Practical Considerations for Implementing the Isolation Forest EWMA Control Chart in Phase II Process Monitoring

T. Perdikis[1],G. Celano[2],S.Chakraborti[3]

[1]University of Piraeus, Greece
[2]Università di Catania, Catania, Italy
[3]University of Alabama, Tuscaloosa, AL, USA

General Introduction

Statistical process control

Statistical process control (SPC), is a valuable process, using statistical methods for monitoring and controlling the process performance.

Control Charts

On-line process monitoring technique whose purpose is to quickly detect shifts in the process.

Applications

SPC techniques are being widely used not only in industrial and manufacturing fields but also in Time series monitoring (see, Sven and Schmid (2008)).

Phase I and Phase II Implementation

Definitions

$$\underbrace{\underbrace{\Big\{\quad \overset{\text{First sample of size } n}{\overset{\downarrow}{\ldots}} \quad\Big\}}_{\text{1-st inspection}}, \underbrace{\Big\{\quad \overset{\text{Second sample of size } n}{\overset{\downarrow}{\ldots}} \quad\Big\}}_{\text{2-nd inspection}}, \ldots, \underbrace{\Big\{\quad \overset{\text{i-th sample of size } n}{\overset{\downarrow}{\ldots}} \quad\Big\}}_{\text{i-th inspection}}, \ldots}_{\text{Production Horizon}}$$

Phase I

- Properly define the chart's design parameters which correspond to the in-control case.

Phase II

- At each sampling point, a subgroup of size n is collected and the charting statistic is computed and plotted.

Phase II implementation

Example: The $\bar{X}$ chart

Design parameters/assumptions

- Normally distributed samples.
- μ_0, σ_0 known.
- $\mathrm{UCL} = \mu_0 + z_{\alpha/2}\frac{\sigma_0}{\sqrt{n}}$.
- $\mathrm{UCL} = \mu_0 - z_{\alpha/2}\frac{\sigma_0}{\sqrt{n}}$.

Plotting statistic

- Compute $\bar{X}_t$.
- Check if $\bar{X}_t \in [\mathrm{UCL}, \mathrm{LCL}]$.

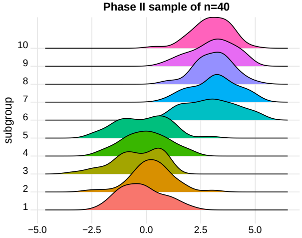

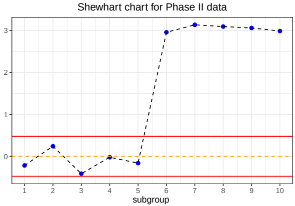

Motivation Example:Monitoring Credit Card Transactions

- We consider the implementation of an online detection method to monitor credit card transactions aimed at detecting potentially fraudulent situations.
- The dataset contains transactions made by credit cards in September 2013 by European cardholders.
- It contains 29 numerical input variables but due to confidentiality issues, the original features names are not provided.
- The dataset comprises N = 10.200 transactions.
- Comprehensive details regarding the data's origin are documented in $https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud$.

Isolation Forest

Motivation

- An unsupervised anomaly detection method proposed by Liu et al. (2008).
- Isolates anomalies by means of an ensemble of decision trees (called ITrees). Iteratively partitions the dataset until all instances are isolated.
- An anomaly, is expected to be rapidly isolated by the tree after a shorter path than a normal instance.
- $h(x)$:path length within an iTree, $c(n)$: normalisation constant.

Isolation Tree

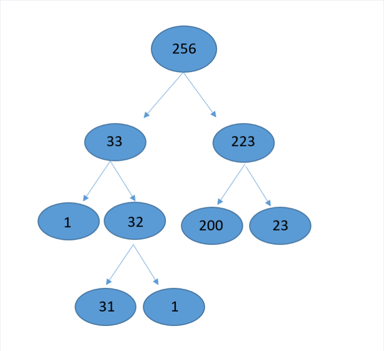

anomaly Scores

$$s(x,n) = 2^{-\frac{E(h(x))}{c(n)}}$$

Isolation Forest

Decision Rule

- Values of the anomaly scores close to 1 indicate a possibly anomalous observation detected by the IF method.

Rule of thumb: Liu et al. (2008)

- An anomaly score larger than 0.5 (a *threshold value*) indicates an anomaly (and classifies the observation as anomalous).
- An anomaly score less than 0.5, there is a strong indication that the corresponding instance (observation) is "normal".
- If all the scores are approximately close to 0.5, then the entire training sample is thought to possess no distinct anomalies.

Isolation Forest

Illustrative Example #1

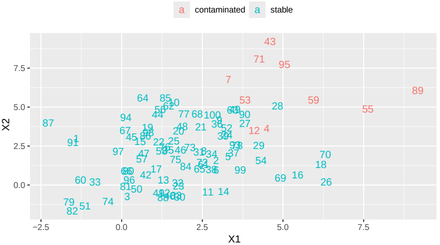

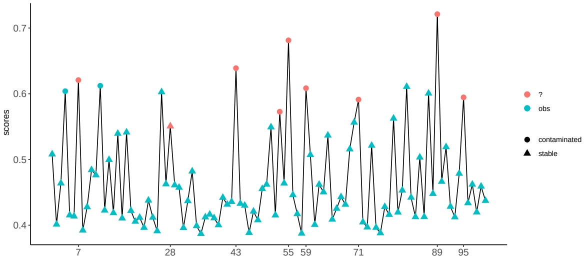

The II IF EWMA control chart

Wang and Liu (2023)

- A stable reference (IC) sample of M multivariate observations $\mathbf{X} = \{\mathbf{X}_1, \mathbf{X}_2, \ldots, \mathbf{X}_M\}$, is assumed to be available.

Phase II implementation

- Set the tuning parameters of the IF model: t (number of ITrees) and S (sub-sample size). Then, train the IF model based on the reference sample $\mathbf{X}$.
- Calculate UCL by means of bootstrapping to meet $\mathrm{MRL} = \mathrm{MRL}_0$.
- Set $\lambda = 0.05$; $Z_0 = 0$ $(Z_0 = E(s(\mathbf{X}|M)))$.
- For a new test observation $\mathbf{Y}_j$ compute $s(\mathbf{Y}_j, \mathbf{X})$, from the trained IF model. The EWMA statistic Z_j is:

$$Z_j = \lambda \times s(\mathbf{Y}_j, \mathbf{X}) + (1 - \lambda) Z_{j-1}.$$

Determination of the UCL

- **Step 1.** Fix $\mathrm{MRL}_0 = 500$, M, $B = 5000$.
- **Step 2.** Set $\lambda = 0.05$ and an initial value for $UCL = L_0$.
- **Step 3.** Set $Z_0 = \frac{1}{M}\sum_{i=1}^{M} s(\mathbf{X_i}|\mathbf{X})$.
- **Step 5.** For $b = 1, \ldots, B$:
  1. Generate $n = 5000$ test observations $\mathbf{Y_j}$, by sampling with replacement from $\mathbf{X}$.
  2. For each $\mathbf{Y_j}$, calculate the score $s(\mathbf{Y_j}|\mathbf{X})$ and compute the corresponding value of the EWMA statistic Z_j.
- **Step 6.** Set the value of run length $\mathrm{RL}(b)$, for $b = 1, \ldots B$, to be $\min(n, (j : Z_j > UCL))$.
- **Step 7.** Compute MRL as the 50^{th} percentile of $\mathrm{RL}(b)$, for $b = 1, \ldots B$
- **Step 8.** If $|\mathrm{MRL} - 500)| < \zeta$, with $\zeta = 2.0$, then the control limit is set to be equal to $\mathrm{UCL} = L_0$ and STOP.

Motivation Example: Phase II implementation

- A reference sample $\mathbf{X}$ of size $M = 10.000$ is considered.
- The UCL is determined targeting an $\mathrm{MRL}_0 = 500$, $\mathrm{UCL} = 0.4695$.
- We display the last $r = 300$ observations from the reference sample $\boldsymbol{X}$ alongside the $n = 200$ test observations Y_i.
- The control chart signals potentially fraudulent activity for observations numbered as 10011 to 10022 and 10024.

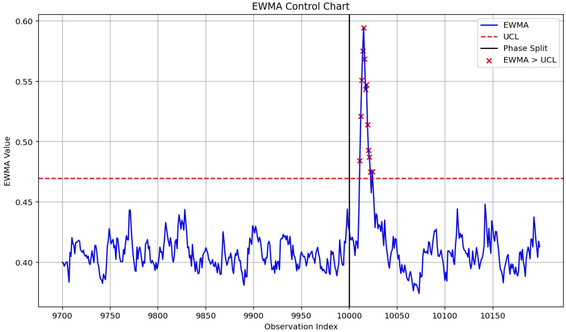

Motivation Example: SHAP values

Extending SHAP explanation approach to Multivariate SPC at the fault identification stage, after a signal is triggered by the control chart, the output $f(\mathbf{Y_i^*})$ of the ML prediction model for an OOC observation $\mathbf{Y_i^*}$ can be explained by:

$$f(\mathbf{Y_i^*}) = \phi_0 + \sum_{r=1}^{P} \phi_{r,i}$$

where:

- $f(\mathbf{Y_i^*})$ is the model prediction;
- $\phi_0 = \frac{1}{M} \times \sum_{m=1}^{M} f(\mathbf{X_m}) = E[f(\mathbf{X})]$;
- $\phi_{r,i}$ is the SHAP value for feature r in OOC observation $f(\mathbf{Y_i^*})$;
- P is the *number* of features.

Motivation Example:Fault identification

- Once the IF EWMA control chart triggers an OOC signal, the SHAP framework is applied for Post-Signal Diagnostics.
- This diagnostic process focuses on the L = 13 transaction records classified as OOC.

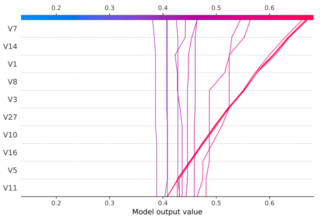

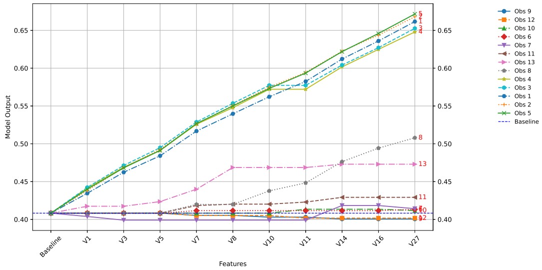

Motivation Example:Fault identification

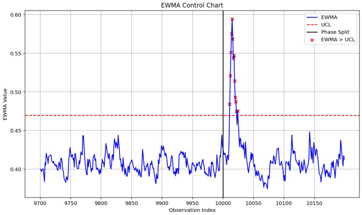

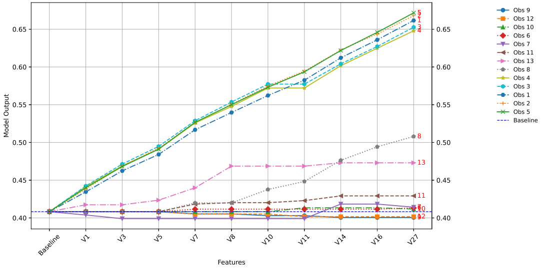

Motivation Example:Fault identification

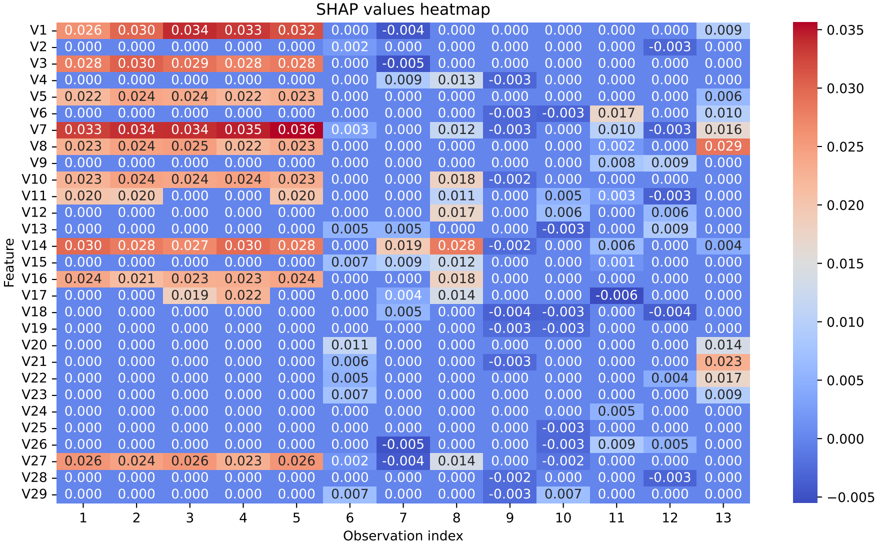

Motivation Example:Fault identification

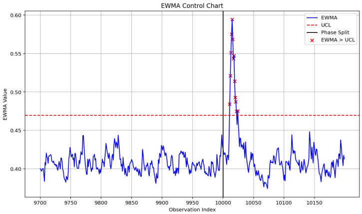

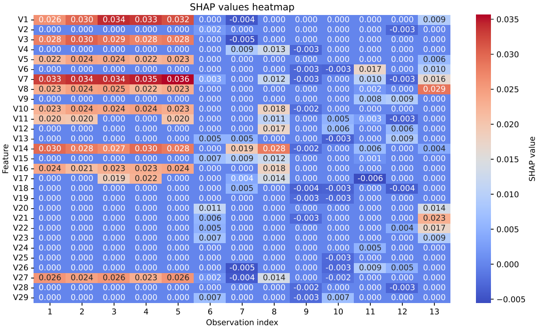

Motivation Example:Phase I implementation

- The fault detection study is completed by investigating the run charts of the most influencing features.
- The run charts for $V7$ and $V14$ reveal a sudden and simultaneous negative shift (spike) for observations suspected to be fraudulent.

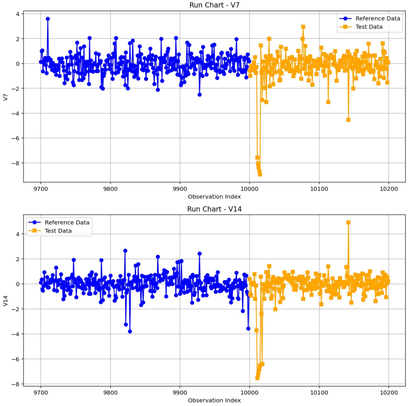

SHAP Values Approximation

For each $\mathbf{Y}_i^*$, the set of *exact* SHAP values ϕ_i, is computed through the marginal contributions of each feature r across all possible combinations of features as:

$$\phi_{r,i} = \sum_{\mathcal{S} \subseteq P \setminus \{r\}} \frac{|\mathcal{S}|!(P - |\mathcal{S}| - 1)!}{P!} \left[f_{\mathcal{S} \cup \{r\}}(\mathbf{Y}^*_{i,\mathcal{S} \cup \{r\}}) - f_{\mathcal{S}}(\mathbf{Y}^*_{i,\mathcal{S}}) \right]$$

where:

- $\mathcal{P}$ is the set of features ("grand coalition");
- $\mathcal{S} \subseteq \mathcal{P} \setminus \{r\}$, all possible subsets of features not containing r;
- $f_{\mathcal{S} \cup \{r\}}(\mathbf{Y}^*_{\mathbf{i},\mathcal{S} \cup \{\mathbf{r}\}})$ is the output for OOC observation $\mathbf{Y}^*_{\mathbf{i},\mathcal{S} \cup \{\mathbf{r}\}}$, predicted by the ML model retrained on $\mathbf{X}_{\mathcal{S} \cup \{\mathbf{r}\}}$;
- $f_{\mathcal{S}}(\mathbf{Y}^*_{\mathbf{i},\mathcal{S}})$ is the output for OOC observation $\mathbf{Y^*}_{i,\mathcal{S}}$, predicted by the ML model retrained on $\mathbf{X}_{\mathcal{S}}$.

SHAP Values Approximation

Exact computation of the SHAP values can be computationally intractable. SHAP values can also be obtained as the solution of the following weighted least squares problem:

$$\min_{\phi_i} \left\{ \sum_{\mathcal{S} \subseteq \mathcal{P}} \left(\nu_i(\mathcal{S}) - \left(\phi_0 + \sum_{j \in \mathcal{S}} \phi_{j,i} \right) \right)^2 k(P, |\mathcal{S}|) \right\}$$

where $k(P, |\mathcal{S}|) = \frac{P-1}{\binom{P}{|\mathcal{S}|} |\mathcal{S}| \, (P - |\mathcal{S}|)}$ are the Shapley kernel weights. The problem can then be reformulated as:

$$\min_{\phi_i} \left\{ (\boldsymbol{\nu}_i - \mathbf{T}\phi_i)^T \, \mathbf{W} \, (\boldsymbol{\nu}_i - \mathbf{T}\phi_i) \right\}$$

and the solution is given by $\phi_i = \left(\mathbf{T}^T \mathbf{W} \mathbf{T} \right)^{-1} \mathbf{T}^T \mathbf{W} \boldsymbol{\nu_i}$, where $\mathbf{T}$ is the $2^P \times (P+1)$ binary matrix of all possible coalitions, $\mathbf{W}$ is the $2^P \times 2^P$ diagonal matrix of kernel weights, ϕ_i is the vector of SHAP values, and $\boldsymbol{\nu}_i$ is the vector of coalition functions $\nu_i(\mathcal{S})$.

SHAP Values Approximation

Kernel SHAP approximation: This method constructs a reduced matrix $\mathbf{T}_{\mathcal{D}}$ by selecting a subset of rows from the full coalition matrix $\mathbf{T}$.

$$\hat{\phi}_{\mathbf{i}} = \left(\mathbf{T}_{\mathcal{D}}^T \mathbf{W}_{\mathcal{D}} \mathbf{T}_{\mathcal{D}}\right)^{-1} \mathbf{T}_{\mathcal{D}}^T \mathbf{W}_{\mathcal{D}} \boldsymbol{\nu}_{\mathbf{i}}$$

where $\nu_i(\mathcal{S}) = E[f(\mathbf{Y}_{\mathbf{i},\mathcal{S}}, \mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}}) \mid (\mathbf{Y}_{\mathbf{i},\mathcal{S}} = \mathbf{Y}^*_{\mathbf{i},\mathcal{S}})]$.

- **Interventional Approach**:

$$\nu_i(\mathcal{S}) = \int f(\mathbf{Y}^*_{\mathbf{i},\mathcal{S}}, \mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}})\ p(\mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}})\ d\mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}}$$

  The integral above can be approximated as:

$$\nu_i(\mathcal{S}) = \frac{1}{Q} \sum_{q=1}^{Q} f(\mathbf{Y}^{\mathbf{q}}_{\mathbf{i},\bar{\mathcal{S}}}, \mathbf{Y}^*_{\mathbf{i},\mathcal{S}})$$

- **Observational Approach**:

$$\nu_i(\mathcal{S}) = \int f(\mathbf{Y}^*_{\mathbf{i},\mathcal{S}}, \mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}})\ p(\mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}} \mid \mathbf{Y}_{\mathbf{i},\mathcal{S}} = \mathbf{Y}^*_{\mathbf{i},\mathcal{S}})\ d\mathbf{Y}_{\mathbf{i},\bar{\mathcal{S}}}$$

Computational analysis

- $M \in \{5000, 10000\}$, $P \in \{10, 30\}$.
- Marginals are defined as $N(2, \sqrt{3})$, Clayton copula with $\tau = 0.5$.
- We simulated an (OOC) signal ($\mu_i = 7$ for $i = 1, \ldots, \frac{P}{2}$).

Q	kNNDD	IF	T^2
100	10.1(48.4)	0.46(1.15)	0.05(0.2)
300	33.8(155.6)	1.62(3.22)	0.07(0.37)
500	54.6(248.2)	2.56(6.06)	0.12(0.53)
1000	115.1(478.3)	4.97(11.83)	0.21(0.8)
2000	235.5()	11.01(23.26)	0.41(1.71)
3000	521.9(-)	16.14(39.56)	0.57(2.96)
Exact	10.15(NA)	0.37(NA)	0.3(NA)

$M = 5000, \; P = 10(30)$

Q	kNNDD	IF	T^2
100	20.1(93.9)	0.45(1.12)	0.04(0.1)
300	69.9(301.2)	1.4(3.02)	0.06(0.3)
500	114.8(485.7)	2.4(5.6)	0.09(0.4)
1000	230.9(893.9)	4.7(10.1)	0.20(0.8)
2000	452.6(1820.2)	10.1(24.1)	0.3(1.7)
3000	659.01(-)	17.2(35.8)	0.4(3.5)
5000	-(-)	27.4(56.8)	0.8(12.5)
Exact	26.1(NA)	0.2545(NA)	0.5994(NA)

$M = 10000, \; P = 10(30)$

Computational analysis

- $\text{MAE}_r = \frac{1}{B}\sum_{b=1}^{B} |\hat{\phi}_r^b - \phi_r|$, for $r = 1, 2, \ldots, P$, $P = 10$, $M = 5000$.
- Marginals are defined as $N(2, \sqrt{3})$, Clayton copula with $\tau = 0.5$.
- We simulated an (OOC) signal ($\mu_i = 7$ for $i = 1, \ldots, \frac{P}{2}$).

	Q=1000	Q=2000	Q=3000	Q=4000	Q=M	Exact SHAP
Feature 1 (OOC)	0.0010	0.0010	0.0009	0.0007	0.0004	0.0484
Feature 2 (OOC)	0.0000	0.0002	0.0001	0.0002	0.0002	0.0304
Feature 3 (OOC)	0.0004	0.0004	0.0004	0.0005	0.0007	0.0315
Feature 4 (OOC)	0.0003	0.0006	0.0006	0.0007	0.0008	0.0425
Feature 5 (OOC)	0.0007	0.0008	0.0007	0.0006	0.0007	0.0424
Feature 6 (IC)	0.0020	0.0020	0.0020	0.0019	0.0018	-0.0047
Feature 7 (IC)	0.0001	0.0002	0.0002	0.0003	0.0004	0.0125
Feature 8 (IC)	0.0000	0.0000	0.0001	0.0000	0.0000	-0.0053
Feature 9 (IC)	0.0000	0.0001	0.0001	0.0001	0.0000	-0.0041
Feature 10 (IC)	0.0000	0.0002	0.0003	0.0003	0.0004	0.0026

Implementation of a Phase II IF EWMA control chart

- Is the IF model training step distribution-free? Is the distribution of score statistics independent of the underlying distribution of the observations?
- How does the selection of an appropriate size M, for the reference sample $\mathbf{X}$, and the dimensionality P, affect the distribution of the score statistics?

Investigation on the IC Scores Distribution

Investigated distributions

Cases	Copula	Marginal Distributions
1	Clayton	$X_i \sim N(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3}), i = 1, \ldots, P$
2	Clayton	$X_i \sim G(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3}), i = 1, \ldots, P$
3	Clayton	$X_i \sim \mathrm{t-Student}(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3}, df = 3), i = 1, \ldots, P$
4	Clayton	$X_1, \ldots X_{P-\lceil P/2 \rceil} \sim N(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3})$ $X_{(P-\lceil P/2 \rceil)+1}, \ldots X_p \sim G(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3})$
5	Clayton	$X_1, \ldots X_{P-\lceil P/2 \rceil} \sim \mathrm{t-Student}(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3}, df = 3)$ $X_{(P-\lceil P/2 \rceil)+1}, \ldots X_P \sim G(\mu_{X_i} = 2, \sigma_{X_i} = \sqrt{3})$

1. $R = 5000$ reference samples $\mathbf{X}^{(r)}$ of size M of observations are simulated from Cases 1-5.
2. For each reference sample the score statistics $s(\mathbf{X_i}^{(r)}, M)$ are calculated.
3. Repeating this procedure R times, the empirical distributions of several score quantiles can be graphically represented.

Effect of the reference sample size M

$m \in \{500, 1000, 1500, 2000, 3000, \ldots, 5000\}$

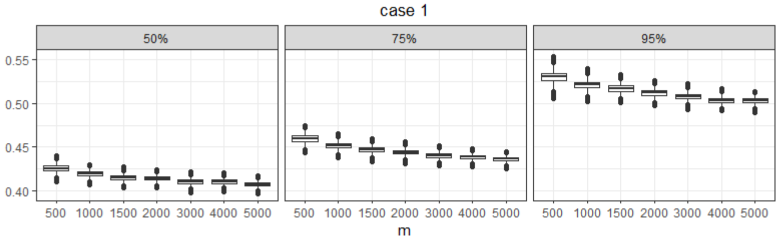

Effect of dimension P

$p \in \{2, 5, 10, 15, 20\}$

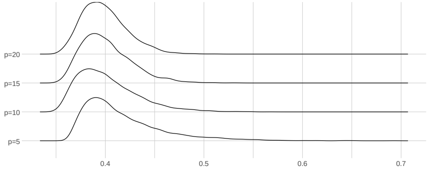

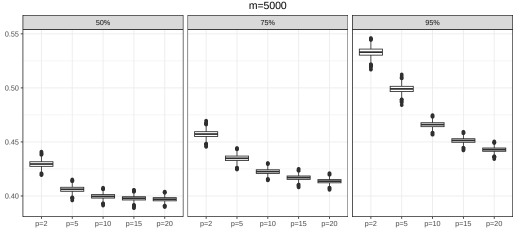

IC Performance: Run length properties

- For each value of M, we get the n simulated test observations with Algorithm 1 and calculate the UCL.
- To investigate the robustness of the UCL design procedure, we calculate *a posteriori* the IC MRL by simulating the Phase II test observations independently of the reference sample, but from the same IC distribution.
- A good IC performance would require an actual IC MRL as close as possible to the nominal value MRL_0 fixed to design UCL.

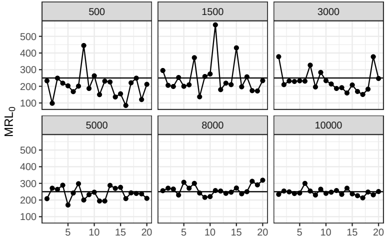

Figure: Case 1, $p = 5$.

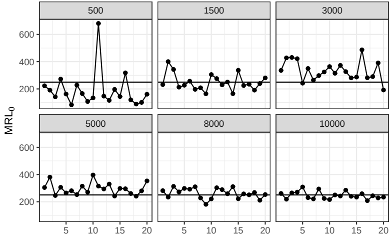

Figure: Case 3, $p = 10$.

Investigation on the IC Scores Distribution

Effect of the Underlying Distribution of Observations

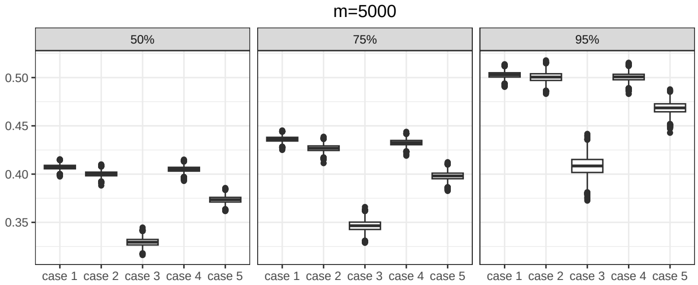

Conclusions

- The IF model scores are *not* distribution-free. Determination UCL is only possible by means of resampling from the Phase I sample.
- The required Phase I sample size must be large, say $M \geq 5000$.
- PSD through SHAP values can be considered as an simple and efficient way to identify which variable caused the shift in the process.

Published Work

- The First part of this presentation has been recently submitted in *Quality Technology & Quantitative Management*: **Perdikis, T.**, Celano, G. and Chakraborti, S., "Post Signal Diagnostics Using SHAP Values for Machine Learning Based Control Charts in Multivariate SPM".
- The Second part of this presentation can be found in **Perdikis, T.**, Celano, G. and Chakraborti, S. (2026), "Practical Considerations for Implementing the Isolation Forest EWMA Control Chart in Phase II Process Monitoring", *Statistics and Computing*, 36(9).

References

Cuentas S., Peñabaena-Niebles R., Garcia, E. Support Vector Machine in Statistical Process Monitoring: a Methodological and Analytical Review. International Journal of Advanced Manufacturing Technology, 2017; 91, 485–500.

Deng H.,Runger G.,Tuv E. System Monitoring with Real-Time Contrasts. Journal of Quality Technology, 2012; 44(1), 9–27

James G., Witten D., Hastie T., Tibshirani R. An Introduction to Statistical Learning with Applications in R, II Edn., 2023; Springer.

Jones-Farmer L.A., Woodall W.H.,Steiner S.H.,Champ C.W. An Overview of Phase I Analysis for Process Improvement and Monitoring. Journal of Quality Technology, 2014; 46(3), 265–280.

Sven K., Wolfgang S. Control charts for time series; Frontiers in statistical quality control 7 (2004): 210–236.

Liu F.T., Ting K.M., Zhou Z-H. Isolation forest. Eight IEEE International Conference on data mining, 2008; 413–422.

Lucas J.M., Saccucci M.S. Exponentially Weighted Moving Average Control Schemes: Properties and Enhancements. Technometrics, 1990; 32(1), 1–12

Wang J., Liu L. A New Multivariate Control Chart Based on the Isolation Forest Algorithm. Quality Engineering, 2023; doi =10.1080/08982112.2023.2220773.

Thanks for your attention! Any questions?