

A Comprehensive Review of Statistical Process Control Techniques and Their Applications

Kavithanjali S¹, Gunasekaran S², Anandhi P³, Srividya M⁴

¹²³⁴Assistant Professor, Department of Mathematics, Sona College of Arts and Science, Tamil Nadu, India

Corresponding Email: kavistat7@gmail.com

Abstract - Statistical Process Control (SPC) provides the statistical foundation for monitoring, controlling, and improving the quality of manufacturing and service processes by distinguishing common cause variation, which is inherent to a process, from special-cause variation, which signals abnormal behaviour requiring intervention. Since Shewhart formulated the original control chart, the field has progressed through successive generations of techniques developed to overcome specific limitations exposed by the methods that preceded them: the insensitivity of basic charts to small and gradual shifts, the difficulty of monitoring several correlated quality characteristics simultaneously, and the restrictive distributional assumptions underlying classical procedures. This narrative review traces that progression, synthesising the literature on univariate charts (Shewhart, CUSUM, EWMA, and their subsequent refinements), adaptive and economically designed monitoring schemes, multivariate control procedures, nonparametric and distribution-free methods, and profile-monitoring approaches. Rather than presenting these contributions as an inventory of isolated studies, the review discusses them as an interconnected line of methodological reasoning in which each advance responds to a gap left by earlier work, and situates the reviewed literature within the broader critical commentary on SPC theory and practice. A brief, clearly separated illustrative example using a simulated manufacturing dataset demonstrates the mechanics of a basic X-bar control chart. The review closes by identifying persistent and emerging gaps in the literature, particularly around high-dimensional and streaming data, real-time and Industry 4.0 enabled monitoring, and the integration of machine learning with classical SPC theory, arguing that these gaps follow directly from the trajectory the field has already taken. The review is intended as a structured reference for researchers, academicians, and quality engineers seeking to understand how Statistical Process Control has developed and where its unresolved problems remain.

Key Words: Statistical Process Control, Shewhart Control Chart; CUSUM, EWMA, Multivariate SPC, Nonparametric Control Charts, Process Monitoring.

1.INTRODUCTION

Statistical Process Control is one of the principal tools by which manufacturing and service organisations distinguish ordinary, inherent process variation from variation that signals a genuine change in the underlying process. The distinction matters because it determines how an organisation should respond: common-cause variation is a property of the process as designed and does not warrant intervention, whereas special-cause variation indicates that something has changed and requires investigation. Shewhart's original control chart, developed in the 1920s, formalised this distinction using statistically derived control limits and gave quality engineers a simple, visual criterion for deciding when a process had moved out of statistical control.

For several decades this basic charting logic was sufficient for the coarse-grained monitoring problems it was designed to address, but as manufacturing processes became more tightly tolerated and quality requirements more demanding, its limitations became apparent. A Shewhart chart evaluates each sample independently of the ones before it, which makes it effective at detecting large, sudden shifts but comparatively slow to detect small, sustained shifts that accumulate gradually rather than appearing as an abrupt jump. This single limitation is the starting point for much of the subsequent methodological development in SPC: cumulative and weighted-memory charts were introduced specifically to retain sensitivity to the kind of small, persistent shift that a memoryless chart tends to miss.

As industrial processes grew more complex, a second limitation became equally pressing. Modern manufacturing and service systems rarely produce a single quality characteristic in isolation; instead, several correlated variables must be monitored together, and treating them with separate univariate charts inflates the overall false-alarm rate while ignoring the correlation structure that often carries useful diagnostic information. This motivated the development of multivariate SPC procedures capable of monitoring several characteristics jointly and, where possible, of identifying which variable was responsible for an out-of-control signal. A third limitation concerns the distributional assumptions on which classical control-chart theory rests: most traditional charts assume that the monitored characteristic is normally distributed, an assumption that does not hold for many practical quality measures, which prompted the development of nonparametric and distribution-free monitoring schemes.

More recently, the growing availability of high-frequency and high-dimensional process data, together with the broader shift toward automated and data-driven manufacturing associated with Industry 4.0, has renewed interest in adapting SPC to settings that its classical formulations were never designed for. Researchers have begun exploring the integration of SPC with machine learning and real-time analytics, but this remains an active and comparatively unsettled area rather than a mature branch of the literature. Taken together, these developments mean that the SPC literature is no longer a single coherent method but a family of related techniques, each suited to a different combination of shift size, dimensionality, distributional behaviour, and monitoring environment, and each traceable to a specific limitation identified in earlier work.

This review synthesises that literature as a continuous narrative of methodological development rather than as a sequence of disconnected study summaries. Its purpose is to explain why the field evolved in the direction it did, connecting each major class of technique to the monitoring problem that motivated it, to the statistical idea it introduced, and to the limitation that subsequent research then sought to address. In doing so, the review aims to provide researchers, academicians, and quality engineers with a structured understanding of how classical, adaptive, multivariate, and nonparametric SPC methods relate to one another, and to identify where the reviewed literature leaves genuine gaps for future research.

1.1. Objectives of the Review

This review pursues the following objectives:

1. To trace the methodological evolution of Statistical Process Control from its Shewhart based origins through subsequent univariate, adaptive, multivariate, and nonparametric developments, explaining the motivation behind each advance rather than merely listing its features.
2. To compare the major families of SPC technique Shewhart, CUSUM, EWMA, adaptive, multivariate, and nonparametric/distribution-free charts with respect to their sensitivity to shift size, distributional assumptions, computational complexity, and practical applicability.
3. To examine how SPC has been applied and extended across industrial and process settings, and how it has been positioned within broader quality-management frameworks.
4. To identify, on the basis of the reviewed literature, the research gaps and emerging directions that remain open in Statistical Process Control, and to explain how these gaps arise from the field's

developmental trajectory rather than treating them as an unconnected list of future technologies.

2. REVIEW METHODOLOGY

This paper is a qualitative narrative literature review rather than a systematic review; it does not follow a PRISMA type protocol, and no formal database search strings, screening counts, or inclusion/exclusion thresholds were applied. The literature discussed was selected on the basis of its foundational status, its influence on subsequent SPC research, and its relevance to the review's objectives, spanning classical control-chart theory, adaptive and economically designed schemes, multivariate monitoring, nonparametric methods, profile monitoring, and broader critical commentary on SPC practice.

Once selected, the studies were grouped thematically according to the type of monitoring problem they address for example, detection of small univariate shifts, monitoring of several correlated characteristics, or relaxation of distributional assumptions rather than being treated as a chronological or alphabetical list. Within each theme, the literature is synthesised by identifying the limitation of earlier methods that motivated a given contribution, the statistical mechanism it introduced, and the way later work built on or qualified it. This thematic, developmental synthesis is what allows the review to be organised as a continuous narrative rather than as a set of independent study summaries.

A short illustrative example using a simulated manufacturing dataset and a basic X-bar control chart is included separately from the literature synthesis, in the section that follows the traditional control-chart discussion. This example is intended only to demonstrate the mechanics of control-chart construction; it is not drawn from, and should not be read as, evidence from the reviewed literature.

3. EVOLUTION AND FOUNDATIONS OF STATISTICAL PROCESS CONTROL

The origins of Statistical Process Control lie in Shewhart's work at Bell Telephone Laboratories, where he proposed using statistically derived limits, typically set at three standard deviations from a process centre line, to distinguish routine process variation from variation requiring corrective action (Shewhart, 1924). The essential contribution of this work was conceptual as much as technical: it separated the notion of a process being “in control,” in the sense of exhibiting only common-cause variation, from the notion of a process meeting its specifications, and it gave quality engineers a statistically grounded, visual tool the control chart for making that judgement in real time. This distinction between common-cause and special-cause variation remains the conceptual foundation on which essentially all subsequent SPC methodology is built, and every technique discussed in

this review can be understood as a different statistical strategy for making that same distinction more sensitively, more efficiently, or under weaker assumptions than the original Shewhart chart allowed.

4. TRADITIONAL UNIVARIATE CONTROL CHARTS AND THEIR REFINEMENT

The most direct limitation of the Shewhart chart is also the most consequential for later research: because each plotted point is evaluated using only the current sample, the chart has no memory of the recent pattern of observations, and it is correspondingly slow to detect shifts that are small relative to the natural process variation, even when those shifts persist over many samples. Page (1954) addressed this limitation directly by proposing the Cumulative Sum (CUSUM) control chart, which plots the running sum of deviations of each observation from a target value rather than the observation itself. Because this cumulative statistic carries forward information from every prior sample, small but sustained shifts accumulate visibly over time even though no single sample would trigger a Shewhart-type signal, giving CUSUM charts materially better sensitivity to small process shifts than memoryless alternatives.

A parallel solution to the same underlying problem was proposed by Roberts (1959), whose Exponentially Weighted Moving Average (EWMA) chart plots a weighted average of current and past observations, with the weight assigned to older observations decaying geometrically. Where CUSUM accumulates deviations without discounting them, EWMA achieves a similar memory effect while allowing the analyst to tune, through a single smoothing constant, how strongly the chart favours recent observations over historical ones; a smoothing constant close to one recovers behaviour close to a Shewhart chart, while smaller values yield behaviour closer to CUSUM. This flexibility is part of why EWMA and CUSUM have developed as parallel rather than competing lines of research: both solve the same memory problem, but through different statistical mechanisms with different tuning behaviour.

Once these two memory-based charts existed, subsequent research turned to refining their operating characteristics rather than replacing the underlying idea. Lucas and Crosier (1982) observed that a standard CUSUM chart starts from a “clean” state and is therefore relatively slow to detect a shift that is already present when monitoring begins; their Fast Initial Response (FIR) modification adds a “head start” to the cumulative statistic so that an already-shifted process is flagged sooner, without materially increasing the false-alarm rate once the process is genuinely in control. On the EWMA side, Crowder (1987) and later Knoth (2005) developed more accurate computational methods for evaluating the Average Run Length (ARL) the expected number of samples before a chart signals which is the

standard criterion used to compare the sensitivity and false-alarm behaviour of competing charts; these contributions did not change the EWMA statistic itself but made it possible to select smoothing constants and control limits with much greater precision. Lucas and Saccucci (1990) extended this line of work into a comprehensive study of how the smoothing constant and control-limit width jointly determine EWMA performance, and their recommendations remain a standard reference for practitioners designing an EWMA scheme. Champ and Woodall (1987) pursued a different route to the same general goal of improving sensitivity without abandoning the simplicity of the Shewhart chart, showing that supplementary run rules signalling out-of-control conditions based on patterns of several consecutive points rather than a single extreme point measurably improve sensitivity to moderate shifts while preserving the chart's basic Shewhart structure.

Taken together, this body of work illustrates a recurring pattern in the SPC literature: an initial technique addresses a specific weakness of its predecessor, and a subsequent wave of research then refines that technique's operating characteristics its starting behaviour, its run-length properties, its parameter selection without displacing the underlying statistical idea. CUSUM and EWMA remain, even after this refinement, univariate procedures intended for a single quality characteristic under an approximately normal distribution, and both limitations become the starting point for the developments discussed in the following sections.

5. ADAPTIVE TECHNIQUES AND ECONOMICALLY DESIGNED MONITORING SCHEMES

A separate strand of research addressed not the statistic being plotted but the way sampling itself is conducted. Duncan (1956) was among the first to frame control chart design as an economic optimisation problem, showing that sample size, sampling interval, and control-limit width could be chosen jointly to minimise the combined cost of inspection, production, and undetected defects rather than being set by statistical convention alone; Montgomery (1980) later consolidated and extended this economic-design literature, reinforcing that the choice of monitoring parameters involves an explicit trade-off between inspection cost and the cost of failing to detect a shift promptly.

A related but distinct idea is that the sampling scheme itself need not be fixed: rather than inspecting at a constant interval regardless of process behaviour, Reynolds et al. (1988) proposed adjusting the sampling interval according to the most recent evidence about process state, sampling more frequently when a chart statistic suggests the process may be drifting and less frequently when it appears stable. Reynolds (1996) subsequently compared Shewhart and EWMA charts operated under this kind of variable-sampling-interval scheme and found that adaptive sampling improved detection speed for both, with EWMA performing particularly

well for small shifts a result consistent with the memory-based sensitivity advantage already established for EWMA in the preceding section. Capizzi and Masarotto (2003) generalised this adaptive philosophy further, developing charts whose monitoring parameters adjust automatically in response to observed process behaviour rather than being fixed at the design stage or switched between a small number of preset states, and reported that this automatic adjustment improved sensitivity to small shifts while keeping false-alarm rates controlled.

Other researchers extended the adaptive and efficiency-oriented agenda to different monitoring targets. Khoo (2004) applied a moving-average charting logic, conceptually related to EWMA's weighted-memory idea, to the specific problem of monitoring the fraction of nonconforming units, and found it more sensitive than a standard Shewhart chart for small and moderate shifts in nonconformance rate. Reynolds and Arnold (2001) redirected the CUSUM logic itself toward a different monitoring objective, using cumulative-sum statistics to detect changes in process dispersion rather than process mean, which matters because a shift in variability can occur without any corresponding shift in the mean and would otherwise go undetected by mean-monitoring charts. Reynolds and Stoumbos (2004) examined how inspection effort should be allocated across a monitoring scheme to maximise detection performance for a given sampling budget, extending Duncan's original economic-design question to more general resource-allocation settings, while Mahmoud et al. (2010) addressed a more foundational input to all of these methods, showing that the accuracy with which process standard deviation is estimated has a material effect on the performance of the resulting control limits regardless of which charting technique is used.

Collectively, this literature shows that improving SPC performance is not solely a matter of choosing a more sensitive statistic; the sampling scheme, the allocation of inspection resources, and even the accuracy of preliminary parameter estimates all affect how quickly and reliably a monitoring scheme detects a genuine process change, and each of these dimensions has generated its own line of methodological refinement alongside the CUSUM/EWMA development discussed earlier.

6. MULTIVARIATE STATISTICAL PROCESS CONTROL

The refinements discussed so far all operate on a single quality characteristic, but industrial processes typically generate several correlated measurements simultaneously for instance, multiple dimensions of a machined part, or several chemical properties of a batch process. Monitoring each characteristic with a separate univariate chart is statistically inefficient for two reasons: the overall probability of a false

alarm across many simultaneous charts is much higher than the false-alarm rate of any single chart, and separate charts cannot exploit the correlation between characteristics, which often carries useful information about the process. This is the problem that motivated multivariate SPC, and it is a direct extension of the univariate developments already discussed rather than an unrelated branch of the literature.

Crosier (1988) extended the CUSUM idea to the multivariate setting with the Multivariate CUSUM (MCUSUM) chart, which accumulates a vector-valued deviation statistic and was shown to detect multivariate shifts more effectively than monitoring each characteristic separately, directly mirroring the sensitivity advantage that univariate CUSUM had already demonstrated over Shewhart charts. Lowry et al. (1992) performed an analogous extension of EWMA, proposing the Multivariate EWMA (MEWMA) chart, and their simulation studies showed performance advantages over earlier multivariate schemes, again paralleling the univariate EWMA story. Because both extensions inherit the tuning behaviour of their univariate counterparts, MCUSUM and MEWMA together form the multivariate analogue of the CUSUM/EWMA relationship discussed in the traditional-chart section.

A separate strand of multivariate research addressed interpretation rather than detection sensitivity. Hawkins (1991) proposed monitoring regression-adjusted variables residuals from regressing each characteristic on the others so that a signal reflects a genuine change in the relationship between correlated variables rather than being driven by a shift in one variable that is merely correlated with the others. Mason et al. (1997) addressed a closely related interpretive problem for the classical multivariate statistic, Hotelling's T^2 , by developing a decomposition method that attributes an out-of-control signal to the specific variable or variables responsible for it; this mattered because the aggregate T^2 statistic, while effective at detecting that something had changed, gave quality engineers no direct way to identify what had changed, limiting its practical diagnostic value. Zou et al. (2011) later introduced a likelihood ratio-based approach to multivariate monitoring that improved sensitivity to small multivariate shifts relative to earlier likelihood-free schemes, extending the small-shift-detection agenda that had originally motivated CUSUM and EWMA into the multivariate domain. Bersimis et al. (2007) subsequently reviewed and compared Hotelling's T^2 , MEWMA, MCUSUM, and related multivariate procedures, and their comparison remains a widely cited reference point for situating individual multivariate techniques within the broader family.

The multivariate literature therefore reproduces, in a correlated-variable setting, the same developmental pattern already seen in the univariate literature: an initial extension establishes multivariate analogues of CUSUM and EWMA to regain small-shift sensitivity, and subsequent work then addresses the practical limitations that arise specifically in the

multivariate case principally, the difficulty of diagnosing which variable is responsible for a signal once one has been raised.

7. NONPARAMETRIC AND DISTRIBUTION-FREE APPROACHES TO SPC

A further limitation shared by Shewhart, CUSUM, EWMA, and their multivariate extensions alike is that their control limits are typically derived under an assumption of normality. Many practical quality characteristics count data, skewed measurements, or data from processes with unknown or unstable distributions do not satisfy this assumption, and control limits calculated as though they did can produce false-alarm rates very different from the nominal rate the chart was designed to achieve. This distributional concern motivated a line of nonparametric or distribution-free SPC research that runs alongside, rather than after, the multivariate developments discussed above.

Chakraborti et al. (2004) proposed a class of distribution free control charts based on statistics whose in-control distribution does not depend on the underlying process distribution, providing reliable monitoring performance for non-normal data without requiring the process distribution to be known or estimated. Qiu (2008) extended the nonparametric agenda into the multivariate setting using data depth, a concept that generalises the notion of an order statistic to multiple dimensions without assuming multivariate normality, and demonstrated reliable detection performance across a range of underlying distributions. Li et al. (2014) developed this direction further with a nonparametric multivariate procedure explicitly designed for correlated variables without a specific distributional assumption, reporting favourable performance for non-normal process data relative to methods that presuppose normality. Read together, this literature shows the field addressing its distributional assumption in essentially the same order that it addressed shift sensitivity and dimensionality: first in the univariate case, and then, once the multivariate machinery existed, in the multivariate case as well.

8. PROFILE MONITORING AND COMPLEX PROCESS DATA

A further extension of the multivariate agenda arises when the quality characteristic of interest is not a single number, or even a fixed vector of numbers, but an entire functional relationship for example, a calibration curve or a dose-response profile that must remain stable over time. Noorossana et al. (2008) surveyed statistical methods developed for monitoring such process profiles, arguing that profile monitoring provides a more complete picture of process behaviour than reducing a functional relationship to one or two summary statistics before applying a conventional chart. Zou et al. (2015) proposed a multivariate EWMA-based

approach specifically adapted to profile data, extending the same weighted-memory logic already established for scalar and vector characteristics to functional data, and reported faster detection of profile-level changes than earlier profile-monitoring techniques. Profile monitoring therefore represents the point at which the multivariate line of development, discussed above, meets increasingly complex and high-dimensional data structures the same territory that motivates much of the research-gaps discussion later in this review.

9. ILLUSTRATIVE EXAMPLE: X-BAR CONTROL CHART

The following example is included solely to illustrate the mechanics of a basic univariate Shewhart-type control chart, as described in Section 5 above; it is a computational demonstration using a simulated dataset and should not be read as evidence drawn from the reviewed literature. A sample manufacturing dataset representing the diameter measurements (in millimetres) of twenty machined components was used to construct an X-bar control chart in Python:

data = [49.8, 50.2, 49.9, 50.1, 50.3, 49.7, 50.0, 49.9, 50.2, 50.1, 49.8, 50.0, 50.4, 49.9, 50.1, 49.8, 50.0, 50.2, 49.9, 50.5]

The process mean, sample standard deviation, and the resulting control limits were calculated using the standard formulas below:

$$\bar{X} = \frac{\sum_{i=1}^n X_i}{n} \quad \text{for } i = 1 \text{ to } n$$

$$s = \sqrt{\frac{\sum_{i=1}^n (X_i - \bar{X})^2}{n-1}} \quad \text{for } i = 1 \text{ to } n$$

$$UCL = \bar{X} + 3s \quad LCL = \bar{X} - 3s \quad CL = \bar{X}$$




Fig-1: X-bar control chart

As shown in Figure 1, all twenty sample measurements fall within the Upper Control Limit (UCL) and Lower Control Limit (LCL), indicating that the simulated process remained under statistical control over the sampling period, with no observation signalling the presence of an assignable cause. This example is deliberately simple; it demonstrates only the basic Shewhart-chart logic described in Section 5 and does not illustrate the CUSUM, EWMA, multivariate, or nonparametric procedures discussed elsewhere in this review.

10. INDUSTRIAL APPLICATIONS AND INTEGRATION WITH QUALITY MANAGEMENT

Beyond the methodological developments traced above, a parallel body of literature has examined how SPC is applied in practice and how it relates to broader organisational quality-management efforts. Kolesar (1993) examined the relationship between SPC and Total Quality Management, arguing that control charts function most effectively not as a standalone statistical tool but as part of a wider organisational commitment to continuous improvement, which connects the technical developments discussed above to the managerial context in which they are actually deployed. Shaw (1991) extended the discussion to continuous-process industries such as chemical and petroleum manufacturing, showing that the same charting logic developed originally for discrete-parts manufacturing could be adapted to reduce variability in continuous production, broadening SPC's applicability well beyond its original discrete-manufacturing setting. Porteus and Angelus (1997) connected SPC more explicitly to economic decision-making, proposing ways to integrate inspection-policy optimisation with quality monitoring; this line of argument echoes, in an applied setting, the economic-design agenda already discussed in Section 6, reinforcing that the question of how to monitor a process economically has been a recurring concern across both the theoretical and applied SPC literature.

11. BROADER CRITICAL PERSPECTIVES ON STATISTICAL PROCESS CONTROL

Alongside the technical literature that develops new charts, a smaller but influential body of work has stepped back to assess the field as a whole. Woodall (1985) examined the statistical design choices underlying control charts more generally, emphasising that the selection of control limits and sampling strategy materially affects false-alarm behaviour and chart performance, a concern that runs through nearly every methodological contribution discussed in this review. Woodall and Montgomery (1999) surveyed the state of SPC research at the end of the twentieth century and identified opportunities for closer integration between SPC and modern manufacturing systems, several of which multivariate monitoring, profile monitoring, and economically efficient sampling had already

begun to develop by the time of their review and are discussed in the sections above. Woodall (2000) subsequently examined a set of persistent controversies and misconceptions in SPC practice, cautioning that control-chart interpretation and process-capability analysis are sometimes applied without adequate attention to the statistical assumptions on which they rest, a caution that applies directly to the normality assumption discussed in Section 8. Psarakis and Papaleonida (2007) provided a further consolidating review of SPC developments and applications, reinforcing the increasing role of SPC in manufacturing quality improvement and pointing toward several of the same future directions discussed later in this review. Together, these broader reviews function less as sources of new charting techniques and more as periodic checkpoints on the field, and they are drawn on throughout this review to contextualise the more technical contributions discussed above.

12. COMPARATIVE CRITICAL DISCUSSION OF SPC TECHNIQUES

The literature synthesised above supports a structured comparison across the major families of SPC technique, organised around the practical dimensions that determine which method is appropriate for a given monitoring problem.

On sensitivity to shift size, Shewhart charts remain best suited to detecting large, sudden shifts, while CUSUM and EWMA charts, and their multivariate analogues, are specifically designed for small to moderate, sustained shifts that a memoryless chart would be slow to detect; supplementary run rules narrow but do not eliminate this gap for Shewhart charts. On distributional assumptions, Shewhart, CUSUM, EWMA, MCUSUM, and MEWMA all rely, in their standard forms, on an assumption of normality (or multivariate normality) that nonparametric and distribution-free methods, such as those of Chakraborti et al. (2004), Qiu (2008), and Li et al. (2014), are specifically designed to relax, at some cost in the simplicity of the underlying theory. On dimensionality, univariate charts are appropriate when a single characteristic is of interest, but multivariate methods become necessary once several correlated characteristics must be monitored jointly, both to control the overall false-alarm rate and to exploit correlation structure that univariate charts discard; the trade-off is that multivariate charts are harder to interpret once a signal occurs, which is precisely the problem that decomposition methods such as that of Mason et al. (1997) were developed to address.

On computational complexity, Shewhart charts are the simplest to implement and interpret, EWMA and CUSUM require only modest additional computation and a small number of tuning parameters, while multivariate, nonparametric, and profile-monitoring methods are progressively more demanding, both computationally and in

terms of the statistical expertise needed to implement and interpret them correctly. On robustness, distribution-free and data-depth-based methods are, by construction, more robust to departures from normality than their parametric counterparts, though this robustness typically comes at some cost in efficiency when the normality assumption does in fact hold. On practical and industrial applicability, the literature reviewed above suggests that simpler univariate charts remain the default choice in many industrial settings precisely because of their interpretability and low computational burden, while more advanced multivariate, adaptive, and nonparametric methods tend to be adopted selectively, in settings where their added complexity is justified by a genuine need for higher sensitivity, joint monitoring of correlated characteristics, or robustness to non-normal data.

No single technique dominates across all of these dimensions simultaneously, and the reviewed literature does not report a consistent numerical ranking across studies that would justify a more quantitative comparison here; the appropriate choice of method depends on the specific combination of shift size, dimensionality, distributional behaviour, and computational constraints present in a given monitoring problem.

13. RESEARCH GAPS AND EMERGING DIRECTIONS

The gaps identified in this section follow directly from the developmental trajectory traced above rather than constituting an unrelated list of future technologies. Much of the multivariate and profile-monitoring literature discussed above including the work of Lowry et al. (1992), Zou et al. (2011), Noorossana et al. (2008), and Zou et al. (2015) was developed for settings with a modest, fixed number of correlated variables or a single functional profile; as manufacturing systems increasingly generate high-dimensional and streaming sensor data, the computational and statistical assumptions underlying these methods are placed under considerably greater strain, and the reviewed literature does not yet offer a settled answer for monitoring very high-dimensional or continuously streaming quality characteristics.

Similarly, the economic-design and adaptive-sampling literature discussed above Duncan (1956), Reynolds et al. (1988), Reynolds (1996), and Capizzi and Masarotto (2003) was developed for discrete, periodic sampling schemes rather than for the continuous, real-time data streams increasingly available under Industry 4.0-style automated manufacturing; adapting the economic logic of this literature to a setting where inspection is effectively continuous, rather than scheduled, is a natural but not yet fully resolved extension of that earlier work.

Finally, the broader reviews discussed above Woodall and Montgomery (1999), Woodall (2000), and Psarakis and Papaleonida (2007) anticipated, in general terms, a future in which SPC would need to integrate with more

sophisticated data-analytic techniques, and recent interest in combining SPC with machine learning reflects this anticipated direction. However, the reviewed literature in this area remains comparatively exploratory: it is not yet clear how classical performance criteria such as Average Run Length, which underpin nearly every comparison discussed in this review, should be adapted for monitoring schemes built on machine-learning models rather than explicit statistical distributions, and this represents an open methodological question rather than a solved problem.

14. FUTURE RESEARCH DIRECTIONS

On the basis of the gaps identified above, several directions for future research follow naturally from the reviewed literature. First, extending multivariate and profile-monitoring methods to genuinely high-dimensional and streaming settings, building on but substantially adapting the work of Lowry et al. (1992), Zou et al. (2011), and Zou et al. (2015), is a natural continuation of the multivariate line of development traced in this review. Second, adapting the economic-design and adaptive-sampling logic established by Duncan (1956) and later authors to continuous, real-time monitoring environments would extend that literature into settings more representative of automated modern manufacturing. Third, developing performance criteria and theoretical guarantees for SPC schemes that incorporate machine-learning components, comparable in rigour to the Average Run Length framework that underlies classical chart comparison, would help resolve the open methodological question identified in the previous section. Finally, further nonparametric and distribution-free developments, extending the work of Chakraborti et al. (2004), Qiu (2008), and Li et al. (2014), remain relevant as quality data from increasingly diverse and automated sources continues to depart from the normality assumptions on which much classical SPC theory still rests.

15. CONCLUSIONS

The literature reviewed in this paper shows Statistical Process Control developing not as a single fixed method but as a continuing sequence of responses to specific limitations exposed by earlier work. Shewhart's original chart established the conceptual distinction between common-cause and special-cause variation but was slow to detect small, sustained shifts; CUSUM and EWMA were developed specifically to address that limitation, and a subsequent wave of research refined their initial response, run-length behaviour, and parameter selection. As monitoring problems grew to involve multiple correlated characteristics, the same memory-based logic was extended into the multivariate domain, generating its own diagnostic challenge identifying which variable was responsible for a signal which later decomposition and regression-adjusted methods addressed. In parallel, the normality assumption underlying nearly all of these methods motivated a distinct nonparametric line of development,

extended in turn to the multivariate case, while economic-design and adaptive-sampling research addressed the practical question of how and when to sample rather than what statistic to plot.

Across this trajectory, traditional univariate charts remain the appropriate choice where interpretability and simplicity are paramount and shifts are expected to be moderate or large; memory-based univariate charts are preferable when small, sustained shifts matter; multivariate methods become necessary once several correlated characteristics must be jointly monitored; and nonparametric methods are warranted when normality cannot reasonably be assumed. No technique in the reviewed literature dominates across all of these considerations simultaneously, which is itself the review's central finding: the appropriate SPC method is a function of the specific monitoring problem, not a single universally superior chart.

The gaps that remain high-dimensional and streaming data, real-time and Industry 4.0-enabled monitoring, and the integration of machine learning with classical SPC theory are not a disconnected list of fashionable topics but a direct continuation of the same developmental logic traced throughout this review: each responds to a limitation of the methods that currently exist, in much the same way that CUSUM responded to the limitations of the Shewhart chart nearly a century ago. Future research that extends the reviewed literature along these lines, while preserving the statistical rigour that has characterised the field since Shewhart's original work, is likely to determine how effectively Statistical Process Control adapts to the increasingly complex and automated production environments it is now being asked to monitor.

REFERENCES

1. Bersimis, S., Psarakis, S., & Panaretos, J. (2007). Multivariate statistical process control charts: An overview. *Quality and Reliability Engineering International*, 23(5), 517–543.
2. Capizzi, G., & Masarotto, G. (2003). Adaptive control charts for monitoring processes. *Journal of Quality Technology*, 35(3), 285–298.
3. Chakraborti, S., van der Laan, P., & Van de Wiel, M. A. (2004). A class of distribution-free control charts. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 53(3), 443–462.
4. Champ, C. W., & Woodall, W. H. (1987). Exact results for Shewhart control charts with supplementary runs rules. *Technometrics*, 29(4), 393–399.
5. Crosier, R. B. (1988). Multivariate generalizations of cumulative sum quality-control schemes. *Technometrics*, 30(3), 291–303.
6. Crowder, S. V. (1987). A simple method for studying run-length distributions of exponentially weighted moving average charts. *Technometrics*, 29(4), 401–407.
7. Duncan, A. J. (1956). The economic design of $\bar{X}$ control charts used to maintain current control of a process. *Journal of the American Statistical Association*, 51(274), 228–242.
8. Hawkins, D. M. (1991). Multivariate quality control based on regression-adjusted variables. *Technometrics*, 33(1), 61–75.
9. Khoo, M. B. C. (2004). Moving average control charts for monitoring the fraction nonconforming. *Quality Engineering*, 16(3), 399–408.
10. Knoth, S. (2005). Accurate ARL computation for EWMA control charts. *Journal of Quality Technology*, 37(1), 78–89.
11. Kolesar, P. J. (1993). The relevance of research on statistical process control to the total quality movement. *Journal of Engineering and Technology Management*, 10(4), 317–338.
12. Li, Z., Qiu, P., & Hawkins, D. M. (2014). Nonparametric multivariate statistical process control. *Technometrics*, 56(3), 344–354.
13. Lowry, C. A., Woodall, W. H., Champ, C. W., & Rigdon, S. E. (1992). A multivariate exponentially weighted moving average control chart. *Technometrics*, 34(1), 46–53.
14. Lucas, J. M., & Crosier, R. B. (1982). Fast initial response for CUSUM quality-control schemes. *Technometrics*, 24(3), 199–205.
15. Lucas, J. M., & Saccucci, M. S. (1990). Exponentially weighted moving average control schemes: Properties and enhancements. *Technometrics*, 32(1), 1–12.
16. Mahmoud, M. A., Henderson, G. R., Epprecht, E. K., & Woodall, W. H. (2010). Estimating the standard deviation in quality-control applications. *Journal of Quality Technology*, 42(4), 348–357.
17. Mason, R. L., Tracy, N. D., & Young, J. C. (1997). Decomposition of T^2 for multivariate control chart interpretation. *Journal of Quality Technology*, 29(2), 109–122.
18. Montgomery, D. C. (1980). The economic design of control charts: A review and literature survey. *Journal of Quality Technology*, 12(2), 75–87.
19. Noorossana, R., Saghaei, A., & Amiri, A. (2008). Statistical analysis of profile monitoring. *Computers & Industrial Engineering*, 54(3), 630–641.
20. Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, 41(1–2), 100–115.
21. Porteus, E. L., & Angelus, A. (1997). Opportunities for improved statistical process control. *Management Science*, 43(9), 1214–1228.

22. Psarakis, S., & Papaleonida, G. (2007). SPC procedures and recent applications. *Quality Technology & Quantitative Management*, 4(2), 181–204.
23. Qiu, P. (2008). Distribution-free multivariate process control based on data depth. *Technometrics*, 50(1), 72–81.
24. Reynolds, M. R. (1996). Shewhart and EWMA control charts with variable sampling intervals. *Sequential Analysis*, 15(4), 387–407.
25. Reynolds, M. R., Amin, R. W., Arnold, J. C., & Nachlas, J. A. (1988). $\bar{X}$ charts with variable sampling intervals. *Technometrics*, 30(2), 181–192.
26. Reynolds, M. R., & Arnold, J. C. (2001). CUSUM charts for monitoring process dispersion. *Journal of Quality Technology*, 33(3), 315–326.
27. Reynolds, M. R., & Stoumbos, Z. G. (2004). Control charts and the efficient allocation of sampling resources. *Technometrics*, 46(2), 200–214.
28. Roberts, S. W. (1959). Control chart tests based on geometric moving averages. *Technometrics*, 1(3), 239–250.
29. Shaw, J. A. (1991). Statistical process control for the process industries. *ISA Transactions*, 30(1), 99–106.
30. Shewhart, W. A. (1924). Some applications of statistical methods to the analysis of physical and engineering data. *Bell System Technical Journal*, 3(1), 43–87.
31. Woodall, W. H. (1985). The statistical design of quality control charts. *The Statistician*, 34(2), 155–160.
32. Woodall, W. H. (2000). Controversies and contradictions in statistical process control. *Journal of Quality Technology*, 32(4), 341–350.
33. Woodall, W. H., & Montgomery, D. C. (1999). Research issues and ideas in statistical process control. *Journal of Quality Technology*, 31(4), 376–386.
34. Zou, C., Tsung, F., & Wang, Z. (2015). Monitoring general profiles using multivariate EWMA charts. *Technometrics*, 57(2), 232–244.
35. Zou, C., Wang, Z., & Tsung, F. (2011). Monitoring multivariate processes using likelihood ratio statistics. *Journal of Quality Technology*, 43(4), 318–330.