When Randomness Changes:

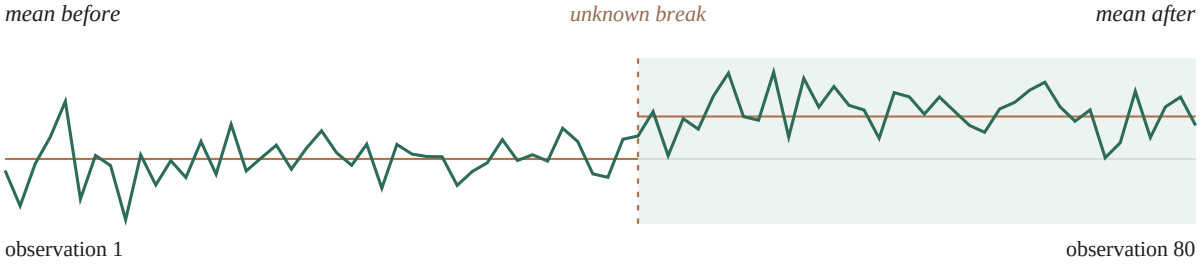
Detecting Change Points

BAMS project - Mathematical Statistics

Abstract.

A machine produces measurements that fluctuate around a stable level. At some unknown moment its behaviour changes: the readings begin to cluster around a new value. Looking only at the noisy sequence, can we decide whether a change occurred, and locate it?

This project studies **change-point detection** for a sequence of observations. Starting with a Gaussian model in which the mean can change once at an unknown time, we derive a likelihood-based change-point estimator and connect it to cumulative-sum (CUSUM) statistics. We then investigate how reliably these procedures distinguish a genuine shift from random fluctuations.



A simulated sequence with one hidden shift in the mean. The observations fluctuate on both sides of the change.

Prerequisites / Difficulty.

Probability and Mathematical Statistics: expectations, variance, Gaussian distributions, likelihood estimation and hypothesis testing. Basic Python or R is useful for simulations. No previous time-series course is required.

Intermediate. The core problem involves one unknown change in the mean. Changes in variance, sequential monitoring and multiple change points offer more advanced directions.

Starting materials.

D. V. Hinkley (1970), *Inference about the change-point in a sequence of random variables*. Begin with the model and likelihood formulation.

C. Truong, L. Oudre and N. Vayatis (2020), *Selective review of offline change point detection methods*. Use the introduction and the single-change framework.

E. S. Page (1954), *Continuous Inspection Schemes*. A starting point for the optional sequential-monitoring extension.

What to do.

The common goal is to detect and estimate **one unknown change in the mean**. The project develops the likelihood calculation, a CUSUM representation and experiments testing when the procedures succeed.

1. Set up the change-point model. Let $X_1, ..., X_n$ be independent Gaussian observations with common variance. Under the null there is one mean; under the alternative there are two means separated by an unknown index k. Derive the maximum-likelihood estimates of the means for a fixed split and explain what cannot be identified under the null.

2. Find the most likely split. Write the likelihood under both models. Show that maximizing the likelihood improvement is equivalent to examining a weighted squared difference between the means on the two sides of a split. Search over admissible values of k and discuss why the maximum over many candidate dates is not an ordinary fixed-parameter test.

3. Connect likelihoods to cumulative sums. Define $C_k = \Sigma_{i \le k}(X_i - \bar{X})$. Derive the identity connecting C_k to the difference of segment means, and show how standardized CUSUM scores relate to the likelihood comparison. Which splits receive more weight, and why?

4. Evaluate detection in simulations. Generate samples with no change and with shifts of different magnitudes and locations. Use simulations under the null to calibrate detection thresholds. Compare false-alarm rates, power and the error in estimating the change point as sample size or shift size varies.

5. Apply the method to an observed series. If time permits, examine a dataset such as sensor measurements. Contrast the fitted segments with the original sequence and discuss what the model assumes. A change detected in a real series does not, by itself, identify its cause.

Further directions.

- **Changing volatility.** Adapt the model to a change in variance and compare the outcome with the mean-shift detector; financial returns provide a possible application.
- **Multiple changes.** Find several change points while penalizing unnecessary splits.
- **Online monitoring.** Study Page's CUSUM for detecting changes as new observations arrive.

Expected outcome.

A short illustrated mathematical report and a reproducible notebook containing the likelihood and CUSUM derivations, simulation comparisons, estimated change locations and examples of both correct detections and false alarms. The theoretical and computational parts can be divided among participants.