
The International Journal of Flexible Manufacturing Systems, 16, 45–64, 2004

Diagnosing Manufacturing Variation Using Second-Order and Fourth-Order Statistics

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

DANIEL W. APLEY* apley@northwestern.edu
Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208-3119

Abstract. This article discusses a method that can aid in diagnosing root causes of product and process variability in complex manufacturing processes, when large amounts of multivariate in-process measurement data are available. A linear structured model, similar to the standard factor analysis model, is used to generically represent the variation patterns that result from the root causes. Blind source separation techniques form the basis for identifying the precise characteristics of each individual variation pattern in order to facilitate the identification of their root causes. The second-order and fourth-order statistics that are used in various blind separation algorithms are combined in an optimal manner to form a more effective and black-box method with wider applicability.

Key Words: blind source separation, factor rotation, manufacturing variation, multivariate statistical process control, principal components analysis

1. Introduction

A common characteristic of many modern complex manufacturing processes is the availability of large amounts of multivariate in-process measurement and inspection data for quality control purposes. One example is automobile body assembly, in which laser-optical measurement stations are built into the assembly line at various stages, typically immediately after major subassemblies are completed (Apley and Shi, 2001). In each measurement station, well over 100 key dimensional characteristics distributed over the subassembly may be measured. Moreover, 100% of the autobodies produced are measured. Another prime example is printed circuit board (PCB) assembly. In PCB assembly, laser-optical measurement is also commonly used to obtain detailed dimensional characteristics of the wet solder paste, after it is deposited onto the board during the screen printing stage. After the electronic components are placed in position on the board and the solder is cured in the reflow oven, additional dimensional characteristics of each cured solder joint are obtained via X-ray laminography (Glass and Thomsen, 1993). In state-of-the-art PCB assembly operations, the solder paste is measured in-process for 100% of the boards produced, and the cured solder joints are measured for nearly 100% of the boards.

The measurement data may contain a wealth of buried diagnostic information concerning the numerous variation sources that contribute to overall levels of process variability. Each

*To whom correspondence should be addressed.

variation source will typically result in a distinct variation pattern in the data. The patterns will have "spatial" characteristics that indicate how a variation source causes different measured variables or features to interact, as well as "temporal" characteristics that indicate how a variation source evolves over time. As in any data mining application, a primary objective is to extract concise, relevant information from the raw data in a form that can be clearly presented to a human operator. For the purpose of understanding and reducing process variation, relevant information refers to the precise nature of the variation pattern caused by each source. The presumption is that the operators, when provided a clearer understanding of the nature of each variation pattern, will be better equipped to identify and eliminate the underlying root causes of process variation.

The objective of identifying spatial variation patterns in a sample of multivariate data is similar to that of principal components analysis (PCA) and factor analysis (Jackson, 1980, 1981; Johnson and Wichern, 1998). When there are numerous variation sources present, however, PCA and factor analysis methods cannot produce unique estimates of the variation patterns. To overcome this nonuniqueness problem, Apley and Lee (2003) proposed using a technique referred to as blind source separation. Blind source separation methods originated in the fields of radar and sonar signal processing (Monzingo and Miller, 1980) for the purpose of identifying the spatial location of signal sources. As noted in Apley and Lee (2003), the underlying model used to represent radar and sonar signals is of surprisingly similar structure to a model that can be used to represent manufacturing variation patterns. Consequently, blind separation methods can be applied directly to the problem of diagnosing manufacturing variation.

Two main classes of blind separation techniques are fourth-order and second-order methods, which utilize fourth-order and second-order statistics, respectively, to estimate the variation patterns. Blind separation methods are able to produce unique estimates of the variation patterns by imposing additional conditions on the statistical nature of the variation sources. Fourth-order methods require that no more than one of the variation sources follows a Gaussian distribution. Second-order methods require that the variation sources are temporally autocorrelated and have different autocorrelation functions. Apley and Lee (2003) demonstrated that fourth-order and second-order methods are effective when their respective uniqueness conditions are satisfied. If their conditions are violated, however, they are unable to produce unique estimates of the variation patterns, much like PCA and factor analysis. In practice, it is difficult to know a priori which set of conditions will be better satisfied and, consequently, which method will be more effective. This article develops a new method for estimating the variation patterns by combining the second-order and fourth-order statistics. The uniqueness conditions of the new method are less restrictive versions of the conditions for the individual second-order and fourth-order methods. The result is a more black-box method with wider applicability, in which the end user is not burdened with verifying which set of conditions are better satisfied.

The remainder of the article is organized as follows. Section 2 provides background information on the model used to represent variation patterns, existing methods for identifying the variation patterns, and the limitations of existing methods. Section 3 introduces the new method for combining second-order and fourth-order statistics, which is intended to overcome some of the limitations of existing methods. Section 4 derives the theoretical

criteria under which the variation patterns can be uniquely identified using the combined method. Section 5 investigates the performance of the combined method, relative to existing methods.

2. Representing and identifying spatial variation patterns

This section provides background information on the model we use to represent variation patterns and existing methods (PCA, factor rotation, and blind source separation) for identifying the variation patterns. The capabilities and limitations of the existing methods are discussed as motivation for the combined method that will be presented in the following section.

2.1. Representing and interpreting variation patterns

Let $\boldsymbol{x} = [x_1, x_2, \ldots, x_n]'$ be an $n \times 1$ random vector that represents a set of n measured characteristics from the product or process. Let $\boldsymbol{x}_i$, $i = 1, 2, \ldots, N$, be a sample of N observations of $\boldsymbol{x}$. In autobody assembly, for example, $\boldsymbol{x}$ could represent the vector of all measured features across a given autobody, and N would be the number of autobodies in a sample. It is assumed that $\boldsymbol{x}$ obeys the model

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

where $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_p]$ is an $n \times p$ constant matrix with full rank p. The vector $\boldsymbol{v} = [v_1, v_2, \ldots, v_p]'$ is a $p \times 1$ zero-mean random vector representing p distinct, statistically independent "variation sources." Each source $v_i, i = 1, 2, \ldots, p$, is scaled to have unit variance without loss of generality. The vector $\boldsymbol{w}$ is an $n \times 1$ zero-mean random vector, independent of $\boldsymbol{v}$, that represents the aggregated effects of measurement noise and any inherent unmodeled variation. For notational convenience, it will be assumed that the covariance matrix of $\boldsymbol{w}$ is $\Sigma_w = \sigma^2\mathbf{I}$, where $\mathbf{I}$ is the identity matrix, and σ^2 is a scalar constant. This can be assumed without loss of generality, providing that a reasonable estimate of $\boldsymbol{\Sigma}_w$ is available. Typically, Σ_w would be estimated from a sample of "in-control" data collected when there were no variation sources present, in which case $\boldsymbol{x} = \boldsymbol{w}$. In this situation, the methods in this article would be applied to data that are first transformed by the inverse of a square root of Σ_w, as discussed in Apley and Lee (2003). Since the focus is on identifying sources of variation, as opposed to mean shifts, we have assumed all random variables are zero mean. If not, the mean of $\boldsymbol{x}$ must be estimated and subtracted from the data.

The model structure may be interpreted as follows. Each variation source is assumed to have a linear effect on $\boldsymbol{x}$ that is represented by the corresponding column of $\mathbf{C}$. Together, $\mathbf{c}_i v_i$ describes the effect of the ith source on $\boldsymbol{x}$. The pattern vector $\mathbf{c}_i$ indicates the spatial nature of the variation due to the ith source. Since each element v_i is scaled to have unit variance, $\mathbf{c}_i$ also indicates the magnitude or severity of the variation due to the ith source. The model (1) is nearly identical to the standard linear orthogonal factor analysis model (Johnson and Wichern, 1998), in which the components of $\boldsymbol{v}$ are the common factors, the components of $\boldsymbol{w}$ are the specific factors, and $\mathbf{C}$ is the factor loading matrix.

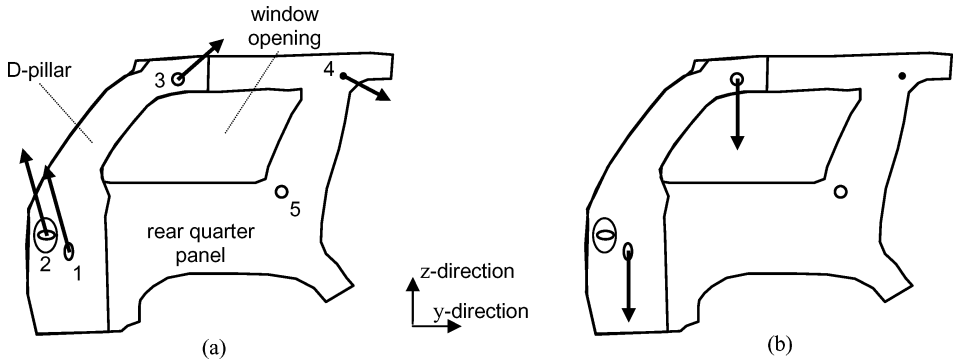

Figure 1. Illustration of two linear spatial variation patterns in autobody assembly: (a) a rotation of the quarter panel subassembly about locating hole 5 and (b) a translation of the D-pillar in the z-direction.

The objective is to estimate the pattern vectors $\{\mathbf{c}_i : i = 1, 2, \ldots, p\}$ based on a sample of data and then to use the estimates to gain insight into the root causes of the variation patterns. To illustrate this, consider the following example from autobody assembly. Figure 1 shows a rear quarter panel subassembly of an autobody. The measurement vector $\boldsymbol{x}$ for the quarter panel subassembly consists of $n = 10$ y/z-plane coordinates of five separate features that are numbered 1 through 5 in Figure 1(a). The measurements are taken after the quarter panel subassembly is joined to the bodyside, which is not shown in the figure. For a more detailed description of the assembly process, we refer the reader to Apley and Shi (1998) or Apley and Shi (2001).

In a sample of $N = 200$ measured autobodies, there were two major variation patterns present, and estimates of $\mathbf{c}_1$ and $\mathbf{c}_2$ were obtained using the method to be described in Section 3. The estimates of $\mathbf{c}_1$ and $\mathbf{c}_2$ are illustrated on a figure of the autobody in Figures 1(a) and (b), respectively. The elements of each variation pattern vector have been plotted as arrows at the locations of the features to which they correspond. The y/z coordinates of each feature have been combined into a single arrow. The estimate of $\mathbf{c}_1$ shown in Figure 1(a) appears to be a rotation of the entire subassembly about feature 5. The source signal v_1 would be a random variable that is proportional to the angle of rotation (clockwise on some autobodies; counterclockwise on others) of each quarter panel subassembly. The estimate of $\mathbf{c}_2$ shown in Figure 1(b) appears to be a z-direction (up/down) translation of the D-pillar with respect to the rest of the quarter panel subassembly. The source signal v_2 would be a random variable that is proportional to the amount of translation (up on some autobodies; down on others) of each D-pillar.

Graphical displays of the estimated variation pattern vectors, such as in Figure 1, generally facilitate identification of the root causes of variability. The root cause of the first variation pattern was found to be a loose locating element that failed to properly constrain the quarter panel subassembly when it was placed into a fixture and welded to the bodyside. The geometry of the fixture and the position of the loose locating element were such that the quarter panel subassembly was free to rotate by small amounts about feature 5 (a hole that mates with a pin rigidly attached to the fixture). The root cause of the second variation

pattern was found to be an elongated hole in the D-pillar due to improper stamping. The hole mated with a pin whose purpose was to constrain the D-pillar when it was welded to the quarter panel. The elongated hole allowed the D-pillar to translate by small amounts in the z-direction, relative to the rest of the quarter panel.

2.2. *Principal component analysis and factor rotation*

Since the model (1) is essentially the standard linear orthogonal factor analysis model (Johnson and Wichern, 1998), one may consider using factor analysis methods to estimate $\mathbf{C}$. Let $\mathbf{\Sigma}_x$ denote the covariance matrix of $\boldsymbol{x}$. Most factor analysis methods are based on PCA, which involves analyzing the eigenvectors and eigenvalues of $\mathbf{\Sigma}_x$. Let $\{\lambda_i : i = 1, 2, \ldots, n\}$ denote the eigenvalues of $\mathbf{\Sigma}_x$ arranged in descending order, and let $\{\boldsymbol{z}_i : i = 1, 2, \ldots, n\}$ denote the corresponding eigenvectors, which we take to be an orthonormal set. Define $\mathbf{Z}_p = [\mathbf{z}_1, \mathbf{z}_2, \ldots, \mathbf{z}_p]$ and $\mathbf{\Lambda}_p = \text{diag}\{\lambda_1, \lambda_2, \ldots, \lambda_p\}$, which are constructed from the p largest eigenvalues and their associated eigenvectors. In order to be consistent with the covariance structure, we know (Apley and Shi, 2001) that any estimate of $\mathbf{C}$ must be of the form $\mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{U}$, where $\mathbf{U}$ is an arbitrary $p \times p$ orthogonal matrix.

We note that the methods in this article are presented in the context that the true covariance matrix $\mathbf{\Sigma}_x$ and other distribution parameters are known. To implement the methods, all parameters are replaced by their estimates, obtained from the sample of data. The number of "dominant" eigenvalues of the sample covariance matrix serves as an estimate of p, and the average of the remaining $n–p$ eigenvalues serves as an estimate of σ^2. Estimates of $\mathbf{Z}_p$ and $\mathbf{\Lambda}_p$ are constructed from the eigenvectors and eigenvalues of the sample covariance matrix. Apley and Shi (2001) discuss this in more detail, including statistical tests for determining how many eigenvalues are dominant.

Since p, σ^2, $\mathbf{Z}_p$, and $\mathbf{\Lambda}_p$ are all available from PCA, the problem reduces to finding an appropriate $p \times p$ orthogonal (rotation) matrix $\mathbf{U}$ and then using $\mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{U}$ as an estimate of $\mathbf{C}$. This is referred to as factor rotation (Jackson, 1981; Johnson and Wichern, 1998). Let $\mathbf{Q}$ denote the value of $\mathbf{U}$ that yields the true $\mathbf{C}$. In other words, $\mathbf{C} = \mathbf{Z}_p[\mathbf{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{Q}$. An underlying premise of this article (refer to Apley and Shi (2001) for a more detailed discussion) is that there is some "true" $\mathbf{C}$ whose structure is dictated entirely by the physics of the process. Each column of the true $\mathbf{C}$ represents a distinct variation pattern with a distinct physical root cause. In the autobody assembly example of Section 2.2, the root cause of the first pattern was a loose tooling element that caused a rotation of the entire quarter panel subassembly about feature 5. The first column $\mathbf{c}_1$ (illustrated in Figure 1(a)) represents that rotation. The root cause of the second pattern was an elongated hole that caused a translation of the D-pillar, and the second column $\mathbf{c}_2$ (illustrated in Figure 1(b)) represents that translation.

Standard factor rotation techniques are not intended to produce a $\mathbf{U}$ that is necessarily close to $\mathbf{Q}$. Rather, they produce a $\mathbf{U}$ that optimizes a somewhat artificial interpretability criterion such as the varimax criterion (Johnson and Wichern, 1998). Consequently, they do not necessarily produce an estimate of $\mathbf{C}$ that is close to the true $\mathbf{C}$. In terms of understanding the root causes of variability, however, the most effective interpretation of an estimate of $\mathbf{C}$

will surely result when it equals the true $\mathbf{C}$. The methods discussed in the remainder of this article are intended to accomplish this.

2.3. *Second- and fourth-order blind source separation methods*

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 separate signal-emitting sources (Cardoso, 1998; Haykin, 2000). Applications include radar and sonar signal processing, biomedical (e.g., EEG, EKG, and fetal heartbeat) and geophysical signal monitoring, and wireless communications. *Blind* separation refers to the situation where information on the individual sources must be determined from only a sample of observed data $\{\boldsymbol{x}_t : t = 1, 2, \ldots, N\}$, with no prior knowledge of the relationship between the sources and observed data other than the assumed linear structure of (1). To accomplish this, it is necessary to first estimate $\mathbf{C}$ "blindly" from the data, which is precisely the objective considered in this article. Given the similarities, it is not surprising that blind source separation methods can often be applied to identifying manufacturing variation patterns. Apley and Lee (2003) have investigated two primary blind separation methods, termed second-order and fourth-order methods. The basic concepts behind these methods are briefly described in this section. For a more detailed discussion, refer to Apley and Lee (2003) and the references therein.

Both the second-order and fourth-order methods work with a transformed version of the data with spatially white (uncorrelated) components. The whitened data vector is defined as $\boldsymbol{y} = \mathbf{W}\boldsymbol{x}$, where $\mathbf{W} = [\boldsymbol{\Lambda}_p - \sigma^2\mathbf{I}]^{-1/2}\mathbf{Z}_p'$ is the $p \times n$ "whitening matrix." Premultiplying (1) by $\mathbf{W}$ and using the relationship $\mathbf{C} = \mathbf{Z}_p[\boldsymbol{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{Q}$, the transformed data vector becomes

$$\boldsymbol{y} = \mathbf{W}\boldsymbol{x} = \mathbf{W}[\mathbf{C}\boldsymbol{v} + \boldsymbol{w}] = \mathbf{Q}\boldsymbol{v} + \mathbf{W}\boldsymbol{w}. \tag{2}$$

To illustrate the concepts behind the second-order method, let $\boldsymbol{v}_t = [v_{1,t}, v_{2,t}, \ldots, v_{p,t}]'$ denote the source vector at time t. The source autocovariance matrix at lag $\tau \geq 0$ is $\boldsymbol{\Sigma}_{v,\tau} = E[\boldsymbol{v}_t\boldsymbol{v}_{t+\tau}'] = \text{diag}\{\rho_{1,\tau}, \rho_{2,\tau}, \ldots, \rho_{p,\tau}\}$, where $\rho_{i,\tau}$ denotes the autocorrelation (the sources have unit variance) function of the ith source. Consider the autocovariance matrix $\boldsymbol{\Sigma}_{y,\tau}$ of the whitened data $\boldsymbol{y}$. Belouchrani, Abed-Meraim, Cardoso, and Moulines (1997) have shown that $\mathbf{Q}$ is the orthogonal matrix that jointly diagonalizes the entire set $\{\boldsymbol{\Sigma}_{y,\tau} : \tau > 0\}$ in the sense that for any $\tau > 0$, $\mathbf{Q}'\boldsymbol{\Sigma}_{y,\tau}\mathbf{Q}$ is the diagonal matrix $\boldsymbol{\Sigma}_{v,\tau}$. Under certain conditions that will be discussed in Section 2.4, this yields a unique estimate of $\mathbf{Q}$.

A slightly different strategy must be used to implement the second-order method using estimates of $\boldsymbol{\Sigma}_{y,\tau}$, since no orthogonal matrix will diagonalize the entire set of sample autocovariance matrices. Belouchrani et al. (1997) have suggested taking $\mathbf{Q}$ to be the orthogonal matrix $\mathbf{U}$ that minimizes the sum of the squares of the off-diagonal elements of the set $\{\mathbf{U}'\boldsymbol{\Sigma}_{y,\tau}\mathbf{U} : \tau = 1, 2, \ldots, T\}$ for some specified integer T (e.g., $T = 10$). A computationally efficient algorithm for performing this "joint approximate diagonalization", proposed by Cardoso and Souloumiac (1994), is outlined in the Appendix.

To illustrate the concepts behind the fourth-order method, consider the set of fourth-order cumulants of the (zero-mean) random vector $\boldsymbol{y} = [y_1, y_2, \ldots, y_p]'$. The fourth-order cumulant of its ith, jth, kth, and lth elements, $1 \leq i, j, k, l \leq p$, is defined as

$$C_{i,j,k,l}(\boldsymbol{y}) = E[y_i y_j y_k y_l] - E[y_i y_j]E[y_k y_l] - E[y_i y_k]E[y_j y_l] - E[y_i y_l]E[y_j y_k]. \tag{3}$$

Also define the set of p^2 cumulant matrix $\{\mathbf{M}(i, j): 1 \leq i, j \leq p\}$ so that the kth-row, lth-column element of the $p \times p$ matrix $\mathbf{M}(i, j)$ is $C_{i,j,k,l}(\boldsymbol{y})$. Cardoso and Souloumiac (1994) have shown that $\mathbf{Q}$ jointly diagonalizes the entire set of cumulant matrices in the sense that $\mathbf{Q}'\mathbf{M}(i, j)\mathbf{Q} = \text{diag}\{C_{1,1,1,1}(\boldsymbol{v})q_{i,1}q_{j,1}, C_{2,2,2,2}(\boldsymbol{v})q_{i,2}q_{j,2}, \ldots, C_{p,p,p,p}(\boldsymbol{v})q_{i,p}q_{j,p}\}$ for $1 \leq i, j \leq p$, where $q_{k,l}$ denotes the kth-row, lth-column element of $\mathbf{Q}$. This also yields a unique estimate of $\mathbf{Q}$, under certain conditions that will be discussed in Section 2.4.

As in the second-order method, no single orthogonal matrix will diagonalize the entire set of cumulant matrices formed from the sample cumulants (i.e., the sample averages of (3)). Consequently, a similar joint approximate diagonalization criterion must be used, in which $\mathbf{Q}$ is taken to be the orthogonal matrix $\mathbf{U}$ that minimizes the sum of the squares of the off-diagonal elements of the set $\{\mathbf{U}'\mathbf{M}(i, j)\mathbf{U}: 1 \leq i, j \leq p\}$. The algorithm discussed in the Appendix can be used to perform the minimization.

2.4. Uniqueness conditions for second-order and fourth-order methods

In order to resolve the ambiguity in selecting $\mathbf{U}$ in factor rotation, blind separation methods utilize additional information beyond what is contained in the covariance matrix $\boldsymbol{\Sigma}_x$. Second-order methods use autocovariance information, and fourth-order methods use fourth-order cumulant information. The sources must satisfy certain conditions, however, if this additional information is to result in a unique estimate of $\mathbf{Q}$ and therefore a unique estimate of the true $\mathbf{C}$. The uniqueness condition for the second-order method is that no pair of sources share the exact same autocorrelation function (Belouchrani et al., 1997). The uniqueness condition for the fourth-order method is that no more than one of the sources follows a Gaussian distribution (more precisely, has zero kurtosis) and that the noise is either Gaussian or negligible (Cardoso and Souloumiac, 1994).

Note that a minimum requirement for the second-order uniqueness condition to be satisfied is that at least $p-1$ of the p sources have some level of autocorrelation. Otherwise, there exists a pair of sources with the exact same (zero) autocorrelation function $\rho_\tau = 0$ for $\tau > 0$. Autocorrelated data are increasingly common in industrial environments where automated in-process measurement produces large volumes of measurement data. Source autocorrelation is only a necessary condition for uniqueness, however, and not a sufficient condition. In addition, the sources must also have *different* autocorrelation functions. To illustrate, suppose there are two variation sources, and they both follow first-order autoregressive (AR) models of the form $v_{i,t} = \phi v_{i,t-1} + a_t$. Following standard time series notation, ϕ is referred to as the AR parameter, and the a_t's represent a sequence of zero-mean, temporally uncorrelated random shocks (Box, Jenkins, and Reinsel, 1994). The autocorrelation function of a first-order AR process is $\rho_\tau = \phi^\tau$ $(\tau > 0)$. Figure 2 shows typical sets of 200 observations

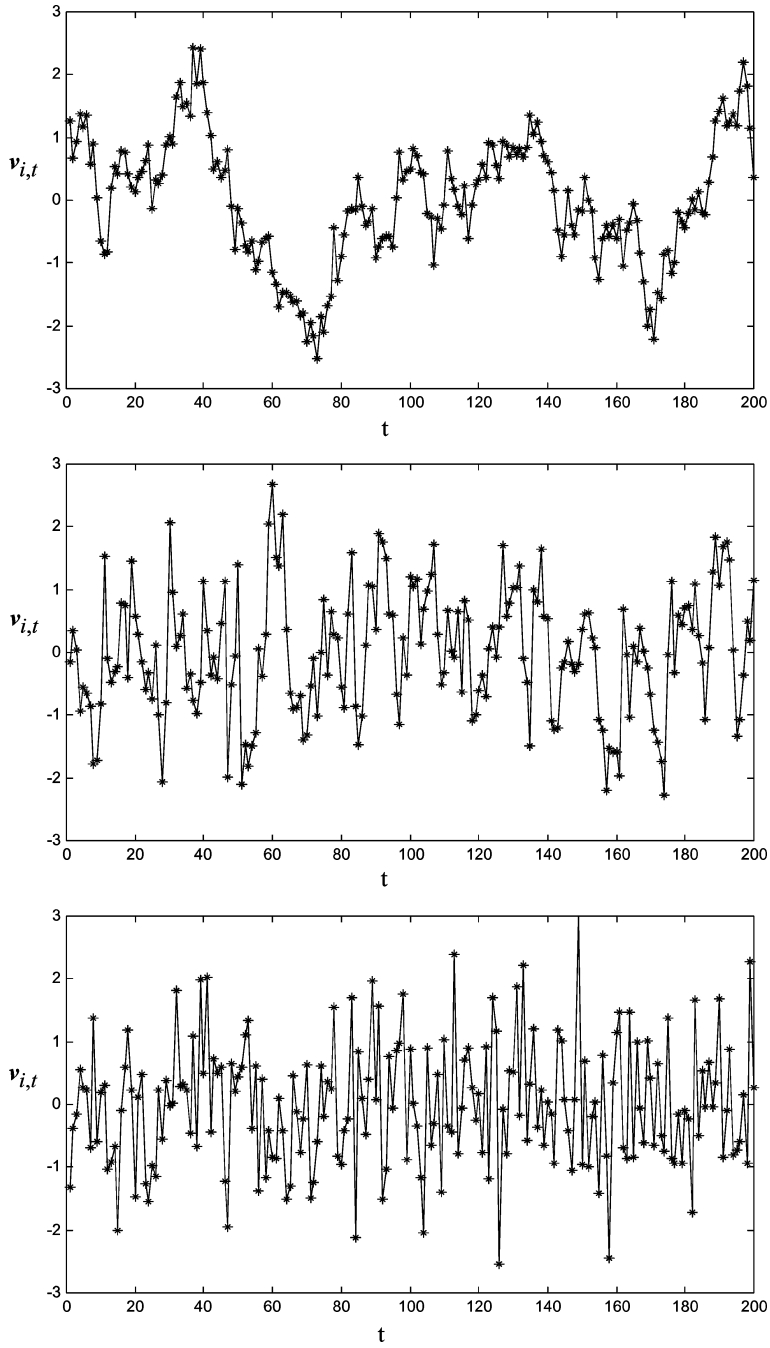

Figure 2. Illustration of first-order AR source signals with different levels of autocorrelation: top—high autocorrelation ($\phi = 0.9$); middle—moderate autocorrelation ($\phi = 0.6$); bottom—no autocorrelation ($\phi = 0$).

of three different first-order AR processes, each with a different level of autocorrelation. The top, middle, and bottom panels of Figure 2 represent high autocorrelation ($\phi = 0.9$), moderate autocorrelation ($\phi = 0.6$), and no autocorrelation ($\phi = 0$), respectively. If the two sources that are present follow the same AR model (with $\phi = 0.9$, say), they will have identical autocorrelation functions, and the second-order uniqueness condition will be violated. In this case, the second-order method cannot produce a unique estimate of $\mathbf{C}$. If on the other hand the two sources follow different AR models (the first with $\phi = 0.9$ and the second with $\phi = 0.6$, say), their autocorrelation functions will differ, and the second-order uniqueness condition will be satisfied.

Apley and Lee (2003) discuss in more detail the physical meaning of the uniqueness conditions and examples of manufacturing scenarios in which they apply. They also describe how to check whether the uniqueness conditions are satisfied based on the data sample. After $\mathbf{C}$ is estimated, the source signals over the data sample can be estimated via the least squares equation $[\mathbf{C}'\mathbf{C}]^{-1}\mathbf{C}'\boldsymbol{x}$. Histograms of the estimated source signals can then be used to ascertain whether more than one source is Gaussian, in which case we cannot uniquely estimate $\mathbf{C}$ using the fourth-order method. Sample autocorrelation plots of the estimated source signals can be used to ascertain whether a pair of sources shares the exact same autocorrelation function, in which case we cannot estimate $\mathbf{C}$ uniquely using the second-order method.

Although there are many manufacturing situations in which we could expect the uniqueness conditions to be satisfied, in particular when the number of sources are small, there are also many situations in which the conditions would not be fully satisfied. As the number of sources increases, it becomes increasingly likely that a pair will have approximately the same autocorrelation function, in which case the second-order method will not produce a robust estimate of $\mathbf{C}$. It also becomes increasingly likely that two or more sources will follow approximate Gaussian distributions, in which case the fourth-order method will not produce a robust estimate of $\mathbf{C}$. The remainder of this article investigates a method of combining the second-order and fourth-order criteria in order to relax the uniqueness conditions.

3. Combining second-order and fourth-order statistics

Although derived from completely different perspectives, the second-order and fourth-order methods share a commonality. Both methods seek to jointly approximately diagonalize a set of matrices. A natural means of combining the second-order and fourth-order criteria is to jointly approximately diagonalize the cumulant matrices $\{\mathbf{M}(i, j): 1 \leq i, j \leq p\}$ together with the autocovariance matrices $\{\boldsymbol{\Sigma}_{y,\tau}: \tau = 1, 2, \ldots, T\}$. More specifically, consider the set of $K = T + p^2$ matrices $\{\boldsymbol{A}_k: k = 1, 2, \ldots, K\}$ defined as follows. The first T matrices $\{\boldsymbol{A}_k: k = 1, 2, \ldots, T\}$ are the set of scaled autocovariance matrices $\{\boldsymbol{\Sigma}_{y,\tau}/s_1: \tau = 1, 2, \ldots, T\}$, where $s_1 = p^{-1}\sum_{i=1}^{p} E[y_i^2]$ is the average second moment of the whitened data. The remaining p^2 matrices $\{\boldsymbol{A}_k: k = T + 1, 2, \ldots, K\}$ are the set of scaled cumulant matrices $\{\mathbf{M}(i, j)/s_2: 1 \leq i, j \leq p\}$, where $s_2 = p^{-1}\sum_{i=1}^{p} E[y_i^4]$ is the average fourth moment of the whitened data. For example, if $T = 3$ and $p = 2$, we have $K = 7$ and $\{\boldsymbol{A}_k: k = 1, 2, \ldots, K\} = \{\boldsymbol{\Sigma}_{y,1}/s_1, \boldsymbol{\Sigma}_{y,2}/s_1, \boldsymbol{\Sigma}_{y,3}/s_1, \mathbf{M}(1,1)/s_2, \mathbf{M}(2,1)/s_2, \mathbf{M}(1,2)/s_2, \mathbf{M}(2,2)/s_2\}$.

Let $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \ldots, \alpha_K]'$ be a vector of nonnegative weighting coefficients, to be determined. We propose to take the estimate of $\mathbf{Q}$ to be the orthogonal matrix $\mathbf{U}$ that minimizes

$$\sum_{k=1}^{K} \alpha_k \text{ off } [\mathbf{U}'\mathbf{A}_k\mathbf{U}], \tag{4}$$

where off $[\cdot]$ denotes the sum of the squares of the off-diagonal elements of a matrix. This can be viewed as jointly approximately diagonalizing the weighted set of matrices $\{\alpha_k^{1/2}\mathbf{A}_k : k = 1, 2, \ldots, K\}$. The algorithm in the Appendix can be used to perform the joint approximate diagonalization.

We seek a means of "optimally" selecting the weighting coefficients $\boldsymbol{\alpha}$. Some of the matrices in the set $\{\boldsymbol{A}_k : k = 1, 2, \ldots, K\}$ will contain more information than others and should therefore be given higher weight. For example, suppose that after a certain time lag ($\tau \geq 3$, say) the autocorrelation functions for each source become identical, but that at time lags one and two the autocorrelation functions differ substantially. Then we would like to use larger weighting coefficients for the two matrices $\boldsymbol{\Sigma}_{y,1}/s_1$ and $\boldsymbol{\Sigma}_{y,2}/s_1$ and smaller coefficients for $\{\boldsymbol{\Sigma}_{y,\tau}/s_1 : \tau \geq 3\}$, since the latter set contains very little information that can aid in separating the sources. As an other example, consider an extreme case where the fourth-order uniqueness condition happens to be satisfied, but all sources are temporally uncorrelated so that the matrices $\{\boldsymbol{\Sigma}_{y,\tau}/s_1 : \tau > 0\}$ contain no useful information. In this case we would like to assign all of the weight to the cumulant matrices and no weight to the autocovariance matrices. In the opposite extreme, if all sources are Gaussian but the second-order uniqueness conditions are satisfied, we would like to assign all of the weight to the autocovariance matrices and no weight to the cumulant matrices. Using large coefficients for matrices that contain no information only adds noise to the estimation problem. The method proposed below for assigning the weighting coefficients results in weights that are commensurate with the information contained in each matrix.

In the following, we use the overscore symbol "$\hat{}$" to distinguish between theoretical quantities and their estimates from the sample data. Theoretically, every matrix in $\{\boldsymbol{A}_k : k = 1, 2, \ldots, K\}$ is exactly diagonalized by $\mathbf{Q}$, so that the criterion (4) is zero for any choice of $\boldsymbol{\alpha}$. The joint diagonalizer $\hat{\mathbf{Q}}$ that minimizes (4) with $\mathbf{A}_k$ replaced by $\hat{\boldsymbol{A}}_k$ will depend on $\boldsymbol{\alpha}$ and will generally differ from $\mathbf{Q}$. An attractive strategy for selecting $\boldsymbol{\alpha}$ is to attempt to minimize some measure of the error between $\mathbf{Q}$ and $\hat{\mathbf{Q}}$.

Represent the error between $\hat{\mathbf{Q}}$ and $\mathbf{Q}$ via $\delta\mathbf{E}$, defined such that $\hat{\mathbf{Q}} = \mathbf{Q}(\mathbf{I} + \delta\mathbf{E})$. It is straightforward to adapt a result from Cardoso (1994) to our criterion (4), which gives the following approximate expression for $\delta\mathbf{E}$. The ith row, jth column element of $\delta\mathbf{E}$ is $(1 \leq i \neq j \leq p)$

$$\delta\mathbf{E}_{i,j} \approx \frac{\sum_k \alpha_k (d_k(j) - d_k(i)) \mathbf{q}_i' \hat{\mathbf{A}}_k \mathbf{q}_j}{\sum_k \alpha_k (d_k(j) - d_k(i))^2}, \tag{5}$$

where $\mathbf{q}_i$ is the ith column of $\mathbf{Q}$, and $d_k(i)$ is the ith diagonal element of $\mathbf{Q}'\mathbf{A}_k\mathbf{Q}$. The error $\delta\mathbf{E}_{i,j}$ is affected by three types of quantities: $\boldsymbol{\alpha}$, $\{d_k(j) - d_k(i) : k = 1, 2, \ldots, K\}$, and

$\mathbf{q}_i'\hat{\mathbf{A}}_k\mathbf{q}_j$. Although $\mathbf{q}_i'\mathbf{A}_k\mathbf{q}_j = 0$ for $1 \le i \ne j \le p$ ($\mathbf{Q}'\mathbf{A}_k\mathbf{Q}$ is exactly diagonal), $\mathbf{q}_i'\hat{\mathbf{A}}_k\mathbf{q}_j$ will generally differ from zero. Writing $\mathbf{q}_i'\hat{\mathbf{A}}_k\mathbf{q}_j = \mathbf{q}_i'(\hat{\mathbf{A}}_k - \mathbf{A}_k)\mathbf{q}_j$, we see that the term $\mathbf{q}'_i\hat{\mathbf{A}}_k\mathbf{q}_j$ in (5) stems from the error in estimating $\mathbf{A}_k$ with a finite set of sample data. The quantities $\{d_k(j) - d_k(i): k = 1, 2, \ldots, K\}$ relate closely to the second-order and fourth-order uniqueness conditions. Suppose that two Gaussian sources (the ith and jth sources, say) have very similar autocorrelation functions. Since $d_k(j)$–$d_k(i)$ will be close to zero for all k in this case, we expect to have large error $\delta\mathbf{E}_{i,j}$ and to have difficulty separating the ith and jth sources. This is consistent with the theoretical uniqueness condition derived in Section 4 for the combined method, since two Gaussian sources with identical autocorrelation functions violate the uniqueness condition.

While (5) provides some insight into the factors that affect accuracy, selecting $\boldsymbol{\alpha}$ in an attempt to minimize it directly would not be straightforward. Rather, we propose the suboptimal approach of selecting $\boldsymbol{\alpha}$ to maximize the denominator terms in (5), which will generally result in smaller values of the $\delta\mathbf{E}_{i,j}$ terms in (5). Specifically, we recommend selecting $\boldsymbol{\alpha}$ to maximize

$$\sum_{i\ne j}\sum_{k}\alpha_k(d_k(j) - d_k(i))^2 = \sum_{k}\alpha_k\sum_{i\ne j}(d_k(j) - d_k(i))^2 \tag{6}$$

under some equality constraint on the norm of $\boldsymbol{\alpha}$ ($\boldsymbol{\alpha}'\boldsymbol{\alpha}$ equals some constant value). Since rescaling $\boldsymbol{\alpha}$ does not affect the solution to (4), it does not matter what value is specified for the norm of $\boldsymbol{\alpha}$. Equation (6) is maximized by selecting

$$\alpha_k \propto \sum_{j\ne i}(d_k(j) - d_k(i))^2. \tag{7}$$

In other words, the weight assigned to $\mathbf{A}_k$ is proportional to $\sum_{j\ne i}(d_k(j) - d_k(i))^2$, which is closely related to the information contained in $\mathbf{A}_k$. If $\sum_{j\ne i}(d_k(j) - d_k(i))^2 = 0$ for some k, there is no information in $\mathbf{A}_k$ that can be used to separate the sources.

We cannot implement (7) exactly, since the d_k(i) terms are the diagonal entries of $\mathbf{Q}'\mathbf{A}_k\mathbf{Q}$, and we do not know the true values of $\mathbf{Q}$ and $\{\mathbf{A}_k: k = 1, 2, \ldots, K\}$. In light of this, we recommend first obtaining an initial estimate $\hat{\mathbf{Q}}_\mathbf{0}$ of $\mathbf{Q}$ by minimizing (4) with equal weighting ($\alpha_k = 1 : k = 1, 2, \ldots, K$). The ith diagonal element of $\hat{\mathbf{Q}}_\mathbf{0}'\hat{\mathbf{A}}_k\hat{\mathbf{Q}}_\mathbf{0}$ can then be substituted for $d_k(i)$ in (7) to obtain the final weighting coefficients. The entire procedure for the combined method is summarized as follows. For notational convenience, we again omit the "^" symbol on all quantities, which are meant to be estimated values from the sample data. The algorithm in the Appendix is used for steps 5 and 8.

(1) From the data sample $\{\boldsymbol{x}_i: i = 1, 2, \ldots, N\}$, calculate the sample covariance matrix $\boldsymbol{\Sigma}_x = N^{-1}\sum_{i=1}^{N}(\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})'$, where $\bar{\mathbf{x}} = N^{-1}\sum_{i=1}^{N}\boldsymbol{x}_i$.
(2) Based on a PCA decomposition of $\boldsymbol{\Sigma}_x$, calculate the whitening matrix $\mathbf{W} = [\boldsymbol{\Lambda}_p - \sigma^2\mathbf{I}]^{-1/2}\mathbf{Z}'_p$ and the whitened data $\{\boldsymbol{y}_i = \mathbf{W}\boldsymbol{x}_i: i = 1, 2, \ldots, N\}$.

(3) Select T and calculate the sets of cumulant and autocovariance matrices $\{\mathbf{M}(i, j): 1 \leq i, j \leq p\}$ and $\{\boldsymbol{\Sigma}_{y,\tau}: \tau = 1, 2, \ldots, T\}$ of the whitened data.
(4) Calculate the average second and fourth moments s_1 and s_2 of the whitened data and form the set of matrices $\{\mathbf{A}_k: k = 1, 2, \ldots, K\}$.
(5) Find the joint approximate diagonalizer $\mathbf{Q}_0$ of the set $\{\mathbf{A}_k: k = 1, 2, \ldots, K\}$.
(6) Set $d_k(i)$ equal to the ith diagonal element of $\mathbf{Q}'_0\mathbf{A}_k\mathbf{Q}_0$ $(i = 1, 2, \ldots, p; k = 1, 2, \ldots, K)$.
(7) Select optimal weights $\boldsymbol{\alpha}$ according to (7).
(8) Find the joint approximate diagonalizer $\mathbf{Q}$ of the set $\{\alpha_k^{1/2}\mathbf{A}_k: k = 1, 2, \ldots, K\}$.
(9) Take the estimate of $\mathbf{C}$ to be $\mathbf{Z}_p[\boldsymbol{\Lambda}_p - \sigma^2\mathbf{I}]^{1/2}\mathbf{Q}$.

In step 3, we recommend using a relatively large value for T (e.g., $T = 10$ to 20). The reason is that if there happens to be little information contained in the autocovariance matrices at larger time lags, the weighting scheme (7) tends to automatically discount them by assigning them small weights. For example, suppose the source autocovariance decays to very small values for time lags greater than 10, but we have chosen $T = 20$. Consider a value of k such that $\mathbf{A}_k$ corresponds to $\boldsymbol{\Sigma}_{y,\tau}$ for some time lag between 10 and 20. Since $d_k(i)$ is the autocorrelation function (scaled by s_1) for the ith source at that time lag, $d_k(i)$ will be close to zero for $i = 1, 2, \ldots, p$. Equation (7) would then select a small value for α_k, and the autocovariance matrices for time lags greater than 10 would be largely ignored in the joint diagonalization.

4. Uniqueness condition for the combined method

This section derives the uniqueness condition for the combined method, which is a less restrictive version of the uniqueness conditions for the individual second-order and fourth-order methods. The derivation parallels that presented in Belouchrani et al. (1997) for the second-order method. The uniqueness condition is derived in the context that the second and fourth moments of the data are known, which is equivalent to having an infinitely large sample of data. For the finite sample sizes one must work with in practice, the effectiveness of the methods depends on the extent to which the conditions are satisfied, as discussed in Apley and Lee (2003).

Since $\mathbf{Q}$ diagonalizes the entire set $\{\mathbf{A}_k: k = 1, 2, \ldots, K\}$, and $\mathbf{Q}$ is orthogonal, we can write $\mathbf{A}_k = \mathbf{Q}\mathbf{D}_k\mathbf{Q}'$ with $\mathbf{D}_k = \text{diag}\{d_k(1), \ldots, d_k(p)\}$. Recall that $\mathbf{A}_k$ is either of the form $\boldsymbol{\Sigma}_{y,\tau}/s_1$ for some τ or of the form $\mathbf{M}(s, m)/s_2$ for some s and m. In the former case $d_k(i) = \rho_{i,\tau}/s_1$, and in the latter case $d_k(i) = C_{i,i,i,i}(\boldsymbol{v})q_{s,i}q_{m,i}/s_2$.

Proposition 1. *Suppose that no pair of Gaussian sources share the exact same autocorrelation function, and consider two arbitrary sources v_i and v_j $(i \neq j)$. Then there exists an index k such that $d_k(i) \neq d_k(j)$.*

Proof: There are three possibilities: both v_i and v_j are Gaussian (case 1); one source is Gaussian and the other source is not Gaussian (case 2); and neither source is Gaussian (case 3). We separate the proof by case.

Case 1. By assumption, the two Gaussian sources must have different autocorrelation function. Thus, there exists a time lag τ such that $\rho_{i,\tau} \neq \rho_{j,\tau}$, which completes the proof for case 1.

Case 2. Let v_i denote the Gaussian source and v_j denote the non-Gaussian source. Since v_i is Gaussian, $C_{i,i,i,i}(\boldsymbol{v})$ is zero, which implies that $C_{i,i,i,i}(\boldsymbol{v})q_{s,i}q_{m,i} = 0$ for all $1 \leq s, m \leq p$. Since v_j is non-Gaussian, $C_{j,j,j,j}(\boldsymbol{v}) \neq 0$, and there exists an s and m such that $C_{j,j,j,j}(\boldsymbol{v})q_{s,j}q_{m,j} \neq 0 = C_{i,i,i,i}(\boldsymbol{v})q_{s,i}q_{m,i}$. This completes the proof for case 2.

Case 3. Since both sources are non-Gaussian, $C_{i,i,i,i}(\boldsymbol{v}) \neq 0$ and $C_{j,j,j,j}(\boldsymbol{v}) \neq 0$. Since $\mathbf{q}_i$ and $\mathbf{q}_j$ are orthogonal and have unit norm, there exist an s and m (possibly equal) such that $C_{i,i,i,i}(\boldsymbol{v})q_{s,i}q_{m,i} \neq C_{j,j,j,j}(\boldsymbol{v})q_{s,j}q_{m,j}$, which completes the proof for case 3.

□

Theorem 1 (Uniqueness condition for the combined method)**.** *If no pair of Gaussian sources share the exact same autocorrelation function, then the orthogonal joint diagonalizer of the set* $\{\alpha_k^{1/2}\mathbf{A}_k\colon k = 1, 2, \ldots, K\}$ *is unique and equal to* $\boldsymbol{Q}$ *(up to an interchange of its columns).*

Proof: Let $\mathbf{U}$ be any orthogonal joint diagonalizer of the set $\{\alpha_k^{1/2}\mathbf{A}_k\colon k = 1, 2, \ldots, K\}$, and let $\mathbf{u}$ denote any column of $\mathbf{U}$. Since $\{\mathbf{q}_s\colon s = 1, 2, \ldots, p\}$ are orthogonal, we can represent $\mathbf{u}$ as some linear combination $\mathbf{u} = \sum_{s=1}^{p} \beta_s \mathbf{q}_s$, where at least one of the coefficients (say β_i) is nonzero. For any $1 \leq j \leq p$ with $j \neq i$ there exists an index k such that $d_k(i) \neq d_k(j)$ by Proposition 1. Equation (7) implies that $\alpha_k \neq 0$. Since $\mathbf{U}$ diagonalizes $\alpha_{\mathbf{k}}^{1/2}\mathbf{A}_{\mathbf{k}}$, $\mathbf{u}$ is an eigenvector of $\mathbf{A}_k$. If γ denotes the corresponding eigenvalue, we have $\gamma\mathbf{u} = \mathbf{A}_k\mathbf{u}$. Substituting $\mathbf{u} = \sum_{s=1}^{p} \beta_s \mathbf{q}_s$ gives $\sum_{s=1}^{p} \beta_s \gamma \mathbf{q}_s = \sum_{s=1}^{p} \beta_s \mathbf{A}_k \mathbf{q}_s = \Sigma_{s=1}^{p} \beta_s d_k(s)\mathbf{q}_s$, where the last equality follows from the fact that $\mathbf{A}_k = \mathbf{Q}\mathbf{D}_k\mathbf{Q}'$. Equating the coefficients of each $\mathbf{q}_s$ implies that $\beta_s[d_k(s) - \gamma] = 0$ for $s = 1, \ldots, p$. Since $\beta_i \neq 0$, it follows that $\gamma = d_k(i) \neq d_k(j)$, which in turn implies that $\beta_j = 0$. Since $j(\neq i)$ was arbitrary, we have $\beta_j = 0$ for any $j \neq i$, and thus $\mathbf{u} = \mathbf{q}_i$. Repeating for each column of $\mathbf{U}$ completes the proof. □

Theorem 1 states that the uniqueness condition for the combined method is that no pair of *Gaussian* sources share the same autocorrelation function. In other words, multiple Gaussian sources are allowed as long as the subset of Gaussian sources have different autocorrelation functions. The autocorrelation functions of the non-Gaussian sources are irrelevant. Suppose we divide the sources into subsets that have the same autocorrelation function (e.g., one subset of uncorrelated sources, a second subset of moderately correlated sources, and a third subset of highly correlated sources). An equivalent statement of the uniqueness condition is that within each subset having the same autocorrelation function there is at most one Gaussian source. Hence, multiple sources with the same autocorrelation function are also allowed. The uniqueness condition for the combined method is therefore less restrictive than the individual uniqueness conditions for the second-order and fourth-order methods. The following section provides an example in which the uniqueness condition for the combined method is satisfied but the conditions for the individual methods are not.

5. Performance comparison

This section compares the performance of the second-order method, the fourth-order method, and the combined method in a simulation example. The performance of the methods in situations where their uniqueness conditions are satisfied and in situations where they are violated are both considered. The following "baseline" example represents the situation where the uniqueness conditions are satisfied for all three methods. In the baseline example, a simple beam is used to represent the part being manufactured, and $n = 20$ measurements are distributed across the beam. This is illustrated in Figure 3. There are three variation sources in the baseline example, with $\mathbf{c}_1$, $\mathbf{c}_2$, and $\mathbf{c}_3$ illustrated in Figures 3(a)–(c), respectively. The beam could be considered a subcomponent of a larger assembly, in which case the variation patterns could represent assembly variation. The first pattern would represent a rigid vertical translation of the beam, the second pattern would represent a rigid rotation about the midpoint of the beam, and the third pattern could represent a bending of the beam about the midpoint. Alternatively, the beam could be considered a separate part, in which case the variation patterns could represent fabrication (e.g., extrusion) variation. The first pattern may represent variation in the thickness of the beam that occurs uniformly across its length. The second pattern may represent variation in the thickness that, when larger on one end of the beam, is smaller on the other end. The third pattern may represent variation in the thickness that, when larger on the ends of the beam, is smaller in the center.

The pattern vectors were scaled so that $\mathbf{c}'_1\mathbf{c}_1 = \mathbf{c}'_2\mathbf{c}_2 = \mathbf{c}'_3\mathbf{c}_3 = n\sigma^2$. Since the component of $\boldsymbol{x}$ due to the ith variation source is just $\mathbf{c}_i v_i$, the total variation due to the ith source is $E[(\mathbf{c}_i v_i)'(\mathbf{c}_i v_i)] = \mathbf{c}'_i\mathbf{c}_i$. Likewise, the total variance due to the noise is $E[\boldsymbol{w}'\boldsymbol{w}] = n\sigma^2$.

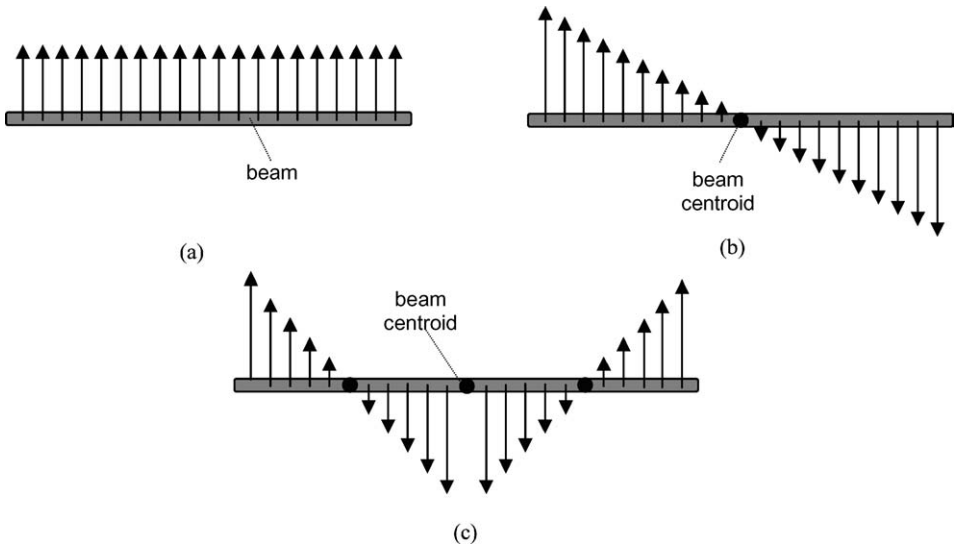

Figure 3. Illustration of the two variation patterns in the baseline example: (a) $\mathbf{c}_1$, representing a beam translation, (b) $\mathbf{c}_2$, representing a beam rotation, and (c) $\mathbf{c}_3$, representing a beam bending.

Consequently, the scaling of the pattern vectors in the baseline example represents the situation where the variance due to each source is equal to the variance due to the noise. The sample size was $N = 200$ in all cases. For autobody assembly, this represents roughly 3 h of production.

The source distributions for the baseline example were chosen so that the uniqueness conditions of both the second-order and the fourth-order methods were satisfied. The third source $\{v_{3,t}: t = 1, 2, \ldots\}$ was generated via the first-order AR model $v_{3,t} = \phi v_{3,t-1} + a_t$, with AR parameter $\phi = 0.9$ and a_t an independent sequence of zero-mean Gaussian random shocks with variance $(1 - \phi^2)$. The starting value $v_{3,1}$ was drawn from the standard normal distribution. Thus, $v_{3,t}$ is a stationary process having a marginal Gaussian distribution with zero mean and unit variance. Its autocorrelation function is $\rho_{3,\tau} = \phi^\tau = 0.9^\tau$ (Box et al., 1994). The first source $\{v_{1,t} : t = 1, 2, \ldots\}$ was generated as an independent sequence of discrete random variables taking on values of ± 1 with equal probability 0.5, which we refer to as the (scaled and shifted) Bernoulli distribution. The second source $\{v_{2,t} : t = 1, 2, \ldots\}$ was generated via

$$v_{2,t} = \begin{cases} v_{2,t-1}\text{: with probability } q \\ -v_{2,t-1}\text{: with probability } 1 - q \end{cases} \tag{8}$$

with $q = 0.05$ and starting value $v_{2,1}$ drawn from the Bernoulli distribution. It can be shown that $v_{2,t}$ is a stationary process with a marginal Bernoulli distribution and autocorrelation function $\rho_{2,\tau} = (2q - 1)^\tau = (-0.9)^\tau$. The parameter q can be used to control the autocorrelation of the Bernoulli sources just like ϕ can be used to control the autocorrelation of the Gaussian sources. For $q < 0.5$, the lag-one autocorrelation will be negative. A typical realization of 100 observations for $q = 0.05$ is shown in Figure 4(a). This could represent the effects of multiple tooling, in which the same operation is performed in two stations in order to increase throughput in bottleneck processes. If the stations are not calibrated perfectly, parts produced in each station may follow different distributions. When mixed together in an approximately alternating manner, the result would appear as the source signal in Figure 4(a) with noise $\boldsymbol{w}$ added. For $q > 0.5$, the source autocorrelation will be positive. A typical realization of 100 observations for $q = 0.95$ is shown in Figure 4(b). This could represent the situation where the process mean jumps back and forth between two different values at occasional random times. For $q = 0.5$, the source will have no temporal autocorrelation, which is illustrated in Figure 4(c). We can view the uncorrelated $v_{1,t}$ in the baseline example as being generated via the same model (8) with $q = 0.5$.

The following performance measure was used to evaluate the methods. Let $\hat{\mathbf{c}}_i$ denote the estimate of the ith pattern vector, and define $\hat{\mathbf{b}}_i = \hat{\mathbf{c}}_i \|\hat{\mathbf{c}}_i\|^{-1}$ and $\mathbf{b}_i = \mathbf{c}_i \|\mathbf{c}_i\|^{-1}$ to be the directions of $\hat{\mathbf{c}}_i$ and $\mathbf{c}_i$. The performance measure for estimating the ith pattern vector is defined as $J_i = E[\|\hat{\mathbf{b}}_i - \mathbf{b}_i\|]$, which is the expected norm of the error in estimating the direction of $\mathbf{c}_i$. Alternatively, we could have used the expected norm of the actual error $E[\|\hat{\mathbf{c}}_i - \mathbf{c}_i\|]\|\mathbf{c}_i\|^{-1}$ as the measure of performance. Since this was quite consistent with the directional error in the simulations, we report only the directional error values. For the purpose of interpreting the nature of the ith variation pattern, the direction of $\mathbf{c}_i$ is more important than the magnitude. The magnitude of $\mathbf{c}_i$ is related to the severity of the pattern.

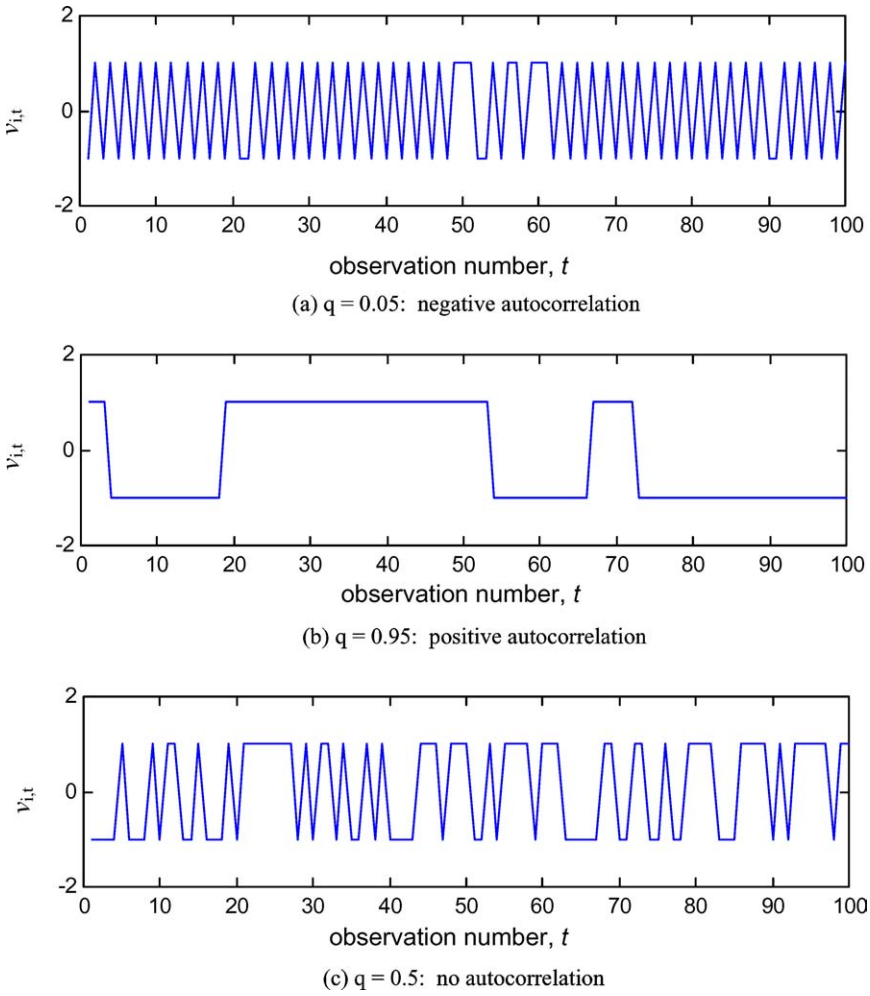

Figure 4. Illustration of autocorrelated Bernoulli sources with different levels of autocorrelation (different values of q).

The average value of $\|\hat{\mathbf{b}}_i - \mathbf{b}_i\|$ over a Monte Carlo simulation with 10,000 replicates was used to approximate J_i in all cases. $T = 20$ was chosen for the second-order and combined methods.

The Monte Carlo results for the three methods in the baseline example are shown in the first three rows of Table 1. Since the autocorrelation functions of the three sources differ $[\rho_{1,\tau} = 0, \rho_{2,\tau} = (-0.9)^{\tau}, \rho_{3,\tau} = 0.9^{\tau}]$, and only one source is Gaussian, the uniqueness conditions for both the second-order and the fourth-order methods are satisfied in the baseline example. As might be expected, Table 1 demonstrates that the second-order and fourth-order methods both perform quite well in this situation.

Table 1. Summary of the Monte Carlo simulation results comparing the performance of the three methods.

Source properties	Method	J_1	J_2	J_3
Baseline: no conditions violated	Fourth-order method	0.102	0.083	0.076
	Second-order method	0.101	0.107	0.107
	Combined method	0.083	0.096	0.096
Fourth-order conditions violated	Fourth-order method	0.599	0.569	0.571
	Second-order method	0.137	0.127	0.128
	Combined method	0.105	0.112	0.113
Second-order conditions violated	Fourth-order method	0.105	0.105	0.104
	Second-order method	0.637	0.606	0.608
	Combined method	0.110	0.110	0.110
Second-order and fourth-order conditions violated	Fourth-order method	0.121	0.347	0.348
	Second-order method	0.468	0.468	0.135
	Combined method	0.107	0.088	0.121

To investigate the effects of violating the fourth-order uniqueness conditions, we next generated the three sources as Gaussian first-order AR processes via $v_{i,t} = \phi v_{i,t-1} + a_t$. AR parameter values of 0, -0.7, and 0.7 were used for sources one, two, and three, respectively. The source autocorrelation functions were therefore $\rho_{1,\tau} = 0$, $\rho_{2,\tau} = (-0.7)^\tau$, $\rho_{3,\tau} = 0.7^\tau$, and the uniqueness condition for the second-order method was satisfied. Since all three sources were Gaussian, the uniqueness condition for the fourth-order method was violated. Rows 4 through 6 of Table 1 show the Monte Carlo results for this situation. The performance of the fourth-order method is substantially worse than in the baseline case, whereas the second-order method still performed well. The combined method was slightly more effective than the second-order method, most likely because the combined method automatically assigned less weight to the autocovariance matrices at larger time lags.

To investigate the effects of violating the second-order uniqueness condition, we next generated each of the three sources as an independent sequence of uniformly distributed random variables with zero mean and unit variance. Since all three sources are temporally uncorrelated, the second-order uniqueness conditions are violated. The fourth-order uniqueness conditions are still satisfied, however, because all three sources are non-Gaussian. Rows 7 through 9 of Table 1 show the Monte Carlo results for this situation. The performance of the second-order method is now substantially worse than in the baseline case. The performances of the fourth-order method and the combined method are roughly the same and are both reasonably good.

The final simulation represents the situation where the uniqueness conditions of both the second-order and the fourth-order methods are violated, but the condition of the combined method is satisfied. Everything is as in the baseline example, except that the second source is changed from an autocorrelated Bernoulli random variable with $q = 0.05$ to a temporally uncorrelated Gaussian random variable. We now have two sources (v_2 and v_3) that follow a Gaussian distribution, so the uniqueness condition of the fourth-order method is violated.

We also now have two sources with the same autocorrelation function ($\rho_{1,\tau} = \rho_{2,\tau} = 0$ for all $\tau > 0$), so the uniqueness condition of the second-order method is violated. Since the two Gaussian sources have different autocorrelation functions, however, the uniqueness condition of the combined method is satisfied. The last three rows of Table 1 show that the combined method performs quite well, whereas the second-order and fourth-order methods do not. It is interesting to note that the fourth-order method is able to separate the first source reasonably well ($J_1 = 0.121$) but is unable to separate the two Gaussian sources ($J_2 = 0.347$ and $J_3 = 0.348$). In contrast, the second-order method is able to separate the third source reasonably well ($J_3 = 0.135$) but is unable to separate the two temporally uncorrelated sources ($J_1 = 0.468$ and $J_2 = 0.468$). It is essentially this "block separability" characteristic of the blind source separation methods that allows the combined method to fully separate the sources.

This section has only investigated the effects of violating the uniqueness conditions. There are a number of other factors that affect the overall accuracy of the methods, including the number of sources p, the sample size N, the dimension n of the measurement vector, the magnitude of the noise variance relative to the severity of the variation patterns, and how close the pattern vectors are to being linearly independent. Additional simulation results indicated that these factors affect all three methods by roughly the same degree. Since the focus of this article is on the relative performance of the combined method versus the individual second-order and fourth-order methods, we omit the details of these simulations. One observation was that the performance measure was roughly inversely proportional to the square root of the sample size N and the square root of the signal-to-noise ratio $\mathbf{c}'_i\mathbf{c}_i n^{-1}\sigma^{-2}$, which agrees with the asymptotic results of Cardoso (1998). The reader is referred to Apley and Lee (2003) for additional discussion of the effects of these factors and other performance issues.

6. Conclusions

This article investigated a method to aid in diagnosing root causes of product and process variability in complex manufacturing processes in which there are large amounts of multivariate in-process measurement data. In order to identify the nature of each individual variation pattern when multiple variation sources are present, blind source separation concepts have been applied. The second-order and fourth-order statistics that are used in two different classes of blind source separation algorithms have been combined into a single algorithm. A method for suboptimally assigning weights to the second-order and fourth-order criteria was developed, and the conditions under which the patterns can be uniquely identified using the combined method were derived. The uniqueness condition is a less restrictive version of the conditions for the individual second-order and the fourth-order methods, which gives the combined method broader applicability than either of the two individual methods. This also lessens the burden of verifying which set of uniqueness conditions are better satisfied, resulting in a more black-box method that requires less statistical expertise on the part of the user.

It should be emphasized that the method is not universally applicable. Certain conditions regarding the source distributions (as stated in Theorem 1) must still be satisfied.

Moreover, the linear model (1) must provide an adequate representation of the variation patterns. Although the resulting algorithms would almost certainly be much more complicated, extension of the blind separation concepts to nonlinear variation patterns would further broaden the applicability of the approach.

Appendix

This appendix summarizes the algorithm of Cardoso and Souloumiac (1994), to be used for jointly approximately diagonalizing a set of matrices $\{\mathbf{A}_k: k = 1, 2, \ldots, K\}$. The algorithm involves representing the orthogonal matrix $\mathbf{U}$ as the product of a series of Givens rotation matrices (Golub and Van Loan, 1989). For a specified rotation angle θ and a pair of indices i and j with $1 \leq i \neq j \leq p$, the Givens rotation matrix $\mathbf{U}(i, j, \theta)$ is defined as a slightly modified version of the identity matrix. All elements of $\mathbf{U}(i, j, \theta)$ are the same as the identity matrix, except that $u_{ii}(i, j, \theta) = \cos(\theta)$, $u_{ij}(i, j, \theta) = -\sin(\theta)$, $u_{ji}(i, j, \theta) = \sin(\theta)$, $u_{jj}(i, j, \theta) = \cos(\theta)$, where $u_{kl}(i, j, \theta)$ denotes the kth-row, lth-column element of $\mathbf{U}(i, j, \theta)$. As the algorithm iterates over i and j, the θ that determines each $\mathbf{U}(i, j, \theta)$ is chosen to minimize

$$\sum_{k=1}^{K} \text{off}[\mathbf{U}'(i, j, \theta)\mathbf{A}_k\mathbf{U}(i, j, \theta)], \tag{9}$$

which is a function of only a single parameter θ. The matrices $\{\mathbf{A}_k: k = 1, 2, \ldots, K\}$ are updated via $\mathbf{A}_k \rightarrow \mathbf{U}'(i, j, \theta)\mathbf{A}_k\mathbf{U}(i, j,\theta)$ after each new $\mathbf{U}(i, j, \theta)$ is found.

For a fixed i and j, the θ that minimizes (9) is easily determined as follows. Define the 2×2 symmetric matrix

$$\mathbf{H} = \sum_{k=1}^{K} \mathbf{h}'(\mathbf{A}_k)\mathbf{h}(\mathbf{A}_k), \tag{10}$$

where $\mathbf{h}(\mathbf{A})$ is defined as $[a_{ii} - a_{jj}\ a_{ij} + a_{ji}]$ with a_{ij} denoting the ith-row jth-column element of a matrix $\mathbf{A}$. The θ that minimizes (9) is given by (see Cardoso and Souloumiac, 1994)

$$\cos\theta = \sqrt{\frac{z_1 + 1}{2}} \quad \text{and} \quad \sin\theta = \frac{z_2}{\sqrt{2(z_1 + 1)}}, \tag{11}$$

where $[z_1\ z_2]'$ denotes the eigenvector of $\mathbf{H}$ that corresponds to its largest eigenvalue.

The algorithm is iterated over i and j values with $1 \leq i \neq j \leq p$ until (9) converges. The final $\mathbf{Q}$ matrix is then taken to be the product (from left to right) of the individual $\mathbf{U}(i, j, \theta)$ matrices at each iteration. Matlab code for performing the joint diagonalization is available from the authors upon request.

Acknowledgments

This work was supported by the State of Texas Advanced Technology Program under grant 000512-0289-1999, the National Science Foundation under grant DMI-0093580, and Ford Motor Company.

References

Apley, D. W. and Lee, H. Y., "Identifying Spatial Variation Patterns in Multivariate Manufacturing Processes: A Blind Separation Approach," *Technometrics*, Vol. 45, pp. 220–234 (2003).

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

Apley, D. W. and Shi, J., "A Factor-Analysis Method for Diagnosing Variability in Multivariate Manufacturing Processes," *Technometrics*, Vol. 43, pp. 84–95 (2001).

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

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

Cardoso, J. F., "Perturbation of Joint Diagonalizers," Technical Report Ref. #94d027, Telecom Paris (1994).

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

Cardoso, J. F. and Souloumiac, A., "Blind Beamforming for Non-Gaussian Signals," *IEE Proceedings*, Vol. 140, pp. 362–370 (1994).

Glass, S. and Thomsen, J., "How SMT Boards are Assembled," *Printed Circuit Fabrication*, Vol. 16, pp. 42–47 (1993).

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

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

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

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

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

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