Published in final edited form as:

Biometrika. 2014 September 1; 101(3): 553–566. doi:10.1093/biomet/asu016.

Statistical inference methods for recurrent event processes with shape and size parameters

MEI-CHENG WANG and

Department of Biostatistics, Bloomberg School of Public Health, Johns Hopkins University, 615 N. Wolfe Street, Baltimore, Maryland 21205, U.S.A

CHIUNG-YU HUANG

Division of Biostatistics and Bioinformatics, Sidney Kimmel Comprehensive Cancer Center, Johns Hopkins University, 550 N. Broadway, Baltimore, Maryland 21205, U.S.A

MEI-CHENG WANG: mcwang@jhu.edu; CHIUNG-YU HUANG: cyhuang@jhmi.edu

Summary

This paper proposes a unified framework to characterize the rate function of a recurrent event process through shape and size parameters. In contrast to the intensity function, which is the event occurrence rate conditional on the event history, the rate function is the occurrence rate unconditional on the event history, and thus it can be interpreted as a population-averaged count of events in unit time. In this paper, shape and size parameters are introduced and used to characterize the association between the rate function $\lambda(\cdot)$ and a random variable X . Measures of association between X and $\lambda(\cdot)$ are defined via shape- and size-based coefficients. Rate-independence of X and $\lambda(\cdot)$ is studied through tests of shape-independence and size-independence, where the shape- and size-based test statistics can be used separately or in combination. These tests can be applied when X is a covariable possibly correlated with the recurrent event process through $\lambda(\cdot)$ or, in the one-sample setting, when X is the censoring time at which the observation of $N(\cdot)$ is terminated. The proposed tests are shape- and size-based, so when a null hypothesis is rejected, the test results can serve to distinguish the source of violation.

Keywords

Intensity function; Point process; Poisson process; Rate-independence; Rate function; Shape-independence; Size-independence

1. Introduction and background

Recurrent events are natural outcomes in many longitudinal studies. Conventionally, recurrent events of a subject are modelled as the realization of an underlying point process. Statistical models for analysing recurrent event data are usually formulated based on either the intensity function or the rate function of the point process. Many authors, including Gail

Supplementary material

Supplementary material available at *Biometrika* online includes the proofs of equation (1) and Theorems 2 and 3, as well as Q-Q plots for examining the null performance of the proposed tests.

et al. (1980), Prentice et al. (1981) and Andersen & Gill (1982), have considered modelling the occurrence probability of events given the preceding covariates and event history. Others, including Nelson (1988) and Lawless & Nadeau (1995), have developed nonparametric estimation methods to estimate the rate function of a recurrent event process. Lin et al. (2000) studied regression of recurrent events unconditional on event history. Readers are referred to Therneau & Grambsch (2000), Kalbfleisch & Prentice (2002) and Cook & Lawless (2007) for reviews of existing models and methods. Examples of econometric applications can be found in Lancaster (1990) and references therein.

Let $N(t)$ denote the number of recurrent events occurring before or at time t , where $t \in [0, \tau]$, with the constant $\tau > 0$ determined from the knowledge that recurrent events could potentially be observed up to time τ . Here we assume that the recurrent events are of the same type. Let Y denote the censoring time at which observation of $N(\cdot)$ is terminated, so that only $\{N(t) : 0 \leq t \leq Y\}$ is observable. Let X be a random variable which is possibly correlated with $N(\cdot)$. The variable X can be a covariate, such as a treatment indicator or health measurement, or it could be the censoring time; in the latter case, $X = Y$.

The intensity function for the point process $N(\cdot)$ is defined as the event occurrence rate conditional on the event history,

$$\lambda\{t|\mathcal{H}(t)\} = \lim_{\Delta \rightarrow 0^+} \frac{\text{pr}[N\{(t+\Delta)^-\} - N(t) > 0 | \mathcal{H}(t)]}{\Delta},$$

where $\mathcal{H}(t) = \{N(u) : 0 \leq u < t\}$ represents the history of recurrent events up to t , for $t \in (0, \tau]$. The intensity function uniquely determines the probability structure of the point process, provided that $\text{pr}[N\{(t+\Delta)^-\} - N(t) > 1] = o(\Delta)$ (Cox & Isham, 1980). The rate function $\lambda(t)$ is defined as the occurrence rate unconditional on the event history; that is,

$$\lambda(t) = \lim_{\Delta \rightarrow 0^+} \frac{\text{pr}[N\{(t+\Delta)^-\} - N(t) > 0]}{\Delta}.$$

The rate function $\lambda(t)$ may be interpreted as the instantaneous population-averaged count of recurrent events in unit time. Unlike the intensity function, the rate function does not fully determine the probability structure of a point process, and different intensity functions could share the same rate function. As illustrated by the following examples, the rate function is conceptually and quantitatively different from the intensity function.

Example 1—Suppose that $N(\cdot)$ is a nonstationary Poisson process with intensity function $\lambda(t)$. Because of the independent-increments property of Poisson processes, the intensity function is memoryless and coincides with the rate function; that is, $\lambda\{t|\mathcal{H}(t)\} = \lambda(t)$.

Example 2—Suppose that there exists a nonnegative random variable Z such that, given $Z = z$, $N(\cdot)$ is a nonstationary Poisson process with subject-specific intensity function $z\lambda_0(t)$, where $\lambda_0(t)$ is a nonnegative-valued function. The variable Z is assumed to be latent in the

model and is used to account for population heterogeneity. Let G denote the distribution function of Z . Conditioning on $Z = z$, the intensity function $z\lambda_0(t)$ is also the rate function. Unconditional on Z , the rate function is $\lambda(t) = \mu_Z \lambda_0(t)$, with $\mu_Z = E(Z)$. Note that $\lambda(t)$ is the same for those distribution functions G with the same mean.

For a fixed $t \in [0, \tau]$, let $t_1 < \dots < t_{N(t)}$ denote the ordered event times before or at t . Let $\Lambda_0(t) = \int_0^t \lambda_0(u) du$. For simplicity of notation, let $p(\cdot|\cdot)$ denote a general conditional probability density function and $P(\cdot|\cdot)$ a general conditional distribution function. The probability density function of the event history prior to t , given $Z = z$, is

$$\begin{aligned} p\{\mathcal{H}(t)|z\} &= p\{N(t)|z\} \times p\{t_1, t_2, \dots, t_{N(t)}|z, N(t)\} \\ &= \left[\frac{\exp\{-z\Lambda_0(t)\} \{z\Lambda_0(t)\}^{N(t)}}{N(t)!} \right] \left\{ N(t)! \prod_{j=1}^{N(t)} \frac{\lambda_0(t_j)}{\Lambda_0(t)} \right\}. \end{aligned}$$

By Bayes' theorem, we obtain that

$$dP\{z|\mathcal{H}(t)\} = \frac{p\{\mathcal{H}(t)|z\}dG(z)}{\int p\{\mathcal{H}(t)|z'\}dG(z')},$$

The memoryless property of the Poisson process implies that $\lambda\{t|z, \mathcal{H}(t)\} = z\lambda_0(t)$, so the intensity function can be expressed as

$$\lambda\{t|\mathcal{H}(t)\} = \int \lambda\{t|z, \mathcal{H}(t)\} dP\{z|\mathcal{H}(t)\} = \frac{E[Z^{N(t)+1} \exp\{-Z\Lambda_0(t)\}]}{E[Z^{N(t)} \exp\{-Z\Lambda_0(t)\}]} \lambda_0(t).$$

Thus, the intensity function $\lambda\{t|\mathcal{H}(t)\}$ is a transformation of $\{N(t), \Lambda_0(t), G\}$ and depends on the event history only through $N(t)$. In the special case where Z is a gamma random variable with mean γ/β and variance γ/β^2 , the rate function is $(\gamma/\beta) \lambda_0(t)$ and the intensity function can be expressed as

$$\lambda\{t|\mathcal{H}(t)\} = \left\{ \frac{N(t)+\gamma}{\Lambda_0(t)+\beta} \right\} \lambda_0(t). \quad (1)$$

The detailed calculation can be found in the Supplementary Material. Clearly, in this case, the intensity function $\lambda\{t|\mathcal{H}(t)\}$ depends on the event history through the total number of events, $N(t)$, and is structurally different from the rate function $\lambda(t)$.

2. Size, shape and rate

Assume that the rate function $\lambda(t)$ is continuous in t for $t \in [0, \tau]$. The cumulative rate function, defined as $\Lambda(t) = \int_0^t \lambda(u) du$, is nondecreasing and continuous in t , with $E\{N(t)\} = \Lambda(t)$ for $t \in [0, \tau]$. In this section we introduce the shape and size parameters of the rate function. These two parameters serve as the basis for defining shape- and size-

independence between X and $\lambda(\cdot)$. Shape- and size-based measures will be constructed to quantify the association between X and $\lambda(\cdot)$.

We first note that the rate function can be characterized by shape and size parameters, where the size parameter is defined to be the cumulative rate function evaluated at τ , i.e., $\Lambda(\tau)$, and the shape parameter is defined by

$$f(t) = \frac{\lambda(t)}{\Lambda(\tau)} \quad (t \in [0, \tau]).$$

Thus the rate function can be decomposed into the product of the shape and size parameters: $\lambda(t) = f(t) \times \Lambda(\tau)$ for $t \in [0, \tau]$. Clearly, $f(t)$ defines a proper probability density function on $[0, \tau]$, and $\Lambda(\tau)$ can be interpreted as the total magnitude of the rate function or, more precisely, the expected number of events occurring in the prespecified interval $[0, \tau]$. The function $f(t)$ is essentially the standardized rate function over the interval $[0, \tau]$. It is called the shape parameter because it characterizes the shape of the marginal occurrence rate over $[0, \tau]$, adjusted by the magnitude of the rate function. In contrast, the size parameter quantifies the magnitude of the rate function without containing additional information on the pattern of the occurrence rate over $[0, \tau]$. For given $X = x$, let the conditional rate function be denoted by $\lambda(t|x)$. Based on the shape and size parameters, we define rate-, shape- and size-independence between X and $\lambda(\cdot)$ as follows.

Definition 1—The rate function $\lambda(\cdot)$ is rate-independent of X if $\lambda(t|x)$ depends on (t, x) only through t or, equivalently, $\lambda(t|x) = \lambda(t)$ for $(t, x) \in [0, \tau] \times R_X$, where R_X is the support of the distribution of X .

Definition 2—The rate function $\lambda(\cdot)$ is shape-independent of X if, for $(t, x) \in [0, \tau] \times R_X$, its shape function

$$f(t|x) = \frac{\lambda(t|x)}{\int I(0 \leq u \leq \tau) \lambda(u|x) du}$$

depends on (t, x) only through t or, equivalently, $\lambda(t|x) = a(x)f(t)$ where $a(\cdot)$ is a nonnegative-valued function of x .

Definition 3—The rate function $\lambda(\cdot)$ is size-independent of X if, for $x \in R_X$, $\Lambda(\tau|x)$ depends on (τ, x) only through τ .

It can be shown that rate-independence holds if and only if both shape-independence and size-independence hold. In general, rate-independence is weaker than independence between $N(\cdot)$ and X . We now consider a number of rate-, shape- and size-independent models.

Example 3—Suppose that X coincides with Y , the censoring time when observation of $N(\cdot)$ is terminated. An independent censoring condition usually refers to the assumption that X is

independent of $N(\cdot)$, although a weaker condition of rate-independence would be sufficient for nonparametric estimation of the rate function (Nelson, 1988, 1995; Lawless & Nadeau, 1995).

Example 4—Poisson count models are commonly used for modelling the number of events occurring in a fixed time interval $[0, \tau]$ in the absence of censoring. Such a model focuses only on the total event count at τ , i.e., $N(\tau)$, although it may implicitly impose a Poisson process assumption on the point process $N(\cdot)$. A typical question of interest is the association between a variable X and $N(\tau)$. In such a model, the count variable follows a Poisson distribution with mean parameter $\Lambda(\tau | x)$, where size-independence holds if $\Lambda(\tau | x) = \Lambda(\tau)$ for $x \in R_X$.

Example 5—The popular proportional rate model considered by Pepe & Cai (1993), Lin et al. (2000) and Schaubel & Cai (2005), among others, is given by $\lambda(t | x) = f(t) \exp(\alpha + \beta x)$ for $t \in [0, \tau]$, where α and β are unknown parameters and $f(t)$ is the shape function. This is a shape-independent model with $a(x) = \exp(\alpha + \beta x)$. In the proportional rate model, the effect of X on $N(\cdot)$ is modelled only through the size parameter of the rate function.

Example 6—The shared-frailty proportional rate model is a commonly adopted model, in which the recurrent event process and censoring share the same subject-specific latent variable, Z . Given $Z = z$ and $X = x$, the rate function is modelled as $\lambda(t | z, x) = z f(t) \exp(\alpha + \beta x)$ for $t \in [0, \tau]$, where α and β are unknown parameters and $f(t)$ is the shape function. This model has been considered by many authors, including Lancaster & Intrator (1998), Wang et al. (2001), Huang & Wang (2004), Liu et al. (2004) and Ye et al. (2007), and it is a shape-independent model with $a(x) = E(Z | x) \exp(\alpha + \beta x)$. Similar to Example 5, the regression effect of X on $N(\cdot)$ is modelled only through the size parameter of the event process.

3. Association measures

When studying the relationship between X and $\lambda(\cdot)$, a question of interest is how to quantify the dependence of $\lambda(\cdot)$ on X . In a general model setting, it is useful to characterize the dependence relationship through the shape and size components of the rate function. In this section we consider new measures, based on Pearson's correlation coefficient and Kendall's tau, for quantifying shape- and size-based associations between X and $\lambda(\cdot)$.

Conceptually, the shape parameter coincides with the population-averaged distributional pattern of event times in $[0, \tau]$, and the size parameter is the expected value of the total number of events, $K = N(\tau)$, occurring in $[0, \tau]$. Define the processes

$$N^*(t) = N(t) / \Lambda(\tau), \quad S(t) = \int_0^t u \, dN^*(u) \quad (0 \leq t \leq \tau).$$

Heuristically, $E\{dN^*(t)\} = f(t) \, dt$, $dS(t) = t \, dN^*(t)$ and $E\{dS(t)\} = t f(t) \, dt$. The weighted process $N^*(t)$ identifies the shape of $N(t)$ for $t \in [0, \tau]$ and therefore can be regarded as a shape process. Further, $S(t)$ is a marked point process with mark t tracing the location of $dN^*(t) > 0$. The processes $N^*(t)$ and $S(t)$ can be used together to define the shape-based

association measures. In what follows, we present two approaches to quantifying associations.

Pearson's correlation coefficient can readily be used to define the size-based association measure

$$\rho_{sz} = \frac{E\{(X - \mu_X)(K - \mu_K)\}}{\sigma_X \sigma_K},$$

where μ_X and μ_K are the means and σ_X and σ_K the standard deviations of X and K . To define the shape-based association measure, we extend Pearson's correlation coefficient to quantify the association between variable X and process $dS(\cdot)$ as follows. The covariance between X and $dS(\cdot)$ in the interval $[0, \tau]$ can be defined as

$\int_0^\tau E[(X - \mu_X)\{dS(t) - t f(t) dt\}] = E[(X - \mu_X)\{S(\tau) - \mu_S\}]$, where $\mu_S = E\{S(\tau)\} = \int_0^\tau t f(t) dt$ is the mean of the shape function f . The shape-based association measure is thus defined as the correlation coefficient between X and the variable $S(\tau)$:

$$\rho_{sh} = \frac{E[(X - \mu_X)\{S(\tau) - \mu_S\}]}{\sigma_X \sigma_S},$$

where σ_S is the standard deviation of $S(\tau)$. Because $dN(t)$ takes values only from $\{0, 1\}$, one can show that

$$\sigma_S^2 = \Lambda(\tau)^{-1} \int_0^\tau u^2 f(u) du - \left\{ \int_0^\tau u f(u) du \right\}^2.$$

Interestingly, σ_S^2 coincides with, is smaller than, or is larger than the variance of the shape function f , which is a probability density function, when $\Lambda(\tau)$ equals, is greater than, or is less than 1, respectively. This is intuitively correct in view of the fact that, with $\Lambda(\tau)$ being the expectation of $N(\tau)$ and $S(\tau) = \int_0^\tau u dN(u) / \Lambda(\tau)$ representing a type of mean occurrence time of recurrent events, a larger value of $\Lambda(\tau)$ implies a smaller variance of $S(\tau)$ and vice versa.

An alternative approach is to define the size-based association based on Kendall's tau:

$$\pi_{sz} = 2 \Pr\{(X_1 - X_2)(K_1 - K_2) > 0\} + \Pr\{(X_1 - X_2)(K_1 - K_2) = 0\} - 1,$$

where $\{X_1, N_1(\cdot)\}$ and $\{X_2, N_2(\cdot)\}$ are assumed to be independent and identically distributed. In an analogous way, Kendall's tau can be extended to define the shape-based association between the random variable X and the shape process $N^*(\cdot)$:

$$\pi_{sh} = 2 \int_0^\tau \int_0^\tau E[I\{(X_1 - X_2)(u - v) > 0\}] dN_1^*(u) dN_2^*(v) + \int_0^\tau \int_0^\tau E[I\{(X_1 - X_2)(u - v) = 0\}] dN_1^*(u) dN_2^*(v) - 1.$$

Clearly, both π_{sz} and π_{sh} take values between -1 and 1 . This property holds for π_{sh} in view of the fact that $E\{dN_j^*(t)\} = f(t) dt$ ($j=1, 2$) is essentially a probability density function.

Theorem 1

- i. Shape-independence between X and $\lambda(\cdot)$ implies $\rho_{sh} = 0$. Size-independence between X and $\lambda(\cdot)$ implies $\rho_{sz} = 0$.
- ii. Shape-independence between X and $\lambda(\cdot)$ implies $\pi_{sh} = 0$. Size-independence between X and $\lambda(\cdot)$ implies $\pi_{sz} = 0$.

Theorem 1 quantifies the relationship between shape- or size-independence and the association measures. Further, since rate-independence holds if and only if shape- and size-independence both hold, it follows that rate-independence implies $(\rho_{sh}, \rho_{sz}) = (0, 0)$ and $(\pi_{sh}, \pi_{sz}) = (0, 0)$.

In general, it can be shown that the association measures ρ_{sh} , ρ_{sz} , π_{sh} and π_{sz} are invariant under linear transformations of the time scale t . In nonparametric models with censoring, the association measures ρ_{sh} , ρ_{sz} , π_{sh} and π_{sz} are neither identifiable nor estimable because of incompleteness in the observed samples. However, despite the nonidentifiability of these measures, nonparametric test statistics can still be constructed to test shape- and size-independence between X and $\lambda(\cdot)$ in the observable region of samples. These tests will be introduced in the next section.

4. Shape- and size-based tests

4.1. Test statistics

In this section we present a formal procedure for testing rate-independence between $\lambda(\cdot)$ and X for recurrent event data with censoring, where the null hypothesis is formulated as $H: \lambda(t|x) = \lambda(t)$ for $(t, x) \in [0, \tau] \times R_X$. Assume that the observed data $\{X_i, Y_i, \mathcal{H}_i(Y_i)\}$ ($i = 1, \dots, n$) are independent and identically distributed copies of $\{X, Y, \mathcal{H}(Y)\}$. The proposed statistical test is a two-step procedure, where a shape-independence test is performed in step 1, followed by a size-independence test in step 2. By performing the shape- and size-independence tests, the dependence between X and $\lambda(\cdot)$ can be clearly explained and better characterized. As the appropriateness of the size-independence test depends on the validity of the shape-independence hypothesis, the two tests are performed in succession. These test statistics can be used in the following two cases.

Case 1: A covariable X is possibly correlated with $N(\cdot)$. In this case, we require an additional assumption that the censoring time Y is rate-independent of $\lambda(\cdot)$ given X .

Case 2: The censoring time Y coincides with X , i.e., $X = Y$. In this case, testing the rate-independence of $\lambda(\cdot)$ and Y is of interest.

Tests for rate-independence in Case 1 can be used to examine the association between a potential risk factor and the recurrent event process. Tests in Case 2 are used to examine the independent censoring assumption that is typically required for recurrent event data analysis in one-sample settings. In applications, study subjects are often censored because of loss to follow-up or end of the study; yet it is also possible that a terminal event, such as death, is a

component of the censoring mechanism. Let D denote the time to death and C the time to loss to follow-up or end of study. Then the censoring time is given by $Y = \min(D, C)$. When postulating statistical models, the recurrent event process, $N(t)$, is defined for $t \in [0, \tau]$ and could be considered latent for $D < t - \tau$, should death occur before τ . As will be shown later in this section, the proposed shape- and size-based tests are constructed based on recurrent events in the observable region $\{N(t) : t \in [0, Y]\}$. Therefore, in Case 2, the proposed tests can also be viewed as quasi-independence tests between Y and $\lambda(t)$, for $t \in [0, Y]$, to avoid the controversy of allowing the occurrence of recurrent events after death under the model assumption (Ghosh & Lin, 2003). In Case 2, our test statistics are designed to test shape- and size-independence between Y and $\lambda(\cdot)$. Thus, if the dependence between D and $\lambda(\cdot)$ is of interest, the proposed tests may not be optimal for data analysis, even though the dependence between D and $\lambda(\cdot)$ is partially carried over to the dependence between Y and $\lambda(\cdot)$.

Consider testing the null hypothesis that X is shape-independent of $N(\cdot)$, that is, $H_{\text{sh}}: \lambda(t | x) = a(x)f(t)$ for $(t, x) \in [0, \tau] \times R_X$. We propose a U-statistic to test the hypothesis based on the location of event times:

$$W_{\text{sh}} = \frac{2}{n(n-1)} \sum_{i < j} w_{ij}^{\text{sh}} \text{sgn}(X_i - X_j) \int \Delta_{ij}(u, t) \text{sgn}(u - t) dN_i(u) dN_j(t). \quad (2)$$

In (2), the sum is over all selections (i, j) from $(1, \dots, n)$ such that $i < j$,

$w_{ij}^{\text{sh}} = w^{\text{sh}}(X_i, X_j, Y_i, Y_j)$ is a nonnegative-valued weight function of (X_i, X_j, Y_i, Y_j) satisfying $w_{ij}^{\text{sh}} = w_{ji}^{\text{sh}}$, $\text{sgn}(r) = 1, -1$ or 0 if r is positive, negative or zero, respectively, and $\Delta_{ij}(u, t) = I\{\max(u, t) \leq \min(Y_i, Y_j)\}$. Define the kernel function

$$h_{\text{sh}}\{(X_i, N_i), (X_j, N_j)\} = w_{ij}^{\text{sh}} \text{sgn}(X_i - X_j) \int \Delta_{ij}(u, t) \text{sgn}(u - t) dN_i(u) dN_j(t),$$

where N_i and N_j represent the recurrent event process in $[0, \tau]$ for subject i and subject j , respectively. Then

$$W_{\text{sh}} = \frac{2}{n(n-1)} \sum_{i < j} h_{\text{sh}}\{(X_i, N_i), (X_j, N_j)\}$$

is a U-statistic with kernel $h_{\text{sh}}(\cdot, \cdot)$, which is symmetric in its arguments. By the pairwise comparison of event times, W_{sh} is designed to test whether the shape of the recurrent event process is associated with X . The bivariate indicator Δ_{ij} in (2) ensures the comparability of event times from subjects i and j .

Next, we consider testing for shape-independence. In general, when censoring is present, the size-independence between X and $\lambda(\cdot)$ cannot be tested nonparametrically due to nonidentifiability of $\lambda(\tau)$. Nevertheless, if shape-independence holds, a test statistic can be

constructed to test the size-independence between X and $\lambda(\cdot)$ in the presence of censoring. Under shape-independence, the null hypothesis H_{sz} : $\Lambda(\tau | x) = \Lambda(\tau)$ for $x \in R_X$ is equivalent to H_{sz} : $\lambda(t | x) = \Lambda(\tau) f(t)$ for $(t, x) \in [0, \tau] \times R_X$.

To test H_{sz} , an initial idea might be to construct a test statistic in the spirit of the log-rank test,

$$\sum_{i=1}^n \int_0^\infty \left\{ X_i - \frac{\sum_{j=1}^n X_j I(Y_j > u)}{\sum_{j=1}^n I(Y_j > u)} \right\} dN_i(u).$$

However, the log-rank statistic requires independent censoring to guarantee the representativeness of the risk set under both null and alternative hypotheses. In Case 1, the censoring time is independent of the recurrent event process given X , so the log-rank statistic is a proper test. In Case 2, the censoring time Y is equal to X and the risk sets under the alternative hypotheses fail to be representative. Thus, the log-rank statistic can be applied to Case 1 but not Case 2. Here we consider a size-independence test statistic which can be used for both Case 1 and Case 2:

$$W_{sz} = \frac{2}{n(n-1)} \sum_{i < j} w_{ij}^{sz} \operatorname{sgn}(X_i - X_j) \int \Delta_{ij}(t, t) \{dN_i(t) - dN_j(t)\}, \quad (3)$$

where $w_{ij}^{sz} = w^{sz}(X_i, X_j, Y_i, Y_j)$ is a nonnegative-valued weight function of (X_i, X_j, Y_i, Y_j) satisfying $w_{ij}^{sz} = w_{ji}^{sz}$, and $\Delta_{ij}(t, t) = I\{t \leq \min(Y_i, Y_j)\}$. Define the kernel function

$$h_{sz}\{(X_i, N_i), (X_j, N_j)\} = w_{ij}^{sz} \operatorname{sgn}(X_i - X_j) \int \Delta_{ij}(t, t) \{dN_i(t) - dN_j(t)\}.$$

Then

$$W_{sz} = \frac{2}{n(n-1)} \sum_{i < j} h_{sz}\{(X_i, N_i), (X_j, N_j)\}$$

is a U-statistic with symmetric kernel $h_{sz}(\cdot, \cdot)$.

Both W_{sh} and W_{sz} can be considered Kendall's-tau-type statistics, and can handle censored and truncated data. The indicator function $\Delta_{ij}(t, t)$ in (3) determines the time interval in which counts from subjects i and j are comparable, so the test statistic W_{sz} tests the size-independence of X and $\lambda(\cdot)$ based on recurrent event data from a comparable interval $[0, Y_i \wedge Y_j]$. This test strategy is appropriate only when the shape of $N(\cdot)$ is unaffected by X , as in this case the difference of sizes can be detected by comparing proportional counts.

In either Case 1 or Case 2, the rate-independence under H is accepted if the null hypotheses of shape-independence and size-independence are both accepted. Rate-independence is

rejected if either H_{sh} is rejected in step 1, or H_{sh} is accepted in step 1 but H_{sz} is rejected in step 2.

4.2. Properties of the test statistics

This section establishes some properties of the proposed test statistics. The discussion will not distinguish between Cases 1 and 2, but one should keep in mind that Case 1 requires the additional assumption that Y is rate-independent of $\lambda(\cdot)$ given X .

Theorem 2—Define $\mu_{sh} = E[h_{sh}\{(X_i, N_i), (X_j, N_j)\}]$ and assume $E(h_{sh}^2) < \infty$. Then:

- i. $n^{1/2}(W_{sh} - \mu_{sh})$ converges weakly to the normal distribution with mean zero and variance σ_{sh}^0 , where

$$\sigma_{sh}^2 = 4 \text{cov}[h_{sh}\{(X_1, N_1), (X_2, N_2)\}, h_{sh}\{(X_2, N_2), (X_3, N_3)\}];$$
- ii. under the null hypothesis H_{sh} , $n^{1/2}W_{sh}$ converges weakly to the normal distribution with mean zero and variance σ_{sh}^2 , where

$$\sigma_{sh}^2 = 4E[h_{sh}\{(X_1, N_1), (X_2, N_2)\} h_{sh}\{(X_2, N_2), (X_3, N_3)\}].$$

Theorem 3—Define $\mu_{sz} = E[h_{sz}\{(X_i, N_i), (X_j, N_j)\}]$. Assume that the null hypothesis H_{sh} holds and that $E(h_{sz}^2) < \infty$. Then:

- i. $n^{1/2}(W_{sz} - \mu_{sz})$ converges weakly to the normal distribution with mean zero and variance σ_{sz}^2 , where $\sigma_{sz}^2 = 4\text{cov}[h_{sz}\{(X_1, N_1), (X_2, N_2)\}, h_{sz}\{(X_2, N_2), (X_3, N_3)\}]$;
- ii. under H_{sz} , $n^{1/2}W_{sz}$ converges weakly to the normal distribution with mean zero and variance σ_{sz}^2 , where

$$\sigma_{sz}^2 = 4E[h_{sz}\{(X_1, N_1), (X_2, N_2)\} h_{sz}\{(X_2, N_2), (X_3, N_3)\}].$$

The proofs of Theorems 2 and 3 are given in the Appendix and the Supplementary Material. To conclude this section, we remark that an alternative test statistic can be constructed by replacing the sign function sgn in W_{sh} and W_{sz} with a function $\varphi(u, t)$ that satisfies $\varphi(u, t) = -\varphi(t, u)$, such as $\varphi(u, t) = (u - t)^k$ where k is an odd integer. Finally, the asymptotic variances of the test statistics W_{sh} and W_{sz} can be consistently estimated by the corresponding sample U-statistics

$$\hat{\sigma}_{sh}^2 = \frac{4}{n} \left[\frac{1}{n(n-1)(n-2)} \sum_{i < j < k} h_{sh}\{(X_i, N_i), (X_j, N_j)\} h_{sh}\{(X_i, N_i), (X_k, N_k)\} - W_{sh}^2 \right],$$

$$\hat{\sigma}_{sz}^2 = \frac{4}{n} \left[\frac{1}{n(n-1)(n-2)} \sum_{i < j < k} h_{sz}\{(X_i, N_i), (X_j, N_j)\} h_{sz}\{(X_i, N_i), (X_k, N_k)\} - W_{sz}^2 \right].$$

5. Numerical studies

Two simulation studies were conducted to examine the performance of the proposed testing procedures in Cases 1 and 2 described in § 4.1. In all simulations, 2000 datasets were generated, each with n independent subjects, where we took $n = 100$ and 200. For Case 1, a

binary time-independent covariate X was generated from the Bernoulli distribution with probability 0.5. A subject-specific frailty Z was simulated from a gamma random distribution with mean 1 and variance 0.25 to account for population heterogeneity. Given $X = x$ and $Z = z$, the recurrent event times of a subject were generated from a nonstationary Poisson process $\{N(t) : t \in [0, 10]\}$ with intensity function

$$\lambda(t|z, x) = 2z \exp(-2^{-1}t + \beta_0 x + \gamma_0 x t). \quad (4)$$

By integrating out the frailty variable, the rate function of $N(\cdot)$, conditional on $X = x$, is $\lambda(t|x) = 2 \exp(-2^{-1}t + \beta_0 x + \gamma_0 x t)$. Thus the shape parameter is

$$f(t|x) = \frac{(\gamma_0 x - 2^{-1}) \exp\{t(\gamma_0 x - 2^{-1})\}}{\exp\{10(\gamma_0 x - 2^{-1})\} - 1} \quad (t \in [0, 10]),$$

and the size parameter is

$$\Lambda(10|x) = \frac{2 \exp(\beta_0 x) [\exp\{10(\gamma_0 x - 2^{-1})\} - 1]}{\gamma_0 x - 2^{-1}}.$$

It is easy to see that $\gamma_0 = 0$ implies that X is shape-independent of $\lambda(\cdot)$, while $\gamma_0 = \beta_0 = 0$ implies that X is rate-independent of $\lambda(\cdot)$. Finally, the censoring time Y was independently generated from a uniform distribution on $[0, 10]$. We considered combinations of simulation parameters $\gamma_0 = 0, 0.2$ and $\beta_0 = 0, -0.3$. The average number of recurrent events per subject ranges from 2.8 to 3.9 across all scenarios.

Table 1 summarizes the estimated powers of the shape- and size-independence tests used separately and in combination. The significance levels were set to 0.05 if tests were applied separately. To control the overall Type I error rate so that it is lower than 0.05, the significance level of each individual test was set to 0.025 when test statistics W_{sh} and W_{sz} were combined to test rate-independence. A unit weight function was used for both test statistics unless otherwise noted; that is, $w_{ij}^{sh} = w_{ij}^{sz} = 1$ for all i, j . When the shape parameter is independent of X , i.e., $\gamma_0 = 0$, the probability of the null hypothesis H_{sh} being falsely rejected is very close to the nominal level, 0.05. When $\gamma_0 = 0$ and $\beta_0 = 0$, both the size- and shape-independence tests reject the corresponding null hypotheses H_{sh} and H_{sz} approximately 5% of the time, demonstrating satisfactory performance of the proposed tests at the 5% level. When W_{sh} and W_{sz} are used in conjunction to test shape-independence, the estimated overall power is approximately 0.05 when the null hypothesis holds. Moreover, as illustrated in the Supplementary Material, the distributions of the p -values derived from the proposed test statistics under the null hypotheses are very close to uniform, suggesting that the test statistics perform well at all significance levels.

Next, we evaluated the performance of the test statistics in Case 2. We generated X from a uniform distribution on $[0, 1]$ and set the follow-up time to be $Y = 10X$. Thus Y has a

uniform distribution on $[0, 10]$. Moreover, we simulated recurrent events from the Poisson process model (4) with a subject-specific frailty Z simulated from a gamma random variable with mean 1 and variance 0.25. Thus $\gamma_0 = 0$ implies that Y is shape-independent of $\lambda(\cdot)$, while $\gamma_0 = \beta_0 = 0$ implies rate-independence. As shown in Table 1, when shape-independence holds for the follow-up time, the estimated power of the shape-independence test statistic W_{sh} is very close to the nominal level, 0.05. When both the size and shape parameters are independent of the follow-up time Y , the proportion of false rejections with combined use of the two test statistics is approximately 0.05. Q-Q plots in the Supplementary Material suggest that the test statistics under the null perform well at all significance levels.

6. Data example

We applied the proposed testing procedures to data collected by the Danish Psychiatric Central Register. This case register records data from all admissions to psychiatric hospitals in the entire nation of Denmark, and participation is mandatory for all Danish psychiatric hospitals, relevant clinical departments and units treating patients with schizophrenia. The register covers psychiatric inpatient services since 1969, as well as outpatient contacts from 1994 onward (Munk-Jørgensen & Mortensen, 1997). We analysed data from a cohort of 1235 patients who had their first contact with psychiatric services between 1 April 1970 and 1 April 1973. Of these patients, 810 were male and 226 had early-onset schizophrenia, defined as onset before 20 years of age. To demonstrate the use of the proposed test statistics in Case 1, we artificially censored the follow-up for all individuals on 1 April 1973 so that the censoring time is expected to be conditionally rate-independent of the recurrent event process. As a result, the median follow-up was 1.3 years and the average number of inpatient psychiatric admissions per patient was 1.2. A total of 28 deaths was recorded during the three-year study period.

We first examined the assumption of shape-independence of the censoring time Y and the counting process of repeated hospitalizations, $N(\cdot)$, within each gender group. The shape- and size-independence tests yield p -values of 0.56 and 0.06, respectively, for males, and 0.70 and 0.052 for females. Thus the null hypothesis that Y is rate-independent of $\lambda(\cdot)$ conditional on gender is not rejected. Next, we applied the shape- and size-independence tests to examine the association between recurrent inpatient psychiatric admissions and gender. The test for shape-independence yields a p -value of 0.70, and the test for size-independence yields a p -value of 0.0002. Hence gender is not associated with the shape parameter of $\lambda(\cdot)$ but is significantly associated with the size parameter. Thus, the proportional rate model (Pepe & Cai, 1993; Lin et al., 2000) is a natural choice for evaluating the effect of gender on rehospitalization risk.

Similarly, we checked the assumption of rate-independence of Y and $\lambda(\cdot)$ for individuals with early- and late-onset schizophrenia. The p -values given by the shape- and size-independence tests are 0.12 and 0.06, respectively, for the early-onset group, and 0.83 and 0.16 for the later-onset group. Thus, the conditional rate-independence assumption of Y and $\lambda(\cdot)$ is not violated for the first three years since initial hospitalization. Moreover, the shape-independence test for the association between $N(\cdot)$ and age of onset yields a p -value of 0.31,

indicating that the shape parameter is shape-independent of onset age. On the other hand, the size-independence test for the association between $N(\cdot)$ and age of onset yields a p -value of 0.0002; thus the age of onset is significantly associated with the size parameter. The results suggest that the proportional rate model can be used to estimate the effect of the age of onset on the risk of repeated hospitalization.

The associate editor suggested conducting the tests with the two groups combined and the group effect accounted for. One can construct stratified tests in the presence of categorical covariates based on the weighted average of the test statistics from the subgroups, where the weights are inversely proportional to the sample sizes of the subgroups. The p -values of the stratified tests, stratified by gender and age of onset, are 0.21 for shape-independence and 0.39 for size-independence. Hence the rate-independence assumption of Y and $\lambda(\cdot)$ is not rejected after controlling for gender and age of onset.

7. Discussion

When studying a recurrent event process, the rate function is often of interest due to its interpretation as a population average. By decomposing the rate function into shape and size parameters, the probability structure of the recurrent event process can be better characterized and analysed. In either Case 1 or Case 2, the statistic W_{sh} is designed to test shape-independence between X and $\lambda(\cdot)$. If H_{sh} is accepted, the statistic W_{sz} is then used to test the size-independence hypothesis H_{sz} . The rate-independence between X and $\lambda(\cdot)$ is rejected if either H_{sh} or H_{sz} is rejected.

As the shape- and size-based statistics W_{sh} and W_{sz} are constructed in the spirit of Kendall's tau, their use would be most appropriate when, for example, the model under the null and alternative hypotheses possesses a stochastic ordering property, i.e., for the shape-independence test, that the underlying model $F(t | x)$ is monotone in x . Similarly, with validity of shape-independence established, a stochastic ordering model where $A(\tau | x)$ is monotone in x would be appropriate for the size-based test.

Along the lines of using shape and size parameters to characterize the recurrent event process, it would also be of interest to quantify the relationship between the size or shape of a process and the covariates in a regression model. Such statistical models and the accompanying methods are currently under investigation and the results will be reported in a future paper.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Thanks are due to the four reviewers for suggestions. The authors also thank Professors Preben Bo Mortensen and William Eaton for providing the anonymous Danish schizophrenia data. This research was supported by the U.S. National Institutes of Health.

References

- Andersen PK, Gill RD. Cox's regression model for counting processes: A large sample study. *Ann Statist.* 1982; 10:1100–20.
- Cook, RJ.; Lawless, JF. *The Statistical Analysis of Recurrent Events*. New York: Springer; 2007.
- Cox, DR.; Isham, V. *Point Processes*. London: Chapman & Hall; 1980.
- Gail MH, Santner TJ, Brown CC. An analysis of comparative carcinogenesis experiments based on multiple times to tumor. *Biometrics.* 1980; 36:255–66. [PubMed: 7407314]
- Ghosh D, Lin DY. Semiparametric analysis of recurrent events data in the presence of dependent censoring. *Biometrics.* 2003; 59:877–85. [PubMed: 14969466]
- Hoeffding W. A class of statistics with asymptotically normal distribution. *Ann Math Statist.* 1948; 19:293–325.
- Huang CY, Wang MC. Joint modeling and estimation for recurrent event processes and failure time data. *J Am Statist Assoc.* 2004; 99:1153–65.
- Kalbfleisch, JD.; Prentice, RL. *The Statistical Analysis of Failure Time Data*. 2. Hoboken, New Jersey: John Wiley & Sons; 2002.
- Lancaster, T. *The Econometric Analysis of Transition Data*. Cambridge: Cambridge University Press; 1990.
- Lancaster T, Intrator O. Panel data with survival: Hospitalization of HIV-positive patients. *J Am Statist Assoc.* 1998; 93:46–53.
- Lawless JF, Nadeau C. Some simple robust methods for the analysis of recurrent events. *Technometrics.* 1995; 37:158–68.
- Lin DY, Wei LJ, Yang I, Ying Z. Semiparametric regression for the mean and rate functions of recurrent events. *J R Statist Soc B.* 2000; 62:711–30.
- Liu L, Wolfe RA, Huang X. Shared frailty models for recurrent events and a terminal event. *Biometrics.* 2004; 60:747–56. [PubMed: 15339298]
- Munk-Jørgensen P, Mortensen PB. The Danish psychiatric central register. *Dan Med Bull.* 1997; 44:82–4. [PubMed: 9062767]
- Nelson W. Graphical analysis of system repair data. *J Qual Technol.* 1988; 20:24–35.
- Nelson W. Confidence limits for recurrence data—applied to cost or number of product repairs. *Technometrics.* 1995; 37:147–57.
- Pepe MS, Cai J. Some graphical displays and marginal regression analyses for recurrent failure times and time dependent covariates. *J Am Statist Assoc.* 1993; 88:811–20.
- Prentice RL, Williams BJ, Peterson AV. On the regression analysis of multivariate failure time data. *Biometrika.* 1981; 68:373–9.
- Schaubel DE, Cai J. Semiparametric methods for clustered recurrent event data. *Lifetime Data Anal.* 2005; 11:405–25. [PubMed: 16133887]
- Therneau, TM.; Grambsch, PM. *Modeling Survival Data: Extending the Cox Model*. New York: Springer; 2000.
- Wang MC, Qin J, Chiang CT. Analyzing recurrent event data with informative censoring. *J Am Statist Assoc.* 2001; 96:1057–65.
- Ye Y, Kalbfleisch JD, Schaubel DE. Semiparametric analysis of correlated recurrent and terminal events. *Biometrics.* 2007; 63:78–87. [PubMed: 17447932]

Appendix

Proof of Theorem 1

If $\lambda(\cdot)$ is shape-independent of X , then $dN(t \mid X = x)$ has rate $\lambda(t \mid X = x) = a(x)f(t)$. This implies that $E\{dS(t \mid X = x)\} = tf(t) dt$ and

$$E\{XS(\tau)\} = \int_0^\tau E\{X dS(t)\} = \int_0^\tau \mu_X t f(t) dt = \mu_X \mu_S,$$

and hence $\rho_{sh} = 0$. Size-independence between X and $\lambda(\cdot)$ implies that

$$E(XK) = E\{XE(K|X)\} = E\{X\Lambda(\tau|X)\} = E(X)\Lambda(\tau) = \mu_X \mu_K,$$

and thus $\rho_{sz} = 0$.

Part (ii) deals with the association measures π_{sh} and π_{sz} . If X is size-independent of $\lambda(\cdot)$, then $\pi_{sz} = 0$ because of the independence between X_j and K_j ($j = 1, 2$). If X is shape-independent of $\lambda(\cdot)$, then the independent and identically distributed structure of $\{X_1, N_1^*(\cdot)\}$ and $\{X_2, N_2^*(\cdot)\}$ together with the fact that $E\{dN_j^*(t)|X_j=x\} = tf(t) dt$ ($j=1, 2$) imply that

$$\begin{aligned} & \int_0^\tau \int_0^\tau E[I\{(X_1 - X_2)(u-v) > 0\} dN_1^*(u) dN_2^*(v)] \\ &= \int_0^\tau \int_0^\tau E[I\{(X_2 - X_1)(u-v) > 0\} dN_1^*(u) dN_2^*(v)], \end{aligned}$$

and so $\pi_{sh} = 0$ follows by symmetry of probabilities.

Proof of Theorem 2

Theorem 2 can be established by verifying the required conditions for the asymptotic normality of U-statistics. If $E[h_{sh}^2\{(X_1, N_1), (X_2, N_2)\}] < \infty$ holds, then W_{sh} converges to μ_{sh} in probability, and the asymptotic normality of $n^{1/2}(W_{sh} - \mu_{sh})$ in (i) follows from the theory of U-statistics (Hoeffding, 1948).

The test statistic W_{sh} is designed to test shape-independence between X and $\lambda(\cdot)$. Let $g(u, t, Y_i, Y_j) = ij(u, t)\text{sgn}(t-u)f(u)f(t)$. With $ij(u, t) = ij(t, u)$, it is easy to see that $g(u, t, y_i, y_j) = -g(t, u, y_i, y_j)$ and $\int g(u, t, y_i, y_j) du dt = \int g(t, u, y_i, y_j) du dt$. Thus, we derive the equality $\int g(u, t, y_i, y_j) du dt = -\int g(u, t, y_i, y_j) du dt$, which implies that $\int g(u, t, y_i, y_j) du dt = 0$.

Under the hypothesis of shape-independence, $H_{sh}: \lambda(t|x) = a(x)f(t)$, one can prove that $\mu_{sh} = E[h_{sh}\{(X_i, N_i), (X_j, N_j)\}] = 0$. The asymptotic normality of $n^{1/2}W_{sh}$ under H_{sh} follows by a standard argument for U-statistics: $n^{1/2}W_{sh}$ converges weakly to the normal distribution $N(2, \sigma_{sh}^2)$ where $\sigma_{sh}^2 = 4E[h_{sh}\{(X_1, N_1), (X_2, N_2)\}h_{sh}\{(X_2, N_2), (X_3, N_3)\}]$ provided that $E(h_{sh}^2) < \infty$.

The assumption $E(h_{sh}^2) < \infty$ in Theorem 2 is implied by the simple moment assumptions

$E\{(w_{ij}^{sh})^4\} < \infty$ for $i \neq j$ and $E\{N_i(\tau)^4\} < \infty$. Details are given in the Supplementary Material.

Proof of Theorem 3

Assertion (i) is derived directly from the theory of U-statistics. To prove (ii), we further examine the zero-unbiasedness of W_{sz} under the additional hypothesis H_{sz} . Size-independence and shape-independence together imply rate-independence. Thus, when both H_{sh} and H_{sz} hold, we have

$$\begin{aligned} E(W_{sz}) &= E[E\{W_{sz} | (X_i, X_j, Y_i, Y_j)\}] \\ &= \frac{2}{n(n-1)} \sum_{i < j} E \left[w_{ij}^{sz} \operatorname{sgn}(X_i - X_j) \int \Delta_{ij}(t, t) E\{dN_i(t) - dN_j(t) | (X_i, X_j, Y_i, Y_j)\} \right] \\ &= \frac{2}{n(n-1)} \sum_{i < j} E \left[w_{ij}^{sz} \operatorname{sgn}(X_i - X_j) \int \Delta_{ij}(t, t) E\{dN_i(t) - dN_j(t) | (X_i, X_j)\} \right] \\ &= \frac{2}{n(n-1)} \sum_{i < j} E \left[w_{ij}^{sz} \operatorname{sgn}(X_i - X_j) \int \Delta_{ij}(t, t) \{\lambda(t) - \lambda(t)\} dt \right] = 0. \end{aligned}$$

One can also prove that $E(h_{sh}^2) < \infty$ is implied by simpler moment constraints, namely $E\{(w_{ij}^{sh})^4\} < \infty$ for $i \neq j$ and $E\{N_i(\tau)^4\} < \infty$. These constraints are sufficient to validate the asymptotic normality of $n^{-1/2}W_{sz}$ in Theorem 3, although they may be stronger than necessary. Details can be found in the Supplementary Material.

Table 1

Estimated power (%) for the shape-, size- and rate-independence tests, with nominal level 5%. Results are the percentages of rejected null hypotheses H_{sh} , H_{sz} and H among 2000 replications

		$n = 100$						$n = 200$					
	γ_0	β_0	m^-	Shape	Size	Rate	Shape	Size	Rate	Shape	Size	Rate	
Case 1	0	0	3.2	4	6	5	5	6	5				
	0	-0.3	2.8	4	47	37	5	75	66				
	0.2	0	3.9	62	49	69	90	78	97				
	0.2	-0.3	3.3	53	6	42	87	5	79				
Case 2	0	0	3.2	5	5	4	5	5	5				
	0	0.5	4.3	4	42	32	5	71	62				
	0.4	0	5.6	62	55	69	92	85	96				
	0.4	0.5	7.8	78	98	98	98	100	100				

m , average number of observed events per subject.