Statistics
in Medicine

Tutorial in Biostatistics

Received 30 September 2009, Accepted 5 April 2011 Published online 22 July 2011 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.4277

Frailty models: Applications to biomedical and genetic studies

Usha S. Govindarajulu,[a*†] Haiqun Lin,[b] Kathryn L. Lunetta[c‡] and R. B. D'Agostino, Sr.[d]

In survival analysis, frailty models are potential choices for modeling unexplained heterogeneity in a population. This tutorial presents an overview and general framework of frailty modeling and estimation for multiplicative hazards models in the context of biomedical and genetic studies. Other topics in frailty models, such as diagnostic methods for model adequacy and inference in frailty models, are also discussed. Examples of analyses using multivariate frailty models in a non-parametric hazards setting on biomedical datasets are provided, and the implications of choosing to use frailty and relevance to genetic applications are discussed.

Keywords: frailty; Cox model; tutorial; correlated; genetic

1. Introduction

Survival analysis methods are used to model time-to-event data. Either time-to-event or hazard of the event occurrence can be modeled as a response variable. In a hazards model, the hazard can be specified in either multiplicative or additive form. The Cox proportional hazards regression model (Cox model) is still the most popular survival model because it has broad versatility. The standard Cox model does not take into account the unmeasured variability among subjects beyond that of the measured covariates. If heterogeneity across study subjects beyond the measured covariates is suspected, a model that can account for this variability is preferred. Frailty models are one possible extension to the Cox model or other survival models that allow for such dependence.

Frailty models are essentially survival models with both fixed and random effect terms. Although the fixed effects comprise the explained or observed portion of the model, the random effect term accounts for the unexplained variability of the model. In other words, the random effect, or frailty, models the unexplained heterogeneity in the model, which the heterogeneity among hazard rates beyond recorded covariates shown in a population could also exhibit. The frailty term [1] was initially developed to describe heterogeneity at the individual level but was expanded to describe heterogeneity among groups of individuals or within an individual. These can be considered levels of clustering, where an individual is one level of clustering [2].

In this tutorial, we explore frailty models for use with modeling biomedical data and data for individuals who are correlated because of genetic relationships (family data). We discuss existing frailty models, model diagnostics, and inference for frailty models, as well as choosing between frailty and no-frailty models for actual biomedical and family data in a non-parametric hazards setting.

[a]*Department of Medicine, Brigham and Women's Hospital, Boston, MA, 02115*
[b]*Division of Biostatistics, Yale School of Public Health, New Haven, CT, 06510*
[c]*Department of Biostatistics, Boston University, Boston, MA, 02118*
[d]*Department of Math and Statistics, Boston University, Boston, MA, 02215*
**Correspondence to: Usha S. Govindarajulu, Department of Medicine, Brigham and Women's Hospital, Boston, MA, 02115.*
[†]*E-mail: usha@alum.bu.edu*
[‡]*We acknowledge the funding source for the noted author: NIA grant no. 5R01AG29451.*

2. Background and overview

2.1. Overall model

The following form of the hazard function with proportionality in covariate effects depending on $c(.)$ represents the multiplicative hazards model, parametric or semi-parametric:

$$h(t|x) = h_0(t)c(\beta' x) \tag{1}$$

where t represents survival time, x represents covariates, β contains the regression coefficients for the covariates, $h_0(t)$ is the unspecified baseline hazard, and $c(\beta' x)$ is the link function, which assumes a non-negative form [3], depending upon the specific model. In the parametric survival model, a parametric distribution (i.e., Weibull, exponential, etc.) models the baseline hazard function. In the semi-parametric survival model (i.e., Cox model), the baseline hazard function is unspecified, the function of covariates are assumed to be parametric, and the proportionality of the covariate effects on the hazards is assumed.

Frailty models are essentially survival models with a random effect term, and similar to proportional hazards models, may be parametric or semi-parametric. In general, the hazard function for the frailty model, either semi-parametric or parametric, can be written by adding a frailty to Equation (1) as

$$h_{ij}(t|x,z) = h_0(t)c(\beta' x_{ij} + z'_{ij}\mathbf{w}_{ij}) \tag{2}$$

where h_{ij} is the hazard function for the jth observation from the ith cluster, where the cluster can be an individual or a group, and $\mathbf{w}_{ij}$ is a vector of random effects associated with the covariates vector z_{ij} for cluster i observation j, which usually includes an intercept. The random effects account for the heterogeneity effects of the z_{ij}. When z_{ij} includes only an intercept, the model (2) is a simple multiplicative frailty model. The $\mathbf{w}_{ij}$ accounts for the correlation among individuals within a cluster. The frailty model in (2) has proportional hazards conditionally, by integrating out the frailty according to the method of Bremaud [4]. Vaida and Xu [5] utilized the generalized mixed-effect model framework specified in Equation (2) with a Cox proportional hazards model. This generalization enables that frailty as the random effects in the log relative risk not only act on the baseline hazards through the intercept in z_{ij} but also can include additional frailty terms as random coefficients associated with respective covariates in z_{ij}. Each component of the frailty represents a cluster-specific effect with the respective covariate. This type of model allows inference of cluster-specific random effects and potentially greater comprehension and interpretation of variability in the data.

2.2. Univariate and multivariate frailty

In univariate survival data, each cluster has only one individual with only one survival outcome. Kosorok *et al.* [6] addressed that univariate survival data can only have univariate frailty under certain conditions. Multivariate survival data consists of a cluster of more than one individual. The clusters may be multiple survival outcomes for a single individual or one or more survival outcomes for multiple correlated individuals, such as relatives. Multivariate survival data can have either univariate or multivariate frailties.

The frailties across different clusters are assumed to have a distribution; they account for unexplained heterogeneity at the cluster level. The frailty term, $\mathbf{w}_{ij,}$ can be univariate or multivariate. A univariate frailty model contains a single frailty per cluster with $\mathbf{w}_{ij}$ being a scalar and the same across all different j's within the same cluster i, z_{ij} being an intercept of one. Individuals within a cluster share a common frailty (also termed shared frailty). In a multivariate frailty model, each cluster has two or more frailties, and $\mathbf{w}_{ij}$'s are therefore different across different j's even within the same cluster i—the frailties $\mathbf{w}_{ij}$ can come from a multivariate distribution or have more than one term in it (in the case of correlated frailty described in the following text). The frailties among different observations within a single cluster are marginally correlated. In the case of correlated frailty model, two additive terms within $\mathbf{w}_{ij}$ consist of a final $\mathbf{w}_{ij}$, one term is the common frailty term shared among all the observations from the cluster i, and the other is the individual frailty not shared; the two frailty terms have two different distributions. For example, Hsu *et al.* [7] found that using a correlated frailty model was helpful in the context of a case-control family study because having risk heterogeneity due to unmeasured risk factors among family members complicated these designs.

Gutierrez [8] noted that a frailty model can be used with univariate data or multivariate survival data when there is a common group effect, which is latent, and the univariate frailty is modeled among groups, hence a shared frailty model. Gutierrez [8] noted that choosing to use a shared frailty model is only a consequence of the way in which the data are organized (univariate or multivariate survival data) and not reflective as to the source of the frailty. According to Hougaard [9], the shared frailty model includes the univariate model, so fundamental probability results are common between the two cases, just the source of heterogeneity is viewed differently, consequently, the interpretation, identifiability of parameters, and the estimation can be quite different.

A further discussion on univariate and multivariate frailty models is noteworthy. Hougaard [9, 10] described these models in detail. He observed the idea of frailty models for survival data as considering the variability in life times as coming from two separate sources: simple randomness (described by the hazard function) and a random effect (the true frailty). In the case of the latter, the frailty generates dependence between the individuals in the group.

Cui and Sun [11] had adopted a marginal survival model for multivariate survival data and imposed a shared univariate gamma frailty correlation structure to the correlated data. They developed a test for checking the adequacy of the gamma frailty distribution. Rondeau *et al.* [12] showed an example of using penalized likelihood estimation in a gamma frailty model among correlated data.

The frailties or random effects among individuals or families could be modeled as shared or correlated. According to Costigan and Klein [13], the shared random effect can represent the common environmental effect and/or some genetic trait the subjects shared, making this a random genetic or batch effect. Furthermore, Aslanidou and Dey [14] noted that the frailty term in a model represents the common unobservable or neglected covariates.

The kinship frailty model described by Pankratz *et al.* [15] and implemented in the kinship package of the R computing environment [16–18] utilizes a correlated frailty in a Cox proportional hazards model specifically designed to deal with genetically related individuals. For this model, in Equation (2), $\mathbf{w}_{ij}$ is the vector of random effects, which are distributed normally with mean 0 and variance Σ, where the elements of Σ are two times the elements of the kinship matrix for the individuals. A kinship matrix defines the genetic relationships between all pairs of individuals. The kinship coefficient is the probability that random alleles drawn from each of two individuals at a specific locus are copies of the same ancestral allele. For this model, only individuals within a cluster (pedigree) are correlated: the kinship coefficient between individuals in two different pedigrees is 0. The model can be extended to a multivariate frailty model by allowing two or more frailty terms. For example, one shared frailty for all family members, and a second, correlated frailty that accounts for individual correlations because of genetic relationships. The model, as implemented in the kinship package, uses normal random effects as described by Ripatti and Palmgren [19] because this model is easily generalized to arbitrary covariance matrices.

2.3. *Frailty distribution*

Hougaard [9, 20] discussed choice of frailty distribution, mainly for the shared frailty case. Hougaard [9, 10] described how the theoretical properties of the various models and the distributions impacting the fit of the model are used in selection of the frailty distribution. The gamma distribution has typically been used to fit the frailty random effect. According to Hougaard [9], the gamma distribution was chosen for mathematical reasons; there are no known biological reasons motivating the choice of the gamma distribution. If one chooses a gamma distribution for the frailty, the advantages are that they have simple densities for which parameters are easily obtained through likelihood estimation. The frailty distribution is gamma with the same shape parameter and different scale parameter for the survivors at a given age, and the frailty distribution is also gamma with shape parameter increased by one and a scale parameter, which is a function of age at death, for the persons dying at any age [9]. Other than the gamma distribution, Hougaard [9] discussed how the choice of frailty distribution was extended to the natural exponential family, where the gamma distributions are the simplest family of this kind. Examples of these distributions include the inverse Gaussian and the positive stable distributions. Hanagal [21] discussed the use of different frailty distributions in a Weibull extended bivariate exponential regression model.

Identifiability refers to being able to uniquely estimate both the parameters of the hazard function as well as of the frailty distribution in univariate data [22]. In univariate frailty models (generally parametric), because of the identifiability problem, the distribution specified for the frailty is often given a prespecified fixed mean of 1 for multiplicative (or proportional) hazards model to identify the

frailty distribution. For example, for a gamma distribution with shape parameter, α, this would become $\Gamma(\alpha, 1/\alpha)$ so that only one parameter would need to be estimated in the model. Heckman and Singer [23] showed that identifiability is maintained for a Weibull hazard function even without regressors, as long as the frailty distribution has a finite mean. Barker and Henderson [24] have shown a method of dealing with estimating the univariate Cox proportional hazards model to be discussed further in Section 2.4. Multivariate frailty models can have either fixed or infinite mean frailty distributions [22]. For a fixed mean frailty distribution, the mean can be restricted to 1, as in the case of the univariate frailty model. For the infinite mean frailty distribution, the scale parameter must be restricted in some other way.

2.4. Methods of estimation

Various methods of estimation of the frailty model exist and are currently used in some software packages. Some of these methods are as follows: likelihood estimation using Newton–Raphson, expectation–maximization (EM) algorithm, penalized likelihood, and Bayesian methods [i.e., Monte Carlo Markov Chain (MCMC)], but this list is not exhaustive. These methods essentially seek to solve the log-likelihood in which the frailty, w, will be integrated out, using counting process notation:

$$\sum_i \sum_j \left[\int_0^\infty Y_{ij}(t) \left[\log(\lambda_0(t) + X_{ij}\beta + Z_i w_{i_j} \right] \mathrm{d}N_{ij}(t) - \left[\int_0^\infty Y_{ij}(t) \exp\left(X_{ij}\beta + Z_i w_{i_j} \right) \lambda_0(t) \mathrm{d}t \right] + \log p(w_{i_j}; \theta) \right] \tag{3}$$

Likelihood requiring the frailty to be integrated out is referred to as mixture likelihood because the frailty serves as a mixing distribution [22]. Depending upon the particular distribution specified for the frailty random effect, the frailty is integrated out of the joint likelihood to obtain an observed likelihood as a function of survival time given the covariates. For a univariate frailty model with gamma frailty, the previously mentioned integration has a closed form. For most other distributions, it does not have a closed form and can only be integrated out numerically. Lancaster [25] provided a tractable description of using joint likelihood for a univariate frailty model. After obtaining a likelihood with the frailty terms integrated out, a numerical procedure like Newton–Raphson is used for maximum likelihood parameter estimation.

The EM algorithm approach for the Cox frailty model (a similar approach exists for parametric frailty) begins with the full likelihood being a function of the observed event times and the unobservable frailties. In the two steps of the EM Algorithm, the E-step involves computing the expectation of the full likelihood with respect to the observable data, and the M-step invokes a partial likelihood being constructed for estimation of the covariate effects using a profile likelihood technique. One then iterates between the E-step and the M-step until convergence is obtained. A more detailed explanation of the EM algorithm can be obtained from Klein [26] or Nielsen *et al.* [27].

Another method of estimation is the use of penalized likelihood estimation. This estimation, which uses penalized regression, has been adapted for estimation of Cox frailty models [28]. The frailty terms are treated as additional regression coefficients, which are then constrained by a penalty function added to the log-likelihood. This method is supposedly much faster than the EM algorithm approach. Therneau and Grambsch [28] provided a detailed description.

Use of MCMC methods in frailty analysis has gained a foothold in the literature as well because of its ability to handle computationally intensive modeling. Chen *et al.* [29] developed a cure model for multivariate data, which incorporates multivariate failure time data for populations with a cure fraction and a frailty term, which creates a correlation structure between the failure times. MCMC is then used to sample from the posterior distribution of the parameters. Even for frailty models with doubly censored data, when both initial and final times are censored, Jones [30] developed a novel way to use MCMC for estimation of this type of model. Modeling the frailty effect as not only a function of unobserved heterogeneity but also observed covariates, Govindarajulu *et al.* [31] also employed MCMC to sample from the posterior distribution of this complex model.

Certainly other techniques exist. Xue [32] has utilized quasi-likelihood estimating equations to estimate parameters and avoid specifications of the frailty distribution in which a parametric formulation is typically given. Barker and Henderson [24] discussed estimation of a univariate Cox proportional hazards model with gamma frailty via the EM algorithm, in which a local likelihood form of the

baseline hazard replaced the non-parametric Breslow estimator. This method of estimation aids in avoiding the underestimated frailty variance that the usual EM algorithm estimation of the univariate Cox frailty model incurred.

3. Diagnostic methods for frailty models

Various ideas for assessing diagnostics in frailty models exist in the literature and are a growing area of research. Approaches for checking distributional assumptions among models have pervaded. Given a shared gamma frailty model, Glidden [33] provided numerical and graphical techniques for checking model adequacy using martingale residuals without specifying the marginal distributions. The counting process martingale residuals were rederived in the shared gamma frailty context. These can be plotted against time to assess model adequacy. Viswanathan and Manatunga [34] also presented visual diagnostic plots for ascertaining the frailty distribution for multivariate survival data in a correlated frailty model. The plot uses kernel regression smoothing on the time scale with bandwidth chosen by cross-validation to differentiate between the gamma and positive stable frailty models.

Chen and Bandeen-Roche [35] have proposed diagnostic methods for fit of certain clustered failure time distributions, Archimedean copulas (AC), for bivariate survival models. They presented a visual display rather than a formal test of fit for a given AC model. They first defined the conditional hazard ratio, which describes the statistical association between the failure times. A plot of the conditional hazard ratio against the estimated survival function approximates the functional relationship between them and can demonstrate if the AC distribution is gamma or positive stable.

There is currently no readily available software for any of the aforementioned diagnostic techniques. Therefore, few of these techniques are employed in standard practice. However, it is useful to be aware of these methods.

4. Inference for frailty models

Although research in inference for frailty models has been limited, it is worth mentioning. Murphy [36, 37] has provided asymptotic theory for the shared Cox gamma frailty model in which they derived asymptotic distributions of the integrated baseline hazard and the frailty variance. They have also demonstrated consistency and asymptotic normality of the non-parametric maximum likelihood estimator (NPMLE). Parner [38] has extended this work but for a correlated gamma frailty model with covariates. They demonstrated consistency of the NPMLE and the frailty variance.

Kosorok *et al.* [6] has discussed robust inference for univariate proportional hazards frailty regression models for univariate survival data. The frailty, unique to each individual, is assumed to have a mean of 1 within a well-defined one-parameter family of distributions. Inference in this context is based on NPMLE, and the behavior of the model is studied when the fitted model is misspecified. Asymptotic theory in regards to efficiency, uniform consistency, and weak convergence are discussed.

5. Choosing between frailty and non-frailty models

It is important to ascertain if there are unmeasured covariates or other possible reasons for heterogeneity in the population, which is important enough to require the use of a frailty model. In a univariate data setting, we could ascertain that there was a variation or mixture of hazards in the population because of the differential survival pattern of members of the population. We could potentially observe this by plotting the hazard function across time for the population and observing that the hazard rate was not constant but instead exhibits a mixture of hazards. If a mixture of hazards or other heterogeneity is noted, one might decide to fit a frailty model. Outside of the univariate data setting, a multivariate frailty model would be a logical choice if we wanted to model the heterogeneity among groups of individuals or within groups of individuals. In this case, we would need to have an idea *a priori* that there is underlying unobserved heterogeneity between groups or among individuals.

One can incorporate a frailty term into a parametric or semi-parametric model, for example, and the model can either be univariate or bivariate or correlated/multivariate. If the frailty term is not significant in the model, it implies there is no significant frailty effect and, therefore, a frailty model for the data is not considered to be distinctly different from a non-frailty model being fit to the same data. In the case

that the frailty term is significant in the model, one could proceed further using model validation to compare how well the frailty model compares with a non-frailty model, perhaps employing the diagnostic methods discussed in Section 3.

Several methods exist for assessing the frailty effect. In terms of the shared frailty model, Klein [26] described the frailty as a common random effect where individuals in a subgroup share an unobservable, random covariate (frailty), acting multiplicatively on the hazard rate of each group member, and he used a one-parameter gamma distribution for the frailty, z, whose density is

$$g(z) = \frac{z^{(1/\theta-1)} \exp(-z/\theta)}{\Gamma(1/\theta)\theta^{1/\theta}}, \quad \theta \geqslant 0 \tag{4}$$

and the mean value of the frailty is 1, and variance of the frailty is θ, which may be used as a measure of association. Large values of θ reflect greater heterogeneity between subgroups and a stronger association among members of a subgroup. In addition, Kendall's τ can measure the association between pairs of individuals in the ith group, which is $(\theta/\theta+2)$, so the strength of the association is monotone increasing in θ, with $\theta = 0$ corresponding to independence between group members [26].

Another valid issue in the literature concerns whether or not a likelihood ratio test (LRT) for frailty is valid because the test involves the boundary of the parameter space. The LRT for the frailty is computed as twice the difference between the log partial-likelihood with the frailty terms integrated out and the log partial-likelihood of a no-frailty model; it is a one-degree-of-freedom test [28]. Therneau and Grambsch [28] noted that the LRT for a significant frailty involves the boundary of the parameter space, but Nielsen *et al.* [27] have shown that the one-degree-of-freedom chi-square approximation is valid. Wald tests are not relevant for testing a boundary hypothesis and would not work in this context.

6. Application

Examples of using non-parametric frailty models with increasing complexity are demonstrated using two datasets. The first dataset comes from a retrospective survey that the Instituto de Nutrición de Centroamérica y Panamá and the RAND Corporation conducted in 1974–1976, collecting information on child mortality in Guatemala. The data contain information on multiple children per family with the mother defining the family. There are a total of $n = 3120$ children with full covariate information, and there were 403 deaths. We wish to model the time to death of each child as a function of various predictors of the outcome in a shared frailty model, where family is the grouping factor. The predictors of interest are mother's age (linear and quadratic), birth interval, birth order, indicator of whether the previous child before birth of the current child in the family birth order died (only for immediately previous birth, not all previous births), and length of previous birth interval (where first birth serves as the reference level). We used the same thresholds for birth interval as utilized by Guo and Rodriguez [39].

Guo and Rodriguez [39] and Pebley and Stupp [40] used dummy variables for piecewise exponential modeling of the hazard for these data. Here, we applied a shared Cox gamma frailty model using the `coxph` function in the R software package [16]. We first fit a Cox model without frailty and then a model with gamma frailty to test if the frailty effect was significant (Table I). Using the `frailty` function within `coxph` to fit a shared Cox frailty model, an example of the code is given:

```
sharedfrail <- coxph(Surv(time,death) ~ mage + mage2 + f0011 + f1223 + borde + pdead + p0014
 + p1523 + p2435 + p36up + frailty(momid,dist='gamma'), data=pebfull, na.action=na.omit)
```

The `frailty` function is fit to the IDs of the mothers, which is used as the family ID for fitting a shared frailty model. To test if the frailty effect is significant, we created an LRT as described in Section 5. The LRT statistic is $2(3104.017 - 3101.9) = 4.23$. The test statistic is statistically significant at a 0.05 significance level, thus the shared frailty effect is significant. This frailty effect was observed for those variables that were statistically significant; coefficient estimates and standard errors increased slightly (Table I). We also observed that the optimal mother's age was 28 years, given by our calculation using the linear and quadratic terms for mother's age: $-0.17842/2^*(0.00319)$ from the fitted Cox model. Therefore, there appeared to be significant heterogeneity between the families in the study, and modeling this heterogeneity in the time-to-death model had a significant effect.

Table I. Cox model estimates from Pebley and Stupp [40] dataset.

	Cox model			Cox model with frailty		
Variable	Coefficient	SE (coefficient)	*p*-value	Coefficient	SE (coefficient)	*p*-value
Mother's age (linear)	−0.178	0.059	0.0026	−0.184	0.062	0.0027
Mother's age (quadratic)	0.003	0.001	0.0024	0.003	0.001	0.0026
Birth interval (length under a year)	2.219	0.192	<0.0001	2.291	0.204	<0.0001
Birth interval (length of 1–2 years)	0.876	0.107	<0.0001	0.897	0.109	<0.0001
Birth order	0.062	0.033	0.0620	0.065	0.036	0.0730
Previous child in family died	0.122	0.149	0.4100	−0.073	0.153	0.6300
Length of previous birth interval (0–14 months)	0.638	0.211	0.0025	0.690	0.217	0.0015
Length of previous birth interval (15–23 months)	−0.064	0.187	0.7300	−0.030	0.190	0.8800
Length of previous birth interval (24–35 months)	−0.103	0.187	0.5800	−0.098	0.191	0.6100
Length of previous birth interval (36 or more months)	−0.234	0.210	0.2600	−0.225	0.214	0.2900

In a second analysis, we examined age-at-death data from the original cohort [41] of the Framingham Heart Study, the longest running cohort study in the USA. There were a total of $n = 5205$ individuals with age-at-death or last contact data, coming from 2601 nuclear and extended families. Of the 5205 individuals, 4653 had died, and 552 were censored at age of last contact for this analysis. Families ranged in size from having a single individual to 25 individuals with age data. Most families (96%) contributed five or fewer members, and 38% of families contributed two individuals. Among the 5205 individuals, there are 1307 first-degree relative pairs (sibling and parent–offspring pairs), 61 second-degree relative pairs (avuncular and half-sibling pairs), and 11 third-degree relative pairs (cousins).

We used Cox proportional hazards models with a frailty component to determine whether there are significant family-specific frailty effects and individual frailty effects correlated according to degree of relationship between individuals. The correlated frailty can be considered an extension of the additive polygenic model often used to estimate heritability for quantitative phenotypes (see for example [42]). A family-specific frailty model assumes that everyone within the same family shares the same increased/decreased risk. Pedigree ID number defines family. Thus, individuals that are not genetically related but are related through marriage will share the same family-specific effect. This model might be reasonable when the increased risk is due to some environmental exposure that affects all family members. For the frailty model described in Equation (2), the random effect, $\mathbf{w}_{ij}$, is the same for all individuals (j) in pedigree i. In contrast, the correlated frailty model specifies the degree of genetic relationship between individuals. Individuals that are more closely related are expected to have more similar frailty, whereas more distant relatives and unrelated individuals have less similar frailty. This model might be appropriate when individuals in a family are correlated because of unmeasured genotypic effects.

We used the R software [16] `kinship` package [15, 17, 18] to fit the models. The `kinship` package `coxme` function fits mixed-effect Cox models, including those with correlated frailties. In addition, the `kinship` package uses sparse block diagonal matrices to compute and represent kinship matrices. These matrices can be manipulated in a manner similar to ordinary matrices, but the total memory use is typically much smaller.

We fit four models with increasing complexity to the age-at-death data. For each model, we used as the outcome age at death or for censored observations, age at last contact (survage), left truncated at age at examination 1 of the study (ageentry) . Model 1 was a standard Cox model with covariates sex and birth year. This model was fit using the `coxph` function of the R `survival` package:

```
model1<-
coxph(Surv(ageentry,survage,dead)~sex+birth-year,data=death.dat).
```

Model 2 was the first mixed-effects (shared frailty) model. In this model, each family has a family-specific frailty, drawn from a Gaussian model. To fit this model, the R `kinship` package function `coxme` is used. The variable 'famid1' is a unique identifier for each family:

```
model2<-
coxme(Surv(ageentry,survage,dead)~sex+birth-year,data=death.dat1,
random=~1|famid1)
```

Model 3 is the correlated individual frailty model described earlier. The individual random effects within a family are correlated between pairs of individuals, with correlation equal to the relationship coefficient, which is twice the kinship coefficient of the pair. To fit this model, we must first create the kinship matrix. The R `kinship` package has a function that does this. The kinship matrix must be determined using the full pedigree, including all individuals who define the relationships between individuals with phenotypic data. The individuals in the full pedigree with no phenotypic data are 'linkers' who are required to correctly specify kinship coefficients. After the kinship matrix is defined, the subset of the kinship matrix that includes the individuals with phenotypic data in the analysis is created:

```
##kinship matrix for whole study:
deathdat.kmat<-
makekinship(famid=death.dat$famid1,id=death.dat$id,
father.id=death.dat$faid,mother.id=death.dat$moid)
##omit linker individuals with no phenotypic data, used only to
##properly specify the pedigree when calculating kinship
##coefficints:
death.dat1<-subset(death.dat,!is.na(survage))
##temp is an array with the indices of the individuals with
##survages:
all<-dimnames(deathdat.kmat)[[1]]
temp<-which(!is.na(death.dat$survage[match(all,death.dat$id)]))
##kinship matrix for the subset with phenotype survage:
deathdat1.kmat<-deathdat.kmat[temp,temp]

model3<-
coxme(Surv(ageentry,survage,dead)~sex+birth-year,data=death.dat1,
random=~1|id,varlist=deathdat1.kmat)
```

Model 4 is a multivariate frailty model, with both a shared frailty effect for family and the correlated individual frailty effect. To fit this model, one must create a block diagonal matrix indicating family membership that has the same dimensions as the kinship matrix:

```
famblockf.mat<-bdsBlock(death.dat1$id,death.dat1$famid1)
```

Then, to specify the model, we provide the `coxme` function with a list of the two variance matrices described earlier, `deathdat1.kmat` and `famblockf.mat`, that specify the frailty correlations for the two frailty effects:

```
model4<-
coxme(Surv(ageentry,survage,dead)~sex+birth-year,data=death.dat1,
random=~1|id,varlist=list(deathdat1.kmat,famblockf.mat),pdcheck=FALSE)
```

Table II contains the results for models 1–4. We determine the significance of the frailty effects using likelihood ratio tests. First, we test whether the family frailty in model 2 is significant by comparing the model 2 integrated likelihood with the model 1 likelihood:

```
pchisq(-2*(model1$loglik[2]-model2$loglik[2]),1,lower.tail=FALSE)
  Integrated
2.962689e-07
```

The `coxme` code returns a likelihood vector of three elements: the penalized partial likelihood for the case, where the frailty effect is 0 and the integrated likelihood and the penalized partial likelihood for the final solution. We compare the integrated likelihoods of the models in our tests. The family-specific frailty component σ_f^2 is very significant; the p-value of the test is 3×10^{-7}.

Table II. The four models for Framingham age-at-death data.

Model	Effect	Estimate	*p*-value*
Model 1. No random effects	Sex	−0.532	<0.0001
	Birth year	−0.006	0.0018
Model 2. Shared family	Sex	−0.574	<0.0001
	Birth year	−0.005	0.005
	σ_f^2	0.095	3e−07
Model 3. Correlated frailty	Sex	−0.631	<0.0001
	Birth year	−0.005	0.014
	σ_p^2	0.264	1e−07
Model 4. Shared family and correlated frailty	Sex	−0.628	<0.0001
	Birth year	−0.005	0.014
	σ_f^2	0.076	0.003
	σ_p^2	0.179	0.001

*Likelihood ratio test *p*-values (see text).

When we compare the model 3 correlated frailty (because of polygenes) likelihood with the model 1 (no frailty) likelihood, the likelihood ratio test p-value is 1.1×10^{-7}. Thus, the correlated frailty component σ_p^2 is significantly different from 0.

To test the significance of the correlated frailty in a model that already has a family frailty, we compare the model with both a family-specific and a correlated frailty (model 4) with the model with just the family frailty (model 2) and find that the likelihood ratio test p-value is 0.001. Thus, the correlated frailty component is significant even after accounting for family-specific frailty. Similarly, the likelihood ratio test of the shared family component when the correlated frailty component is already in the model is also significant ($p = 0.003$). We can conclude that both the family frailty and the correlated individual frailty components are significant. However, both frailty component variances are smaller in model 4 than in the single frailty component models (models 2 and 3). Thus, correlated individual frailties (model 3) can explain some of the shared family frailty observed in model 2 and vice versa.

7. Interpretation

The variance components are modeled on the log-hazard scale. Pankratz *et al.* [15] described that exponentiating the square root of the variance component provides information concerning the relative risk of the outcome that corresponds to the random effect. The estimated variances for the frailty effects in model 4 are $\sigma_f^2 = 0.076$ and $\sigma_p^2 = 0.179$. Because we expect the majority of families to fall within 1 SD of the norm, most individual families will have risk of death at a given age up to 32% larger or smaller than the average risk of death ($\exp(\sqrt{0.076}) = 1.32$). Similarly, $\exp(\sqrt{0.179}) = 1.53$, so most individual relative risks of dying at a given point in time are up to 1.53 times larger or smaller than the overall risk; within families, these individual relative risks are correlated, with the degree of correlation between pairs of individuals dictated by the relationship coefficient (twice the kinship coefficient). Both sex and birth year are significant covariates. The effect of birth year does not change substantially when frailties are modeled. The effect of age does change somewhat. The log hazard ratio for women versus men increases by almost 10% when the frailties are included in the model compared with the model with no frailties.

Ultimately, the models aid in exploring the underlying heterogeneity between families that is not accounted for by measured covariates. The analysis of both datasets shows that there are significant differences in frailty between the families. Each analysis can be thought of as assessing the correlation of age at death among related individuals in a genetic framework.

8. Summary

This paper provides a tutorial on frailty models applied in a biomedical and genetic study context. A background and overview of frailty models is presented in detail along with a discussion of more current topics in diagnostics and inference. Finally, we applied the methodology to real data, showing examples

of how to choose between frailty or no-frailty models. A highlight of the applications presents the use of frailty models in assessing genetic variability, which can be a potentially useful application of frailty models. Even Jonker *et al.* [43] posed the use of a gamma frailty model for linkage analysis. This tutorial hopefully contains a useful overview and application to frailty models in these contexts.

References

1. Vaupel JW, Manton KG, Stallard E. The impact of heterogeneity in individual frailty on the dynamics of mortality. *Demography* 1979; **16**(3):439–454.
2. Duchateau L, Janssen P. *The Frailty Model*. Springer: New York, 2008.
3. Klein JP, Moeschberger ML. *Survival Analysis: Techniques for Censored and Truncated Data*. Springer: New York, 1997.
4. Bremaud P. *Point Processes and Queues: Martingale Dynamics*. Springer-Verlag: New York, 1981.
5. Vaida F, Xu R. Proportional hazards model with random effects. *Statistics in Medicine* 2000; **19**:3309–3324.
6. Kosorok MR, Lee B-L, Fine JP. Robust inference for univariate proportional hazards frailty regression models. *The Annals of Statistics* 2004; **32**(4):1448–1491.
7. Hsu L, Chen L, Gorfine M, Malone K. Semiparametric estimation of marginal hazard function from case-control family studies. *Biometrics* 2004; **60**:936–944.
8. Gutierrez RG. Parametric frailty and shared frailty survival models. *The Stata Journal* 2002; **2**(1):22–44.
9. Hougaard P. Frailty models for survival data. *Lifetime Data Analysis* 1995; **1**:255–273.
10. Hougaard P. *Analysis of Multivariate Survival Data*. Springer: New York, 2000.
11. Cui S, Sun Y. Checking for the gamma frailty distribution under the marginal proportional hazards frailty model. *Statistica Sinica* 2004; **14**:249–267.
12. Rondeau V, Commenges D, Joly P. Maximum penalized likelihood estimation in a gamma-frailty model. *Lifetime Data Analysis* 2003; **9**:139–153.
13. Costigan TM, Klein JP. Multivariate survival analysis based on frailty models. In *Advances in Reliability*, AP Basu (ed.). North Holland: New York, 1993; 43–58.
14. Aslanidou H, Dey DK. Model determination for multivariate survival data under gamma frailty models. *ASA Proceedings of the Biometrics Section*, 1996; 292–297.
15. Pankratz VS, de Andrade M, Therneau TM. Random-effects Cox proportional hazard model: general variance components methods for time-to-event data. *Genetic Epidemiology* 2005; **28**:97–109.
16. R Development Core Team. *R: A language and environment for statistical computing*. R Foundation for Statistical Computing: Vienna, Austria, 2008.
17. Therneau T, Lumley T. *Survival: Survival analysis, including penalized likelihood*. R package version 2.34-1, 2008.
18. Therneau T. *Kinship: Mixed-effects Cox models, sparse matrices, and modeling data from large pedigrees.* R package version 1.1.0-19, 2008.
19. Ripatti S, Palmgren J. Estimation of multivariate frailty models using penalized partial likelihood. *Biometrics* 2000; **56**:1016–1022.
20. Hougaard P. Life table methods for heterogeneous populations: distributions describing the heterogeneity. *Biometrika* 1984; **71**(1):75–83.
21. Hanagal DD. Weibull extension of bivariate exponential regression model with different frailty distributions. *Statistical Papers* 2009; **50**(1):29–49.
22. Pickles A, Crouchley R. A comparison of frailty models for multivariate survival data. *Statistics in Medicine* 1995; **14**:1447–1461.
23. Heckman JJ, Singer B. A method for minimizing the impact of distributional assumptions in econometric models for duration data. *Econometrica* 1984; **52**:271–320.
24. Barker P, Henderson R. Small sample bias in the gamma frailty model for univariate survival. *Lifetime Data Analysis* 2005; **11**:265–284.
25. Lancaster T. Econometric methods for the duration of unemployment. *Econometrica* 1979; **47**(4):939–956.
26. Klein JP. Semiparametric estimation of random effects using the Cox model based on the EM algorithm. *Biometrics* 1992; **48**:795–806.
27. Nielsen GG, Gill RD, Andersen PK, Sørensen TI. A counting process approach to maximum likelihood estimation of frailty models. *Scandinavian Journal of Statistics* 1992; **19**:25–43.
28. Therneau TM, Grambsch PM. *Modeling Survival Data: Extending the Cox Model*. Springer-Verlag: New York, 2000.
29. Chen M-H, Ibrahim JG, Sinha D. Bayesian inference for multivariate survival data with a cure fraction. *Journal of Multivariate Analysis* 2002; **80**(1):101–126.
30. Jones G. Markov chain Monte Carlo methods for inference in frailty models with doubly-censored data. *Journal of Data Science* 2004; **2**:33–47.
31. Govindarajulu US, Glickman ME, D'Agostino RB. Modeling frailty as a function of observed covariates. *Journal of Statistical Theory and Practice* 2007; **1**(1):117–135.
32. Xue X. Multivariate survival data under bivariate frailty: an estimating equation approach. *Biometrics* 1998; **54**(4):1631–1637.
33. Glidden DV. Checking the adequacy of the gamma frailty model for multivariate failure times. *Biometrika* 1999; **86**:381–393.
34. Viswanathan B, Manatunga AK. Diagnostic plots for assessing the frailty distribution in multivariate survival data. *Lifetime Data Analysis* 2001; **7**:143–155.
35. Chen M-C, Bandeen-Roche K. A diagnostic for association in bivariate survival models. *Lifetime Data Analysis* 2005; **11**:245–264.

36. Murphy S. Consistency in a proportional hazards model incorporating a random effect. *Annals of Statistics* 1994; **22**:712–731.
37. Murphy S. Asymptotic theory for the frailty model. *Annals of Statistics* 1995; **23**(1):182–198.
38. Parner E. Asymptotic theory for the correlated gamma-frailty model. *Annals of Statistics* 1998; **26**(1):183–214.
39. Guo G, Rodriguez G. Estimating a multivariate proportional hazards model for clustered data using the EM Algorithm, with an application to child survival in Guatemala. *Journal of the American Statistical Association* 1992; **87**(420):969–976.
40. Pebley A, Stupp P. Reproductive patterns and child mortality in Guatemala. *Demography* 1987; **24**:43–60.
41. Dawber TR, Meadors GF, Moore Jr. FE. Epidemiological approaches to heart disease: the Framingham Study. *American Journal of Public Health and the Nation's Health* 1951; **41**(3):279–81.
42. Almasy L, Blangero J. Multipoint quantitative-trait linkage analysis in general pedigrees. *American Journal of Human Genetics* 1998; **62**(5):1198–1211.
43. Jonker MA, Bhulai S, Boomsma DI, Ligthart RSL, Posthuma D, Van Der Vaart AW. Gamma frailty model for linkage analysis with application to interval-censored migraine data. *Biostatistics* 2009; **10**(1):187–200.