

SURvival Control Chart ESTimation Software in R: the success package

Daniel Gomon
Leiden University

Marta Fiocco
Leiden University

Hein Putter
Leiden University

Mirko Signorelli
Leiden University

Abstract

Monitoring the quality of statistical processes has been of great importance, mostly in industrial applications. Control charts are widely used for this purpose, but often lack the possibility to monitor survival outcomes. Recently, inspecting survival outcomes has become of interest, especially in medical settings where outcomes often depend on risk factors of patients. For this reason many new survival control charts have been devised and existing ones have been extended to incorporate survival outcomes. The R package **success** allows users to construct risk-adjusted control charts for survival data. Functions to determine control chart parameters are included, which can be used even without expert knowledge on the subject of control charts. The package allows to create static as well as interactive charts, which are built using **ggplot2** (Wickham 2016) and **plotly** (Sievert 2020).

Keywords: CUSUM, Control charts, Quality control, Survival analysis, funnel plot, R.

1. Introduction

Inspecting the quality of a survival process is of great importance, especially in the medical field. Many of the methods currently used to inspect the quality of survival processes in a medical setting, such as funnel plots (Spiegelhalter 2005), Bernoulli Cumulative sum (CUSUM) charts (Steiner, Cook, Farewell, and Treasure 2000) and exponentially weighted moving average (EWMA) charts (Cook, Coory, and Webster 2011) work only with binary outcomes, and are thus not appropriate for survival outcomes. These charts require the continuous time outcome to be dichotomized, often leading to delays when trying to detect problems in the quality of care. To overcome this limitation, Biswas and Kalbfleisch (2008) developed a continuous time CUSUM procedure (which we call the BK-CUSUM), that can be used to inspect survival outcomes in real time. Gomon, Putter, Nelissen, and van der Pas (2022) proposed a generalization of the BK-CUSUM chart called the Continuous Time Generalized Rapid Response CUSUM (CGR-CUSUM). The CGR-CUSUM allows to estimate some of the parameters involved in the construction of the chart, overcoming the need for the user to correctly specify parameters that the BK-CUSUM procedure relies on. Recently other procedures for the continuous time inspection of survival outcomes have been developed, such as the improved Bernoulli CUSUM (Keefe, Loda, Elhabashy, and Woodall 2017), uEWMA chart (Steiner and Jones 2009), STRAND chart (Grigg 2018) and funnel plot for survival outcomes (Putter, Eikema, de Wreede, McGrath, Sánchez-Ortega, Saccardi, Snowden, and van Zwet 2022). To the best of our knowledge there are no publicly available software im-

plementations of these methods, therefore the focus of this article is on survival CUSUM charts.

When constructing control charts in continuous time, not only the time to failure of a subject is of interest, but also the information provided by the survival up until current time is crucial. Many quality control methods cannot incorporate continuous time (survival) outcomes, requiring the continuous time outcome to be dichotomized (e.g. 30-day survival). The resulting binary data is called discrete time data. We provide an overview of some of the existing R (R Core Team 2021) packages which can be used for constructing control charts. The package **qcc** (Scrucca 2004) for discrete time data contains functions to construct many types of Shewhart, binary CUSUM, EWMA and other charts. The packages **qcr** (Flores 2021), **qicharts** (Anhoej 2021) and **ggQC** (Grey 2018) also allow for the construction of discrete time control charts, but differ in their graphical possibilities and their intended application area such as medicine, industry and economics. It is possible to assess some of the discrete time control charts by means of their zero/steady state average run length using the package **spc** (Knoth 2021). The package **funnelR** (Kumar 2018) allows for the construction of funnel plots for proportion data, a method often used in medical statistics to visualise the difference in proportion over a time frame. The package **cusum** (Hubig 2019) can be used to monitor hospital performance using a Bernoulli CUSUM, also allowing users to easily determine control limits for continuously inspecting hospitals. The package however requires the input to be presented in a binary format.

Whereas many packages exist allowing for the construction of quality control charts on discrete time data, to the best of our knowledge there are currently few statistical software packages allowing for the construction of quality control charts on survival data and no R packages allowing for the continuous time inspection of survival outcomes.

The main contribution of this article is to present the R (R Core Team 2021) package **success** (SURvival Control Chart ESTimation Software), a tool for constructing quality control charts on survival data. With this package, we aim to fill the gap in available open source software for the construction of control charts on survival data. The primary goal of **success** is to allow users to easily construct the BK- and CGR-CUSUM; moreover, **success** can also be used to construct the discrete time funnel plot (Spiegelhalter 2005) and the Bernoulli CUSUM chart (Steiner *et al.* 2000) on survival data without manually dichotomizing the outcomes. This way, users can determine the possible gain in detection speed by using continuous time quality control methods over some popular discrete time methods. Our aim is to have **success** become the reference package for users looking to compute control charts for survival outcomes.

The article is structured as follows. In Section 2, we briefly describe the theory that underlines the control charts presented in the package. Section 3 then explains how to prepare the input data and describes the functions and their arguments. For the reader not interested in technical details, only Subsection 3.2 is of interest. In Section 4, the control charts in the package are applied to data based on a clinical trial for breast cancer. Finally, the article ends with a discussion about the methods presented.

2. Theory and Models

Throughout this article, we are interested in comparing institutional (hospital) performance for survival after a specific medical procedure (surgery). Even though our focus is on medical

applications, the methods can be applied to any data set containing survival outcomes, ranging from production lines to customer satisfaction inspection.

In this section we introduce the funnel plot, Bernoulli CUSUM, BK-CUSUM and CGR-CUSUM implemented in the package **success**. The goal of each of the methods is to detect a deviation from a certain target performance measure and discover hospitals with increased mortality rates as quickly as possible.

Reading guide

This section consists of two parts. Subsection 2.1 summarizes the minimal required knowledge in layman's terms. Subsection 2.2 delves into the mathematical details and assumptions. Further Subsections introduce the mathematical theory of the methods available in **success**.

2.1. Terminology

We assume that hospitals/units have a constant and steady stream of patients/subjects coming in for a treatment of interest (e.g. surgery). In survival quality control, we are interested in determining whether failure/death rates after treatment at a certain hospital deviate significantly from a certain target measure. A target measure defines the acceptable failure rate. This measure can be set or estimated from a large set of (historical) data. Most of the time, we are only interested in detecting an increase in failure rate as this indicates that the hospital in question is performing worse than expected, and that corrective interventions may be necessary to improve the quality of care at such a hospital.

The process is defined to be *in-control* when failures occur according to the target level. It is *out-of-control* if failures happen at a higher rate than expected. Hospitals may have in-control periods followed by out-of-control periods. The goal is to detect when observations at a hospital start going out-of-control as soon as possible, so action can be taken quickly.

To continuously inspect the quality of the process, we construct a control chart to monitor the process failure rate from the start of the study. Some charts only change value when an outcome is observed (discrete time), while others change value at each time point (continuous time). When the control chart exceeds a pre-defined *control limit*, a signal is produced, indicating that the hospital in question is performing worse/better than the target measure. The time it takes for a control chart to exceed the control limit is called the *run length* of the chart. For some control charts it is necessary to fix the expected increase in failure rate in advance. This is done by specifying a parameter θ , where e^θ indicates the expected increase in the failure odds (discrete time) or expected increase in the failure rate (continuous time).

Not all patients have an equal probability of failure at any given point in time. For example, people who smoke may have a larger risk of failure than non-smokers. Therefore, patient's risk can be incorporated into the control charts by using a risk-adjustment model. Relevant characteristics (called covariates) are then used to determine the increase/decrease in the risk of failure for each patient.

2.2. Mathematical notation

Consider a single hospital. For each patient ($i = 1, 2, \dots$) let S_i and X_i be the chronological entry time and survival/failure time from the time of entry respectively. The chronological time of failure is then given by $T_i = S_i + X_i$. Assume that patients arrive (enter the study)

at the hospital according to a Poisson process with rate ψ , and that each patient i has a set of p covariates \mathbf{Z}_i . Let $h_i(x) = h_0(x)e^{\mathbf{Z}_i\beta}$ be the subject-specific hazard rate obtained from the Cox (1972) proportional hazards model. Let $Y_i(t) = \mathbb{1}\{S_i \leq t \leq \min(T_i, R_i)\}$ indicate whether a patient is *at risk* (of failure) at time t , where R_i denotes the right censoring time of patient i .

Define by $N_i(t)$ the counting process indicating whether patient i experiences a failure at or before time t . Let $N(t) = \sum_{i \geq 1} N_i(t)$ be the total number of observed failures at the hospital at or before time t . Define the cumulative intensity of patient i as $\Lambda_i(t) = \int_0^t Y_i(u)h_i(u)du$ and $\lambda_i(t) = Y_i(t)h_i(t)$. Note that $\lambda_i(t)$ is equal to zero when patient i is not at risk. Similarly, let $\Lambda(t) = \sum_{i \geq 1} \Lambda_i(t)$ be the total cumulative intensity at the hospital at time t .

For some methods, we will be interested in detecting a fixed increase in the cumulative intensity at the hospital from $\Lambda(t)$ to $\Lambda^\theta(t) := e^\theta \sum_{i \geq 1} \Lambda_i(t)$. In a similar fashion, denote $h_i^\theta := e^\theta h_i(t)$. The corresponding density and distribution functions are denoted by f_i^θ and F_i^θ . We call e^θ the *true hazard ratio*. When $e^\theta = 1$ (similarly $\theta = 0$), we say that the failure rate is *in-control*. Alternatively, when $e^\theta > 1$ we say that the failure rate is *out-of-control*.

For a chart $C(t)$ with changing values over time define the *average run length* for a given control limit h as $\mathbb{E}[\tau_h]$, where $\tau_h = \inf\{t > 0 : C(t) \geq h\}$.

2.3. Funnel plot

The risk-adjusted funnel plot (Spiegelhalter 2005) is a graphical method used to compare performance between hospitals over a fixed period of time. The general structure of the data is as follows: there are k centers/hospitals ($j=1 \dots k$) with n_j treated patients in hospital j . For each patient we observe a binary variable $X_{i,j}$:

$$X_{i,j} = \begin{cases} 1, & \text{if patient } i \text{ at hospital } j \text{ had an adverse event within } C \text{ days,} \\ 0, & \text{if patient } i \text{ at hospital } j \text{ otherwise.} \end{cases} \quad (1)$$

Let $X_{i,j} \sim \text{Ber}(p_j)$, with p_j the probability of failure at hospital j within C days. Consider the hypotheses:

$$H_0 : p_j = p_0 \qquad H_1 : p_j \neq p_0 \quad (2)$$

with p_0 some baseline (in-control) failure proportion. The proportion of failures observed at hospital j is then given by $\gamma_j = \frac{\sum_{i=1}^{n_j} X_{i,j}}{n_j}$. The asymptotic distribution of γ_j under H_0 is

$$\gamma_j|_{H_0} \sim \mathcal{N}\left(p_0, \frac{p_0(1-p_0)}{n_j}\right).$$

We can then signal an increase or decrease in the failure proportion of hospital j with confidence level $1 - 2\alpha$ when

$$\gamma_j \notin \left[p_0 + \xi_\alpha \sqrt{\frac{p_0(1-p_0)}{n_j}}, p_0 - \xi_\alpha \sqrt{\frac{p_0(1-p_0)}{n_j}} \right] \quad (3)$$

with ξ_α the α -th quantile of the standard normal distribution.

It is often desirable to determine patient specific failure probabilities using some of their characteristics. A risk-adjusted funnel plot procedure can then be performed by modelling the patient specific failure probability using a logistic regression model:

$$p_i = \frac{1}{1 + e^{-\beta_0 + \mathbf{Z}_i \beta}} \quad (4)$$

where \mathbf{Z}_i is the vector of p covariates for patient i . The expected number of failures at hospital j is then given by

$$E_j = \mathbb{E} \left[\sum_{i=1}^{n_j} X_{j,i} \right] = \sum_{i=1}^{n_j} p_i = \sum_{i=1}^{n_j} \frac{1}{1 + e^{-\beta_0 + \mathbf{Z}_i \beta}}.$$

Let O_j be the observed number of failures at hospital j , the risk-adjusted proportion of failures at hospital j is as follows:

$$\gamma_j^{\text{RA}} = \frac{O_j}{E_j} \cdot p_0.$$

The quantity γ_j^{RA} is then used in Equation (3) instead of γ_j .

The funnel plot can be used to compare hospital performance over a fixed time period. The funnel plot is often used for monitoring the quality of a process by repeatedly constructing funnel plots over different time intervals. We advocate against such an inspection scheme, as it introduces an increased probability of a type I error due to multiple testing. We recommend to only use the funnel plot as a graphical tool to visually inspect the proportion of failures at all hospitals over a time frame. The funnel plot is a discrete time method and can therefore only be used to compare overall performance over a time span. To determine whether a hospital was performing poorly during the time of interest, one of the following CUSUM charts should be used.

2.4. Bernoulli cumulative sum (CUSUM) chart

The Bernoulli CUSUM chart (Steiner *et al.* 2000) can be used to sequentially test whether the failure rate of patients at a single hospital has changed starting from some patient $\nu \geq 1$. Consider a hospital with patients $i = 1, \dots, \nu, \dots$ and a binary outcome:

$$X_i = \begin{cases} 1, & \text{if patient } i \text{ had an undesirable outcome within } C \text{ days;} \\ 0, & \text{if patient } i \text{ had a desirable outcome within } C \text{ days;} \end{cases}$$

where $X_i \sim \text{Ber}(p_i)$ with p_i the failure probability within C days for patient i . The Bernoulli CUSUM can be used to test the hypotheses of an increased failure rate starting from some patient ν :

$$H_0 : X_1, X_2, \dots \sim \text{Ber}(p_0) \quad H_1 : \begin{matrix} X_1, \dots, X_{\nu-1} \sim \text{Ber}(p_0) \\ X_{\nu}, X_{\nu+1}, \dots \sim \text{Ber}(p_1) \end{matrix} \quad (5)$$

where $\nu \geq 1$ is not known in advance, $p_0 < p_1$, and patient outcomes are ordered according to the time of entry into the study S_i .

The Bernoulli CUSUM statistic is the likelihood ratio test associated with the hypotheses in (5). The chart is then given by:

$$S_n = \max(0, S_{n-1} + W_n) \quad (6)$$

with $W_n = X_n \ln \left(\frac{p_1(1-p_0)}{p_0(1-p_1)} \right) + \ln \left(\frac{1-p_1}{1-p_0} \right)$. Alternatively, it is possible to reformulate the chart in terms of the Odds Ratio $OR = \frac{p_1(1-p_0)}{p_0(1-p_1)} =: e^\theta$. In that case, $W_n = X_n \ln \left(e^\theta \right) + \ln \left(\frac{1-p_1}{1-p_0+e^\theta p_0} \right)$. The null hypothesis is rejected when the value of the chart exceeds a control limit h .

A risk-adjusted procedure may be performed by modelling patient-specific failure probability $(p_{0,i})$ using a logistic regression model similarly to (4). The risk-adjusted Bernoulli CUSUM can be used as a sequential quality control method for binary outcomes. When survival outcomes are available a proper method is lacking. Dichotomizing the outcome (survival time) can lead to delays in detection. When survival outcomes are available, it is recommended to construct one of the CUSUM charts described in Sections 2.5 - 2.6.

2.5. Biswas and Kalbfleisch CUSUM (BK-CUSUM)

The BK-CUSUM chart can be used to continuously test whether the failure rate of patients at the hospital has changed at some point in time. Consider a hospital with patients $i = 1, 2, \dots$ and assume that the patient specific hazard rate is given by $h_i(x) = h_0(x)e^{\mathbf{Z}_i\beta}$. The BK-CUSUM chart can be used to test the hypotheses that the baseline hazard rate of all active patients has increased from $h_0(x)$ to $h_0(x)e^{\theta_1}$ at some point in time $s > 0$ after the start of the study:

$$H_0 : X_i \sim \Lambda_i(t), i = 1, 2, \dots \quad H_1 : \begin{array}{l} X_i \sim \Lambda_i(t) | t < s, i = 1, 2, \dots \\ X_i \sim \Lambda_i^{\theta_1}(t) | t \geq s, i = 1, 2, \dots \end{array} \quad (7)$$

where $\theta_1 > 0$ is the user's estimate of the true hazard ratio θ and $s > 0$ is the unknown time of change in hazard rate.

The chart displays the likelihood ratio test for the test of hypotheses (7) and is given by:

$$BK(t) = \max_{0 \leq s \leq t} \left\{ \theta_1 N(s, t) - (e^{\theta_1} - 1) \Lambda(s, t) \right\}, \quad (8)$$

where the estimated hazard ratio $e^{\theta_1} > 1$ has to be prespecified and

$$N(s, t) = N(t) - N(s) \text{ and } \Lambda(s, t) = \Lambda(t) - \Lambda(s). \quad (9)$$

The null hypothesis is rejected when the value of the chart exceeds the control limit.

The BK-CUSUM chart can lead to faster detection speeds than the Bernoulli CUSUM chart as the hypotheses can be tested at any point in time, rather than just at the (dichotomized) times of outcome. Unfortunately the chart requires users to specify θ_1 to estimate θ , which is not known a priori in most practical applications. Misspecifying this parameter can lead to large delays in detection (Gomon *et al.* 2022).

2.6. Continuous time Generalized Rapid response CUSUM (CGR-CUSUM)

The CGR-CUSUM chart can be used to test the following hypotheses:

$$H_0 : X_i \sim \Lambda_i, i = 1, 2, \dots \quad H_1 : \begin{array}{l} X_i \sim \Lambda_i, i = 1, 2, \dots, \nu - 1 \\ X_i \sim \Lambda_i^\theta, i = \nu, \nu + 1, \dots \end{array} \quad (10)$$

where e^θ and ν do not need to be prespecified. The CGR-CUSUM chart is then given by

$$CGR(t) = \max_{1 \leq \nu \leq n} \left\{ \hat{\theta}_{\geq \nu}(t) N_{\geq \nu}(t) - \left(\exp \left(\hat{\theta}_{\geq \nu}(t) \right) - 1 \right) \Lambda_{\geq \nu}(t) \right\}, \quad (11)$$

where the subscript “ $\geq \nu$ ” stands for all subjects after the potential change point ν :

$$\begin{aligned} N_{\geq \nu}(t) &= \sum_{i \geq \nu} N_i(t) \quad \text{and} \quad \hat{\theta}_{\geq \nu}(t) = \max \left(0, \log \left(\frac{N_{\geq \nu}(t)}{\Lambda_{\geq \nu}(t)} \right) \right). \end{aligned} \quad (12)$$

The null hypothesis is rejected when the value of the chart exceeds the control limit.

In contrast to the BK-CUSUM where an estimate of e^θ had to be specified in advance, the CGR-CUSUM uses the maximum likelihood estimate \hat{e}^θ to estimate the true hazard ratio e^θ from the data. This means that when e^{θ_1} is misspecified in the BK-CUSUM, the CGR-CUSUM can lead to quicker detections. In practice the real hazard ratio e^θ is never known in advance and may vary over time. The maximum likelihood estimator can alleviate this problem, therefore the CGR-CUSUM is generally the preferred chart. The major difference between the two charts is that the BK-CUSUM tests for a sudden change in the failure rate of all patients currently at risk of failure, while the CGR-CUSUM tests for a sudden change in the failure rate of all patients currently at risk of failure who have entered the hospital after a certain time. A drawback of the CGR-CUSUM is that the computation of the MLE $\hat{\theta}$ requires sufficient information (in the form of survival times/failures) to converge to the true value. This means that the chart can be unstable at the beginning of the study and might not provide reliable values for hospitals with low volumes of patients.

2.7. Choosing control limits

For the funnel plot in Section 2.3, it is sufficient to choose a confidence level to determine which hospitals are performing worse/better than the baseline. For the CUSUM charts, instead, it is necessary to choose a control limit h , so that a signal is produced when the value of the chart exceeds h . The most common ways to choose this control limit are to either restrict the in-control average run length (ARL) of the chart, or to restrict the type I error over a certain time period. With the first method, one could choose to restrict the in-control ARL to approximately 5 years, so that on average we would expect a hospital which performs according to the target to produce a false signal (detection) once every 5 years. Using the second method, one could choose the control limit such that at most a proportion α of the in-control hospitals yields a signal (false detection) within a period of 5 years.

The R package **success** contains functions to determine control limits by restricting the type I error probability α over a chosen time frame.

3. The R package success

The package **success** can be used by both laymen and experts in the field of quality control charts. In Section 3.1 we describe the general data structure to be used for constructing control charts by means of an example data set. In Section 3.2 we describe the `parameter_assist()` function which can be used to determine all necessary parameters for the construction of control charts for the user not interested in technical details. In Sections 3.3 - 3.7 we present

the functions that can be used to compute the control charts described in Section 2.

3.1. Input data

All methods in this package require the user to supply a `data.frame` for the construction of control charts as well as the determination of a benchmark failure rate. We show how to use the **success** package by means of an enclosed data set.

Example data set

The data frame `surgerdat` contains 32529 survival times, censoring times and covariates (age, BMI and sex) of patients from 45 hospitals with 15 small, medium and large hospitals (0.5, 1 and 1.5 patients per day). Patients enter the hospitals for 2 years (730 days) after the start of the study. Survival times were generated using a risk-adjusted Cox proportional hazards model following the procedure in Austin (2012) with coefficient vector $\beta = (\text{age} = 0.003, \text{BMI} = 0.02, \text{sexmale} = 0.2)$ and exponential baseline hazard rate $h_0(t, \lambda = 0.01)e^\theta$. Hospitals are numbered from 1 to 45 with hazard ratio θ sampled from a normal distribution with mean 0 and standard deviation 0.4. This means that some hospitals are performing better than baseline ($\theta < 0$) and some are performing worse ($\theta > 0$).

The `surgerdat` data set can be loaded using:

```
R> library(success)
R> data("surgerdat", package = "success")
R> head(surgerdat)
```

	entrytime	survtime	ensorid	unit	exptheta	psival	age	sex	BMI
1	5	21	0	1	1.887204	0.5	70	male	29.52
2	9	19	1	1	1.887204	0.5	68	male	24.06
3	10	64	1	1	1.887204	0.5	101	female	20.72
4	20	64	1	1	1.887204	0.5	67	female	24.72
5	21	0	1	1	1.887204	0.5	44	male	27.15
6	24	64	1	1	1.887204	0.5	54	male	26.28

Each row represents a patient. The first 2 columns (`entrytime` and `survtime`) are crucial for the construction of control charts. These columns have to be present in the data. The time scale of `entrytime` and `survtime` must be the same. They can both only be supplied in a "numeric" format (dates are not allowed). In the example data, the time scale is in days; `entrytime` is the number of days since the start of the study and `survtime` is the survival time since the time of entry. The column `ensorid` is a censoring indicator, with 0 indicating right-censoring and 1 that the event has occurred. If `ensorid` is missing, a column of 1's will automatically be created, indicating that no observations were right-censored. The column `unit` (indicating the hospital number) is required for the construction of a funnel plot, but not for the CUSUM charts. This is because CUSUM charts are constructed separately for every unit, requiring the user to manually subset the data for each unit. The columns `exptheta` and `psival` indicate the parameters e^θ and ψ used to create the simulated data set. The last three columns `age`, `sex` and `BMI` are the covariates of each individual. In a user supplied

`data.frame` these can of course take any desired name.

3.2. The `parameter-assist` function

Readers not interested in technical details can use the `parameter_assist()` function to determine most of the parameters required for constructing the control charts in this package. The `parameter_assist()` function can be used to determine control limits for all control charts described in Section 2. The function guides users through the following 3 steps:

Step 1: Specify arguments to `parameter_assist()`.

Step 2: Determine control limit(s) for the required control chart(s) by feeding the output of Step 1 to one of the `*_control_limit()` functions.

Step 3: Construct the desired control chart by feeding the output of Step 1 to the function, as well as the control limit from Step 2.

Step 2 can be skipped for the `funnel_plot()` function, as funnel plots do not require control limits.

The `parameter_assist()` function has the following arguments:

- **baseline_data (required):** a `data.frame` in the format described in Section 3.1. It should contain the data to be used to determine the target performance metric (both for discrete and continuous time charts);
- **data (required):** a `data.frame` in the format described in Section 3.1. It should contain the data used to construct control charts. For example: the data from one hospital;
- **formula (optional):** a formula indicating in which way the linear predictor of the risk-adjustment models should be constructed. Only the right-hand side of the formula will be used. If left empty, no risk-adjustment procedure will be performed;
- **followup (required for discrete time charts):** a numeric value which has to be in the same unit as `entrytime` and `survtime` and should specify how long after `entrytime` we consider the binary outcome (failure/non-failure) of the patient. This argument can be left empty when the user does not want to construct discrete time charts;
- **theta (recommended for Bernoulli and BK-CUSUM):** the expected increase in the log-odds of failure/hazard rate. Default is $\log(2)$, meaning the goal will be a detection of a doubling in the failure rate;
- **time (recommended):** time interval for type I error to be determined. Default is the largest `entrytime` of subjects in `baseline_data`;
- **alpha (recommended):** required type I error to control over the time frame specified in `time`. Default is 0.05.

The first two arguments should always be specified. Depending on the number of additional arguments specified, different functions in the package can be used. Most arguments have default values, but these may not always be suitable for the desired inspection scheme. Risk-adjusted procedures can only be constructed if at least `formula` is specified.

An example where all arguments are specified is provided below, but specifying only `baseline_data` and `data` is sufficient to construct a CGR-CUSUM without risk-adjustment.

```
R> assisted_parameters <- parameter_assist(
+   baseline_data = subset(surgerydat, entrytime < 365),
+   data = subset(surgerydat, entrytime >= 365 & unit == 1),
+   formula = formula("~age + sex + BMI"),
+   followup = 30,
+   theta = log(2),
+   time = 365,
+   alpha = 0.05)
```

We use data on patients arriving in the first year (`entrytime < 365`) to determine the target performance measure. We then construct control charts on the first hospital (`unit == 1`) using information on all patients arriving after the first year. Risk-adjustment is performed using the 3 available covariates. We choose to consider patient followup 30 days after surgery and want to detect a doubling of failure rate. The last two arguments were chosen such that the type I error of the procedure is restricted to 0.05 within 1 year (on average 1 in 20 hospitals performing according to baseline will be detected within 1 year).

The `parameter_assist()` then returns a list of arguments to supply to other functions in this package:

```
R> names(assisted_parameters)

[1] "call"           "data"           "baseline_data"  "glmmod"
[5] "coxphmod"       "theta"          "psi"            "time"
[9] "alpha"          "maxtheta"       "followup"       "p0"
```

The user can manually feed the determined parameters to other functions in this package. Conversely, it is possible to feed the output of the `parameter_assist()` function to the following functions directly:

- `funnel_plot()`
- `bernoulli_control_limit()` and `bernoulli_cusum()`
- `bk_control_limit()` and `bk_cusum()`
- `cgr_control_limit()` and `cgr_cusum()`

Step 1 was performed in the code above. For Step 2 we feed the output of `parameter_assist()` to determine control limits for the Bernoulli, BK- and CGR-CUSUM charts.




Figure 1: Bernoulli, BK- and CGR-CUSUM charts for hospital 1 in the surgery data set starting from 1 year after the start of the study.

```
R> bernoulli_control <- bernoulli_control_limit(assist = assisted_parameters)
R> bk_control <- bk_control_limit(assist = assisted_parameters)
R> cgr_control <- cgr_control_limit(assist = assisted_parameters)
```

The determined control limits can then be fed to the control chart functions to finish Step 3.

```
R> bernoulli_assist <- bernoulli_cusum(assist = assisted_parameters,
+                                   h = bernoulli_control$h)
R> bk_assist <- bk_cusum(assist = assisted_parameters, h = bk_control$h)
R> cgr_assist <- cgr_cusum(assist = assisted_parameters, h = cgr_control$h)
```

We plot the control charts using the `plot()` function (see Figure 1).

```
R> plot(bernoulli_assist) + ggtitle("Bernoulli CUSUM of hospital 1") +
+   ylim(c(0, 6.5))
R> plot(bk_assist) + ggtitle("BK-CUSUM of hospital 1") + ylim(c(0, 6.5))
R> plot(cgr_assist) + ggtitle("CGR-CUSUM of hospital 1") + ylim(c(0, 6.5))
```

The run length of the charts (time until control limit is reached) can then be determined using the `runlength()` function.

```
R> runlength(chart = bernoulli_assist, h = bernoulli_control$h)
```

```
[1] 83
```

```
R> runlength(chart = bk_assist, h = bk_control$h)
```

```
[1] 46
```

```
R> runlength(chart = cgr_assist, h = cgr_assist$h)
```

```
[1] 48
```




Figure 2: Funnel plot for hospitals in the surgery data set starting from 1 year after the start of the study.

Note that the run length of the charts are determined from the smallest time of entry of subjects into specified `data`. The study therefore starts at the moment the first subject has a surgery (in this case, at day 375).

The funnel plot does not require control limits, therefore steps 2 and 3 can be skipped.

```
R> funnel_assist <- funnel_plot(assist = assisted_parameters)
R> plot(funnel_assist) + ggtitle("Funnel plot of surgerydat")
```

The resulting plot can be seen in Figure 2. The blue dashed lines indicate the confidence intervals. Each dot represents a hospital, with the colour representing the confidence level at which this hospital would be signalled.

3.3. Manual risk-adjustment

Risk-adjustment models should be estimated on a data set known to have in-control failures, as this allows the coefficients to be determined as precisely as possible. In real life applications it is not known in advance which hospitals have had in-control failures. It is therefore common practice to use all available data to determine risk-adjustment models.

We consider a logistic model to use for risk-adjustment in the discrete time methods (funnel plot and Bernoulli CUSUM), using all available data of patients with surgeries in the first year of the study. We use 30 days mortality as outcome for these charts.

```
R> baseline_data <- subset(surgerydat, entrytime <= 365)
R> followup <- 30
R> formula_glm <- as.formula(
+   "(survtime <= followup) & (censorid == 1) ~ age + sex + BMI")
R> glm_risk_model <- glm(formula_glm, data = baseline_data, family = binomial)
```

Then we estimate a Cox proportional hazards model to use for risk-adjustment in the continuous time BK- and CGR-CUSUM, using the same baseline data. For this we use the

function `Surv()` and `coxph()` from the package **survival** (Terry M. Therneau and Patricia M. Grambsch 2000).

```
R> require(survival)
R> formula_coxph <- as.formula("Surv(survtime, censorid) ~ age + sex + BMI")
R> coxph_risk_model <- coxph(formula_coxph, data = baseline_data)
```

Conversely, we can manually specify a risk-adjustment model:

```
R> RA_manual <- list(formula = formula("~ age + sex + BMI"),
+                    coefficients = c(age = 0.003, BMI = 0.02,
+                                    sexmale = 0.2))
```

This is useful for users who do not want to use the packages **stats** and **survival** for the estimation of the models.

3.4. The funnel plot function

The `funnel_plot()` function can be used to construct the funnel plot described in Section 2.3. The code below constructs a funnel plot over the first 1 year (`ctime = 365`) on the simulated data set. By not specifying `ctime` the funnel plot is constructed over all data (2 years). By leaving the parameter `p0` empty, the average failure proportion within 30 days is used as baseline failure probability.

```
R> funnel <- funnel_plot(data = surgerydat, ctime = 365,
+                        glmmod = glm_risk_model, followup = 30)
```

The function creates an object of class ‘`funnelplot`’. The `plot()` function can be used on classes in the **success** package. This creates a ‘`gg`’ object, whose graphical parameters can further be edited by the user.

```
R> plot(funnel) + ggtitle("Funnel plot of surgery data at 1 year")
```

The resulting plot can be seen in Figure 3. By default, the `funnel_plot()` function will display 95% and 99% confidence levels. This can be changed using the `conflev` argument.

A summary of the funnel plot can be obtained by using the `summary()` function.

```
R> head(summary(funnel))
```

	unit	observed	expected	numtotal	p	0.95	0.99
1	1	105	74.49454	155	0.6724872	worse	worse
2	2	71	80.60038	168	0.4202816	in-control	in-control
3	3	50	68.25903	143	0.3494854	better	better
4	4	76	80.64515	167	0.4496292	in-control	in-control
5	5	70	87.22662	182	0.3828848	better	better
6	6	210	166.23430	348	0.6027230	worse	worse




Figure 3: Funnel plots of first year of simulated data.

The resulting table displays all the relevant statistics for the funnel plot.

3.5. The Bernoulli CUSUM function

The `bernoulli_cusum()` function can be used to construct the Bernoulli CUSUM detailed in Section 2.4. The Bernoulli CUSUM uses the same dichotomized outcome as the funnel plot. For this reason, the syntax of `bernoulli_cusum()` is quite similar to that of `funnel_plot()`. In this section we will construct a Bernoulli CUSUM for the ninth hospital in the simulated data set, again using 30 day post operative survival as outcome and aiming to detect an increase of the odds ratio to 2.

Determining control limits

The Bernoulli CUSUM produces a signal when the value of the chart exceeds a value h called the control limit. The `bernoulli_control_limit()` function can be used to determine a control limit such that the type I error of the Bernoulli CUSUM procedure is restricted over some desired time frame. Suppose we want to restrict the type I error of the procedure to 0.05 over the time frame of 1 year at a hospital with an average of 1 patient per day undergoing surgery. We determine the control limit as follows:

```
R> bern_control <- bernoulli_control_limit(
+   time = 365, alpha = 0.05, followup = 30, psi = 1,
+   glmmod = glm_risk_model, baseline_data = surgerydat, theta = log(2))
```

The determined control limit h can then be retrieved by:

```
R> bern_control$h
```

```
[1] 5.56
```

By default, the control limit is determined using 200 simulated units (hospitals). If higher precision is required, the amount of simulated units can be changed using the argument `n_sim`.




Figure 4: Bernoulli CUSUM of first year of simulated data for hospital 9.

The control limit will be determined up to 2 significant digits. This can be changed using the `h_precision` argument.

Constructing the chart

After determining the control limit, we can construct the control chart. The `bernoulli_cusum()` function requires the user to specify one of the following combinations of parameters:

- a) `glmmod` & `theta`
- b) `p0` & `theta`
- c) `p0` & `p1`

Only the first option allows for risk-adjustment. The difference between these parametrisations is described in Section 2.4. We construct the Bernoulli CUSUM on the data of the ninth simulated hospital, aiming to detect whether the odds ratio of failure for patients is 2 (`theta = log(2)`). Using the `plot()` function we obtain a 'gg' object (see Figure 4).

```
R> followup <- 30
R> Bernoulli <- bernoulli_cusum(
+   data = subset(surgerydat, unit == 9),
+   glmmod = glm_risk_model, followup = followup, theta = log(2))
R> plot(Bernoulli, h = bern_control$h) +
+   ggtitle("Bernoulli CUSUM of Hospital 9")
```

The run length of the chart, in Figure 4 visible as the time at which the chart first crosses the red line, can be found using the `runlength()` function:

```
R> runlength(Bernoulli, h = bern_control$h)
```

```
[1] 211
```


3.6. The BK-CUSUM function

The `bk_cusum()` function can be used to construct the BK-CUSUM chart presented in Section 2.5. The chart is no longer constructed using dichotomized outcomes, therefore leading to faster detections on survival data than discrete time methods.

Determining control limits

The BK-CUSUM produces a signal when the value of the chart exceeds a value h called the control limit. The `bk_control_limit()` function can be used to determine a control limit such that the type I error of the BK-CUSUM procedure is restricted over some desired time frame. Suppose we want to restrict the type I error of the procedure to 0.05 over the time frame of 1 year for a hospital with an average of 1 patient per day undergoing surgery. The control limit is determined as follows:

```
R> BK_control <- bk_control_limit(
+   time = 365, alpha = 0.05, psi = 1, coxphmod = coxph_risk_model,
+   baseline_data = surgerydat, theta = log(2))
```

The determined control limit h can then be retrieved by:

```
R> BK_control$h
```

```
[1] 7.13
```

Constructing the chart

We construct the BK-CUSUM on the data of the ninth hospital aiming to detect a doubling in the hazard rate of patients ($\theta = \log(2)$).

```
R> BK <- bk_cusum(data = subset(surgerydat, unit == 9), theta = log(2),
+                 coxphmod = coxph_risk_model)
R> plot(BK, h = BK_control$h) + ggtitle("BK-CUSUM of hospital 9")
```

The resulting plot is presented in Figure 5. The run length of the chart, in Figure 5 visible as the time at which the chart first crosses the red line, can be found using the `runlength()` function:

```
R> runlength(BK, h = BK_control$h)
```

```
[1] 180
```

When the cumulative baseline hazard is not specified through the argument `cbaseh`, but a Cox Risk-adjustment model `coxphmod` as obtained from `survival::coxph()` is provided, the cumulative baseline hazard will automatically be determined from this Cox model. When the argument `ctimes` is left empty, the chart will only be determined at the times of patient failures, as this is sufficient for detection purposes and saves computation time. A control




Figure 5: BK-CUSUM of the first year of simulated data from hospital 9.

limit `h` can be specified, so that the chart is only constructed until the value of the chart exceeds the value of the control limit. This is very convenient when monitoring the quality of care at a hospital. Sometimes it is desirable to only construct the chart up until a certain time point, for this the argument `stoptime` can be used. The argument `C` can be used to only consider patient outcomes up until C time units after their surgery. [Biswas and Kalbfleisch \(2008\)](#) use $C = 365$, considering patient outcomes only until 1 year post surgery. Finally, a progress bar can be added using the argument `pb`.

Suppose a decrease in the failure rate is of interest, we can then construct a lower-sided BK-CUSUM (see left side of Figure 6) by specifying a `theta` value smaller than 0.

```
R> BK_control_lower <- bk_control_limit(
+   time = 365, alpha = 0.05, psi = 1, coxphmod = coxph_risk_model,
+   baseline_data = surgerydat, theta = -log(2))

R> BK_control_lower <- bk_control_limit(
+   time = 365, alpha = 0.05, psi = 1, coxphmod = coxph_risk_model,
+   baseline_data = surgerydat, theta = -log(2))
R> BKlower <- bk_cusum(data = subset(surgerydat, unit == 9),
+   theta = -log(2), coxphmod = coxph_risk_model)
R> plot(BKlower, h = BK_control_lower$h) +
+   ggtitle("BK-CUSUM of hospital 9 (lower sided)")
```

Similarly, when both an increase and decrease of the failure rate are of interest the argument `twosided = TRUE` can be used. This produces a two-sided BK-CUSUM (see right side of Figure 6). For the lower-sided BK-CUSUM the control limit must be determined separately.

```
R> BKtwosided <- bk_cusum(data = subset(surgerydat, unit == 9),
+   theta = log(2), coxphmod = coxph_risk_model, twosided = TRUE)
R> plot(BKtwosided, h = c(BK_control_lower$h, BK_control$h)) +
+   ggtitle("BK-CUSUM of hospital 9 (two-sided)")
```




Figure 6: BK-CUSUM of the first year of simulated data. Lower sided (left) and two-sided (right).

3.7. The CGR-CUSUM function

The `cgr_cusum()` function can be used to construct the CGR-CUSUM detailed in Section 2.6. This function has almost the same syntax as the `bk_cusum()` function. The difference is that the procedure estimates a suitable value for `theta` through maximum likelihood estimation, instead of requiring the users to specify such a value a priori.

The maximum likelihood estimate in the CGR-CUSUM can be unstable at early time points, when not much information is available about subject failure. For this reason, the value of the maximum likelihood estimate is restricted to $e^{\theta} \leq 6$ by default. This comes down to believing that the true hazard ratio at any hospital is always smaller or equal than 6 times the baseline. To change this belief, the user can supply the `maxtheta` parameter to the `cgr_cusum()` and `cgr_control_limit()` functions.

Determining control limits

Similarly to the `bk_control_limit()` function used for the BK-CUSUM, the `cgr_control_limit()` function can be used to determine the control limit for the CGR-CUSUM chart as follows:

```
R> CGR_control <- cgr_control_limit(
+   time = 365, alpha = 0.05, psi = 1, coxphmod = coxph_risk_model,
+   baseline_data = surgerydat)
```

The determined control limit h can then be retrieved by:

```
R> CGR_control$h
```

```
[1] 8.52
```

By default the control limit for the CGR-chart is determined on only 20 simulated samples (due to the computational intensity of the procedure), but we recommend to increase the




Figure 7: CGR-CUSUM of the first year of simulated data from hospital 9.

number of samples by using the argument `n_sim` to get a more accurate control limit.

Constructing the chart

The CGR-CUSUM for the ninth hospital is created with:

```
R> CGR <- cgr_cusum(data = subset(surgerydat, unit == 9),
+                   coxphmod = coxph_risk_model)
R> plot(CGR, h = CGR_control$h) + ggtitle("CGR-CUSUM of hospital 9")
```

The resulting plot can be seen in Figure 7. To determine the run length of the procedure (time at which chart crosses the control limit), use the function `runlength()`.

```
R> runlength(CGR, h = CGR_control$h)
```

```
[1] 214
```

To construct the control chart up until the time of detection, the parameter `h` can be specified in the `cgr_cusum()` function. This allows for continuous inspection, as well as reducing computation time.

Since the CGR-CUSUM is time consuming when the value has to be computed at many time points, we recommend to leave `ctimes` unspecified, so that the CGR-CUSUM will only be determined at the times necessary for detection purposes.

To reduce computing time, we allow users to parallelize the computations across multiple cores. This can be easily done through the `ncores` argument. The calculation of the CGR-CUSUM proceeds through 2 steps. First, the contributions to the cumulative intensity of each subject are determined at every time point of interest and are stored in a matrix. Afterwards, the value of the chart is computed by performing matrix operations. When `ncores` > 1, both steps are automatically parallelized using functions from the **pbapply** package by [Solymos and Zawadzki \(2021\)](#). When a value for the control limit has been specified, only the first step can be parallelized. For small data sets and/or short runs `cmethode` = "CPU" can be chosen,

thereby recalculating the value of the chart at every desired time point but not requiring a lot of initialization. For small hospitals and/or short detection times it could be the preferred method of construction. As the CGR-CUSUM can take long to construct, it is recommended to display a progress bar by specifying `pb = TRUE`.

```
R> CGR_multicore <- cgr_cusum(data = subset(surgerydat, unit == 9),
+                             coxphmod = coxph_risk_model, ncores = 3, pb = TRUE)
```

3.8. The interactive plot function

The `interactive_plot()` function can be used to plot multiple CUSUM charts together in one figure, while allowing the user to interact with the plot. This is achieved by using the package **plotly** by Sievert (2020). We show how to use these features by plotting some of the CUSUM charts from the previous Sections together in one figure. We first combine all CUSUM charts into a list, together with the control limits.

```
R> Bernoulli$h <- bernoulli_control$h
R> BK$h <- BK_control$h
R> CGR$h <- CGR_control$h
R> cusum_list <- list(Bernoulli, BK, CGR)
R> interactive_plot(cusum_list, unit_names = rep("Hosp 9", 3))
```

The resulting plot can be seen on the left side of Figure 8. As each chart has a different control limit, the separate control limits will not be displayed. We can scale each chart with respect to their control limit by choosing `scale = TRUE` to obtain the right side of Figure 8.

```
R> interactive_plot(cusum_list, unit_names = rep("Hosp 9", 3), scale = TRUE)
```

After scaling, the control limit will be $h = 1$ for all CUSUM charts. By choosing `highlight = TRUE`, the user can highlight CUSUM charts by hovering over them. The **plotly** package allows for many many interactive capabilities with the plot.

4. Application

In Section 3 we employed a simulated data set to show how to use the **success** package. In this Section, we illustrate the use of the **success** package on a data set based on a clinical trial for breast cancer conducted by the European Organisation for Research and Treatment of Cancer (EORTC). Covariates for 2663 patients over 15 treatment centres are available, with patients having surgery over a span of 61 time units. In addition, the chronological time of surgery and time since surgery until a combined endpoint are known.

To analyse the data with the **success** package, we first arrange them in the format presented in Section 3. The outcome of interest is event-free survival. For patients who did not experience an event during the study period, the observations were censored at the last time the patients were known to be event-free. We consider the start of the study as the time when the first patient had surgery. The resulting data was stored in a `data.frame` called `breast`, which can be loaded as follows:




Figure 8: Combined plot for Bernoulli, BK- and CGR-CUSUM charts of hospital 9. Original values (left) and scaled with respect to control limits (right).

```
R> data(breast, package = "success")
```

To determine the risk-adjustment models, we consider 36 time units post surgery followup as outcome. We fit a logistic model to be used for risk adjustment in the funnel plot and Bernoulli CUSUM, and a Cox model for risk adjustment in the BK- and CGR-CUSUM:

```
R> glmmodEORTC <- glm((survtime <= 36) & (censorid == 1) ~ var1 + var2 +
+                   var3 + var4 + var5 + var6 + var7,
+                   data = breast, family = binomial)
R> phmodEORTC <- coxph(Surv(survtime, censorid) ~ var1 + var2 +
+                   var3 + var4 + var5 + var6 + var7, data = breast)
```

We then construct the funnel plot and Bernoulli, BK- and CGR-CUSUM for each of the 15 centres in the data. To make the continuous time charts visually interesting, we determine their values at every time unit from the start of the study using the argument `ctimes`.

```
R> EORTC_funnel <- funnel_plot(data = breast, glmmod = glmmodEORTC,
+                             followup = 36)
R> EORTC_charts <- vector("list", length = 15)
R> for(i in 1:15){
+   max_time <- ceiling(max(subset(breast, unit == i)$entrytime +
+                             subset(breast, unit == i)$survtime))
+   EORTC_charts[[i]]$ber <- bernoulli_cusum(data = subset(breast, unit == i),
+                                             followup = 36, glmmod = glmmodEORTC,
+                                             theta = log(2))
+   EORTC_charts[[i]]$bk <- bk_cusum(data = subset(breast, unit == i),
+                                     coxphmod = phmodEORTC,
+                                     ctimes = seq(0, max_time, 1), theta = log(2))
+ }
```

```
+ EORTC_charts[[i]]$cgr <- cgr_cusum(data = subset(breast, unit == i),
+                                     coxphmod = phmodEORTC,
+                                     ctimes = seq(0, max_time, 1))
+ }
```

We then estimate the arrival rate for each center in the data using the `arrival_rate()` function.

```
R> arrival_rate <- arrival_rate(breast)
R> arrival_rate
```

1	2	3	5	4	6	7
0.1568693	0.7127849	0.7922508	0.9657104	0.9676291	1.3067964	1.3870253
8	9	10	11	12	13	14
1.4589466	1.6650555	3.0589682	3.6197173	6.8752701	7.5733827	11.0912301
15						
18.2107083						

Based on the estimated arrival rate $\hat{\psi}$ (per time unit), we group the centres into 3 categories:

Small: Centres 1-5: $\hat{\psi} \approx 0.9$;

Medium: Centres 6-11: $\hat{\psi} \approx 2.1$;

Large: Centres 12-15: $\hat{\psi} \approx 11$;

For each category, we determine the control limits to use in the CUSUM charts, using the `*_control_limit()` functions. For this, we restrict the simulated type I error over 60 time units to 0.05.

```
R> h <- matrix(0, nrow = 3, ncol = 3,
+             dimnames = list(c(0.9, 2.1, 11), c("Ber", "BK", "CGR")))
R> psi <- c(0.9, 2.1, 11)
R> for(i in 1:3){
+   h[i,1] <- bernoulli_control_limit(time = 60, alpha = 0.05, followup = 36,
+                                     psi = psi[i], n_sim = 300, theta = log(2), glmmmod = glmmmodEORTC,
+                                     baseline_data = breast)$h
+   h[i,2] <- bk_control_limit(time = 60, alpha = 0.05, psi = psi[i], n_sim = 300,
+                               theta = log(2), coxphmod = phmodEORTC, baseline_data = breast)$h
+   h[i,3] <- cgr_control_limit(time = 60, alpha = 0.05, psi = psi[i],
+                               n_sim = 300, coxphmod = phmodEORTC, baseline_data = breast)$h
+ }
R> print(h)
```

	Ber	BK	CGR
0.9	2.29	3.23	4.90
2.1	2.93	3.89	5.51
11	4.71	6.43	6.49

The columns represent the chart and the rows represent the estimated arrival rate. Using these control limits, we determine the times of detection for the 15 centres using the `runlength()` function.

```
R> times_detection <- matrix(0, nrow = 3, ncol = 15,
+                             dimnames = list(c("Ber", "BK", "CGR"), 1:15))
R> for(i in 1:5){
+   times_detection[1,i] <- runlength(EORTC_charts[[i]]$ber, h = h[1,1])
+   times_detection[2,i] <- runlength(EORTC_charts[[i]]$bk, h = h[1,2])
+   times_detection[3,i] <- runlength(EORTC_charts[[i]]$cgr, h = h[1,3])
+ }
R> for(i in 6:11){
+   times_detection[1,i] <- runlength(EORTC_charts[[i]]$ber, h = h[2,1])
+   times_detection[2,i] <- runlength(EORTC_charts[[i]]$bk, h = h[2,2])
+   times_detection[3,i] <- runlength(EORTC_charts[[i]]$cgr, h = h[2,3])
+ }
R> for(i in 12:15){
+   times_detection[1,i] <- runlength(EORTC_charts[[i]]$ber, h = h[3,1])
+   times_detection[2,i] <- runlength(EORTC_charts[[i]]$bk, h = h[3,2])
+   times_detection[3,i] <- runlength(EORTC_charts[[i]]$cgr, h = h[3,3])
+ }
```

We determine the centres which were detected by any of the charts, and compare their detection times.

```
R> ceiling(times_detection[,colSums(is.infinite(times_detection)) != 3])
```

	3	5	9	10	11	14
Ber	50	Inf	41	45	47	Inf
BK	Inf	92	Inf	21	25	127
CGR	21	112	Inf	16	23	Inf

The columns represent the centre numbers, while the rows represent the CUSUM charts. We find that the detections by the considered charts do not coincide perfectly with the centres detected by the funnel plot (10, 11) at a 5 percent significance level. Comparing the continuous time methods, we see that centre 5 is detected faster by the BK-CUSUM, while centres 10 and 11 are signaled faster by the CGR-CUSUM. Centre 14 is only detected by the BK-CUSUM while centre 9 is only detected by the Bernoulli CUSUM.

An important consideration when comparing detection times between discrete and continuous time methods is that the discrete time chart inspect the 36 time units survival probability of patients, while the continuous time charts inspect overall survival. This means that the Bernoulli CUSUM might detect a centre with high post operative failure proportions in the 36 time units after surgery. However, it does not mean that patients necessarily experience failures faster than expected at this centre as the Bernoulli CUSUM makes no distinction between a patient who has failed 1 time unit or 10 time units post treatment. This could explain why only the Bernoulli CUSUM detects centre 9.

For the continuous time charts it is important to keep in mind that multiple consecutive failures cause the BK-CUSUM to jump up by $\log(2)$ for every failure, independent of the probability of failure of the patients at that point in time. This can lead to fast detections when many failures are clustered. The CGR-CUSUM can make smaller or larger jumps, depending on the failure probability for each patient at the time of death. This could explain why only the BK-CUSUM detects Centre 14. For centres with low volumes of patients, the maximum likelihood estimate in the CGR-CUSUM might not converge quickly to an appropriate value therefore causing a delay in detection times. In contrast, a wrong choice of θ_1 in the BK-CUSUM may negatively influence detection times (Gomon *et al.* 2022).

We take a closer look at the disparities between detection times by visualising the funnel plot as well as CUSUM charts for all centres in the EORTC data.

```
R> unnames <- paste(rep("Centre", 15), 1:15)
R> ber_EORTC <- lapply(EORTC_charts, FUN = function(x) x$ber)
R> bk_EORTC <- lapply(EORTC_charts, FUN = function(x) x$bk)
R> cgr_EORTC <- lapply(EORTC_charts, FUN = function(x) x$cgr)
R> for(i in 1:5){
+   ber_EORTC[[i]]$h <- h[1,1]
+   bk_EORTC[[i]]$h <- h[1,2]
+   cgr_EORTC[[i]]$h <- h[1,3]
+ }
R> for(i in 6:11){
+   ber_EORTC[[i]]$h <- h[2,1]
+   bk_EORTC[[i]]$h <- h[2,2]
+   cgr_EORTC[[i]]$h <- h[2,3]
+ }
R> for(i in 12:15){
+   ber_EORTC[[i]]$h <- h[3,1]
+   bk_EORTC[[i]]$h <- h[3,2]
+   cgr_EORTC[[i]]$h <- h[3,3]
+ }
R> cols <- brewer.pal(n = 6, "Set2")
R> col_manual <- rep("lightgrey", 15)
R> col_manual[c(3,5,9,10,11,14)] <- cols
R> t1 <- interactive_plot(ber_EORTC, unit_names = unnames,
+   scale = TRUE, group_by = "type",
+   manual_colors = col_manual)
R> t2 <- interactive_plot(bk_EORTC, unit_names = unnames,
+   scale = TRUE, group_by = "type",
+   manual_colors = col_manual)
R> t3 <- layout(interactive_plot(cgr_EORTC, unit_names = unnames,
+   scale = TRUE, group_by = "type",
+   manual_colors = col_manual))
R> t0 <- ggplotly(plot(EORTC_funnel))
R> layout(subplot(t0, style(t1, showlegend = FALSE),
+   style(t2, showlegend = FALSE), t3, nrows = 2),
+   autosize = FALSE, width = 1000, height = 600,
```

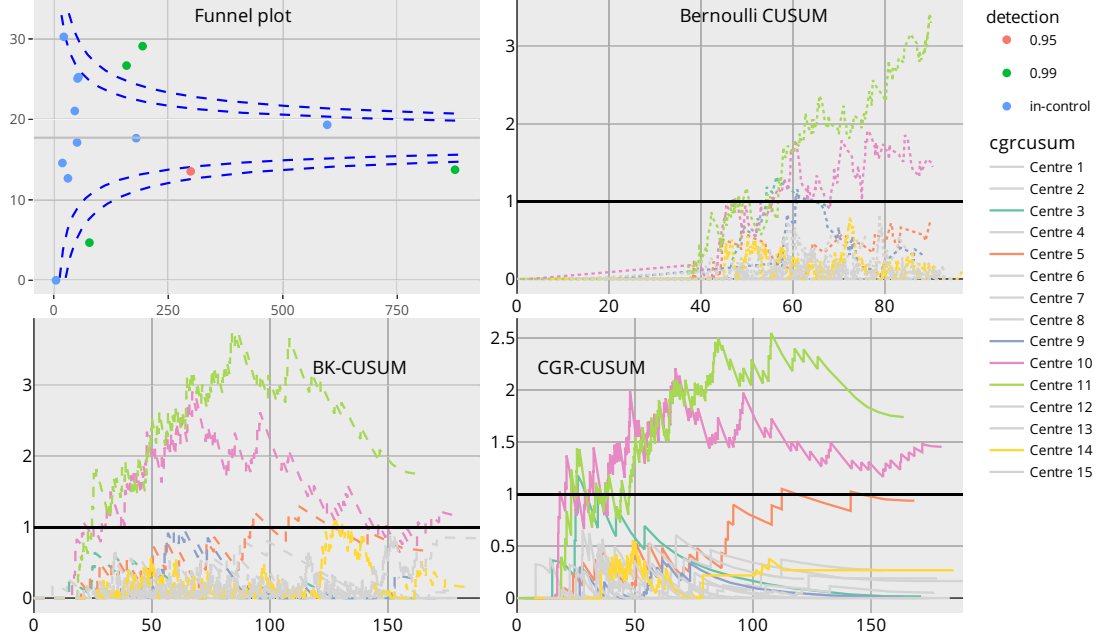


Figure 9: (Top left) Funnel plot of breast data. (Top right) Bernoulli, (Bottom left) BK- and (Bottom right) CGR-CUSUM charts of all 15 centres. Centres not detected by any of the charts are greyed out.

```
+ annotations = annot)
```

The resulting plots are shown in Figure 9. All hospitals detected during the time of the study by the CUSUM charts are highlighted, the remaining centres are shown in gray. We can clearly see the 36 time unit delay in the Bernoulli CUSUM charts. The CGR-CUSUM charts have high initial spikes when the first failures are observed. This happens due to the instability of the maximum likelihood estimate $\hat{\theta}(t)$ when only few failures have been observed. After a sufficient amount of patients have been observed, the CGR-CUSUM makes a clear distinction between centres with respect to their performance. Part of the centres retain a value close to zero while for others the value increases over time. In contrast, the BK-CUSUM charts appear to be less stable, with centres almost hitting the control limit over the period of the study multiple times. This could indicate that the value of $\theta = \ln(2)$ is not suitable for these centres. Only centres 10 (purple) and 11 (lightgreen) were detected by all charts. From the value of the continuous time charts we can presume that centre 10 had a cluster of failures at the beginning of the study, followed by a period of (slightly) above average failures. The Bernoulli CUSUM does not provide such insights, as failures come




Figure 10: Bernoulli, BK- and CGR-CUSUM of centre 5.

in 36 time units after surgery. It seems that centre 11 had a high rate of failure from the start until the end of the study. Centre 9 (dark blue) was only detected by the Bernoulli CUSUM. This disparity between discrete and continuous time charts mostly happens when 36 time units failure proportions are relatively large, but many patients fail at reasonable times (e.g. around 30 time units post surgery). From a visual inspection of the charts, this seems to be the case for Centre 9. Finally, Centre 5 was only detected by the continuous time CUSUM charts. We display all 3 CUSUM charts for Centre 5 in Figure 10. The Bernoulli CUSUM only incorporates the information provided by survival 36 time units after surgery. Because of this, the Bernoulli CUSUM can only be calculated up to 36 time units after the last patient had surgery (in this case, until the 89th time unit). The continuous time charts can incorporate failures of patients at any point in time, therefore producing signals at later times. While the BK-CUSUM always rises by $\ln(2)$ whenever a failure is observed, the CGR-CUSUM can make jumps of different sizes, depending on the risk of failure of the observed patient. This causes a disparity in the times of detection and also in interpretation of the values of the charts. As no patients had surgery later than 60 time units after the start of the study, the arrival rate after 60 time units is $\psi = 0$ for all centres, meaning detections at that point should be taken with a grain of salt.

5. Discussion

The **success** package implements three CUSUM methods for the inspection of the failure rate in survival data in continuous time and the funnel plot. Using the `parameter_assist()` function, quality control charts can also be constructed by users unfamiliar with control chart and survival theory.

We would like to highlight the different type of outcome and purpose of the control charts in this package. The funnel plot should not be used for the continuous inspection of survival data, as it can only be used to test for a difference in failure proportion at a fixed point in time. The Bernoulli CUSUM is closest to the funnel plot, as it uses the same outcome to determine chart values. The Bernoulli CUSUM is suitable for continuous inspection, but the followup time has a great impact on the resulting conclusions as well as the choice of the expected increase in hazard ratio `theta`. In contrast, the BK-CUSUM can incorporate patient failures at any point in time, but also requires the specification of `theta` a priori. The CGR-CUSUM does not have this downside, and as the increase in failure rate is never known in advance in practical applications, it can lead to quicker detection times. Finally, the CUSUM charts test different hypotheses, nuancing their interpretation even further: the Bernoulli and CGR-CUSUM charts can be used to test for a change in the failure rate starting from some patient, while the BK-CUSUM can be used to test for a sudden change in the failure rate of all patients.

Computationally, the funnel plot, Bernoulli CUSUM and BK-CUSUM do not require a lot of computational power, whereas the computation of the CGR-CUSUM is more sophisticated and can require more computation time. For this reason, we provide the user the option to parallelize the computation of this chart. A key part of CUSUM charts are their control limits, which are mostly determined using simulation studies due to the lack of analytical results. This can be done using the `*_control_limit()` functions in the **success** package. The time required to compute control limits depends on the value of the arrival rate `psi` and the proportion of failures in the data. A higher value of `psi` means more patients have to be accounted for, and a higher failure proportion means chart values need to be calculated more often. For the breast cancer data, determining control limits took approximately 10 minutes in total (using a consumer grade laptop), on a simulated sample of 300 in-control centres.

It is important to determine appropriate control limits for the inspection of survival processes using CUSUM charts. The chosen value of `psi` greatly influences the value of the control limit. Heuristically, this can be compared with the confidence levels in the funnel plot in Figure 9. As the number of outcomes in a centre increases, the confidence levels become narrower. For the CUSUM charts, this is expressed in the control limit.

6. Acknowledgements

The authors thank the European Organization for Research and Treatment of Cancer for permission to apply our methods on data based on an EORTC study. The contents of this publication and methods used are solely the responsibility of the authors and do not necessarily represent the official views of the EORTC.

References

- Anhoej J (2021). *qicharts: Quality Improvement Charts*. R package version 0.5.8, URL <https://CRAN.R-project.org/package=qicharts>.
- Austin PC (2012). “Generating survival times to simulate Cox proportional hazards models with time-varying covariates.” *Statistics in Medicine*, **31**(29), 3946–3958. doi:10.1002/sim.5452.
- Biswas P, Kalbfleisch JD (2008). “A risk-adjusted CUSUM in continuous time based on the Cox model.” *Statistics in Medicine*, **27**, 3452–3452.
- Cook D, Coory M, Webster R (2011). “Exponentially weighted moving average charts to compare observed and expected values for monitoring risk-adjusted hospital indicators.” *BMJ quality and safety*, **20**, 469–74.
- Cox DR (1972). “Regression Models and Life-Tables.” *Journal of the Royal Statistical Society. Series B (Methodological)*, **34**(2), 187–220.
- Flores M (2021). *qcr: Quality Control Review*. R package version 1.3, URL <https://CRAN.R-project.org/package=qcr>.
- Gomon D, Putter H, Nelissen RGHH, van der Pas S (2022). “CGR-CUSUM: A Continuous time Generalized Rapid Response Cumulative Sum chart.” doi:10.48550/ARXIV.2205.07618. URL <https://arxiv.org/abs/2205.07618>.
- Grey K (2018). *ggQC: Quality Control Charts for 'ggplot'*. R package version 0.0.31, URL <https://CRAN.R-project.org/package=ggQC>.
- Grigg OA (2018). “The STRAND Chart: A survival time control chart.” *Statistics in Medicine*, **38**(9), 1651–1661. doi:10.1002/sim.8065.
- Hubig L (2019). *cusum: Cumulative Sum (CUSUM) Charts for Monitoring of Hospital Performance*. R package version 0.4.1, URL <https://CRAN.R-project.org/package=cusum>.
- Keefe MJ, Loda JB, Elhabashy AE, Woodall WH (2017). “Improved implementation of the risk-adjusted Bernoulli CUSUM chart to monitor surgical outcome quality.” *International Journal for Quality in Health Care*, **29**(3), 343–348.
- Knoth S (2021). *spc: Statistical Process Control – Calculation of ARL and Other Control Chart Performance Measures*. R package version 0.6.5, URL <https://CRAN.R-project.org/package=spc>.
- Kumar M (2018). *funnelR: Funnel Plots for Proportion Data*. R package version 0.1.0, URL <https://CRAN.R-project.org/package=funnelR>.
- Putter H, Eikema D, de Wreede L, McGrath E, Sánchez-Ortega I, Saccardi R, Snowden JA, van Zwet E (2022). “Benchmarking survival outcomes: A funnel plot for survival data.” *Statistical Methods in Medical Research*, **0**(0).
- R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.

- Scrucca L (2004). “qcc: an R package for quality control charting and statistical process control.” *R News*, **4/1**, 11–17. URL <https://cran.r-project.org/doc/Rnews/>.
- Sievert C (2020). *Interactive Web-Based Data Visualization with R, plotly, and shiny*. Chapman and Hall/CRC. ISBN 9781138331457. URL <https://plotly-r.com>.
- Solymos P, Zawadzki Z (2021). *pbapply: Adding Progress Bar to '*apply' Functions*. R package version 1.5-0, URL <https://CRAN.R-project.org/package=pbapply>.
- Spiegelhalter DJ (2005). “Funnel plots for comparing institutional performance.” *Statistics in Medicine*, **24**(8), 1185–1202.
- Steiner SH, Cook RJ, Farewell V, Treasure T (2000). “Monitoring surgical performance using risk-adjusted cumulative sum charts.” *Biostatistics*, **1**(4), 441–452.
- Steiner SH, Jones M (2009). “Risk-adjusted survival time monitoring with an updating exponentially weighted moving average (EWMA) control chart.” *Statistics in Medicine*, **29**(4), 444–454.
- Terry M Therneau, Patricia M Grambsch (2000). *Modeling Survival Data: Extending the Cox Model*. Springer, New York. ISBN 0-387-98784-3.
- Wickham H (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. ISBN 978-3-319-24277-4. URL <https://ggplot2.tidyverse.org>.

Affiliation:

Daniel Gomon
Leiden University
Mathematical Institute
Niels Bohrweg 1
2333CA Leiden, the Netherlands
E-mail: d.gomon@math.leidenuniv.nl
URL: <https://github.com/d-gomon>