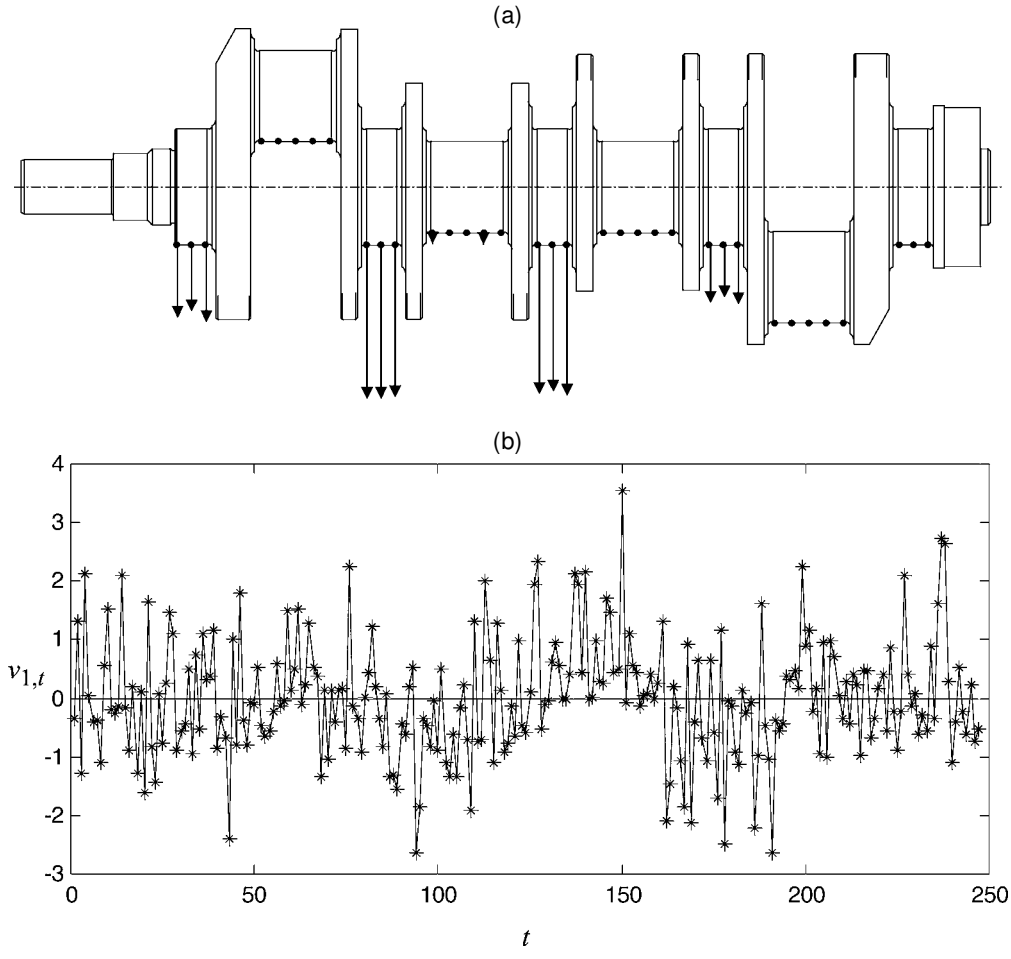

Figure 4. Estimates of the First Pattern Vector $\mathbf{c}_1$ (a) and Source Signal $v_{1,t}$ (b) in the Crankshaft Example. The first pattern affects only the main bearings and appears to be caused by shaft flexibility during rough cutting.

contributions of the three sources were $\mathbf{c}_1'\mathbf{c}_1 = 44{\cdot}5$ (39.3%), $\mathbf{c}_2'\mathbf{c}_2 = 17{\cdot}2$ (15.2%), and $\mathbf{c}_3'\mathbf{c}_3 = 16{\cdot}0$ (14.2%). Together, the three sources account for 68.7% of the total variance of $\mathbf{x}$.

The illustrations in Figures 4–6 might be used by process operators and engineers to aid in diagnosing the major root causes of variation in the bearing diameters. The first variation source, illustrated in Figure 4, has a pronounced effect on all but one of the main bearings and no effect on the pin bearings. Because the arrows on mains 1–4 all point in the same direction, this source causes all of the diameters on these bearings to either increase together or decrease together from crankshaft to crankshaft. The relative lengths of the arrows indicate that the diameters located near the middle of the crankshaft (mains 2 and 3) vary by a much larger amount (roughly 2.5 times larger) than the diameters located nearer to the ends of the crankshaft (mains 1 and 4). The diameters on main 5, which is nearest to the end of the crankshaft, do not vary at all. One possible explanation is that the large cutting forces generated during rough cut machining cause the middle of the crankshaft to flex more than the ends, which are held securely in chucks. We note that the first variation source alone accounted for 83.5% of the total variance of the main 2 and main 3 diameters.

The second variation source, illustrated in Figure 5, affects primarily pins 1 and 2. Main 2, which is located between pins 1 and 2, is slightly affected. The arrows in Figure 5 indicate that the second source causes the diameters on pin 1 to increase/decrease uniformly along its length, while simultaneously causing a taper along pin 2. The estimated source signal also provides insight into the temporal nature of the variation pattern that may aid in identifying its root cause. The plot of $v_{2,t}$ in Figure 5 shows that the second source tends to wander both above and below the zero value for extended periods of time, with a substantial shift occurring around the time of crankshaft 200. The third variation source, illustrated in Figure 6, affects only the pins. If we visually smooth the arrows for this pattern, then it appears that when the diameters near the middle of the crankshaft (pins 2 and 3) increase, the diameters near the ends of the crankshaft (pins 1 and 4) tend to decrease, and vice versa. The plot of $v_{3,t}$ in Figure 6 also reveals an interesting temporal pattern. A large portion of the variation in the third source is due to only a few spikes in the data, occurring at around the times of crankshafts number 60 and 180. By inspection of the plot of $v_{3,t}$, it appears that the third source may come from a mixture of two different distributions during these temporary periods of large variability. The suspected root cause relates to the use of parallel machines to perform the same operation at certain locations in the production line. The mixture distribution observed in $v_{3,t}$ most likely resulted from an intermittent problem experienced by only one of the parallel machines.

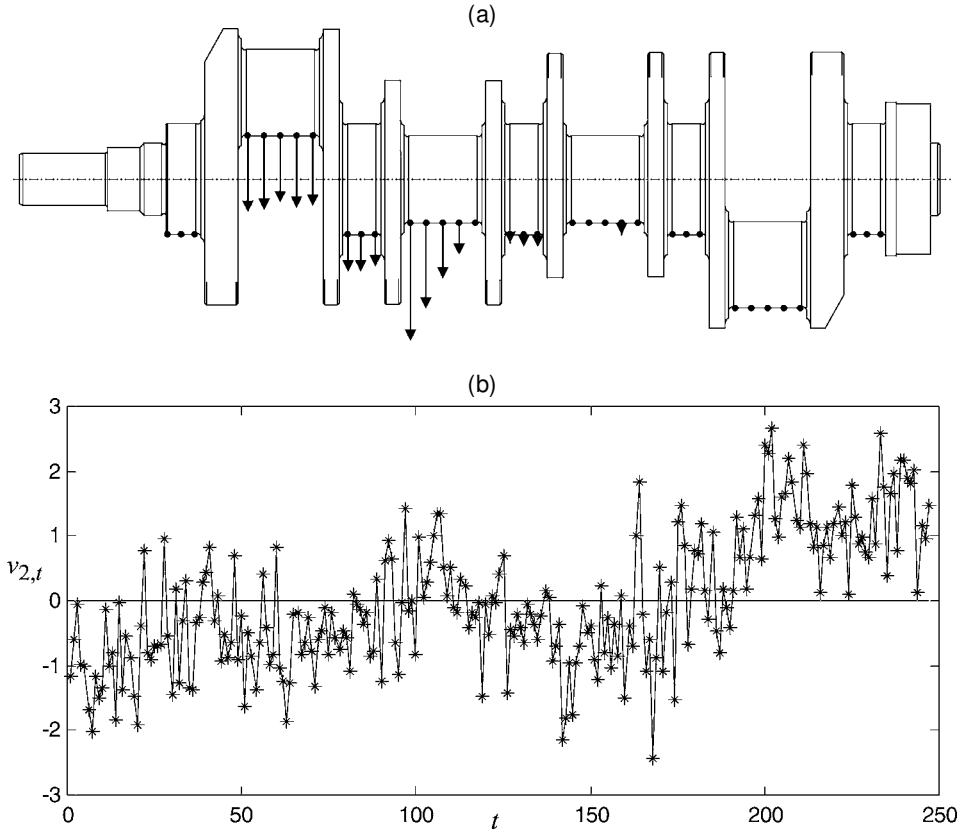

Figure 5. Estimates of the Second Pattern Vector $\boldsymbol{c}_2$ (a) and Source Signal $v_{2,t}$ (b) in the Crankshaft Example. The predominant effects of the second pattern are on pins 1 and 2, and the plot of $v_{2,t}$ indicates a substantial shift around crankshaft number 200.

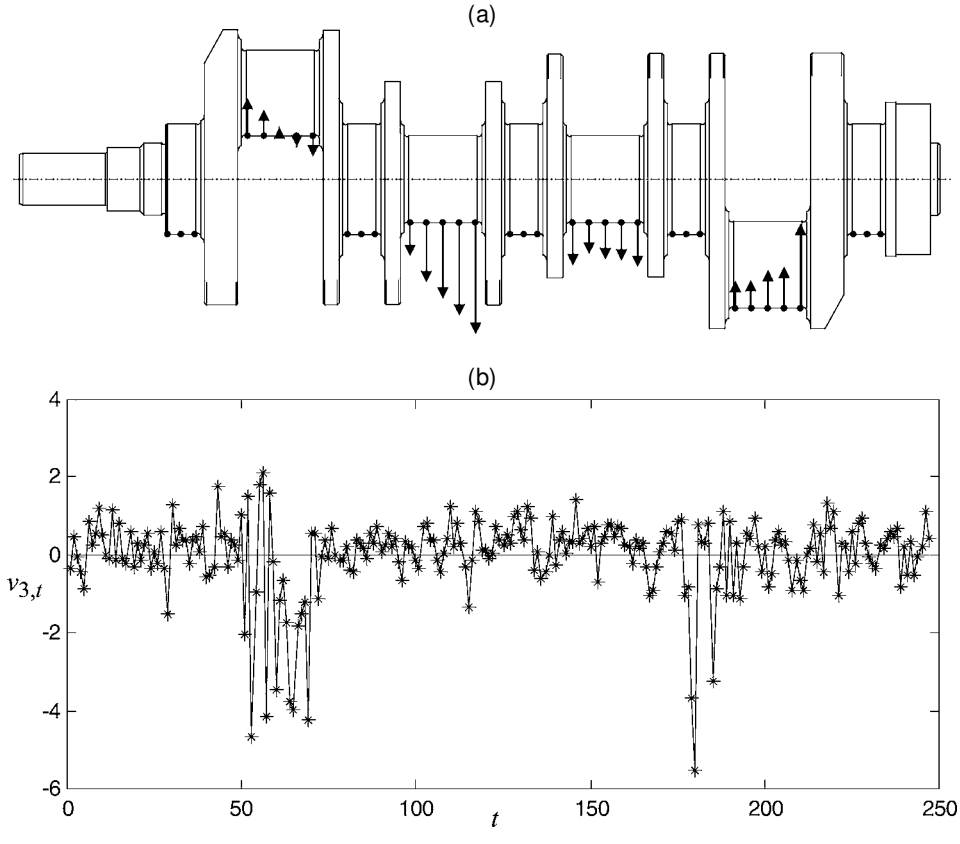

Figure 6. Estimates of the Third Pattern Vector $\boldsymbol{c}_3$ (a) and Source Signal $v_{3,t}$ (b) in the Crankshaft Example. The third pattern affects only the pin bearings, and the plot of $v_{3,t}$ indicates that the temporal nature of the source is intermittent bursts of large variability.

6. DISCUSSION OF THE ASSUMPTIONS

Virtually all methods that attempt to uniquely estimate $\mathbf{C}$ impose explicit or implicit assumptions regarding either the structure of $\mathbf{C}$ or the source distributions. This includes the method of A&S, blind source separation, and the varimax factor rotation method. Which method we could expect to produce a better estimate of $\mathbf{C}$ depends largely on which assumptions are better satisfied. The purpose of this section is to compare the various assumptions and provide guidelines for verifying whether they are satisfied.

6.1 Comparing the Assumptions

The varimax optimization criterion is intended to produce an estimate of $\mathbf{C}$ whose elements are either large in magnitude or small in magnitude, with as few moderate-sized elements as possible (Johnson and Wichern 1998). For any valid estimate of $\mathbf{C}$, the sum of the squares of the elements of any one of its rows is a fixed quantity (equal to the variance of the corresponding element of $\mathbf{x}$, minus σ^2). Consequently, the varimax method seeks an estimate of $\mathbf{C}$ whose structure is as close as possible to what we refer to as the *ideal varimax structure*,

$$\mathbf{C} = \begin{bmatrix} \mathbf{c}_{1,1} & & & \\ & \mathbf{c}_{2,2} & & \\ & & \ddots & \\ & & & \mathbf{c}_{p,p} \end{bmatrix}, \tag{10}$$

where $\mathbf{c}_{i,i}$ is an $n_i \times 1$ vector with $\sum_{i=1}^{p} n_i = n$. This assumes an appropriate reordering of the elements of $\mathbf{x}$. Hence the ideal varimax structure is that the p variation sources affect p *disjoint* subsets of the elements of $\mathbf{x}$. If the true $\mathbf{C}$ does not have this implicitly assumed structure, then the varimax estimate would most likely be inaccurate.

The method of A&S assumes that $\mathbf{C}$ has the ragged lower triangular structure

$$\mathbf{C} = \begin{bmatrix} \mathbf{c}_{1,1} & & & & \\ \mathbf{c}_{2,1} & \mathbf{c}_{2,2} & & & \\ \mathbf{c}_{3,1} & \mathbf{c}_{3,2} & \mathbf{c}_{3,3} & & \\ \vdots & \vdots & \vdots & \ddots & \\ \mathbf{c}_{p,1} & \mathbf{c}_{p,2} & \mathbf{c}_{p,3} & \dots & \mathbf{c}_{p,p} \end{bmatrix}, \tag{11}$$

where $\mathbf{c}_{i,j}$ is an $n_i \times 1$ vector with $\sum_{i=1}^{p} n_i = n$. An additional assumption that each n_i is strictly greater than 1 is required so that the subsets discussed later can be identified. The interpretation of (11) is that there exists a subset of n_1 measurements $\{x_1, x_2, \ldots, x_{n_1}\}$ that are affected by only a single variation source, which we call the *first source*. The effects of the first source on the first measurement subset is represented by $\mathbf{c}_{1,1}$. There must also exist a second subset of n_2 measurements $\{x_{n_1+1}, x_{n_1+2}, \ldots, x_{n_1+n_2}\}$ that are affected by only one of the remaining $p - 1$ variation sources, which we call the *second source*. But the second subset of measurements may also be affected by the first source ($\mathbf{c}_{2,1} \neq \mathbf{0}$), which is a major distinction between the ideal varimax structure and the assumed structure of A&S. There must also exist a third subset of measurements affected by only one of the remaining $p - 2$ sources (although these measurements may also be affected by the first two sources), and so on.

Comparing (10) and (11), it is clear that the ideal varimax structure is a rather restrictive special case of the structure assumed by A&S. Hence their method could be expected to produce a reasonable estimate of $\mathbf{C}$ in many situations where the varimax method would not. Consider the crankshaft example discussed in Section 5, and assume the true $\mathbf{C}$ coincides with the estimates shown in Figures 4–6. Because the second and third sources both have a strong effect on the diameter measurements for pins 1 and 2, $\mathbf{C}$ does not have an ideal varimax structure. In contrast, $\mathbf{C}$ does have the structure of (11). Pins 3 and 4 are affected by only a single source (the source illustrated in Fig. 6). The first measurement subset $\{x_1, x_2, \ldots, x_{n_1}\}$ would consist of the 10 diameter measurements taken on pins 3 and 4. Although pins 1 and 2 are also affected by this source, they are affected by only one of the remaining two sources—the source illustrated in Fig. 5. Thus the second measurement subset $\{x_{n_1+1}, x_{n_1+2}, \ldots, x_{n_1+n_2}\}$ would consist of the 10 diameter measurements taken on pins 1 and 2. The situation is similar for the two variation patterns depicted in Figure 1. Because features 2 and 4 are affected by only one of the two variation sources, $\mathbf{C}$ has the structure of (11), but does not have the ideal varimax structure.

Even when $\mathbf{C}$ has the structure (11) required in the A&S method, there is often ambiguity in selecting the measurement subset affected by a single source, for reasons discussed in Section 6.2. When this measurement subset is selected incorrectly or when $\mathbf{C}$ does not have the required structure, the method of A&S would not be expected to produce accurate estimates. Blind separation methods may still apply in these situations, because they make no assumptions regarding the structure of $\mathbf{C}$. This broader applicability with respect to the structure of $\mathbf{C}$ comes at the expense of narrower applicability with respect to the source distributions. Recall that the fourth-order method requires that no more than one of the p sources follows a Gaussian distribution, and the second-order method requires that no pair of sources shares the same autocorrelation function. The latter is equivalent to requiring that for each pair (i, j) with $1 \leq i \neq j \leq p$, there exists a $\tau = \tau(i, j)$ such that $\rho_{i,\tau} \neq \rho_{j,\tau}$. In other words, the second-order assumptions are satisfied as long as the autocorrelation functions for each pair of sources differ for at least one time lag (providing that the autocovariance matrix at this time lag is included in the set to be jointly diagonalized).

Lee and Apley (2003) developed a method for optimally combining the second-order and fourth-order joint diagonalization criteria to relax the blind separation assumptions required for uniquely identifying $\mathbf{C}$. They showed that the condition for uniquely identifying $\mathbf{C}$ using the combined criteria is that no pair of *Gaussian* sources shares the same autocorrelation function. This is weaker than the assumption in the second-order method that no pair of sources, whether Gaussian or not, shares the same autocorrelation function. It is also weaker than the assumption in the fourth-order method that no more than one source is Gaussian, because two or more Gaussian sources are allowed if their autocorrelation functions differ. Hence the combined method would have broader applicability then either the second-order method or the fourth-order method individually.

It should be noted that the blind separation conditions are theoretical conditions that result in the unique identification of $\mathbf{C}$ in

the hypothetical situation where the theoretical covariances and cumulants (or, equivalently, infinitely large samples) are available. With finite sample sizes, the performance of the methods depends on the extent to which their assumptions are satisfied, as illustrated in Section 7.

6.2 Verifying the Assumptions

Regardless of which method is used, an attempt to verify that its assumptions are satisfied is recommended. In this respect, the blind separation methods have an advantage over the method of A&S. For the fourth-order method, histograms of the estimated sources are useful for determining whether there is more than one Gaussian source. For the second-order method, a plot of the sample autocorrelation functions of the estimated sources is useful for determining whether a pair of sources shares the same autocorrelation function. Figures 7 and 8 show histograms and sample autocorrelation functions for the three source signals from the crankshaft example (with the estimated signals shown in Figs. 4–6). The distributions of the second and third sources appear to be non-Gaussian, which satisfies the assumptions of the fourth-order method. In contrast, the autocorrelation functions for the first and third sources appear to be quite similar, which would violate the assumptions of the second-order method.

When using the method of A&S, it is more difficult to verify whether the structural requirement (11) for $\mathbf{C}$ is satisfied. This relates to how one identifies the measurement subset that is affected by only one source. The strategy described by A&S is as follows. The latent covariance matrix is defined as the portion of $\boldsymbol{\Sigma}_{\mathbf{x}}$ due to the sources. From (3) and (4), the latent covariance matrix is $\mathbf{C}\mathbf{C}' = \mathbf{Z}_p[\boldsymbol{\Lambda}_p - \sigma^2\mathbf{I}]\mathbf{Z}_p'$, which can be determined from the PCA step. The latent correlation matrix is defined in the usual way from the latent covariance matrix. A&S showed that if a subset of measurements is affected by a single variation source, then the (theoretical) latent correlation coefficients of these measurements are all ± 1. The procedure for finding the subsets is to inspect the (sample) latent correlation matrix for a subset of measurements with latent correlation coefficients that are all close to 1 in magnitude.

Although a subset of measurements affected by a single source will have latent correlation coefficients that are all ± 1, the converse is not always true. Consider the situation where $n = 4$, $p = 2$, $\mathbf{c}_1 = [1 \quad 1 \quad 1 \quad 1]'$, and $\mathbf{c}_2 = [1 \quad 1 \quad -1 \quad -1]'$. The latent correlation matrix in this case is

$$\begin{bmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix}.$$

Because two measurement subsets $\{x_1, x_2\}$ and $\{x_3, x_4\}$ have unit magnitude latent correlation coefficients, we might incorrectly conclude that each of these subsets is affected by only

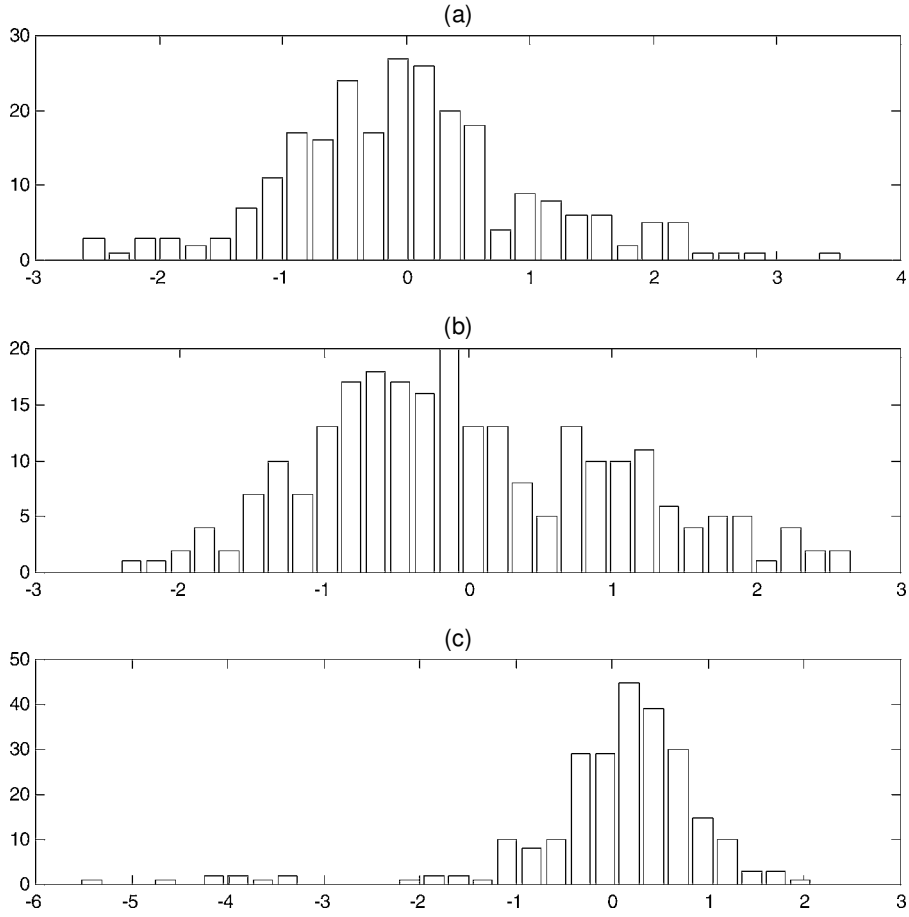

Figure 7. Histograms for the Three Source Signals From the Crankshaft Example: (a) v_1, (b) v_2, (c) v_3.

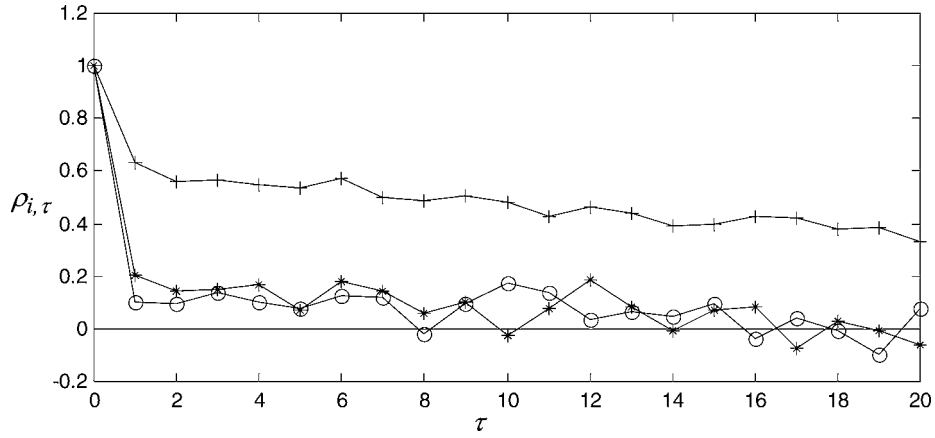

Figure 8. Sample Autocorrelation Functions for the Three Source Signals From the Crankshaft Example ($_\circ$, $\rho_{1,\tau}$; $_+$, $\rho_{2,\tau}$; $_*$, $\rho_{3,\tau}$).

a single source. The reason that these subsets have high latent correlation is that the effects on $\{x_1, x_2\}$ of the first and second sources are $[1 \quad 1]'$ and $[1 \quad 1]'$, which are identical. Likewise, the effects on $\{x_3, x_4\}$ of the first and second sources are $[1 \quad 1]'$ and $[-1 \quad -1]'$, which differ by only a constant scale factor. If the method of A&S were applied using either of these two subsets, then the estimates of $\mathbf{c}_1$ and $\mathbf{c}_2$ would be $[\sqrt{2} \quad \sqrt{2} \quad 0 \quad 0]'$ and $[0 \quad 0 \quad \sqrt{2} \quad \sqrt{2}]'$, which differ substantially from the true pattern vectors. Note that these also coincide with the varimax estimates. In situations like this, it is difficult to verify whether high latent correlation is the result of $\mathbf{C}$ truly having the structure of (11) or the result of two or more sources having exactly the same effect (up to a constant scale factor) on a measurement subset.

Consequently, there is a higher level of subjectivity involved in the A&S method than in blind separation methods. The primary subjectivity in blind separation lies in deciding whether the fourth-order method or the second-order method should be used, which relates to verifying whether the fourth-order assumptions or the second-order assumptions are better satisfied. It is relatively straightforward to do this using histograms and autocorrelation plots, as described earlier. This subjectivity may be reduced further if the fourth-order and second-order criteria are combined as was done by Lee and Apley (2003).

7. THE EFFECTS OF VIOLATING ASSUMPTIONS

The purpose of this section is to provide insight into how the performance of the blind separation methods is affected when their assumptions are violated or close to being violated. We use a simulation example in which a beam represents the part being manufactured and $n = 20$ measurements are distributed uniformly across the beam. There are two variation sources, with $\mathbf{c}_1$ and $\mathbf{c}_2$ as illustrated in Figures 9(a) and 9(b). For simplicity, the number of sources is assumed known, although in practice this must also be estimated. A&S discussed in detail a number of methods for estimating p.

The beam could be considered a subcomponent of a larger assembly, in which case the variation patterns may represent assembly variation. The first pattern would represent a rigid vertical translation of the beam; the second pattern, a rigid rotation about the beam centroid. Alternatively, the beam could be considered a separate part, in which case the variation patterns may represent fabrication (e.g., extrusion) variation. In this case, the first pattern would represent variation in the thickness of the beam that occurs uniformly across its length, and the second pattern would represent thickness variation that when larger on one end of the beam is smaller on the other end.

Both pattern vectors were scaled so that the total variance due to each source ($\mathbf{c}_i'\mathbf{c}_i$) was equal to the total variance due to the noise ($n\sigma^2$). In other words, the signal-to-noise ratio $\mathbf{c}_i'\mathbf{c}_i(n\sigma^2)^{-1}$ was unity for each source. The noise variance σ^2 was also unity. The first source follows a first-order autoregressive (AR) model $v_{1,t} = \phi v_{1,t-1} + a_t$, where $\phi = .9$ and the a_t's are 0-mean independent Gaussian random variables with variance $\sigma_a^2 = 1 - \phi^2$. First-order AR processes are widely encountered in industrial environments (Box, Jenkins, and Reinsel 1994). The variance and autocorrelation function of a first-order AR process are $\sigma_a^2(1 - \phi^2)^{-1}$ and $\rho_\tau = \phi^\tau$ $(\tau = 0, 1, 2, \ldots)$. Thus the marginal distribution of the first source is Gaussian

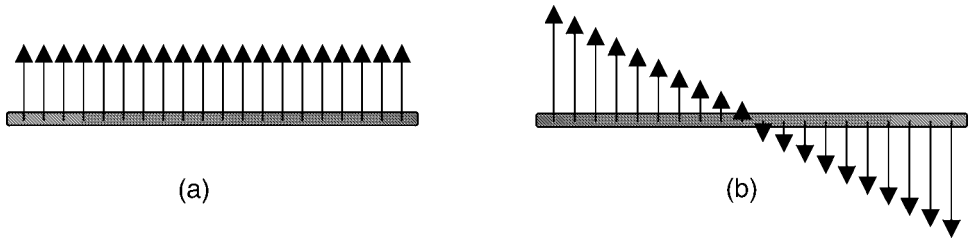

Figure 9. The Two Variation Patterns in the Example of Section 7: (a) $\mathbf{c}_1$, Which Represents a Beam Translation, and (b) $\mathbf{c}_2$, Which Represents a Beam Rotation.

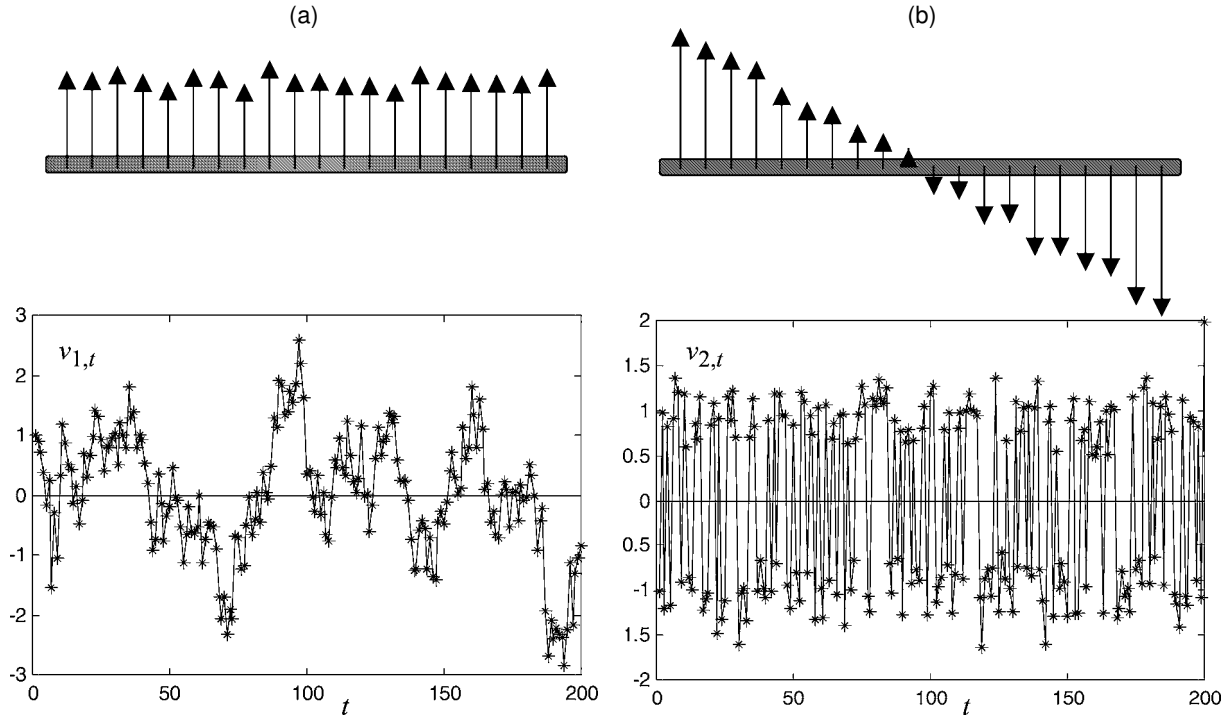

Figure 10. Estimates of the Pattern Vectors and Source Signals in the Simulation Example. (a) Estimates of $\mathbf{c}_1$ and $v_{1,t}$; (b) Estimates of $\mathbf{c}_2$ and $v_{2,t}$.

with 0 mean and unit variance, and its autocorrelation function is $\rho_{1,\tau} = .9^{\tau}$. The second source, $v_{2,t}$, follows a (scaled and shifted) Bernoulli distribution, where the two values ± 1 each occur with equal probability .5. Variation sources of this nature are also commonly observed in manufacturing, the root cause of which may be two parallel machines performing the same operation or the use of components or raw materials from two different suppliers. The second source was temporally uncorrelated. Because only one source is Gaussian and the autocorrelation functions for the two sources differ, the assumptions of both the second-order method and the fourth-order method are satisfied.

Figure 9 shows that $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2]$ has neither the ideal varimax structure (10) nor the structure (11) required in the A&S method. But the structure of $\mathbf{C}$ is close to (11), because the second source has very little effect on the two measurements that lie closest to the beam centroid. Ordering the measurements from left to right, we refer to these two measurements as x_{10} and x_{11}. The elements of $\mathbf{c}_1$ and $\mathbf{c}_2$ associated with $\{x_{10}, x_{11}\}$ are $[1 \quad 1]'$ and $[.087 \quad -.087]'$. It can be shown that the latent correlation coefficient for x_{10} and x_{11} is .985, which is relatively large. When $\{x_{10}, x_{11}\}$ was selected as the measurement subset affected by only a single source, the performance of the A&S method was quite similar to the blind separation performance discussed later. One must use caution when applying the A&S method in this example, however. The elements of $\mathbf{c}_1$ and $\mathbf{c}_2$ associated with the two left-most measurements $\{x_1, x_2\}$ are $[1 \quad 1]'$ and $[1.65 \quad 1.47]'$. It follows that the latent correlation coefficient for x_1 and x_2 is .9986, which is even larger than the latent correlation for x_{10} and x_{11}. When $\{x_1, x_2\}$ was selected as the measurement subset affected by only a single source, the A&S estimates (using the theoretical covariance matrix) of the two pattern vectors were the orthogonal linear combinations $.55\mathbf{c}_1 + .84\mathbf{c}_2$ and $.84\mathbf{c}_1 - .55\mathbf{c}_2$ of the true pattern vectors.

A Monte Carlo simulation with 10,000 replicates was used to compare the second-order and fourth-order blind separation methods. A sample size of $N = 200$ was assumed, and the autocovariance matrices for lags $\tau = 1, 2, \ldots, 6$ were used in the second order-method. Figure 10 shows the estimated pattern vectors and source signals for a typical replicate. The estimated pattern vectors are reasonably close to the true pattern vectors shown in Figure 9 and would most likely be correctly interpreted as a translation and a rotation of the beam. The estimated source signals shown in Figure 10 are noisy versions of the true source signals. The second source is clearly non-Gaussian, which indicates that the fourth-order assumptions were met. Figure 11 shows that the two sources have substantially different sample autocorrelation functions, which indicates that the second-order assumptions were met.

To evaluate the performance of the second-order and fourth-order methods over the entire Monte Carlo simulation, consider the performance measure $J_i = E[\|\mathbf{c}_i - \mathbf{c}_{i\|}]\|\mathbf{c}_i\|^{-1}$, where $\mathbf{c}_i$ denotes an estimate of $\mathbf{c}_i$ $(i = 1, 2)$. The average value of $\|\mathbf{c}_i - \mathbf{c}_i\|\mathbf{c}_i\|^{-1}$ over the 10,000 replicates was used to estimate J_1 and J_2 for both methods. The results, given in the first row of Table 1, indicate that both methods perform similarly for this example.

In the remainder of this section we investigate the performance of the two methods as their assumptions become closer to being violated. While the first source remained Gaussian, the distribution of the second source was modified from the Bernoulli distribution to the uniform, triangular, and Gaussian distributions shown in Figure 12, each having 0 mean and unit

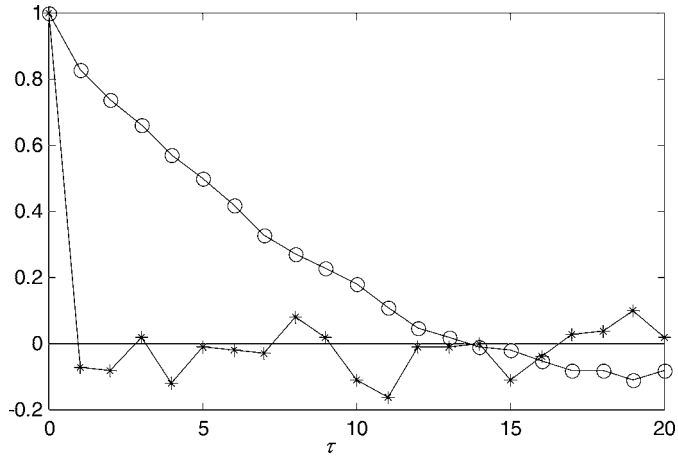

Figure 11. Sample Autocorrelation Functions for the Estimated Source Signals in the Simulation Example ($_\circ$, $\rho_{1,\tau}$; $_*$, $\rho_{2,\tau}$).

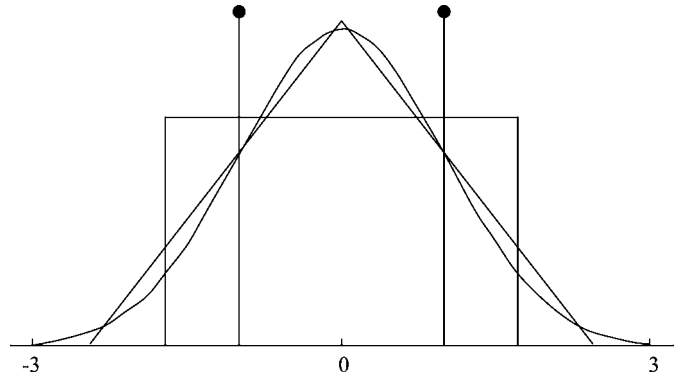

Figure 12. The Bernoulli, Uniform, Triangular, and Gaussian Distributions (With Mean 0 and Unit Variance) Used for the Second Source. The Bernoulli distribution is discrete, whereas the others are continuous distributions.

variance. This represents the second source distribution becoming successively closer to the Gaussian distribution, in terms of its kurtosis. Note that the kurtosis for the Bernoulli, uniform, triangular, and Gaussian distributions are -2, -1.2, $-.6$, and 0. When the second source is exactly Gaussian, the assumptions of the fourth-order method are violated. The results of using 10,000 Monte Carlo replicates to estimate J_1 and J_2 are shown in Table 1. The performance of the fourth-order method clearly deteriorates as the distribution of the second source becomes closer to Gaussian. But its performance is still reasonable ($J_i \approx .1$) when the second source distribution is uniform. The performance of the second-order method is unaffected by the source distribution.

The situation is reversed as the assumptions regarding the source autocorrelation in the second-order method become closer to being violated. While the second source remained temporally uncorrelated, the autocorrelation of the first source was reduced by decreasing the AR parameter ϕ, as illustrated in Figure 13. For $\phi = 0$, the first source is uncorrelated, and the assumptions of the second-order method are violated. Table 1 indicates that the performance of the second-order method deteriorates rapidly for $\phi < .5$, whereas the fourth-order method is unaffected by the source autocorrelation.

The last row of Table 1 gives the results when the assumptions for both methods are violated, in which case neither method performed well. In comparison, the A&S method was still quite effective ($J_1 = .108, J_2 = .085$) in this situation when $\{x_{10}, x_{11}\}$ was selected as the first measurement subset. But for the reasons discussed earlier, it performed poorly ($J_1 = .574$, $J_2 = .573$) when $\{x_1, x_2\}$ was selected as the first measurement subset.

We have also observed that the performance measures for both the second-order and the fourth-order methods are roughly inversely proportional to the square root of the sample size N and the square root of the signal-to-noise ratio. This agrees with the asymptotic results discussed by Cardoso (1998).

Table 1. Performance of the Second-Order and Fourth-Order Methods as Their Assumptions Become Closer to Being Violated

		Fourth-order method		Second-order method	
v_1 autocorrelation	v_2 distribution	J_1	J_2	J_1	J_2
$\phi = .9$	Bernoulli	.075	.097	.103	.082
$\phi = .9$	Uniform	.103	.116	.104	.083
$\phi = .9$	Triangular	.234	.236	.103	.082
$\phi = .9$	Gaussian	.362	.360	.104	.082
$\phi = .7$	Bernoulli	.075	.098	.116	.098
$\phi = .5$	Bernoulli	.075	.098	.144	.132
$\phi = .3$	Bernoulli	.075	.097	.277	.270
$\phi = 0$	Bernoulli	.074	.098	.568	.567
$\phi = 0$	Gaussian	.361	.360	.567	.567

8. ACCOMMODATING OTHER NOISE COVARIANCE STRUCTURES

Throughout this article, we have assumed that $\mathbf{\Sigma_w} = \sigma^2\mathbf{I}$—in other words, the noise variables associated with each element of $\mathbf{x}$ are uncorrelated and have equal, but unknown variance. In applications where the elements of $\mathbf{x}$ are similar entities obtained via similar measurement principles, this often would be a reasonable assumption. In this situation, one may even take the view that the noise variance for each element of $\mathbf{x}$ *should* be equal. If $\mathbf{\Sigma_w} = \sigma^2\mathbf{I}$ were assumed, but a particular element of $\mathbf{w}$ had a much larger variance (because of larger measurement error, for example), then this would appear as an additional variation pattern. The only nonzero element of the associated pattern

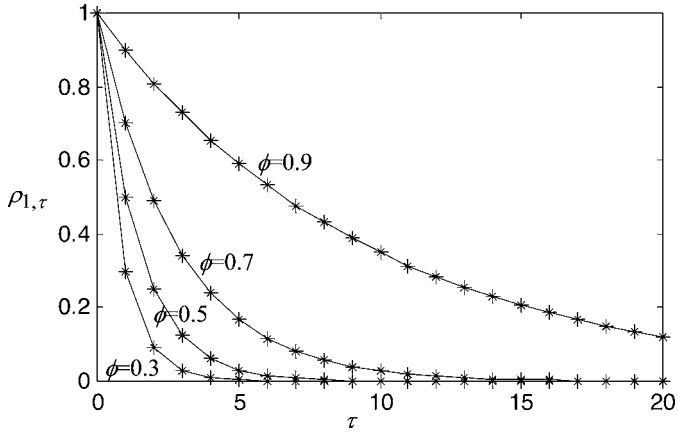

Figure 13. Autocorrelation Function $\rho_{1,\tau} = \phi^\tau$ of the First-Order AR Source, v_1, for Various ϕ.

vector would correspond to the element of $\mathbf{w}$ with larger variance. This may provide an indication that the apparatus used to measure that particular element of $\mathbf{x}$ should be recalibrated or replaced.

Many factor rotation methods assume only that $\boldsymbol{\Sigma}_{\mathbf{w}} = \text{diag}\{\sigma_1^2, \sigma_2^2, \ldots, \sigma_n^2\}$ is diagonal (Johnson and Wichern 1998). When the elements of $\mathbf{x}$ represent different entities measured on different scales, this would be more appropriate than assuming that $\boldsymbol{\Sigma}_{\mathbf{w}} = \sigma^2\mathbf{I}$. The blind separation algorithms may still be applied in this situation, as long as the diagonal $\boldsymbol{\Sigma}_{\mathbf{w}}$ is known (up to multiplication by a scalar constant) or a reasonable estimate is available. Before the algorithms are applied, the data must first be transformed via $\boldsymbol{\Sigma}_{\mathbf{w}}^{-1/2}\mathbf{x} = (\boldsymbol{\Sigma}_{\mathbf{w}}^{-1/2}\mathbf{C})\mathbf{v} + \boldsymbol{\Sigma}_{\mathbf{w}}^{-1/2}\mathbf{w}$, where $\boldsymbol{\Sigma}_{\mathbf{w}}^{-1/2} = \text{diag}\{\sigma_1^{-1}, \sigma_2^{-1}, \ldots, \sigma_n^{-1}\}$. Because the covariance matrix of the transformed noise $\boldsymbol{\Sigma}_{\mathbf{w}}^{-1/2}\mathbf{w}$ is a scalar multiple of the identity matrix, blind separation algorithms can be applied directly to the transformed data to produce an estimate of $\boldsymbol{\Sigma}_{\mathbf{w}}^{-1/2}\mathbf{C}$. This estimate can then be transformed back to $\mathbf{C}$ by premultiplying by $\boldsymbol{\Sigma}_{\mathbf{w}}^{1/2}$.

An estimate of $\boldsymbol{\Sigma}_{\mathbf{w}}$ would often be available in the context of manufacturing statistical process control (SPC). The estimate could be obtained by estimating the noise variances $\{\sigma_1^2, \sigma_2^2, \ldots, \sigma_n^2\}$ from a sample of data collected when the process is known to be in control (i.e., when no variation sources are present, so that $\mathbf{x} = \mathbf{w}$). Gage repeatability and reproducibility studies might also be used to estimate the noise variances.

An alternative to estimating the noise variances is to assume that they are such that $\sigma_i = \alpha T_i$ $(i = 1, 2, \ldots, n)$, where T_i denotes the width of the tolerance interval assigned to x_i and α is some arbitrary scale factor. In other words, the assumption would be that the standard deviation of each x_i is proportional to its tolerance width when no variation sources are present other than the noise. Borrowing SPC terminology, we may view this as *common cause* variability. This assumption translates to equal process capability ratios (the tolerance width divided by six standard deviation units) for all elements of $\mathbf{x}$ when only common cause variability is present. Although there is no statistical validity to this assumption, it has a conceptual appeal that we tend to favor for two reasons. First, tolerances are often assigned and/or manufacturing processes designed so that common cause variability is proportional to the tolerance width. Second, suppose that with only common cause variability present, the process capability ratio for a particular element of $\mathbf{x}$ was substantially smaller than for the other elements. If the blind separation methods were applied under the assumption of equal process capability ratios, then the result would be an additional variation pattern affecting only the variable with the low process capability. This would rightly call attention to the variable in most need of quality improvement efforts.

9. CONCLUSIONS

In this article we have reported on a blind source separation approach to identifying variation patterns in manufacturing measurement data. We considered two main classes (second-order and fourth-order) of blind separation methods and contrasted these with alternative methods for identifying variation patterns. The relative performance of the various methods depends predominantly on whether their specific sets of assumptions are satisfied. Whereas the varimax factor rotation method and the method of A&S assume conditions on the structure of $\mathbf{C}$, blind separation methods assume conditions on the distribution of the sources. Because of this, it is difficult to argue that one method is more broadly applicable than the others for diagnosing manufacturing variation. Nonetheless, blind separation methods have certain advantages over the method of A&S, in that their implementation involves less subjectivity and their assumptions are more straightforward to verify.

As a final comment, we point out that as in-process measurement proliferates, manufacturing data structures become increasingly spatially and temporally dense. The once-clear boundary between typical data structures encountered in manufacturing SPC versus those encountered in signal processing applications is disappearing. As we have illustrated here, signal processing methods may provide very useful tools for diagnosing manufacturing variation.

ACKNOWLEDGMENTS

This work was supported by the State of Texas Advanced Technology Program grant 000512-0289-1999, the National Science Foundation grant DMI-0093580, and Ford Motor Company. The authors are grateful to the anonymous referees, whose suggestions have improved the quality of this article.

[Received April 2001. Revised March 2003.]

REFERENCES

Apley, D. W., and Shi, J. (1998), "Diagnosis of Multiple Fixture Faults in Panel Assembly," *ASME Journal of Manufacturing Science and Engineering*, 120, 793–801.

——— (2001), "A Factor-Analysis Method for Diagnosing Variability in Multivariate Manufacturing Processes," *Technometrics*, 43, 84–95.

Belouchrani, A., Abed-Meraim, K., Cardoso, J. F., and Moulines, E. (1997), "A Blind Source Separation Technique Using Second-Order Statistics," *IEEE Transactions on Signal Processing*, 45, 434–444.

Box, G. E. P., Jenkins, G. M., and Reinsel, G. C. (1994), *Time Series Analysis: Forecasting and Control* (3rd ed.), Englewood Cliffs, NJ: Prentice-Hall.

Cardoso, J. F. (1998), "Blind Signal Separation: Statistical Principles," *Proceedings of the IEEE*, 86, 2009–2025.

Cardoso, J. F., and Souloumiac, A. (1993), "Blind Beamforming for Non-Gaussian Signals," *IEE Proceedings*, 140, 362–370.

Ceglarek, D., and Shi, J. (1996), "Fixture Failure Diagnosis for Autobody Assembly Using Pattern Recognition," *ASME Journal of Engineering for Industry*, 118, 55–66.

Comon, P. (1994), "Independent Component Analysis, a New Concept?" *Signal Processing*, 36, 287–314.

Golub, G. H., and Loan, C. F. V. (1989), *Matrix Computations*, Baltimore: Johns Hopkins University Press.

Haykin, S. (ed.) (2000), *Unsupervised Adaptive Filtering*, Vol. 1, New York: Wiley.

Hyvarinen, A. (1999), "Survey on Independent Component Analysis," *Neural Computing Surveys*, 2, 94–128.

Hyvarinen, A., and Oja, E. (2000), "Independent Component Analysis: Algorithms and Applications," *Neural Networks*, 13, 411–430.

Jackson, J. E. (1980), "Principal Components and Factor Analysis: Part I—Principal Components," *Journal of Quality Technology*, 12, 201–213.

——— (1981), "Principal Components and Factor Analysis: Part II—Additional Topics Related to Principal Components," *Journal of Quality Technology*, 13, 46–58.

Johnson, R. A., and Wichern, D. W. (1998), *Applied Multivariate Statistical Analysis* (4th ed.), Upper Saddle River, NJ: Prentice-Hall.

Lee, H. Y., and Apley, D. W. (2003), "Diagnosing Manufacturing Variation Using Second-Order and Fourth-Order Statistics," *International Journal of Flexible Manufacturing Systems*. To appear.

Monzingo, R. A., and Miller, T. W. (1980), *Introduction to Adaptive Arrays*, New York: Wiley.

Reed, C. W., and Yao, K. (1998), "Performance of Blind Beamforming Algorithms," in *Proceedings of the Ninth IEEE Signal Processing Workshop*, pp. 256–259.

Rosenblatt, M. (1985), *Stationary Sequences and Random Fields*, Boston: Birkhäuser.

Stuart, A., and Ord, J. K. (1987), *Kendall's Advanced Theory of Statistics, Vol. 1, Distribution*, (5th ed.), New York: Wiley.

Tong, L., Soon, V. C., Huang, Y. F., and Liu, R. (1990), "AMUSE: A New Blind Identification Algorithm," *IEEE International Symposium on Circuits and Systems*, 3, 1784–1787.

Wax, M., and Sheinvald, J. (1997), "A Least-Squares Approach to Joint Diagonalization," *IEEE Signal Processing Letters*, 4, 52–53.