

Multivariate process monitoring with Hotelling T^2 and SPE

Read correlated sensors through T^2 and SPE together rather than one statistic, interpret contribution plots, and learn where the method's assumptions break.

In short

- Correlated process variables can move into a wrong combination while each one stays inside its own single-variable limits.
- T^2 measures how extreme a point is within the normal model plane, while SPE measures how far it lies from that plane.
- Both statistics depend on clean normal training data and standardised variables, and both need retraining after regime changes.

01 / The blind spot of single-variable charts

In a pump, flow, pressure and motor current often move together. When flow rises, motor current may rise as well, and pressure depends on system resistance. If you place a control chart on each variable, all three can stay inside their own limits while the pattern they form together has broken. Single-variable charts do not see that wrong combination. Multivariate monitoring models the shared structure of the variables in one model and measures departures from it.

Two separate questions follow. The first asks how extreme the point is within normal variation. The second asks whether the normal relationship between the variables still holds. The two catch different faults, which is why this approach uses two statistics.

02 / The model plane: PCA and standardisation

Principal component analysis (PCA) finds the main directions along which the variables move together. The first component is the direction of largest normal variation, and later components carry the remaining variation in directions perpendicular to the earlier ones. The model keeps k components; the subspace they span is called the model plane. Choose k so that it explains a sufficient share of the normal variation.

Before fitting, standardise the variables: subtract the training mean from each sensor and divide by its standard deviation. Otherwise a sensor with a large numeric range dominates the others. The training data should contain only clean, normal operation. If a fault period enters the training window, the plane treats that fault as part of normal structure.

03 / How T^2 and SPE are calculated

T^2 measures how far a new point lies from the centre of normal variation inside the model plane. It is computed from the point's component scores and reads as a Mahalanobis-type distance:

$$T^2 = \sum (t_i^2 / \lambda_i), \text{ for } i = 1, \dots, k$$

Here t_i is the point's score on component i , and λ_i is the variance of that component's scores in the training data. A large T^2 means the point is unusually extreme within normal variation.

SPE, also called Q , is the squared distance from the point to the model plane. Let $r = x - \hat{x}$ be the residual vector, where $\hat{x}$ is the projection of x onto the plane:

$$SPE = \sum r_j^2$$

A large SPE means the relationship between variables has moved outside the structure seen in training. It can indicate a broken correlation.

04 / A hypothetical calculation: two statistics, two questions

The numbers below are hypothetical, and the variables are assumed to be standardised. The model has $k = 2$ components, and in the training data their score variances are $\lambda_1 = 4$ and $\lambda_2 = 1$. A new point has scores $t_1 = 4$ and $t_2 = 1.5$. Then:

- $T^2 = 4^2 / 4 + 1.5^2 / 1 = 16 / 4 + 2.25 = 4 + 2.25 = 6.25$

Set a hypothetical T^2 limit of 9. The value 6.25 is below it, so the point looks normal inside the plane.

Now let the residual vector be $r = (0.3, -0.4, 0)$. Then $SPE = 0.3^2 + 0.4^2 + 0^2 = 0.09 + 0.16 = 0.25$. With a hypothetical SPE limit of 0.10, the point is flagged on the SPE side.

The conclusion: the point's position within the plane is normal, but it has left the plane. A break of this kind may be invisible when each sensor is viewed on its own chart.

05 / Assumptions and limit values

The method assumes that relationships between variables are largely linear. If they are not, the plane does not represent the real structure, and both T^2 and SPE can raise false alarms. Clean training data is also essential: a fault or a post-maintenance shift inside the training window moves the normal limit with it.

Limits can be set in three ways: an approximation based on the F distribution for T^2 , an approximation based on residual eigenvalues for SPE, or the empirical percentile of the training data. Each rests on assumptions. The empirical percentile is only as representative as the training data. After maintenance, a product change or a change of operating regime, retrain the model.

06 / Where to start on a process

Choose a period of stable operation and record the sensors at the same sampling interval; write the units into the data dictionary. Standardise the variables, and set the number of components and the limits from training data only.

Then replay a past event with the maintenance logs. When did the model signal, and when did the maintenance team act? For each signal, look at the contribution plot and note which variable drove the departure. Interpret those notes together with the engineers who know the process. Finally, record the false-alarm rate and the average time spent reviewing one alarm; these two numbers show directly what the method is worth for your plant.

Try it

In Sentinel's sample sensor table, open one flagged row and write one sentence stating whether T^2 or SPE rose more, and what that suggests about the departure.

References and online guide

[Open the online guide](#)

NIST/SEMATECH e-Handbook: 6.5.4.3. Hotelling T squared

NIST/SEMATECH e-Handbook: 6.5.4.1. Mean Vector and Covariance Matrix

[Hotelling T-squared distribution \(Wikipedia\)](#)

[Principal component analysis \(Wikipedia\)](#)

Proposed pilot steps do not claim field accuracy. Synthetic examples are not measurements from a production site.