PharmaSUG 2025 US – Paper ST-368

Incorporating Frailty into Time-to-Event Analysis: A Practical Approach with R frailtypack

Sunil Kumar Pusarla, Omeros, Seattle, US;
Dr. Avani Alla, Omeros, Frisco, US

ABSTRACT

Mixed models are commonly used in clinical trials with continuous and binary/ordinal endpoints involving repeated measures data, where both fixed and random effects are used to estimate the treatment effect. However, in oncology and rare disease trials with time-to-event endpoints, random effects are not routinely incorporated into the **Cox multivariable regression model**. This omission overlooks unobservable covariates that may influence outcomes. Additionally, correlated outcomes are rarely considered in survival analysis.

Frailty models extend the Cox model by accounting for unmeasured variables, particularly in clustered or recurrent event data. While SAS allows frailty inclusion in **PROC PHREG** via the **RANDOM** statement in a univariate analysis, it cannot handle multivariate joint frailty models. Moreover, it does not provide patient-level predicted frailty scores. Individual predicted frailty scores allow for more precise risk assessment, as population-level estimates alone may obscure individual risk levels, potentially impacting treatment and follow-up strategies.

The **R frailtypack package** addresses this gap by fitting shared, nested, joint, and additive frailty models using penalized likelihood estimation. By generating individual predicted frailty scores, it facilitates the identification of optimal cutoffs for stratifying patients into risk groups such as low, medium, and high. This approach enhances personalized care by ensuring that high-risk patients receive intensive treatment while low-risk patients may benefit from conservative management. Additionally, it improves prognostic accuracy, disease prevention, healthcare efficiency, and overall patient outcomes.

1. INTRODUCTION

This paper presents a comprehensive overview of the R frailtypack package and its core functionalities, emphasizing its application in survival analysis through the incorporation of frailty modeling. We illustrate the practical implementation of the R frailtypack package—using clinical data. Particular attention is given to the package's ability to generate individual-level frailty estimates, which serve as a powerful tool for risk stratification. Through our example, we demonstrate how individual frailty estimates generated by the package can enhance prognostic accuracy, guide treatment strategies, and improve patient outcomes in clinical trials by accounting for unobserved heterogeneity and complex data structures.

2. R FRAILTYPACK

The **R frailtypack** package offers comprehensive methods for estimating frailty models using **penalized likelihood estimation**, enabling improved model stability and convergence. Below are key concepts central to frailty model implementation using this package.

Incorporating Frailty in Survival Models

Frailty models introduce a random effect term, denoted as Z_i, to capture unmeasured heterogeneity. The hazard function for a frailty model is expressed as:

$$h_i(t) = Z_i\, h_0(t)\, \exp(\beta X_i)$$

Where:

- $h_i(t)$ = hazard function for individual i at time t
- Z_i = individual frailty term, typically gamma or log-normal distributed
- $h_0(t)$ = baseline hazard function
- β = regression coefficients for covariates X_i

The frailty term introduces variability among individuals or groups, improving the model's ability to account for unmeasured risk factors.

2.1 FRAILTY MODEL TYPES IN R FRAILTYPACK

The **R frailtypack** package offers several flexible model types:

- **Shared Frailty Model:** A common frailty term is shared among related subjects or groups (e.g., patients within the same hospital or family).
- **Nested Frailty Model:** Accounts for hierarchical clustering by nesting frailty terms within higher-level clusters.
- **Joint Frailty Model:** Simultaneously models recurrent events and terminal events (e.g., disease relapse and death).
- **Additive Frailty Model:** Extends shared frailty models by incorporating additive effects, improving flexibility in modeling complex dependencies.

Each of these models addresses different clinical scenarios, making **R frailtypack** versatile for real-world applications.

2.2 PENALIZED LIKELIHOOD ESTIMATION IN FRAILTYPACK

The **frailtypack** package employs a penalized log-likelihood approach, which balances model fit and complexity by introducing a penalty term. Penalized likelihood estimation improves model convergence and stability, especially in high-dimensional data with sparse observations. The penalized log-likelihood function is given by:

$$\mathcal{PL}(\theta) = \log L\,(\theta) - \frac{1}{2}\lambda\theta^{\top K}$$

where:

- $\log L\,(\theta)$ = log-likelihood of the observed data,
- θ = vector of model parameters,
- λ = smoothing (penalty) parameter, and
- K = penalty matrix, often based on the second derivative of the baseline hazard or the structure of random effects.

This method stabilizes the estimation of baseline hazard functions and frailty terms, enhancing model reliability.

2.3 STEP-BY-STEP IMPLEMENTATION OF R FRAILTYPACK

Step 1: Installing and Loading the Package

```
install.packages("frailtypack")

library(frailtypack)
```

Step 2: Data preparation

Data preparation is an important step for the successful application of frailty models in survival analysis, especially when addressing complex event structures involving recurrent and terminal events. In this study, to account for unobserved heterogeneity among individuals, we employed a joint shared frailty model, assuming the frailty term followed a gamma distribution. The dataset comprised survival time, event indicators, clinically relevant covariates (denoted by rf1 — rf6 in the code snippet), and a clustering variable identifying individual subjects (id). All known risk factors of overall survival (OS) were incorporated as covariates in the model.

We considered two time-to-event outcomes in this analysis:

- Event A
  - Events can be recurrent (e.g. Relapse) or non-recurrent competing (e.g. Acute GVHD)
- Death (mortality)

To accommodate the joint nature of these outcomes, we defined four mutually exclusive event scenarios based on the occurrence or absence of Event A and death using conditional logic:

- **Case 1:** Death occurred without prior occurrence of Event A (**dead = 1**, **eventA = 0**)
- **Case 2:** Death occurred following Event A (**dead = 1**, **eventA = 1**)
- **Case 3:** Neither Event A nor death was observed (**dead = 0**, **eventA = 0**)
- **Case 4:** Event A occurred, but the patient survived (**dead = 0**, **eventA = 1**)

To meet the input structure required by the **frailtypack** package — specifically, a start-stop format with separate indicators for recurrent and terminal events—the dataset was split into subsets corresponding to the four defined cases:

Case 1 (Death only)	Case 2 (Event A followed by death)	Case 3 (No Event A, No death)	Case 4 (Event A only)
A single time interval from start to death	Two consecutive time intervals: 1) From start to the occurrence of Event A and 2) From Event A to death	A single censored time interval	Two time intervals with no terminal event: 1) From start to Event A and 2) From Event A to end of follow-up
t.start = 0; **t.stop = dtxsurv** **tevent = 1;** **event2 = 0;**	**t.start = 0;** **t.stop = dtxeventA;** **tevent = 0;** **event2 = 1;** **t.start = dtxeventA;** **t.stop = dtxsurv;** **tevent = 1;** **event2 = 0;**	**t.start = 0;** **t.stop = dtxsurv;** **tevent = 0;** **event2 = 0;**	**t.start = 0;** **t.stop = dtxeventA;** **tevent = 0;** **event2 = 1;** **t.start = dtxeventA;** **t.stop = dtxsurv;** **tevent = 0;** **event2 = 0;**

*dtxsurv = time to death *dtxeventA = time to eventA

The resulting subsets were concatenated into a single dataset (**bivariate**) and subsequently sorted by patient identifier and time. The dataset was then reviewed to ensure structural consistency and completeness.

```
bivariate <-  read("location of the dataset")
```

Step 3: Fitting a Joint Frailty Model for Recurrent Events and Terminal Events

To simultaneously model recurrent events and a terminal event, a joint frailty model was fitted using the **frailtyPenal function** from the frailtypack package. This approach accounts for the dependence between repeated events and the terminal event through a shared frailty term at the cluster (patient) level. The model was specified with 10 interior knots and a penalization parameter (kappa) set to 1000 to control the smoothness of the baseline hazard. The **Andersen-Gill formulation (recurrentAG = TRUE)** was employed to handle the recurrent event process over time.

Covariates were included separately for the recurrent and terminal components of the model. For the recurrent event process (e.g., EventA), categorical clinical risk factors such as rf1 through rf6 were included. The terminal event component incorporated distinct covariates using the **formula.terminalEvent** argument, modeling the effects of the same or other risk factors (e.g., rf1 to rf6) on the hazard of death. The model allowed for flexible estimation of covariate effects while accounting for unobserved heterogeneity. The joint frailty model was implemented as follows:

```
final  <- frailtyPenal(Surv(t.start, t.stop, eventA) ~ cluster(id) +
as.factor(rf1) + as.factor(rf2), …….. +    as.factor(rf6) +
terminal(tevent),

               formula.terminalEvent = ~ as.factor(rf1) + as.factor(rf2),
…….. + as.factor(rf6),

               data = bivariate,

               n.knots = 10,

               kappa = c(1000),

               recurrentAG = TRUE)

summary(final)

final
```

Running **summary(final)** yields model outputs including hazard ratios (HRs) with 95% confidence intervals (CIs) for each covariate in both the recurrent and terminal components. Shown below are the output values of a few covariates.

Recurrent Events (Event A)	HR (95% CI)	Terminal Event	HR (95% CI)
RF1	0.37 (0.24 0.57)	RF1	0.4 (0.61–0.85)
RF2	1.51 (1.25–1.84)	RF2	1.35 (1.14–2.20)
RF3	0.54 (0.0 -0.7)	RF3	1.13 (0.93–1.33)

Results also provide estimates for key frailty parameters—**theta (θ)**, the variance of random effects, and **alpha (α)**, the association between recurrent and terminal events—alongside likelihood-based fit diagnostics:

Model Fit and Frailty Estimates:

- **Theta (Variance of Frailties):** 0.79 (SE = 0.0276), $p < 1e\text{-}16$
- **Alpha (Frailty-Terminal Association):** 2.20 (SE = 0.2345), $p < 1e\text{-}16$
- **Penalized Marginal Log-Likelihood:** –3.56
- **Likelihood Cross-Validation (LCV):** 0.00313
- **Convergence Criteria Met:** Gradient < 1e-9, Parameters < 1e-6

Both the parameters **theta (θ)** and **alpha (α)** provide an understanding of the frailty effect and dependency between survival and recurrent events. **Theta (θ)** denotes the variance of the frailty distribution, typically assumed to follow a Gamma or log-normal form. It quantifies unobserved heterogeneity in patient risk: higher values indicate substantial variability not captured by observed covariates, while lower values suggest a more homogeneous population. When **θ = 0**, the frailty effect is

negligible, and the model simplifies to a standard survival or recurrent event model without random effects.

Alpha (α) measures the strength of association between the recurrent events and the terminal event. An **α > 1** indicates that individuals with more frequent recurrences are at increased risk of the terminal event, such as mortality. In contrast, **α ≈ 1** suggests independence between the two processes.

Step 4: Predicting Individual Frailty Scores for Risk Stratification

Because the frailty estimates pertain to each subject rather than each event, we first restricted the dataset to one representative record per patient. This dataset (**unique**) was used to generate subject-specific frailty predictions from the fitted joint model using the **$frailty.pred** component**.** These values represent the estimated random effects (**frailty scores**) for each patient and quantify the individual-level deviation from the population baseline hazard.

```
unique <- bivariate[!duplicated(bivariate$ID), ]

  predictions <- unique$frailty.pred
```

2.4 RISK STRATIFICATION

Following the derivation of individual risk scores, patients were stratified into clinically meaningful risk categories to facilitate downstream interpretation and decision-making. The selection of appropriate cut-off points for stratification is highly context-dependent and may be influenced by disease-specific features, patient demographics, clinical profiles, or relevant biomarker distributions. These thresholds can be determined using data-driven techniques such as maximally selected rank statistics (e.g. gamma distribution <1 or >1), receiver operating characteristic (ROC) curve analysis, or established clinical threshold benchmarks.

In the below **figures 1 and 2**, risk categories were defined based on the distribution of predicted individual frailty scores: patients with scores below the mean were classified as low-risk, those between the mean and the third quartile (75th and 80th percentile respectively) as medium-risk, and those above the 75th and 80th percentile respectively as high-risk. This stratification schema was chosen to provide a balanced and interpretable framework for assessing outcome variability across risk groups. To complement the risk-based classification, a frequency distribution plot of stratified patients is presented (see figure below). This visualization aids in understanding the distributional properties of the risk scores and the proportional representation of each group, offering additional insights into the heterogeneity of the study population.

Figure 1	Figure 2
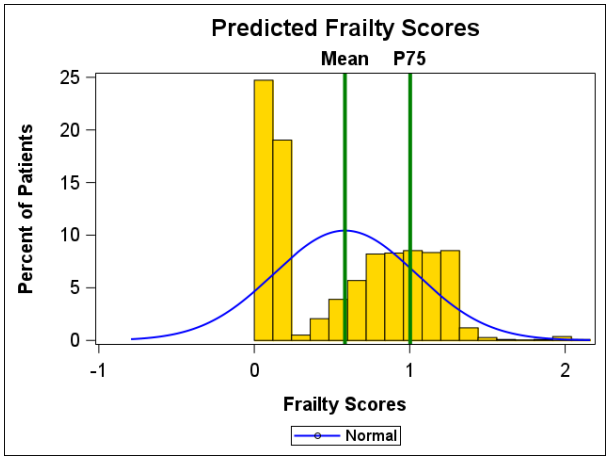	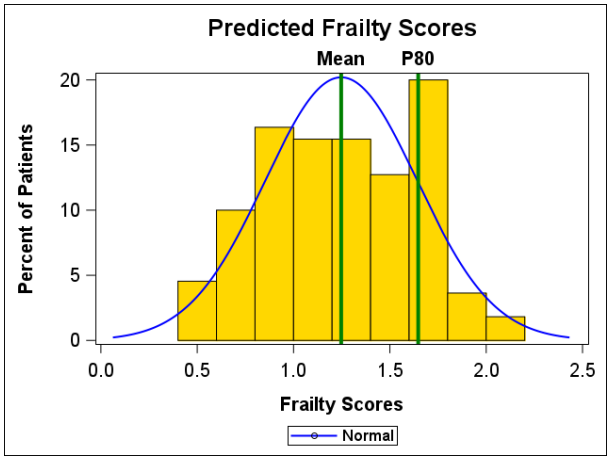

Depending on the clinical or research objectives, patients may be stratified into two (e.g., low vs. high risk), three (e.g., low, intermediate, and high risk), or more groups to capture a spectrum of risk. Such stratification not only supports individualized clinical decision-making and targeted therapeutic interventions but also enhances the efficiency of healthcare resource allocation. Moreover, it enables the identification of high-risk subpopulations that may benefit from enhanced surveillance, early intervention, or enrollment in clinical trials investigating novel treatments.

2.5 SURVIVAL PLOT

To demonstrate the prognostic value of risk stratification, Kaplan–Meier survival curves were generated for each risk group. The plot provides a clear visual comparison of overall survival across strata. As shown in the figure below, patients classified as high-risk exhibited a markedly shorter median survival time relative to those in intermediate- and low-risk groups. This pronounced difference underscores the clinical importance of incorporating risk stratification into treatment decision-making. Uniform treatment approaches may be suboptimal, as the variation in survival outcomes across risk groups suggests the need for tailored therapeutic strategies.

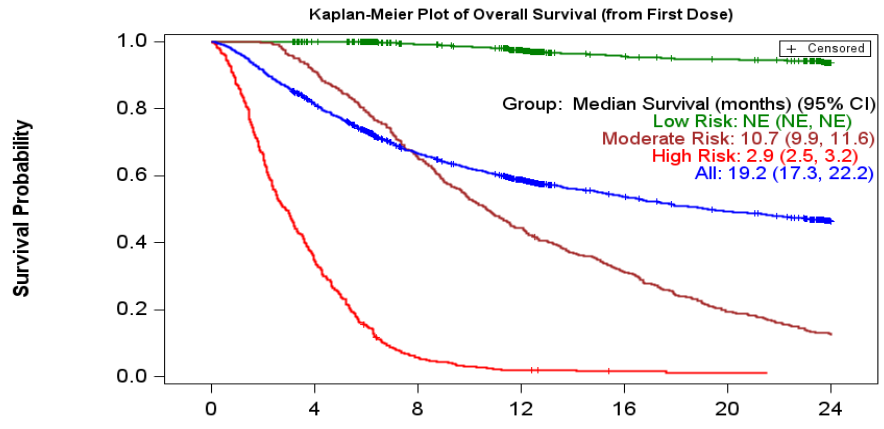

Patients At Risk							
Low Risk	1078	1047	940	858	714	701	639
Moderate Risk	557	508	363	247	174	108	69
High Risk	545	190	31	11	5	4	3
All	2180	1745	1334	1116	893	813	711

Such visualizations not only enhance the interpretability of model outputs but also highlight the potential for risk-guided care. By identifying patients at elevated risk of adverse outcomes, clinicians can make more informed decisions regarding treatment intensity, frequency of follow-up, and inclusion in clinical trials evaluating novel therapies.

3. CLINICAL APPLICATIONS

The R package **frailtypack** has demonstrated versatility across various clinical research domains by facilitating the analysis of correlated survival data through advanced frailty models.

Clinical Applications:

- **Oncology Trials:** In cancer research, frailtypack enables the modeling of recurrent events such as tumor relapses and the assessment of terminal events like patient mortality. By employing joint frailty models, researchers can simultaneously analyze these correlated events, leading to more accurate predictions of recurrence risks and informing post-treatment strategies.
- **Cardiology Studies:** In cardiology, frailtypack assists in evaluating the progression of heart failure by modeling recurrent hospital admissions. The package's ability to handle clustered and recurrent survival times allows for a comprehensive understanding of patient readmission patterns, which is crucial for developing effective intervention strategies.
- **Rare Disease Research:** For rare diseases, where patient cohorts are often small and data may be hierarchically structured, frailtypack's nested frailty models are particularly beneficial. These models account for multiple levels of clustering, enabling the identification of high-risk patients who may require aggressive interventions.

4. CONCLUSION AND FUTURE DIRECTIONS

The R package **frailtypack** provides a robust framework for incorporating frailty effects into survival models, addressing unmeasured heterogeneity that often complicates clinical data analysis. Its support for a wide range of models—including shared, nested, joint, and additive frailty structures—makes it highly adaptable to complex study designs and hierarchical or correlated data.

A key advantage of frailtypack lies in its ability to enhance prognostic accuracy through individual frailty estimates. This facilitates precise risk stratification, enabling clinicians to personalize treatment intensity, follow-up schedules, and eligibility for clinical trials. By tailoring care based on predicted risk, healthcare systems can both improve outcomes for high-risk patients and minimize unnecessary interventions in low-risk groups.

Looking ahead, future enhancements to frailtypack may include integration with machine learning techniques to enable dynamic, data-driven risk prediction. Expanded visualization tools could further improve interpretability of model outputs, and optimization for computational efficiency would facilitate its use with large-scale datasets. Continued development of frailtypack holds substantial promise for advancing personalized medicine and supporting data-informed clinical decision-making across diverse therapeutic areas.

Disclaimer: The opinions expressed in this paper are those of the authors and do not represent the views of Omeros Corporation.

5. REFERENCES

https://cran.r-project.org/web/packages/frailtypack/frailtypack.pdf

https://www.rdocumentation.org/packages/frailtypack/versions/3.6.5/topics/frailtypack-package

Rondeau, V., Mazroui, Y., & Gonzalez, J. R. (2012). frailtypack: An R Package for the Analysis of Correlated Survival Data with Frailty Models. Journal of Statistical Software, 47(4).

6. CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Names: Sunil Kumar Pusarla
Avani Alla
Affiliation: Omeros Corporation
201 Elliott Avenue West, Seattle, WA 98119
Work Phone: 206-676-5000
E-mail: **skumar@omeros.com/sunilpusarla@gmail.com**
aalla@omeros.com/dravareddy@gmail.com