Monitoring Product Size and Edging from Bivariate Profile Data

ROMÁN VIVEROS-AGUILERA

McMaster University, Hamilton, Ontario L8S 4K1, Canada

STEFAN H. STEINER and R. JOCK MACKAY

University of Waterloo, Waterloo, Ontario N2L 3G1, Canada

Profile data consist of the coordinates of points along the edge of the product. Often, several hundred points are involved. Mechanical and automated procedures (e.g., scanning) are used in data gathering. The large data dimensionality presents challenges in the development of control charts to monitor product profiles. The data also show strong cross-correlations between points close to one another. In this article, using the leading principal components of the coordinate covariance matrix, we develop Hotelling's T^2 and upper exponentially weighted moving average (EWMA) charts to monitor product size. The methods are extended to monitor product edging using the angles between the normal vectors of the blueprint and sample profiles. We use a Markov chain approximation to calculate average run length. Through simulations, we assess the performance of the proposed methods and show the upper EWMA chart exhibit good performance in most of the off-target scenarios considered. A comparison with existing methods reveals that the proposed charts are very competitive and require fewer distributional assumptions.

Key Words: Angular Vectors; Average Run Length; EWMA Chart; Hotelling's T^2 Chart; Principal Components; Smoothing; SPC.

1. Introduction

Blueprints or designs of mechanical parts are drawn to represent graphically detailed part specifications. Computer-aided design (CAD) packages are usually used in this process. Often the parts display flat forms, for instance, when obtained by cutting from metal sheets and other flat materials. A turning process using a lathe is another example. If the flat material exhibits sufficiently homogeneous thickness, the main features of the workpiece are determined by its edge. In what follows, we refer to the planar representation of the edge as the *shape* or *profile* of the workpiece.

Shape specifications include dimensional characteristics such as length and width; curvature such as straightness, circularity, and ovality; and edge texture such as roughness and waviness. No matter how tightly a process is run, manufacturers recognize that no two workpieces made from the same design are identical. Every process is bound to exhibit some inherent variation in shape from part to part. The role of statistical process control methods is to spot when the variation in the profile of the manufactured parts shows deviation beyond the natural process variation.

The concepts and methods discussed in this article apply equally to shapes from cross-sections of three-dimensional parts, e.g., the profile of a workpiece at a given height of the part when it rests on a level surface.

Dr. Viveros-Aguilera is Professor in the Department of Mathematics and Statistics. He is a Member of ASQ. His email address is rviveros@ms.mcmaster.ca.

Dr. Steiner is Professor in the Department of Statistics and Actuarial Science. He is a Fellow of ASQ. His email address is shsteiner@uwaterloo.ca.

Dr. MacKay is Adjunct Professor in the Department of Statistics and Actuarial Science. His email address is jock.mackay@uwaterloo.ca.

Profile monitoring initially focused on situations where the profiles are adequately described by a linear regression between a single product quality variable and a single predictor. The primary objective is to uncover when the linear association breaks down in the Phase II process operation (see, e.g. Mahmoud and Woodall (2004) and Sullivan (2002) for more details on Phase I and Phase II analyses). The methods make use of the residuals or coefficient estimates (e.g., Kim et al. (2003), Mahmoud and Woodall (2004), Gupta et al. (2006), among others) or are based on change-point analysis (e.g., Zou et al. (2006)). The methods were then extended to processes where multiple linear or polynomial relationships between a single quality indicator and predictors are appropriate (e.g., Zou et al. (2007), Kazemzadeh et al. (2008), Noorossana et al. (2010)). In the more recent research, attention centers on curved or nonlinear profiles using parametric nonlinear regression (e.g., Williams et al. (2007a, 2007b)) and nonparametric charts (Qiu et al. (2010)).

In the above references, the profile description is univariate. For instance, Qiu et al. (2010) consider data for m profiles, where the ith profile consists of n_i observations, $\{x_{ij}, y_{ij}\}$, $1 \le j \le n_i$, $1 \le i \le m$, described by $y_{ij} = g(x_{ij}) + f_i(x_{ij}) + \epsilon_{ij}$ where $g(x)$ is the target profile, $f_i(x)$ models the variation of the ith profile around the target, and ϵ_{ij} is a random error, the source of which is typically measurement error. Here x is a predictor or covariate, such as time or distance, depending on the process. This is a mixed-effects model and Qiu et al. (2010) fit it nonparametrically to Phase I data.

Colosimo et al. (2008) discuss bivariate profiles where the data are many pairs of (x, y) values corresponding to points along the edge of a product. Unlike the previously discussed profiles, the planar representation of the bivariate profiles are closed-loop curves. Naturally, measurement error can occur in both x and y coordinates. Colosimo et al. (2008) borrow ideas from spatial statistics, specifically spatial correlation, to develop a parametric normal regression model to describe the profiles. After fitting the model to a sample of Phase I profiles, they use the estimated model coefficients to construct a Hotelling's T^2 chart to monitor part size from Phase II sample profiles. Additionally, they provide an informative discussion on many issues in product manufacturing from an engineering perspective that conveys well the context for profile control charts in industrial applications.

We focus on bivariate profiles and develop alternative charts to those of Colosimo et al. (2008) to monitor the size of parts. The methods proposed are nonparametric and straightforward to implement using conventional statistical software such as R, SAS, or MATLAB. We then extend the methods to monitor edge smoothness from the same bivariate profile data.

In the next section, we describe formally the bivariate profiles of interest along with practical aspects of profile data gathering. A short review follows on principal component analysis upon which the construction of the proposed Hotelling's T^2 and exponentially weighted moving average (EWMA) charts for monitoring product size and edging are built. Details on the charts, including average run length (ARL) calculation, are provided. We then focus on simulation of sample profile data yielding rough and smooth sample profiles useful in chart calibration and performance assessment. Next, the results of a numerical study of chart performance is presented using six off-target scenarios. A study that compares numerically the proposed methods with the main competing charts follows. The paper ends with a discussion of the advantages and limitations of the proposed methods.

2. Bivariate Profiles: Target, Sample and Data Gathering

Bivariate profile data consist of the (x, y) coordinates of points sequentially located along the edge of a manufactured item. If n points are sampled, the profile data take the form $\mathbf{p} = \{(x_j, y_j)\}_{j=1}^{n}$. Ideally, the points should be uniformly spread along the edge for slowly changing sections of the product and tighter on fast changing sections. Naturally, keeping track of order is critical.

The target profiles considered here are planar closed loops and thus can adequately be described by bivariate functions indexed over an interval as

$$\mathbf{p}_0(s) = (x_0(s), y_0(s)),\ s \in [a, b]; \quad \mathbf{p}_0(a) = \mathbf{p}_0(b). \tag{1}$$

The components will be continuous and differentiable except for a finite number of points representing "corners" in the items profiled. Representation (1) allows for curved loops of arbitrary shapes. Through a simple linear transformation, one could work with $[a, b] = [0, 1]$.

Most authors working on profile monitoring focus on representations of the regression type. For in-

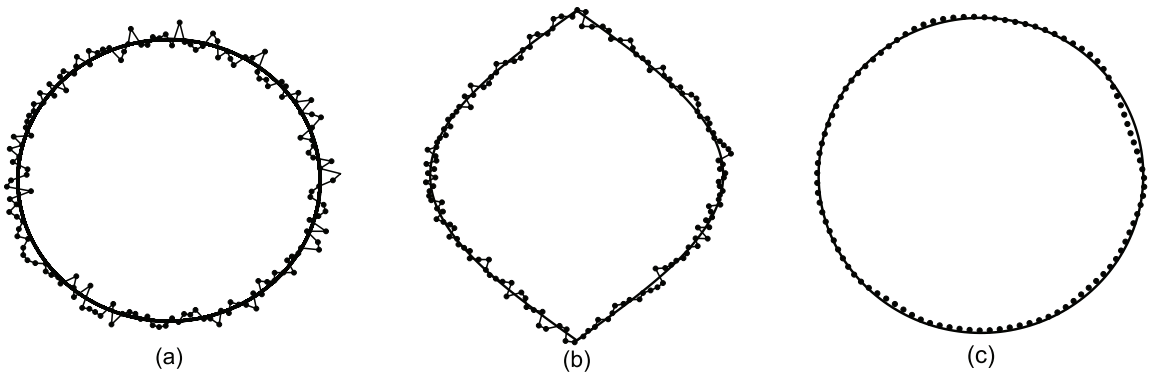

FIGURE 1. Circle and Oval Target Profiles, and Sample Profile Data With and Without Measurement Error.

stance, Qiu et al. (2010) work with

$$(x, g(x)), \quad x \in [a, b], \tag{2}$$

where x is treated as a covariate. Note that equation (2) can be represented as equation (1) by defining

$$(x_0(s), y_0(s)) = (s, g(s)), \quad s \in [a, b]. \tag{3}$$

Profiles described by equation (2) are essentially ordinary univariate functions. Generally they are inadequate to describe profiles shaped as closed loops, as they can represent only sections of such profiles. However, as discussed in the next section, depending on how the data are gathered, there are instances where closed loops can be unfolded, enabling a univariate treatment.

Several target and sample profiles are displayed in Figure 1. In Figures 1(a) and 1(c), the target profile is a circle with functional representation

$$\begin{aligned} \mathbf{p}_0(s) &= (x_0(s), y_0(s)) \\ &= (r\cos(s), r\sin(s)), \quad 0 \le s \le 2\pi, \end{aligned} \tag{4}$$

where r is the radius of the circle. In Figure 1(b), the target profile is a pointed oval with functional form

$$\begin{aligned} \mathbf{p}_0(s) &= (x_0(s), y_0(s)) \\ &= \begin{cases} (1.5\sin(2\pi s), 8s), & 0 \le s \le 0.5; \\ (1.5\sin(2\pi(s-0.5)), \\ \quad 8(1-s)), & 0.5 \le s \le 1. \end{cases} \end{aligned} \tag{5}$$

Sample profiles of 150 points each are displayed in Figure 1. In Figure 1(a) and 1(b), the points are joined by segments to identify the sequence. These cases correspond to situations where substantial measurement error is present. A smooth sample profile is shown in Figure 1(c). Profile data of this type are typically obtained when a high-precision instrument, such as a laser scanner, is used in the data-gathering process.

For planar profiles exhibiting a convex shape (i.e., $\alpha(x_1, y_1)+(1-\alpha)(x_2, y_2)$ lies inside the profile for every pair of points (x_1, y_1) and (x_2, y_2) on the profile and every $0 \le \alpha \le 1$), one may collect the data by the method of angular spanning. First, identify a "center" point for the part. In principle, any point within the profile can be used as center, provided the point is consistently used for all the parts. Second, for a given angle s in radians ($0 \le s \le 2\pi$), trace the segment at angle s starting from the center and extending up to the edge. Last, identify the coordinates $(x(s), y(s))$ of the point of intersection. The process is illustrated in Figure 2(a). One then repeats the process for selected angles $s_1, s_2, \ldots, s_n$. Here $[a, b] = [0, 2\pi]$. The resulting sample data take the form $\{(x(s_j), y(s_j))\}_{j=1}^n$. Colosimo et al. (2008) follow this approach using the radial distances in their analysis.

In other applications, a mechanical tool and object rotation are used to gather data on the shape of objects. The mechanical device has a flexible arm that is capable of retracting or extending as it comes in contact with objects. The method is depicted in Figure 2(b). First, fix the device at some point P. Second, center and fix the product at another point C, where P and C are far enough apart to allow the arm to reach the edge of the object for which we wish to capture the profile. Third, rotate the object around C and record the distance r from the device to the edge of the object. The process may be performed either continuously, for instance on graph paper, or at different points in time $s_1, s_2, \ldots, s_n$ as the object rotates at a constant speed. The data can be converted to coordinates $\{(x(s_j), y(s_j))\}_{j=1}^n$. Alternatively, one

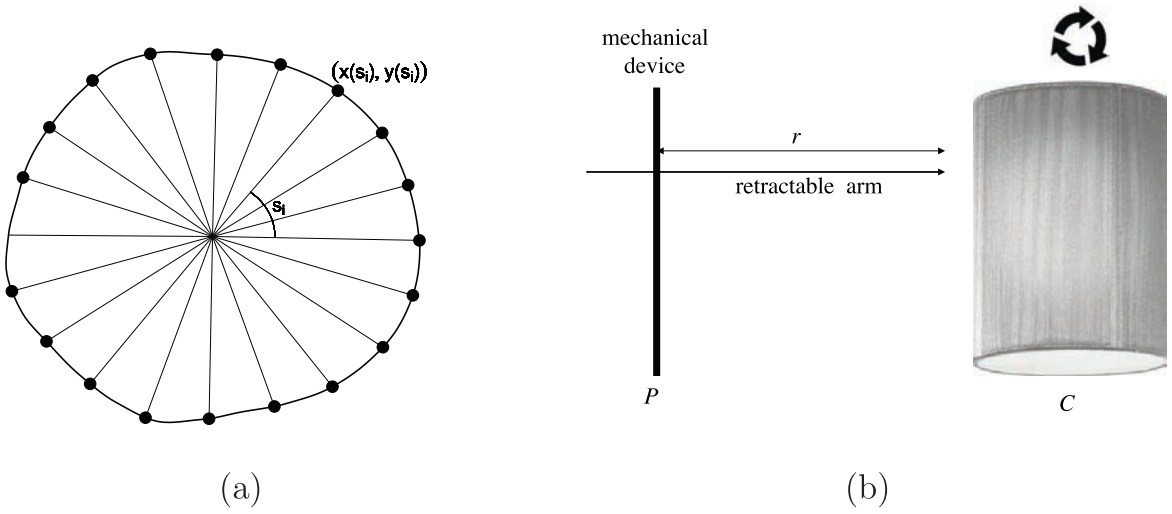

FIGURE 2. Gathering Sample Profile Data by (a) Angular Spanning and (b) Retractable-Armed Device and Object Rotation.

can work with the radial distances from the center C to the edge of the part.

Greater data-gathering efficiency and precision are achieved through the use of a laser scanner to capture the shape of an item. Companies use profile laser scanners routinely in several stages of a manufacturing process, including part inspection, robotic guidance, and shape check. In addition, scanners typically record a larger number of points on the product's edge than do manual and other low-tech methods, thus providing a more detailed description of the shape. Another advantage is their capability of capturing nonconvex shapes, a challenge for manual and mechanical methods.

3. Principal Components Chart for Product Size

3.1 Principal Components of Phase I Sample Profiles

Following the familiar approach for control-chart construction, the starting point is the availability of process data amenable to estimation of the process parameters required in setting up and calibrating the chart. After careful inspection and possibly some data cleaning, we assume that sample profile data from m parts are available when the process is stable. Further, assume that the data were collected systematically so that the indexing of points is the same across the m parts. We refer to this dataset as the Phase I or training sample profiles.

Represent the Phase I profiles as

$$\mathbf{p}_i = \{(x_{ij}, y_{ij}) = (x_i(s_j), y_i(s_j))\}_{j=1}^{n}, \quad i = 1, 2, \ldots, m.$$

Each profile will be treated as a row vector with $2n$ components, namely

$$\mathbf{p}_i = (x_{i1}, x_{i2}, \ldots, x_{in}, y_{i1}, y_{i2}, \ldots, y_{in}), \quad i = 1, 2, \ldots, m.$$

The average Phase I profile is

$$\mathbf{p}_0 = (x_{01}, x_{02}, \ldots, x_{0n}, y_{01}, y_{02}, \ldots, y_{0n}) = \sum_{i=1}^{m} \mathbf{p}_i / m. \tag{6}$$

We further assume that measurement error is either small or has been reduced, for instance, by data smoothing. Thus, the sample profiles should typically look like that in Figure 1(c). Section 4 on profile simulation provides further details.

An estimate of the covariance matrix of the profile coordinates is a key ingredient of the proposed multivariate control chart discussed in the next section. We use the m training profiles to produce this estimate. The main consideration here is that the coordinates of points physically close in the profile of an item likely exhibit high (positive) correlation while points far apart will not. We denote by $\mathbf{S}_0$ the within-profile covariance matrix associated with a profile vector $\mathbf{p}$ for items produced under normal process operation.

We will work with the sample covariance matrix estimate, namely,

$$\mathbf{S}_0 = \left(\begin{array}{c|c} \mathbf{S}_{xx} & \mathbf{S}_{xy} \\ \hline \mathbf{S}'_{xy} & \mathbf{S}_{yy} \end{array}\right),$$

with the $n \times n$ matrices $\mathbf{S}_{xx}$, $\mathbf{S}_{xy}$, and $\mathbf{S}_{yy}$ being the estimates of the within x, between x and y and within y covariance matrices, respectively.

Note that, should there be an underlying structure (e.g., time series or other) in the x and y components for the edge points of the profile, it will be reflected in the above estimates. An attractive feature of the proposed approach is that such a structure does not need to be made explicit. However, if one wishes to investigate the nature and extent of such serial correlations and is able to identify a viable time-series model for the Phase I profile data, the above covariance estimates can be replaced with the corresponding estimates from the time-series model.

Our proposed monitoring approach focuses on the leading principal components that explain the majority of the variability in the Phase I profile data. The main motivation is a reduction in dimension in the monitoring problem by concentrating on the directions that drive the variability in the profile data.

The analysis will be based on the well-known spectral decomposition for quadratic forms (see, e.g., Johnson and Wichern (2007), p. 61; Krzanowski (2008), p. 126). Specifically, if $\lambda_1, \lambda_2, \ldots, \lambda_{2n}$ are the eigenvalues of $\mathbf{S}_0$ (we assume that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_{2n}$) and $\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_{2n}$ the associated normalized column eigenvectors, then

$$\mathbf{S}_0 = \mathbf{U}\mathbf{\Lambda}\mathbf{U}', \tag{8}$$

where $\mathbf{U}$ is the $2n \times 2n$ matrix with corresponding columns given by the eigenvectors $\mathbf{U} = (\mathbf{e}_1|\mathbf{e}_2|\cdots|\mathbf{e}_{2n})$, and $\mathbf{\Lambda}$ is the $2n \times 2n$ diagonal matrix of the eigenvalues. Because the normalized eigenvectors are also orthogonal, then $\mathbf{e}'_i\mathbf{e}_i = 1$ and $\mathbf{e}'_i\mathbf{e}_j = 0$ for $i \neq j$, leading to $\mathbf{U}'\mathbf{U} = \mathbf{U}\mathbf{U}' = \mathbf{I}_{2n}$ where $\mathbf{I}_{2n}$ is the $2n \times 2n$ identity matrix and

$$\mathbf{U}'\mathbf{S}_0\mathbf{U} = \mathbf{\Lambda}.$$

Because $\mathbf{S}_0$ is a (sample) covariance matrix, it is positive semidefinite. As a result, all of its eigenvalues are nonnegative. We chose $\mathbf{U}$ so that the eigenvalues are ordered from largest to smallest.

The principal components for an n-point profile $\mathbf{p} = (x_1, x_2, \ldots, x_n, y_1, y_2, \ldots, y_n)$ are

$$\mathbf{c} = \mathbf{U}'\mathbf{p}',$$

with estimated covariance matrix

$$\mathrm{Var}(\mathbf{c}) = \mathbf{U}'\mathbf{S}_0\mathbf{U} = \mathbf{\Lambda}.$$

A measure of overall variability in a multivariate data set is the *total sample variability* (*tsv*), defined as the sum of all the sample variances for the individual variables,

$$tsv = \mathrm{tr}(\mathbf{S}_0), \tag{9}$$

where $\mathrm{tr}(\mathbf{A})$ denotes the trace of matrix $\mathbf{A}$. Because $\mathrm{tr}(\mathbf{AB}) = \mathrm{tr}(\mathbf{BA})$, it follows from equations (8) and (9) that

$$tsv = \mathrm{tr}(\mathbf{\Lambda}) = \sum_{i=1}^{2n} \lambda_i.$$

Thus, the percentage of the total variability accounted for by the first k eigenvalues is

$$P_k = 100 \times \frac{\sum_{i=1}^{k} \lambda_i}{\sum_{i=1}^{2n} \lambda_i}\%, \quad k = 1, 2, \ldots, 2n. \tag{10}$$

Consider the first k principal components, namely

$$\mathbf{c}_k = \mathbf{U}'_k\mathbf{p}', \tag{11}$$

where $\mathbf{U}_k = (\mathbf{e}_1|\mathbf{e}_2|\ldots|\mathbf{e}_k)$. One can readily show that

$$Var(\mathbf{c}_k) = \mathbf{U}'_k\mathbf{S}_0\mathbf{U}_k = \mathbf{\Lambda}_k, \tag{12}$$

where $\mathbf{\Lambda}_k = \mathrm{diag}(\lambda_1, \lambda_2, \ldots, \lambda_k)$. Thus, the total variability in the first k principal components is $\sum_{i=1}^{k} \lambda_i$ and, hence, the percentage of the total variability in the full data explained by the first k principal components is P_k as given by equation (10).

3.2 Hotelling's T^2 and Upper EWMA Charts for Product Size

Direct application of standard multivariate control charts requires inverting the $2n \times 2n$ sample covariance matrix $\mathbf{S}_0$. With large n and strong cross-correlations, $\mathbf{S}_0$ is nearly singular. To get around this difficulty, we propose to build charts using the scores from the leading principal components. The main motivations of the proposed approach are as follows. In many high-dimensional data sets, a small to moderate number of leading principal components often explain the majority of the variability in the data, thus providing a substantial reduction in dimension. In addition, the principal components associated with the larger eigenvalues tend to be most sensitive to changes in the patterns of multivariate data. Note also that the principal components are nonparametric in the sense that no statistical model

for the multivariate data is needed for their calculation and interpretation. There are many applications of principal components in quality control, particularly on the use of one to three leading components in problems involving a large number of process variables (see, e.g., Jackson (2003), chs. 6 and 7; Montgomery (2008), ch. 11; and the references therein).

A word of caution is necessary here. Looking at the leading principal components helps capture the largest sources of variation in the Phase I profile data. However, these may not always be the best for detecting subtle changes in the process. In other words, there is no reason why some special causes may not act in other directions. Note that this is a lurking feature of any data-reduction–based control chart.

From equations (11) and (12), the Hotelling's T^2 statistic from the leading k_P components for a given $1 \times 2n$ profile vector $\mathbf{p}$ is

$$T^2 = (\mathbf{p} - \mathbf{p}_0)\mathbf{U}_{k_P}\mathbf{\Lambda}_{k_P}^{-1}\mathbf{U}_{k_P}'(\mathbf{p} - \mathbf{p}_0)', \qquad (13)$$

where $\mathbf{p}_0$ is the Phase I sample average (6). The quantity T^2 is a statistical distance between the sample profile $\mathbf{p}$ and the average Phase I profile $\mathbf{p}_0$.

For Phase II profiles, we propose the Hotelling's T^2 chart arising from equation (13). Specifically, one plots T_t^2 vs. t, where T_t^2 is the value from equation (13) obtained for the sample profile $\mathbf{p}$ of an item selected from production at sampling period t. After setting an in-control average run length ARL_0, the (upper) control limit H for the T^2 chart can be obtained as the $(\text{ARL}_0 - 1)/\text{ARL}_0$ percentile of the distribution of T^2 when the process is stable. The percentile can be estimated using the m T^2-values from the Phase I profiles. A better approach is to fit an appropriate statistical distribution to those values and then use the percentile of the fitted distribution to obtain the control limits. Note that, if the Phase II profiles $\mathbf{p}$ are stable and have an approximate multivariate $N_{2n}(\mathbf{p}_0, \mathbf{S}_0)$ distribution, then T^2 follows an approximate χ^2 distribution with k_P degrees of freedom. Experimentally, we found that, in many situations, a gamma distribution, which contains the χ^2 distribution as a particular case, provides a satisfactory approximation to the distribution of T^2 in Phase I. In this case, the required control limit H satisfies

$$P(Z \le H) = \frac{\text{ARL}_0 - 1}{\text{ARL}_0} \qquad (14)$$

where $Z \sim \text{Gamma}(\alpha_0, \beta_0)$ and α_0 and β_0 are the shape and scale parameters of the approximating gamma distribution.

It should be pointed out that, as with every control chart that relies on Phase I sample data to estimate the stable process parameters, the above method provides only approximate average run lengths. This stems from the fact that the process parameters are not set to their true values but to estimates. If the Phase I sample is large, the ensuing estimates have small sampling errors, leading to more accurate ARLs.

It is well known that a more effective chart to detect small persistent changes is the exponentially weighted moving average (EWMA) chart (e.g., Montgomery (2008), ch. 9). Introduced by Roberts (1959), the chart has received much attention, particularly for normally distributed measurements. See, for example, Crowder (1987, 1989) and Lucas and Saccucci (1990). In our situation, the quantity of interest T^2 follows a gamma distribution approximately when the process is in-control.

If T_t^2 is the Hotelling's T^2 value of equation (13) from the Phase II sample profile $\boldsymbol{p}_t$ at sampling period t, the standard EWMA chart statistic is

$$W_t = (1 - \lambda)W_{t-1} + \lambda T_t^2, \qquad (15)$$

where $0 < \lambda \le 1$ is the smoothing parameter with the starting value W_0 set at the expected value for T^2. Assuming T^2 follows a Gamma distribution with shape α_0 and scale β_0 when the process is stable, then we let $W_0 = \alpha_0\beta_0$.

Because the aim is primarily to catch process deterioration, only large values of W_t are of interest. Some authors advocate the use of a reflecting barrier in the EWMA and other charts, aimed at injecting higher efficiency in detecting excursions from the stable condition (see, e.g., Gan (1992, 1994), Zhang and Chen (2004), Knoth (2005), Li et al. (2009)). The idea is to prevent the statistic from reaching values too low that unduly slow its reaction to process changes. The modified resulting chart, termed the upper EWMA, is based on the statistic

$$Z_t = \max\{B, (1 - \lambda)Z_{t-1} + \lambda T_t^2\}, \qquad (16)$$

where B $(0 \le B < \infty)$ is the barrier. Setting $B = 0$ leads to equation (15). A tested and recommended choice for B is $E(T^2)$ for T^2, calculated from on-target profiles. Thus, if T^2 follows a $\text{Gamma}(\alpha_0, \beta_0)$ distribution under stable conditions, then $B = \alpha_0\beta_0$. Further, the initial chart value Z_0 can be taken to be $Z_0 = E(\max\{B, T^2\})$, again, calculated for T^2 from the appropriate gamma distribution. One can show

that, when T^2 follows a Gamma(α_0, β_0) distribution,

$$Z_0 = BG(B; \alpha_0, \beta_0) + \alpha_0\beta_0[1 - G(B; \alpha_0 + 1, \beta_0)], \tag{17}$$

where $G(x; \alpha_0, \beta_0)$ denotes the gamma distribution function. We use the upper EWMA chart throughout our analysis with the above choices for B and Z_0.

Calibrating the upper EWMA chart involves determining the control limit H that yields a specified in-control average run length ARL_0 with the chart signaling when $W_t > H$ for the first time. The Markov chain method works well here. The method is described in detail for the standard EWMA by Lucas and Saccucci (1990) for a normally distributed chart statistic. In our situation, T^2 follows a Gamma(α_0, β_0) distribution. The relevant result is that, given any control limit H, the in-control average run length ARL_0 is approximated by

$$\text{ARL}_0 = \mathbf{v}_0'(\mathbf{I} - \mathbf{P}_0)^{-1}\mathbf{1}, \tag{18}$$

where $\mathbf{P}_0$ is the $N \times N$ transition probability matrix with entries given in equation (19) and $\mathbf{v}_0$ is the $N \times 1$ vector of 0s except for the entry corresponding to the state containing Z_0 from equation (17), where 1 is entered instead. The larger the value of N, the better the approximation. Then H is varied until ARL_0 reaches a desired (large) value. The form of the transition probability matrix $\mathbf{P}_0 = (p(i,j))_{i=1}^{N}{}_{j=1}^{N}$ is

$$p(i,j) = \begin{cases} G\left(\frac{C_j - (1-\lambda)C_i}{\lambda} + \frac{L}{\lambda}; \alpha_0, \beta_0\right), & j = 1; \\ G\left(\frac{C_j - (1-\lambda)C_i}{\lambda} + \frac{L}{\lambda}; \alpha_0, \beta_0\right) & \\ \quad - G\left(\frac{C_j - (1-\lambda)C_i}{\lambda} - \frac{L}{\lambda}; \alpha_0, \beta_0\right), & \\ & 2 \le j \le N; \end{cases} \tag{19}$$

for $1 \le i \le N$, where $L = (UCL - B)/(2m)$, $C_k = B + (2k-1)L$, and $G(x; \alpha_0, \beta_0)$ is the gamma distribution function.

Table 1 contains the control limit H when $\text{ARL}_0 = 400$ for an in-control gamma distributed T^2 for selected values of the shape parameter α_0 with scale fixed at $\beta_0 = 1$. In all cases, $\lambda = 0.1$. If $\beta_0 \neq 1$, the respective control limit is $\beta_0 H$, where H is the control limit from Table 1 for the same α_0. Interpolation can be used for shape values between those in the table. For the H values in Table 1, a total of $N = 1000$ states were used, resulting in a 1000×1000 transition probability matrix.

4. Simulating Profile Data

Many defects in manufactured parts appear in the form of cracks, scratches, or insufficient/excess ma-

TABLE 1. Control Limit (H) for the Upper EWMA Chart for In-Control Average Run Length $\text{ARL}_0 = 400$ and Gamma Distributed Phase I T^2 for Several Shape Parameter Values α_0 and Scale $\beta_0 = 1$ ($\lambda = 0.1$ in All Cases)

α_0	0.1	0.3	0.5	0.7	1.0	2.0	3.0
H	0.4008	0.7652	1.0699	1.3528	1.7560	3.0141	4.2110
α_0	4.0	5.0	6.0	7.0	8.0	9.0	10.0
H	5.3765	6.5223	7.6539	8.7750	9.8877	10.9934	12.0935
α_0	11.0	12.0	13.0	14.0	15.0	16.0	17.0
H	13.1886	14.2796	15.3667	16.4506	17.5315	18.6099	19.6858
α_0	18.0	19.0	20.0	22.0	24.0	26.0	28.0
H	20.7594	21.8311	22.9009	25.0354	27.1639	29.2871	31.4058
α_0	30.0	32.0	34.0	36.0	38.0	40.0	42.0
H	33.5203	35.6309	37.7382	39.8424	41.9438	44.0424	46.1387
α_0	44.0	46.0	48.0	50.0	55.0	60.0	65.0
H	48.2328	50.3245	52.4145	54.5025	59.7151	64.9184	70.1134
α_0	70.0	75.0	80.0	85.0	90.0	95.0	100.0
H	75.3007	80.4817	85.6569	90.8264	95.9912	101.1513	106.3072

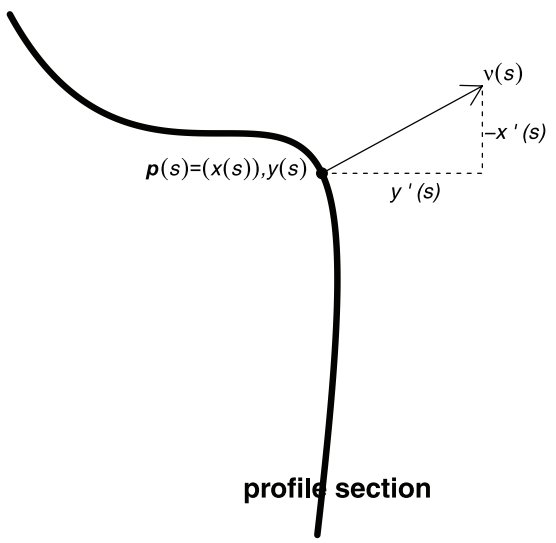

FIGURE 3. Perpendicular (Normal) Profile Vectors.

terial in some areas of the part. It is then useful to gauge disagreement between part and blueprint by measuring how far the edge of the part is from the master piece. In this section, we present a method for generating sample profile data keeping this issue in mind. The method is independent of any monitoring chart. Note that, in assessing the performance of any monitoring scheme theoretically, it is essential to have a procedure to simulate profile data.

Consider a profile $\mathbf{p}$ with functional form $\mathbf{p}(s) = (x(s), y(s))$, $a \le s \le b$, $\mathbf{p}(a) = \mathbf{p}(b)$. From basic vector calculus we know that, for a given $s \in [a, b]$, a vector perpendicular (normal) to the profile at point $\mathbf{p}(s) = (x(s), y(s))$ is any multiple of $\mathbf{v}(s) = (y'(s), -x'(s))$ with origin at $\mathbf{p}(s)$ (see Figure 3).

To illustrate the method, consider the circle target profile discussed in Section 1 with functional form given by equation (4). Without lost of generality, a radius $r = 1$ will be used. We focus on generation of sample Phase I profiles but the method can readily be used for out-of-control profiles with specified nonconformance patterns (see Section 6). Consider generating profiles with n points indexed by equally spaced values, $s_i = 2\pi(i-1)/n$, $i = 1, 2, \ldots, n$. Two types of profiles will be generated; we label them here as "rough" and "smooth".

Rough sample profiles are produced as follows:

(a) For each index s_i, obtain the outside normal vector $\nu(s_i) = (\cos(s_i), \sin(s_i))$ to the circle at point $(x_0(s_i), y_0(s_i))$ and normalize it to get $(y_0'(s_i), -x_0'(s_i))/\sqrt{[y_0'(s_i)]^2 + [-x_0'(s_i)]^2} = (\cos(s_i), \sin(s_i))$.

(b) Generate a normal $N(0, \sigma)$ random value a_i and obtain point $\mathbf{b}_i = (\cos(s_i), \sin(s_i)) + a_i(\cos(s_i), \sin(s_i))$. Here, σ is kept constant for all draws and is expected to be small if the process is tight to the target.

(c) The ensuing rough profile consists of the points $\{\mathbf{b}_i, i = 1, 2, \ldots, n\}$, where the a_i random numbers are generated independently of one another.

The resulting profiles are on target in the sense that the expected (or average) coordinate values for the points generated at each index s_i equal the corresponding blueprint coordinates for the point. Note also that we can replace the normal distribution in (b) with other distributions centered at 0 such as $U(-\sigma, \sigma)$, for example. Furthermore, the a_i values may be correlated, for instance, following a multivariate normal distribution with 0 mean and selected covariance matrix, such as in a circular time series (e.g., see Fisher and Lee (1994)). In all of these cases, the sample profiles will be on target on average. Note also that the ensuing profiles may be spiky, perhaps adequate to model situations with substantial measurement error. This is in contrast with the smooth profiles one typically sees in many finished products.

For the smooth profiles, we go a step further as follows.

(d) Consider a rough profile $\{\mathbf{b}_i = (x_i, y_i),\ i = 1, 2, \ldots, n\}$ from (c) above. Smooth the respective values $(x_1, x_2, \ldots, x_n)$ and $(y_1, y_2, \ldots, y_n)$ over $(s_1, s_2, \ldots, s_n)$, separately. Then, from the smoother, predict corresponding values for $(s_1, s_2, \ldots, s_n)$ resulting in $(x_1^{(s)}, x_2^{(s)}, \ldots, x_n^{(s)})$ and $(y_1^{(s)}, y_2^{(s)}, \ldots, y_n^{(s)})$. Assemble the univariate predictions to form points $\{\mathbf{b}_i^{(s)} = (x_i^{(s)}, y_i^{(s)}),\ i = 1, 2, \ldots, n\}$. These points form the smooth n-point sample profile.

Many smoothers are available in the literature, each offering control parameters to regulate different aspects of the smoothing process. Chief among them is the smoothing parameter that controls the degree of smoothing. We experimented with three smoothers available, either directly or through downloadable libraries, in the R computing package (R Development Core Team (2008)). Most of them involve some form of regression. Perhaps the best known is `loess()` in the core R package based on

fitting local polynomial regressions. Also available in the core R package is `smooth.spline()`, which fits a cubic smoothing spline. A variant of `loess()` is `lpridge()`, which fits a local polynomial regression with ridging, available in library `lpridge` of Seifert (2007). Each of these smoothing methods is nonparametric and produces on-target profiles whenever the original rough data are on target. Also, each method allows predictions. An interesting and important feature of the smooth profiles is that neighboring points will tend to be similar and thus will exhibit high positive coordinate-wise correlation while points far apart will be nearly uncorrelated.

The rough sample profile for the circle shown in Figure 1(a) contains $n = 150$ points generated using the above method, with a_i values independently generated from a normal distribution with mean 0 and standard deviation 0.1. The same process produced the rough sample profile in Figure 1(b) but using the functional form for the oval target profile from equation (5). The smooth sample profile shown in Figure 1(c) comes from a rough sample profile of $n = 100$ points generated with independent $a_i \sim N(0, 0.1)$ values and smoothed using `smooth.spline()` with the smoothing parameter set at spar = 0.6.

For smooth Phase I sample profiles, the role of σ is to reflect the intrinsic leeway of the manufacturing process around the blueprint shape. Simulation of tight processes requires σ to be small, while a large σ recreates processes with substantial wiggling around the target.

Exampe 1: Phase I Profile Sample

The purpose of this example is to illustrate the foregoing methods and to generate a sample of Phase I profiles to be used in calibrating the proposed charts. A sample of $m = 1000$ rough and smooth independent profiles was generated employing the above method with the circle of radius $r = 1$ as the blueprint. The index values were $s_i = 2\pi(i-1)/200$ radians, $i = 1, 2, \ldots, 200$, with the a_i values generated from $N(0, \sigma)$ with $\sigma = 0.1$ for the rough profiles. The rough profiles were then smoothed using `smooth.spline()` in `R` with *spar* $= 0.6$ to obtain the smooth profiles. For the remainder of the article, we consider these sample profiles as typical of profiles seen when the process is stable. Thus, they will form the Phase I sample.

The sample covariance matrix $\mathbf{S}_0$ was obtained for the 400 coordinate values (200 x-coordinates, 200 y-coordinates) for the 1000 smooth profiles. A contour plot for the resulting 400×400 sample correlation matrix was produced and examined. The plot reveals that x-coordinates from neighboring points are highly positively correlated and so are the respective y-coordinates. However, the correlations between x- and y-coordinates for neighboring points show regular sections of very high positive correlations and sections of very high negative correlations as one crosses the quadrants. This feature is explained by the coordinate constraints arising from their circular conformance. Finally, coordinates of points far apart exhibit nearly 0 correlation.

5. Measuring and Monitoring Edging

5.1 Appraising Edging

The question addressed here is: how well does the edge of a Phase II part conform to the edge of the target part? We assess this conformance relative to the conformance of the Phase I profiles when the process operates in control.

The first task is to develop a measure of agreement in edging between a part and the target part. We propose to focus on the angles (in radians) between the outside normal vector of the target profile and the outside normal vector at the edge of the part at points where information is available in the sample profile. Specifically, consider the normal vector $\boldsymbol{\nu}_0(s)$ at point $\mathbf{p}_0(s) = (x_0(s), y_0(s))$ on the target profile, given by $\boldsymbol{\nu}_0(s) = (u_0(s), v_0(s)) = (y_0'(s), -x_0'(s))$ (see Figure 3). Similarly, for a smooth sample profile $\mathbf{p}(s)$ with index values $s_1, s_2, \ldots, s_n$, calculate for each $s = s_i$ the outside normal vector $\boldsymbol{\nu}(s_i) = (u(s_i), v(s_i))$. To do this, one can use the coordinate derivatives provided by the univariate smoothers or use basic numerical approximation of the coordinate derivatives, e.g., $f'(x_0) \doteq [f(x_0+\delta) - f(x_0-\delta)]/(2\delta)$. Next align $\boldsymbol{\nu}_0(s)$ and $\boldsymbol{\nu}(s)$ to have the same origin, for instance using $\mathbf{p}(s)$ as origin. Finally, calculate the signed angle (in radians) between $\boldsymbol{\nu}_0(s)$ and $\boldsymbol{\nu}(s)$,

$$\theta(s) = \arctan\left(\frac{u_0(s)v(s) - v_0(s)u(s)}{u_0(s)u(s) + v_0(s)v(s)}\right). \qquad (20)$$

Repeat the calculation across all the available index points $s_1, s_2, \ldots, s_n$ to obtain the $1 \times n$ angular vector

$$\boldsymbol{\theta} = (\theta_1, \theta_2, \ldots, \theta_n) = (\theta(s_1), \theta(s_2), \ldots, \theta(s_n)).$$

Note that the angle orientation in equation (20) is counterclockwise, that is, $\theta(s)$ is positive if one travels against the clock when going from $\boldsymbol{\nu}_0(s)$ to $\boldsymbol{\nu}(s)$, and negative otherwise.

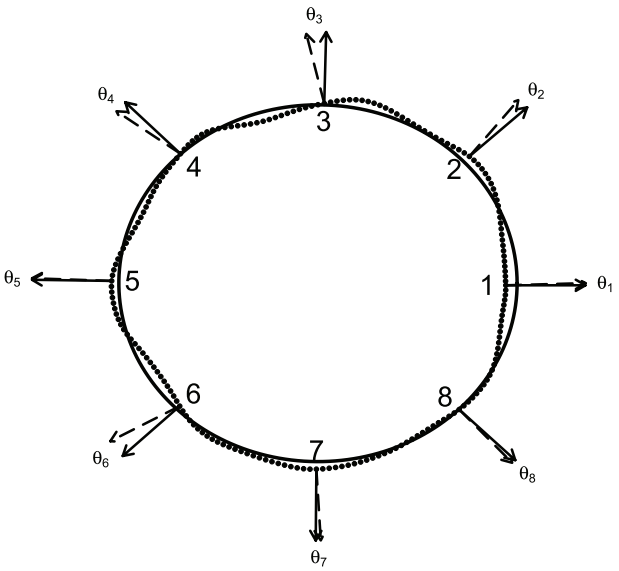

FIGURE 4. Angles Between Outside Normal Vectors for Sample and Target Profiles.

Figure 4 illustrates the angles at eight points from a smooth sample profile of 200 points generated using the method described in Section 4 for the circle target profile of radius $r = 1$. Here, $\sigma = 0.12$ and smoothing was done using `smooth.spline()` with `spar = 0.6`. The numerical values obtained for θ_1 to θ_8 (in radians) were 0.031, 0.151, 0.251, 0.168, −0.021, −0.230, 0.063, and −0.053. Note that the largest disagreement in edging occurs at location 3, even though the two profiles are very close to each other. On the other hand, at locations 1 and 5, the edging is very similar (i.e., small angles) but the target and sample profile points are far apart. These remarks highlight the difference between product size and edging. Note that the angles, including signs, are invariant under positive scaling, translation, and rotation.

5.2 Angular Principal Component Charts for Edging

To construct control charts for edging, one can use the n-point Phase I smooth profiles $\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_m$ discussed earlier to render Phase I angular vectors $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \ldots, \boldsymbol{\theta}_m$ by applying the method just described. We then calculate the sample covariance matrix $\mathbf{A}_0$ of the angular vectors.

Analogous to the charts developed for part size, the leading principal components of $\mathbf{A}_0$ can be used to construct Hotelling's T^2 and EWMA control charts for edging. Denoting by $\mathbf{V}_k$ the $n \times k$ matrix containing as columns the coefficients (loadings) for the k leading principal components of $\mathbf{A}_0$ with associated eigenvalues $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_k$, we then use the scores $\mathbf{V}_k'(\boldsymbol{\theta} - \boldsymbol{\theta}_0)'$ to calculate the Hotelling's T^2 values

$$T^2 = (\boldsymbol{\theta} - \boldsymbol{\theta}_0)\mathbf{V}_k \boldsymbol{\Delta}^{-1} \mathbf{V}_k'(\boldsymbol{\theta} - \boldsymbol{\theta}_0)', \qquad (21)$$

where $\boldsymbol{\theta}$ is the angular vector for a smooth sample profile $\mathbf{p}$, $\boldsymbol{\theta}_0$ is the sample average of the Phase I angular vectors, and $\boldsymbol{\Delta} = \text{diag}(\alpha_1, \alpha_2, \ldots, \alpha_k)$ (see equation (13)). Thus, T^2 is a statistical distance comparing the edging in the sample profile to the average Phase I edging. The T^2 values from the m (in-control) Phase I angular vectors can be used for chart calibration.

In Phase II, the value T_t^2 obtained from the angular vector $\boldsymbol{\theta} = \boldsymbol{\theta}_t$ arising from a part chosen at sampling period t can be used to construct a Hotelling's T^2 chart for edging by plotting T_t^2 vs. t. Further, one can use equation (15) with the angular T_t^2 to obtain an upper EWMA chart for edging.

Example 2: Phase I Angular Profile Sample

Signed angular vectors were obtained by the method just described for the smooth profiles generated in example 1. Each of the $m = 1000$ ensuing vectors has $n = 200$ angles,

$$\begin{aligned}\boldsymbol{\theta}_i &= (\theta_{i1}, \theta_{i2}, \ldots, \theta_{i200}) \\ &= (\theta_i(s_1), \theta_i(s_2), \ldots, \theta_i(s_{200})), \quad i = 1, 2, \ldots, 1000.\end{aligned}$$

We use these data as a Phase I sample of angular vectors.

The 200×200 sample covariance matrix was calculated for the $m = 1000$ Phase I angular vectors. The resulting correlation matrix resembles closely the upper-left and lower-right sections of the correlation matrix for size in example 1. The salient features are a strong positive correlation between angles for profile points close to each other while angles for points far apart are nearly uncorrelated. The underlying cause for the high correlations is the smoothness of the profiles.

6. Phase II Analysis

The methods proposed will be illustrated in this section with the circle profile of equation (4) with radius $r = 1$ as the blueprint. On-target process operation generates smooth profiles as discussed in Sections 4 and 5, with parameters set at $r = 1$, $\sigma = 0.1$, and $spar = 0.6$. Here, r has the biggest say on the size of the part while σ and spar determine the degree of

edge wiggling/smoothness. Changing these parameters produces varying deviations from normal stable operation. The 1000 smooth sample profiles from example 1 and corresponding angular vectors from example 2 are used as the Phase I data.

6.1 Monitoring Size: Chart Calibration

The smooth sample profiles from example 1 were randomly split into two groups of size 500 each. The sample covariance matrix $\mathbf{S}_0$ was calculated for one of the subsets. Then the eigenvalues and eigenvectors of $\mathbf{S}_0$ were obtained. Figure 5(a) displays the percentage of variability explained by the leading principal components as calculated from equation (10) for up to 50 out of the 400 components. The plot reveals that the principal components associated with the largest 15, 18, and 24 eigenvalues account for 90%, 95%, and 99% of the total variability, respectively. We worked with the 24 largest eigenvalues, which were stored in the diagonal matrix $\mathbf{\Lambda}_{24}$, and loadings (coefficients) of the associated components, stored as columns in the 400×24 matrix $\mathbf{U}_{24}$.

Matrices $\mathbf{U}_{24}$ and $\mathbf{\Lambda}_{24}$ were then used to obtain T^2 values, using equation (13), for the 500 Phase I smooth sample profiles in the other subset. If the coordinates for all profile points followed a multivariate normal distribution, the T^2 values would be chi-squared distributed with 24 df, approximately (e.g., see Johnson and Wichern (2007), p. 163). We found that a gamma distribution provides a more satisfactory fit. Using maximum likelihood and the 500 T^2 values, we get parameter estimates $\alpha_0 = 11.5434$ and $\beta_0 = 2.1806$. The fit is very good, as shown by a gamma QQ plot of the T^2 values. See Figure 5(b). Note that the chi-square distribution with 24 df is a member of the gamma family with shape 12 and scale 2.

Control limits were then set for the Hotelling's T^2 and EWMA charts. The on-target average run length was set at $\text{ARL}_0 = 400$ for both charts. For our case, with $\alpha_0 = 11.5451$ and $\beta_0 = 2.1803$, the resulting control limit for the Hotelling's T^2 chart obtained from equation (14) is $H = 50.9655$. As for the EWMA chart, Table 1 was used to calculate the required control limit. Noting that α_0 falls between 11.0 and 12.0, linear interpolation was applied to calculate the control limit, yielding $H = \beta_0[13.1886 + (14.2796 - 13.1886)(\alpha_0 - 11.0)] = 30.0518$. The smoothing parameter in the upper EWMA chart was set to $\lambda = 0.1$ throughout.

6.2 Monitoring Size: Chart Performance

In Phase II, several process departures from the circle profile were selected to test the ability of the

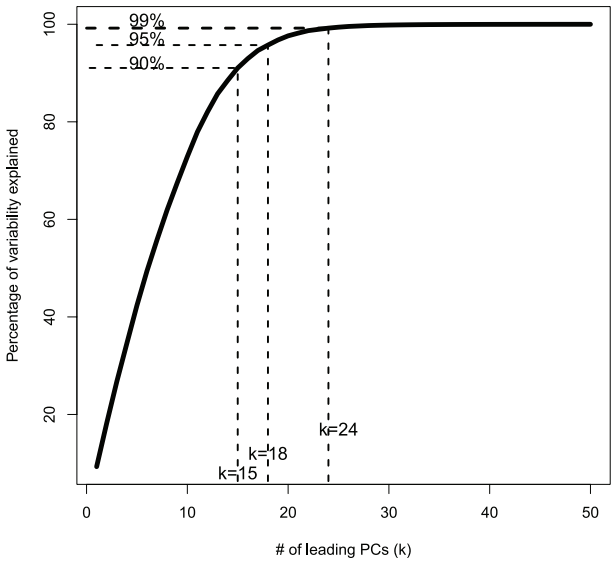

(a)

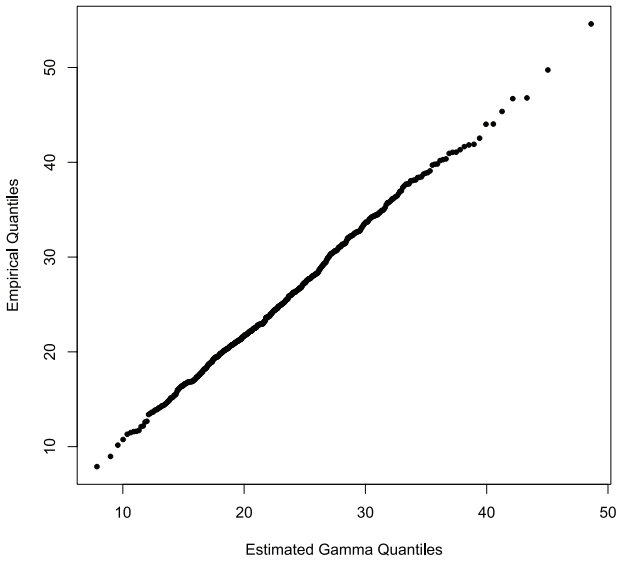

(b)

FIGURE 5. (a) Variability Explained by the Principal Components of the Sample Covariance Matrix for the x- and y-Coordinates of 500 of the Phase I Smooth Profiles. (b) Gamma QQ Plot for Resulting T^2 Values from the 24 Leading Principal Components.

charts to detect changes in size. Naturally, with high-dimensional measurements, there are many ways in which a profile of a manufactured product may deviate from its stable average. Our choices here focus on systematic departures in size, in the degree of edge wiggling/smoothness, and on abnormalities in certain spots of the product.

The following process departures were selected (in all cases $spar = 0.6$):

Case 1: *Undersized Product.* The product's circular profile is systematically of radius $r = 0.95$ (rather than $r = 1.0$), but the degree of wiggling/smoothness remains in control ($\sigma = 0.1$).

Case 2: *Oversized Product.* The radius is systematically $r = 1.03$ (rather than $r = 1.0$) with on-target degree of wiggling/smoothness ($\sigma = 0.1$).

Case 3: *Smoother Product.* The product shows on-target radius ($r = 1.0$) but is systematically smoother than in Phase I with ($\sigma = 0.09$).

Case 4: *Rougher Product.* The degree of wiggling is systematically higher, coming from $\sigma = 0.11$, but the radius is on-target ($r = 1.0$).

Case 5: *Chipped Product.* The product exhibits a dent in a spot but in-control radius ($r = 1.0$) and degree of wiggling ($\sigma = 0.1$).

Case 6: *Knobby Product.* A lump appears systematically in the product but everything else is in control ($r = 1.0$, $\sigma = 0.1$).

Illustrations of these scenarios are shown in Figure 6. Note that the chip is on the southeast section (case 5) and lump in the same place (case 6). The `R` code used to generate chips and lumps is available on request.

The off-target average run length (ARL_1) was then calculated for both the Hotelling's T^2 and upper EWMA charts. Thus, ARL_1 is the number of samples that, on average, is needed to signal a change in the process when a change to a particular case of the six selected scenarios occurred. Although tedious, the most direct approach to accomplish this task is run-length simulation. However, we found that the T^2 values from off-target sample profiles obtained using U_{24} and Λ_{24} also follow closely a gamma distribution but with different shape and scale parameters. Capitalizing on this finding, 2000 off-target smooth profiles were generated and resulting T^2 values calculated. A gamma distribution was fitted and then the average run length was calculated using the proposed Markov chain method with the previously calculated control limit H. Although the ARL_1 values were fairly stable, the method was repeated eight times and the resulting ARL_1 values averaged.

The two columns under "ARL_1 for size" in Table 2 display the results. Recall that both charts were calibrated to $ARL_0 = 400$. When the departure is small (cases 4, 5, and 6), the EWMA chart performs better than the Hotelling's T^2 chart. Note that, when the change results in smoother product (case 3), the T^2 will tend to be smaller than in the on-target case and thus will take much longer to detect the change because we are using only an upper control limit (H).

6.3 Monitoring Edging: Chart Calibration

The 1000 Phase I 1×200 sample angular vectors from example 2 were similarly split into two subsets of 500 vectors each. The 200×200 covariance matrix $\mathbf{A}_0$ was calculated for one subset and the eigenvalue/eigenvector decomposition obtained. The 35 leading principal components, which explain 99% of the total variability, will be used in the charts. The associated eigenvalues were stored in diagonal matrix $\mathbf{\Delta}_{35}$ and the corresponding loadings (eigenvectors) as columns in matrix $\mathbf{V}_{35}$.

Next, the T^2 values for the 500 angular vectors in the other subset were calculated from equation (21). We fit a gamma distribution using maximum likelihood, resulting in parameter estimates $\alpha_0 = 17.5464$ and $\beta_0 = 2.0670$. A gamma QQ plot (not shown) reveals a satisfactory fit. Using equation (14) with $ARL_0 = 400$ yields $H = 65.318$ as the control limit for the Hotelling's T^2 chart. Interpolation in Table 2 gives $H = 41.9037$ as the control limit for the upper EWMA chart.

6.4 Monitoring Edging: Chart Performance

A similar approach to product size was followed to assess the performance of the charts in detecting changes in product edging from the circle blueprint. The six off-target scenarios previously considered were used. The off-target average run lengths obtained are displayed in Table 2 in the columns under "ARL_1 for edging". Some interesting patterns emerge. Somewhat expected, when the product is bigger than the blueprint and the same degree of wiggling occurs, or the product is smoother but similar in size as the blueprint, the outside normal vectors for the blueprint and product will tend to be aligned. As a result, the angles between these vectors will

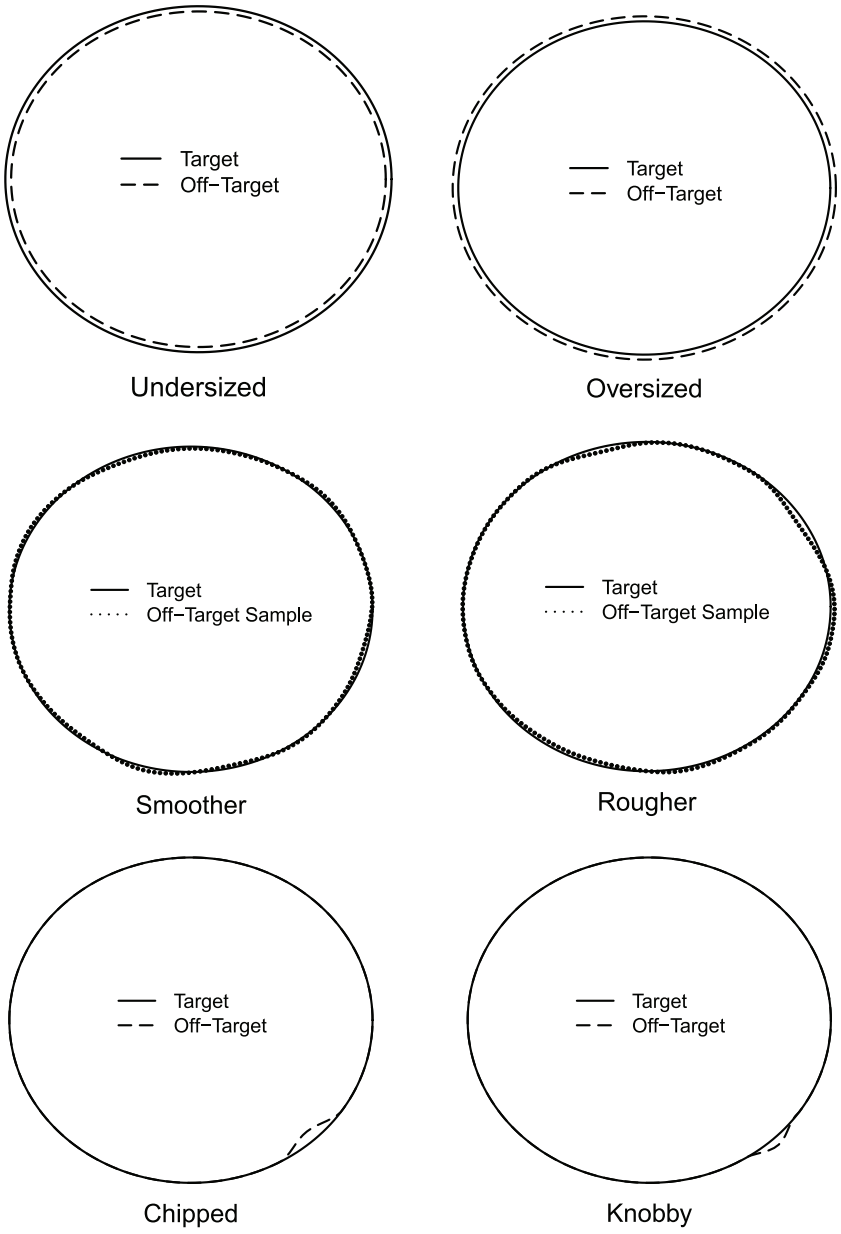

FIGURE 6. Sample Profiles for Each of the Six Off-Target Scenarios Considered.

be smaller and so will be the T^2 values, resulting in very large ARL_1. When the product manufactured is smaller but with the same degree of smoothing, the wiggling will appear larger and thus will be detected faster. For chipped or knobby product, both charts detect the change from on-target product at about the same speed. In the two particular cases shown, the chip and the knob are fairly smooth. However, if they were sharper, i.e., spiky inwards or outwards defects, the edging chart will detect them more rapidly than the chart for size.

7. Comparison with Other Charts

In this section, we compare the performance of our proposed upper EWMA chart for size with the T2+S2 chart proposed by Colosimo et al. (2008). This chart performs the best among the three charts they evaluated. To model their profiles, they consider the spatial autoregressive regression (SAR) model discussed by Cressie (1993, p. 441). This parametric framework has been successfully applied to environmental spatial data, where varying degrees of corre-

TABLE 2. Off-Target Average Run Length (ARL_1) for Detecting a Change in Product Size and Edging for Six Selected Departures from Normal Process Operation with the Hotelling's T^2 and the Upper EWMA Charts Calibrated at $ARL_0 = 400$

Off-target feature	ARL_1 for size		ARL_1 for edging	
	Hotelling's T^2	Upper EWMA	Hotelling's T^2	Upper EWMA
Case 1: Undersized product	1.03	1.02	78.21	19.49
Case 2: Oversized product	1.03	1.02	>2000	>2000
Case 3: Smoother product	>2000	>2000	>2000	>2000
Case 4: Rougher product	45.11	11.26	29.90	7.41
Case 5: Chipped product	9.11	2.97	10.32	3.03
Case 6: Knobby product	19.44	3.97	22.56	4.86

lation as well as short- and long-scale effects need particular consideration.

Colosimo et al. (2008) consider profiles consisting of radial distances from the center to the edge of the product, measured at $N = 748$ angles. They use 100 Phase I real-sample profiles to estimate the SAR model parameters. They select the resulting SAR model of order 2 (denoted SARX(2)) as the baseline model (Phase I) for the profiles. Then they simulate 20,000 Phase I profiles to calibrate their T2+S2 chart, choosing an upper control limit with on-target average run length $ARL_0 = 100$. Their equation (16) describes in detail the generating process for Phase I profiles.

In Phase II, the charts were applied to sample profiles with assignable causes introduced by spindle-motion errors that affect product roundness. Specifically, the off-target scenarios considered were obtained by adding to a baseline (on-target) sample profile a spurious harmonic of a certain frequency and with amplitude directly proportional to a parameter δ. All the values of δ used were between 0 and 1, with larger values producing more severe departures from the baseline profile. Equations (18)–(21) of Colosimo et al. (2008) provide the specific form of the altered shapes. Their Figure 10 shows the baseline profile while Figure 11 depicts several profiles with assignable causes of the type just described. For each off-target scenario, they simulated 1000 run lengths to estimate the off-target average run length ARL_1. Their Figure 12 displays the results.

Our assessment uses the same Phase I and Phase II profile simulation process as in Colosimo et al. (2008). Following our approach, first 20,000 on-target sample radial profiles ($N = 748$ points each) were simulated. Each profile was smoothed using `smooth.spline()` in R with $spar = 0.5$. The 748×748 covariance matrix was calculated for 10,000 of the smooth radial profiles and the principal component analysis performed. The $k = 32$ leading principal components, which explain 99% of the total variability, were retained. Then the T^2 values from equation (13) were calculated for the other 10,000 smooth sample radial profiles. The familiar monotonic transformation $Y = -(T^2)^{-1/3}$ was applied and a QQ normal plot revealed that the transformed values were close to normal. The resulting mean and standard deviation were $\overline{y}_0 = -0.313513$ and $s_0 = 0.051414$, respectively. Then the upper EWMA chart for normal data (using $\overline{y}_0$ and s_0) was calibrated, resulting in an upper control limit of $H = -0.287777$ for an on-target average run length $ARL_0 = 100$. The Markov chain method was used to calculate ARLs.

For each scenario with assignable causes, 1,000 sample profiles were simulated. Using the same coefficients (loadings) for the 32 leading principal components from Phase I, the T^2 values were calculated using equation (13). After transformation, a normal QQ plot was drawn to check for normality, which each plot supported, and the new sample mean ($\overline{x}$) and standard deviation (s) were calculated. Using the Markov chain method with the new $\overline{x}$ and s but the Phase I upper control limit, the off-target ARL was obtained.

Table 3 displays the results of the simulation study. The ARL values for T2+S2 were read from Figure 12 of Colosimo et al. (2008). The smaller ARL of the two methods is highlighted in bold. Both methods take fewer runs on the average to detect the

TABLE 3. Average Run Length Performance of the T2+S2 Chart Proposed by Colosimo et al. (2008) (as Read from Their Figure 12), and of Our Proposed Upper EWMA Chart for Size on the 20 Scenarios with Assignable Causes

Scenario	δ	T2+S2	Upper EWMA
On target		100.0	100.0
Half-frequency spindle-motion error	0.30	38	**11.2**
	0.40	20	**5.2**
	0.50	10	**2.8**
	0.60	6	**2.1**
Bilobe spindle-motion error	0.25	**40**	48.8
	0.50	**12**	17.0
	0.75	**6**	7.9
	1.00	**4**	4.7
Trilobe spindle-motion error	0.25	43	**39.9**
	0.50	**13**	15.5
	0.75	**7**	**7.0**
	1.00	**4**	5.5
Four-lobe spindle-motion error with fixed phase	0.05	36	**1.1**
	0.06	28	**1.0**
	0.07	20	**1.0**
	0.08	14	**1.0**
Four-lobe spindle-motion error with random phase	0.05	39	**1.0**
	0.06	28	**1.0**
	0.07	19	**1.0**
	0.08	16	**1.0**

presence of the assignable cause as δ increases, reflecting the fact that larger values of δ are indicative of greater departure from in-control process operation. It is clear that our proposed chart is competitive compared with the T2+S2 chart in the scenarios simulated with its best showing when a four-lobe spindle-motion error is present with either fixed or random phase.

An important distinction between the two methods is worth noting. The Colosimo et al. (2008) approach is parametric, based on a statistical model for the profile points. Thus, their method would be expected to do well when the model provides a good fit for the profile data. On the other hand, our proposed principal components-based method is approximate but assumes no specific statistical model. We expect it to do well in many types of profile data but perhaps not as well as a chart that is built from the correct statistical model for the profile data. At the same time, such charts can be less robust to model changes.

8. Conclusions

This article presents methods to monitor size and edging from product profile data. The focus is on the leading principal components of the covariance matrix for the sample coordinates of the profile points, which contain information about product size, and the leading principal components of the covariance matrix of certain profile angular vectors, which contain information about product edging. When compared with standard charts for multivariate processes, the most attractive features of the methods proposed are (a) a reduction in dimensionality, (b) accounting for the cross-correlation structure, (c) implementation requiring tools readily available in most

standard statistical software, and (d) the methods are largely nonparametric. A simulation study shows that the methods proposed perform well in detecting a variety of departures from in-control process operation.

In the case of product size, the proposed Hotelling's T^2 chart provides an alternative to the chart developed by Colosimo et al. (2008). The latter assumes a parametric spatial autoregressive regression (SARX) model originally developed for spatial data. Some effort is needed to elicit the covariates to be used because there are no measured covariates in the case of product profile data. It is of interest to more fully compare the performance of the two charts to monitor product size. Also of interest will be to develop an extension of the SARX-based chart to monitor product edging.

The high dimension of product profile data presents other challenges. For instance, a larger number of Phase I multivariate samples (i.e., product profiles) is needed in chart calibration. Another challenge is that departures from normal process operation are not amenable to simple characterization through a few numbers, as is the case with standard multivariate process monitoring. As a result, optimal tuning parameters (e.g., optimal degree of smoothing λ in the case of the EWMA) are not easily found. Note that, in our illustrations in this paper, we worked with the standard value of λ, namely $\lambda = 0.1$.

The proposed methods, which use an upper control limit for the size and edging charts, are designed to spot process worsening, reflected in large values of the respective T_t^2. However, one might be interested in detecting improvements as well. For instance, we might like to detect greater conformance to blueprint arising from reduced wiggling. In this case, we can run two charts, one for catching process deterioration and one for process improvement. Specifically, we can use $W_t = \max\{B, (1-\lambda)W_{t-1} + \lambda T_t^2\}$ with an upper control limit to spot process deterioration (as done in this paper) and $Z_t = \min\{B, (1-\lambda)Z_{t-1} + \lambda T_t^2\}$ with a lower control limit to detect process improvement.

One important follow-up to the methods proposed is the extension to a genuinely multivariate EWMA (e.g., see Lowry et al. (1992)). It entails working with the multivariate chart statistic

$$\mathbf{W}_t = (1-\lambda)\mathbf{W}_{t-1} + \lambda \mathbf{U}_k'(\mathbf{p}_t - \mathbf{p}_0)' \quad t = 1, 2, \ldots,$$

to monitor size and

$$\mathbf{W}_t = (1-\lambda)\mathbf{W}_{t-1} + \lambda \mathbf{V}_k'(\boldsymbol{\theta}_t - \boldsymbol{\theta}_0)', \quad t = 1, 2, \ldots$$

to monitor edging. In either case, an out-of-control signal is issued when $\mathbf{W}_t \boldsymbol{\Sigma}_W^{-1} \mathbf{W}_t' > H$. This chart should have greater power to detect persistent process changes.

Acknowledgments

We are grateful to an anonymous referee whose comments led to substantial improvements in content and presentation. We also thank Professor Bianca Colosimo for kindly lending us the computer programs to generate the radial profiles of Section 7 that facilitated a comparison of our methods with those of Colosimo et al. (2008). This research was supported in part through grants from the Natural Sciences and Engineering Research Council of Canada.

References

Chang, T. C. and Gan, F. F. (1994). "Optimal Designs of One-sided EWMA Charts for Monitoring a Process Variance". *Journal of Statistical Computation and Simulation* 49, pp. 33–48.

Colosimo, B. M.; Semeraro, Q.; and Pacella, M. (2008). "Statistical Process Control for Geometric Specifications: On the Monitoring of Roundness Profiles". *Journal of Quality Technology* 40(1), pp. 1–18.

Cressie, N. A. C. (1993). *Statistics for Spatial Data*, revised edition. New York, NY: John Wiley & Sons.

Crowder, S. V. (1989). "Design of Exponentially Weighted Moving Average Schemes". *Journal of Quality Technology* 21(3), pp. 155–162.

Fisher, N. I. and Lee, A. J. (1994). "Time Series Analysis of Circular Data". *Journal of the Royal Statistical Society, Series B* 56(2), pp. 327–339.

Gan, F. F. (1993). "Exponentially Weighted Moving Average Control Charts with Reflecting Boundaries". *Journal of Statistical Computation and Simulation* 46, pp. 45–67.

Gupta, S.; Montgomery, D. C.; and Woodall, W. H. (2006). "Performance Evaluation of Two Methods for Online Monitoring of Linear Calibration Profiles". *International Journal of Production Research* 44(10), pp. 1927–1942.

Jackson, J. E. (2003) *A User's Guide to Principal Components*. New York, NY: Wiley.

Johnson, R. A. and Wichern, D. W. (2007). *Applied Multivariate Statistical Analysis*, 6th edition. Upper Saddle River, NJ: Pearson Prentice Hall.

Kazemzadeh, R. B.; Noorossana, R.; and Amiri, A. (2008). "Phase I Monitoring of Polynomial Profiles". *Communications in Statistics—Theory and Methods* 37, pp. 1671–1686.

Kim, K.; Mahmoud, M. A.; and Woodall, W. H. (2003). "On the Monitoring of Linear Profiles". *Journal of Quality Technology* 35, pp. 317–328.

Knoth, S. (2005). "Accurate ARL Computation for EWMA–S2 Control Charts". *Statistics and Computing* 15, pp. 341–352.

KRZANOWSKI, W. J. (2008). *Principles of Multivariate Analysis: A User's Perspective*, revised edition. New York, NY: Oxford University Press.

LI, Z.; WANG, Z.; and WU, Z. (2009). "Necessary and Sufficient Conditions for Non-Interaction of a Pair of One-Sided EWMA Schemes with Reflecting Boundaries". *Statistics and Probability Letters* 79(3), pp. 368–374.

LOWRY, C. A.; WOODALL, W. H.; CHAMP, C. W.; and RIDGON, S. E. (1992). "A Multivariate Exponentially Weighted Moving Average Control Chart". *Technometrics* 34(1), pp. 46–53.

LUCAS, J. M. and SACCUCCI, M. S. (1990). "Exponentially Weighted Moving Average Schemes: Properties and Enhancements" (with Discussion). *Technometrics* 32(1), pp. 1–12.

MAHMOUD, M. A. and WOODALL, W. H. (2004). "Phase I Analysis of Linear Profiles with Calibration Applications". *Technometrics* 46, pp. 380–391.

MONTGOMERY, D. C. (2008). *Introduction to Statistical Quality Control,* 6th edition. New York, NY: Wiley.

NOOROSSANA, R.; EYVAZIAN, M.; AMIRI, A.; and MAHMOUD, M. A. (2010). "Statistical Monitoring of Multivariate Multiple Linear Regression Profiles in Phase I with Calibration Application". *Quality and Reliability Engineering International* 26, pp. 291–303.

QIU, P.; ZOU, C.; and WANG, Z. (2010). "Nonparametric Profile Monitoring by Mixed Effects Modeling" (with Discussion). *Technometrics* 53(3), pp. 265–293.

R DEVELOPMENT CORE TEAM (2008). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0. http://www.R-project.org.

ROBERTS, S. W. (1959). "Control Chart Tests Based on Geometric Moving Averages". *Technometrics* 1(3), pp. 239–250.

SEIFERT, B. (2007). *lpridge: Local Polynomial (Ridge) Regression*. http://probability.ca/cran/bin/macosx/universal/contrib/2.8/lpridge_1.0-4.tgz, R package version 1.0-4.

SULLIVAN, J. H. (2002). "Detection of Multiple Change Points from Clustering Individual Observations". *Journal of Quality Technology* 34, pp. 371–383.

WILLIAMS, J. D.; BIRCH, J. B.; WOODALL, W. H.; and FERRY, N. M. (2007a). "Statistical Monitoring of Heteroscedastic Dose Response Profiles from High-Throughput Screening". *Journal of Agricultural, Biological and Environmental Statistics* 12, pp. 216–235.

WILLIAMS, J. D.; WOODALL, W. H.; and BIRCH, J. B. (2007b). "Statistical Monitoring of Nonlinear Product and Process Quality Profiles". *Quality and Reliability Engineering International* 23, pp. 925–941.

ZHANG, L. and CHEN, G. (2004). "EWMA Charts for Monitoring the Mean of Censored Weibull Lifetimes". *Journal of Quality Technology* 36(3), pp. 321–328.

ZHU, L. M.; ZHANG, Y. J.; and WANG, Z. J. (2006). "A Control Chart Based on a Change Point Model for Monitoring Linear Profiles". *IIE Transactions* 38(12), pp. 1093–1103.

ZOU, C.; TSUNG, F.; and WANG, Z. (2007). "Monitoring General Linear Profiles Using Multivariate EWMA Schemes". *Technometrics* 49, pp. 395–408.