CUSUM for Detecting Small Persistent Shifts in a Process Mean

Gary Pai

March 2008

1 Introduction

The CUmulative SUM (CUSUM) chart is a powerful sequential procedure rooted in statistical hypothesis testing. This document provides the theoretical formulation of the standard CUSUM, its parameters, and robust extensions necessary for application in heavy-tailed, non-stationary environments like market risk scoring.

2 Core Theoretical Foundation: The SPRT

The CUSUM procedure, formalized by E. S. Page, is closely related to the Sequential Probability Ratio Test (SPRT) developed by Abraham Wald. The SPRT is the optimal sequential procedure for detecting a change from an "in-control" distribution with parameter θ_0 to a specific "out-of-control" distribution with parameter θ_1.

The CUSUM statistic is an accumulation of the **Log-Likelihood Ratio (LLR)** for each observation X_i:

$$L_t = \sum_{i=1}^{t} \ln\left(\frac{f_{\theta_1}(X_i)}{f_{\theta_0}(X_i)}\right) \tag{1}$$

The detection rule is an **optimal stopping rule**, τ, which minimizes the detection delay for a fixed False Alarm Rate (FAR) or Average Run Length in-control (ARL_0).

3 The Standard CUSUM for Shift in Mean

Assuming a variable X_t that is normally distributed with known variance σ^2, the standard CUSUM chart detects a shift in the mean from the target μ_0 to a potential out-of-control mean $\mu_1 = \mu_0 \pm \delta\sigma$. The procedure uses two one-sided statistics, C_t^+ for an upward shift and C_t^- for a downward shift.

3.1 Upward Shift Statistic (C_t^+)

This statistic accumulates positive deviations that exceed a specified allowance K:

$$C_t^+ = \max\left(0, C_{t-1}^+ + (X_t - (\mu_0 + K))\right) \tag{2}$$

3.2 Downward Shift Statistic (C_t^-)

This statistic accumulates the magnitude of negative deviations that fall below the corresponding allowance (formulated to track positive evidence of a downward shift):

$$C_t^- = \max\left(0, C_{t-1}^- + ((\mu_0 - K) - X_t)\right) \tag{3}$$

3.3 Parameters

The performance is governed by two key parameters:

1. **Reference Value (K):** The standardized shift allowance, often set to half the desired shift magnitude δ:

$$K = \frac{\delta\sigma}{2}$$

2. **Decision Interval (H):** The control limit that determines the sensitivity and the in-control performance (ARL_0). A large H ensures a large ARL_0 (minimizing false positives).

3.4 Alert Rule

An alarm is signaled at time τ if either statistic exceeds the decision interval H:

$$\tau = \min\{t \geq 1 : C_t^+ > H \text{ or } C_t^- > H\} \tag{4}$$

4 Relevance to Robust Regime Classification (Robust CUSUM)

For non-stationary, heavy-tailed data, robust CUSUM methods replace the deviation term $(X_t - \mu_0)$ with a non-parametric statistic.

4.1 Non-Parametric Sign-Test CUSUM

This robust version considers only the sign of the deviation relative to the in-control median (Median_0), making it insensitive to outlier magnitude:

$$C_t^+ = \max\left(0, C_{t-1}^+ + (\mathbf{I}(X_t > \text{Median}_0) - p_0)\right) \tag{5}$$

where $\mathbf{I}(\cdot)$ is the indicator function and p_0 is the expected probability of being above the median (typically 0.5).

4.2 Wilcoxon-Based CUSUM

This approach utilizes the rank statistic of X_t relative to a reference distribution, offering a powerful, robust alternative for detecting location shifts.

These robust sequential procedures directly address the requirements for:

- **Minimizing False Positives:** Controlled by setting a large H (high ARL_0).
- **Tolerating Long Quiet Periods:** Ensured by the accumulation structure, which requires persistent evidence (exceeding K).
- **Robustness under Heavy-Tailed Noise:** Achieved by using non-parametric statistics (ranks or signs).

5 CUSUM Example: Detecting a Shift in a Univariate Risk Score

This example demonstrates the calculation of the CUmulative SUM (CUSUM) statistic for detecting an upward shift in a univariate risk score, X_t. Since the score increases under instability, we focus on the **Upward CUSUM (C_t^+)**.

6 Setup and Parameters

We define the in-control parameters and the control limits for the chart.

- **Target Mean (μ_0):** 10.0 (Stable regime score)
- **Process Standard Deviation (σ):** 1.0
- **Desired Shift to Detect ($\delta\sigma$):** 1.0σ shift
- **Reference Value (Slack, K):** Set to half the desired shift:

$$K = \frac{\delta\sigma}{2} = \frac{1.0}{2} = 0.5$$

- **Decision Interval (Threshold, H):** Set at 5.0σ for a high ARL_0:

$$H = 5.0 \times 1.0 = 5.0$$

6.1 CUSUM Formula (Upward Shift)

The Upward CUSUM statistic accumulates deviations above the target plus the slack ($\mu_0 + K = 10.5$):

$$C_t^+ = \max\left(0, C_{t-1}^+ + (X_t - (\mu_0 + K))\right) \tag{6}$$

7 Calculation and Shift Detection

The table below tracks a sequence of observations where a sustained upward shift begins at $t = 6$. The alarm threshold is $H = 5.0$.

Table 1: CUSUM Calculation for Upward Shift Detection

Time (t)	**Obs.** (X_t)	**Target + Slack**	**Excess Above** ($X_t - 10.5$)	**Previous** $\mathbf{C^+_{t-1}}$	**New** $\mathbf{C^+_t}$	**Signal** ($\mathbf{C^+_t} > \mathbf{5.0}$)
0	—	10.5	—	0.00	—	No
1	10.2	10.5	−0.30	0.00	0.00	No
2	10.6	10.5	0.10	0.00	0.10	No
3	10.1	10.5	−0.40	0.10	0.00	No
4	10.4	10.5	−0.10	0.00	0.00	No
5	11.0	10.5	0.50	0.00	0.50	No
Shift in Mean Occurs Here						
6	11.2	10.5	0.70	0.50	1.20	No
7	11.5	10.5	1.00	1.20	2.20	No
8	11.8	10.5	1.30	2.20	3.50	No
9	**12.0**	10.5	1.50	3.50	**5.00**	**ALARM!**
10	12.1	10.5	1.60	5.00	6.60	ALARM!

8 Conclusion

The accumulation structure of the CUSUM chart ensures that small, sustained deviations ($X_t > 10.5$) are accumulated, while transient noise ($X_t \leq 10.5$) causes the statistic to be reset to zero. The shift, starting at $t = 6$, is successfully detected at $\mathbf{t = 9}$, demonstrating the CUSUM's efficiency in signaling persistent structural instability.