

RESEARCH ARTICLE

Open Access



Flexible semiparametric joint modeling: an application to estimate individual lung function decline and risk of pulmonary exacerbations in cystic fibrosis

Dan Li¹ , Ruth Keogh², John P. Clancy³ and Rhonda D. Szczesniak^{4*}

Abstract

Background: Epidemiologic surveillance of lung function is key to clinical care of individuals with cystic fibrosis, but lung function decline is nonlinear and often impacted by acute respiratory events known as pulmonary exacerbations. Statistical models are needed to simultaneously estimate lung function decline while providing risk estimates for the onset of pulmonary exacerbations, in order to identify relevant predictors of declining lung function and understand how these associations could be used to predict the onset of pulmonary exacerbations.

Methods: Using longitudinal lung function (FEV₁) measurements and time-to-event data on pulmonary exacerbations from individuals in the United States Cystic Fibrosis Registry, we implemented a flexible semiparametric joint model consisting of a mixed-effects submodel with regression splines to fit repeated FEV₁ measurements and a time-to-event submodel for possibly censored data on pulmonary exacerbations. We contrasted this approach with methods currently used in epidemiological studies and highlight clinical implications.

Results: The semiparametric joint model had the best fit of all models examined based on deviance information criterion. Higher starting FEV₁ implied more rapid lung function decline in both separate and joint models; however, individualized risk estimates for pulmonary exacerbation differed depending upon model type. Based on shared parameter estimates from the joint model, which accounts for the nonlinear FEV₁ trajectory, patients with more positive rates of change were less likely to experience a pulmonary exacerbation (HR per one standard deviation increase in FEV₁ rate of change = 0.566, 95% CI 0.516–0.619), and having higher absolute FEV₁ also corresponded to lower risk of having a pulmonary exacerbation (HR per one standard deviation increase in FEV₁ = 0.856, 95% CI 0.781–0.937). At the population level, both submodels indicated significant effects of birth cohort, socioeconomic status and respiratory infections on FEV₁ decline, as well as significant effects of gender, socioeconomic status and birth cohort on pulmonary exacerbation risk.

Conclusions: Through a flexible joint-modeling approach, we provide a means to simultaneously estimate lung function trajectories and the risk of pulmonary exacerbations for individual patients; we demonstrate how this approach offers additional insights into the clinical course of cystic fibrosis that were not possible using conventional approaches.

Keywords: Bayesian analysis, Cystic fibrosis, Functional data analysis, Longitudinal studies, Mixed model analysis, Pulmonary decline, Pulmonary function, Joint modeling, Registry analyses, Spline regression

*Correspondence: rhonda.szczesniak@cchmc.org

⁴Division of Biostatistics and Epidemiology, Cincinnati Children's Hospital Medical Center, MLC 5041, 3333 Burnet Ave, Cincinnati, OH 45229, USA
Full list of author information is available at the end of the article

Background

Maintaining pulmonary function is essential for survival in individuals with cystic fibrosis (CF), a lung disease that currently affects nearly 70,000 individuals worldwide [1]. Lung function is measured primarily as forced expiratory volume in 1 s of percent predicted (hereafter, FEV₁), which is the primary marker of disease severity in individuals with CF. FEV₁ declines in a nonlinear fashion over age and exhibits substantial heterogeneity both between patients and within an individual patient over time (Fig. 1). Most individuals with CF have decreased FEV₁ over age, but initial FEV₁ and rate of decline vary between patients. Linear mixed modeling is an established approach to estimate age-related FEV₁ progression. Historically, CF epidemiologic studies have not accounted for nonlinearity in the FEV₁ trajectory; however, recent approaches have included piecewise polynomials, either in the form of regression splines or change-point models, to estimate decline in FEV₁ [2, 3]. These approaches have shown that FEV₁ decline is variable with maximal loss occurring in adolescence and early adulthood.

Meanwhile, the clinical course of CF is often marked by the occurrence of an acute respiratory event known as a pulmonary exacerbation, which intensifies the severity of this chronic disease. Increased pulmonary symptoms and decreased lung function and weight are common clinical indicators of this event [4]. A study of the United States Cystic Fibrosis Foundation Patient Registry showed that, after having a pulmonary exacerbation, 25% of patients' lung function levels did not recover to baseline (pre-pulmonary exacerbation) levels [5]; this study identified lower socioeconomic status (use of Medicaid insurance), malnutrition, having respiratory infections and pathogens, and being female as risk factors. Another study

identified pulmonary exacerbations as a risk factor for more rapid lung function decline [6].

Although much information has been gleaned, these and other epidemiologic studies in CF have modeled the longitudinal and time-to-event processes of lung function and pulmonary exacerbations separately. Efficiency and additional information about the disease course may be gained by modeling them simultaneously. Joint modeling of longitudinal and event processes have enjoyed a renaissance in recent years among biostatisticians and epidemiologists, largely due to software developments and modern applications demonstrating their clinical utility for the purposes of monitoring disease progression and making predictions [7]. Developments in joint models have also been motivated by data on CF disease progression, focusing on simultaneously fitting a longitudinal submodel of lung function and a survival submodel. For example, Schluchter et al. [8] developed a joint longitudinal-survival model to create prognostic indicators of CF disease severity, such as predicted age at death and other empirical Bayes estimates of parameters in the longitudinal submodel, such as slope, which could be used to estimate rate of decline in FEV₁ adjusted for survivor bias. They combined a linear mixed effects model for longitudinal FEV₁ with a Gaussian model for age at death and applied it to data taken from a local CF center. They later extended this model to account for left truncation, as follow-up does not necessarily begin at age zero for CF patients [9]. Such joint models can be classified as shared parameter approaches, meaning that FEV₁ and the event of interest are assumed to be conditionally independent given a set of random effects.

In this study, we utilize the aforementioned advantages of joint modeling in the shared-parameter framework and combine them with flexible semiparametric regression splines to estimate FEV₁ trajectories in the longitudinal submodel. In this context, joint modeling would account for informative dropout due to pulmonary exacerbation events; not accounting for these events could bias estimates of the FEV₁ trajectories. Our objectives are to identify relevant predictors of lung function trajectories and pulmonary exacerbation events, and to understand how associations found from a joint model could be used to predict the onset of pulmonary exacerbations for individual CF patients.

Methods

Study design and cohort

This longitudinal cohort study consisted of clinical encounter and hospitalization data from the U.S. Cystic Fibrosis Foundation Patient Registry (CFFPR) between January 1, 2003, and December 31, 2011. This registry has been tracking CF demographic and clinical variables for

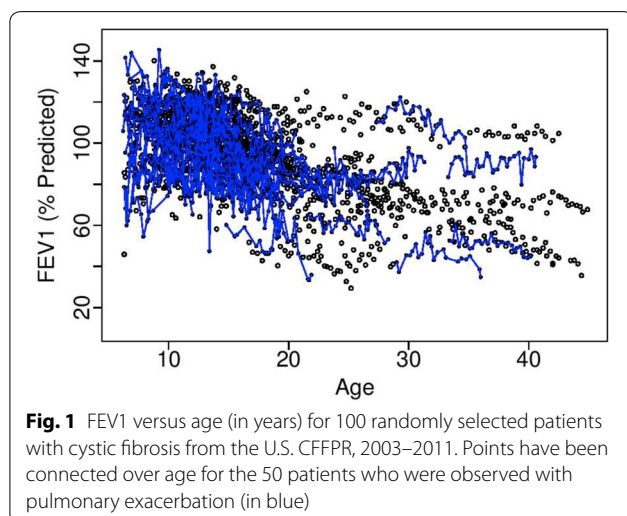


Fig. 1 FEV₁ versus age (in years) for 100 randomly selected patients with cystic fibrosis from the U.S. CFFPR, 2003–2011. Points have been connected over age for the 50 patients who were observed with pulmonary exacerbation (in blue)

more than 40 years, and has been thoroughly described elsewhere [10]. Because the majority of relevant predictors of lung function decline and pulmonary exacerbation events were consistently documented beginning in 2003, we considered data available from this year onward. Patients younger than 6 years of age were excluded due to the potential for unreliable data from pulmonary function testing. Data from those older than 45 years of age were excluded due to the possibility that these individuals have milder phenotypes that are not representative of typical CF disease progression.

The longitudinal outcome of interest, FEV₁ % predicted (hereafter, FEV₁), was obtained at each clinical encounter; % predicted values were calculated using Wang and Hankinson reference equations [11, 12]. The event outcome of interest, onset of first pulmonary exacerbation, was subject to censoring. A pulmonary exacerbation was considered to have occurred if documented in the CFFPR as warranting treatment with intravenous antibiotics. Analyses were restricted to patients with at least ten measurements of FEV₁. Longitudinal data were included up to the time of the first pulmonary exacerbation on record and censored thereafter. Baseline was defined as the time point at which the first FEV₁ measurement was available. Covariates of interest were selected from the literature on FEV₁ [13, 14] and pulmonary exacerbation [5, 6] as separate models, and included age, gender, initial FEV₁ measure and baseline lower socioeconomic status (lower SES, defined as having only state/federal or no insurance; recorded as 1, and 0 otherwise); a birth cohort covariate was used to adjust for potential left truncation bias; time-varying covariates included CF-related diabetes (CFRD, with or without fasting hyperglycemia), positive cultures for *methicillin-resistant staphylococcus aureus* (MRSA), *Burkholderia cepacia* (*B. cepacia*), and *Pseudomonas aeruginosa* (*Pa*).

Joint model formulation and estimation

Each joint model consists of two linked submodels, a mixed model to fit longitudinal FEV₁ and a Weibull model to fit pulmonary exacerbation data.

Longitudinal submodel for FEV₁

Both the separate modeling approach and the joint modeling approach require specification of a model for the longitudinal measure of FEV₁. In this subsection we describe this model before linking it to the time-to-event model in subsequent sections. Suppose there are *N* patients in the CFFPR indexed by *i* = 1, 2, . . . , *N*. Let *FEV*_{1*ij*} represent lung function measured for the *i*th patient at the *j*th time, represented as *age*_{*ij*} (corresponding to patient age, in years), *j* = 1, . . . , *n_i*. The longitudinal submodel consists of a linear mixed model with random effects:

$$FEV_{1ij} = f(\text{age}_{ij}) + m_i(\text{age}_{ij}) + W_{1i}(\text{age}_{ij}) + \varepsilon_{ij}, \tag{1}$$

where the population-level mean response *f*(*age*_{*ij*}) is modeled as a continuous function of time, which combines the time-related fixed effects terms and regression splines¹⁵; *m_i*(*age*_{*ij*}) = **x**'_{1*i*}(*age*_{*ij*})**α**₁ refers to the subject-level fixed effects. The vector **x**_{1*i*}(*age*_{*ij*}) represents static and time-varying covariates defined above, and the vector **α**₁ contains their corresponding regression coefficients. The expression *W*_{1*i*}(*age*_{*ij*}) = **z**'_{1*i*}(*age*_{*ij*})**U**_{*i*} incorporates random effects that describe how subject-specific true FEV₁ levels deviate from their expected behavior, where the vector **U**_{*i*} corresponds to subject-specific random slopes and intercepts. We used the usual Laird-Ware form⁸, *W*_{1*i*}(*age*_{*ij*}) = *U*_{0*i*} + *U*_{1*i*}*age*_{*ij*}, corresponding to **z**_{1*i*}(*age*_{*ij*}) = (1, *age*_{*ij*})'; this specification allows individuals to have varying baseline FEV₁ measurements and different rates of change in FEV₁ over time. Lastly, the term $\varepsilon_{ij} \sim N(0, \sigma_\varepsilon^2)$ denotes zero-mean Gaussian measurement error.

We considered two nonlinear representations of *f*(*age*_{*ij*}), which we hereafter refer to as semiparametric and cubic, and examined both in the joint modeling framework. First, we used penalized splines with a cubic truncated power basis in order to provide smooth estimates of the longitudinal course of FEV₁ measurements. Second, we considered only global cubic polynomials to represent population-level FEV₁ decline by excluding the basis functions in *f*(*age*_{*ij*}). Details of the two formulations are available in the ‘‘Appendix 1’’. Taking the first derivative with respect to age in Eq. (1) yields overall and subject-specific estimates of the rate of change in FEV₁, and has been used previously to model FEV₁ decline [2].

Time-to-event submodel for pulmonary exacerbation

Let *T_i* denote the possibly censored survival time to a pulmonary exacerbation event for the *i*th patient. In a Weibull model, we assume that the survival time follows a Weibull distribution, that is *t_i* ~ Weibull(*k*, *μ_i*(*t*)), where *μ_i*(*t*) = exp{**x**'_{2*i*}(*t*)**α**₂ + *W*_{2*i*}} and *k* > 0. The hazard at time *t* for the *i*th patient is

$$h_i(t) = kt^{(k-1)} \exp(\mathbf{x}'_{2i}(t)\boldsymbol{\alpha}_2 + W_{2i}), \tag{2}$$

which monotonically increases with time if *k* > 1, decreases if *k* < 1, and reduces to the exponential hazard and remains constant if *k* = 1. The vectors **x**_{2*i*}(*t*) and **α**₂ represent possibly time-dependent covariates and the corresponding regression coefficients. Covariates **x**_{2*i*}(*t*) need not have elements in common with **x**_{1*i*}(*age*) in the longitudinal model as shown in Eq. (1). Notice that *W*_{2*i*} could have a time-dependent form [15], but we do not consider it here since this level of complexity is not

required for our data. Similar to the form of $W_{1i}(\text{age})$ in Eq. (1), $W_{2i} = \theta_0 U_{0i} + \theta_1 U_{1i}$ corresponds to patient-specific random intercepts and slopes, but has distinct regression parameter coefficients, θ_0 and θ_1 . If none of the covariates vary over time, W_{2i} reduces to 0 in the absence of random effects. The main idea of this approach is to connect the longitudinal and survival processes with a latent bivariate Gaussian process. We can add a frailty term in W_2 , i.e., $W_{2i} = \theta_0 U_{0i} + \theta_1 U_{1i} + U_{3i}$, where $(U_{0i}, U_{1i})'$ follows a bivariate Gaussian distribution, while $U_{3i} \sim N(0, \sigma_3^2)$, independent of $(U_{0i}, U_{1i})'$.

Submodel links

Pulmonary exacerbation event times were associated with longitudinal FEV₁ measurements through stochastic dependence between W_{1i} and W_{2i} from Eqs. (1)–(2) by assuming:

$$W_{1i}(\text{age}) = U_{0i} + U_{1i} \text{age}; \quad (3)$$

$$W_{2i}(t) = W_{2i} = \theta_0 U_{0i} + \theta_1 U_{1i}. \quad (4)$$

The subject-specific random intercept and slope, depicting varying initial values and rates of FEV₁ decline for each patient after accounting for the nonlinear FEV₁ trajectory, are contained in the vector U_i , which is shared between the longitudinal and time-to-event submodels. Here U_i is assumed to follow a bivariate normal distribution $N(0, \Sigma)$, where Σ has diagonal entries σ_0^2 , σ_1^2 and off-diagonal entries σ_{01} . The joint model with (4) allows both the random intercept U_{0i} and slope U_{1i} , involved in (3), to affect the risk of the event. Thus, deviations of the patient-specific FEV₁ trajectories from the population-level FEV₁ curve enter the pulmonary exacerbation model in the form of random intercepts and slopes. If this type of association exists between the longitudinal FEV₁ and pulmonary exacerbation event processes, then inference from the joint model should be less biased and more efficient, compared to modeling the processes separately [16]. Given that absolute FEV₁ (intercept) and rate of change in FEV₁ (slope) are on different scales, we used the estimated SDs of the random effects to express each corresponding HR as per SD change in the respective covariate. These quantities were obtained from the joint model by estimating $\exp\{\theta_0 * \sigma_0\}$ and $\exp\{\theta_1 * \sigma_1\}$, where θ_0 and θ_1 are the association parameters estimated ordinarily from the HRs $\exp\{\theta_0\}$ and $\exp\{\theta_1\}$, and σ_0 and σ_1 are the respective SDs of intercept U_{0i} and slope U_{1i} as defined previously.

MCMC sampling procedure

Conventional likelihood-based estimation of the parameters requires integration of the two submodels over the distribution of random effects. As the number of random

effects in the model increase, the dimension of integration in the joint model likelihood increases; as a result, parameter estimation typically involves specialized numerical algorithms and becomes computationally burdensome [17]. Alternatively, we employed Markov-Chain Monte-Carlo (MCMC) via Gibbs sampling in WinBUGS [18] to simulate data from the respective posterior distributions under each model. The highest posterior density (HPD) and accompanying 95% credible interval (CI) were used to estimate each parameter of interest. Specification of priors and sampling procedures are detailed in the “Appendix 2”.

Model comparisons

We calculate the deviance information criterion (DIC) to compare the performance of the separate and joint models [19]. The DIC measure balances the fit of a model to the data with its complexity. The components of DIC for the two submodels, denoted as DIC₁ and DIC₂, were also provided to evaluate their relative contributions to the total DIC score. Formulas are provided in the “Appendix 3”. A smaller value of DIC indicates the preferred model.

Research ethics

The Internal Review Board of Cincinnati Children’s Hospital Medical Center (Cincinnati, OH, USA) approved the study.

Results

Study cohort

There were 7672 individuals who met inclusion criteria, yielding a total of 136,051 FEV₁ measurements. Median (IQR) follow-up was 5.8 (4.0–7.7) years; 3349 (43.7%) of the patients experienced a pulmonary exacerbation. Median age at entry was 10.8 (6.9–16.6) years; 4298 (56%) of the patients were male; 2690 (35%) had lower SES (only state/federal or no insurance). In total, 1036 (13.5%), 1176 (15.3%), 2144 (27.9%) and 3316 (43.2%) of the patients were born before 1981, between 1981 and 1988, between 1989 and 1994, and after 1994, respectively. Few patients had infection with *B. cepacia* (2.5%), but 1589 (20.7%) had *Pa* infection at baseline. Mean (SD) FEV₁ at entry was 92.3 (19.6) % predicted.

Longitudinal submodel for FEV₁

Both the separate and joint modeling of longitudinal FEV₁ indicated that decline was nonlinear over age (Table 1, Semiparametric longitudinal submodel). In both models, having higher FEV₁ at entry corresponded to higher FEV₁ over time; lower SES, diagnosis with CFRD and the presence of infections were associated with lower FEV₁ over age. The younger birth cohorts appeared to have similar FEV₁ progression, while being

Table 1 Posterior estimates for lung function decline and pulmonary exacerbation onset based on separate and joint models

Parameter	Separate model (Model I)		Joint model (Model III)	
	Posterior mean	95% HPD CI	Posterior mean	95% HPD CI
<i>Semiparametric longitudinal submodel for FEV₁</i>				
Curve				
β_0 , intercept	18.57	[16.33, 20.68]	16.89	[14.63, 20.39]
β_1 , age	2.846	[2.545, 3.143]	3.121	[2.74, 3.382]
β_2 , age ²	- 0.1295	[- 0.1393, - 0.1197]	- 0.1367	[- 0.1481, - 0.124]
β_3 , age ³	- 0.0012	[- 0.0014, - 0.0009]	- 0.0016	[- 0.0017, - 0.0014]
b_1	0.0058	[0.0044, 0.0077]	0.0073	[0.0058, 0.0084]
b_2	0.0083	[0.0076, 0.0089]	0.0103	[0.0095, 0.0111]
b_3	- 0.0143	[- 0.0187, - 0.0114]	- 0.0206	[- 0.0248, - 0.0167]
b_4	- 0.0023	[- 0.0056, 0.0027]	0.0029	[- 0.001, 0.0085]
b_5	0.0041	[- 0.0015, 0.0101]	0.0015	[- 0.0033, 0.007]
b_6	- 0.0011	[- 0.012, 0.0102]	0.0008	[- 0.0095, 0.0101]
Baseline FEV1	0.6828	[0.6722, 0.6933]	0.6834	[0.6676, 0.6945]
Male	- 0.1665	[- 0.6403, 0.3137]	- 0.1650	[- 0.6054, 0.2971]
Birth cohort				
< 1981	11.37	[9.651, 12.97]	12.39	[10.89, 13.95]
1981–1988	3.766	[2.749, 4.839]	4.936	[3.995, 5.916]
1989–1994	- 0.1912	[- 0.698, 0.3584]	0.0608	[- 0.4059, 0.5908]
> 1994 (reference)				
Lower SES	- 0.6318	[- 1.075, - 0.1686]	- 0.5984	[- 1.062, - 0.1119]
CFRD	- 0.4528	[- 0.765, - 0.1465]	- 0.3695	[- 0.6731, - 0.0633]
MRSA	- 0.8209	[- 1.033, - 0.6121]	- 0.8114	[- 1.016, - 0.6134]
<i>B. cepacia</i>	- 0.8318	[- 1.531, - 0.1367]	- 0.7705	[- 1.441, - 0.1022]
<i>Pa</i>	- 0.5661	[- 0.6963, - 0.4314]	- 0.5673	[- 0.6971, - 0.4373]
Sources of variation				
σ^2 (measurement error)	59.29	[58.83, 59.76]	59.24	[58.77, 59.73]
σ_0^2 (between patients, intercept)	324.9	[307.6, 342.2]	326.8	[311, 343.8]
σ_1^2 (between patients, slope)	2.566	[2.441, 2.694]	2.602	[2.484, 2.727]
σ_{01} (covariance, intercept and slope)	- 25.51	[- 26.82, - 24.16]	- 25.76	[- 27.11, - 24.49]
<i>Weibull event submodel for pulmonary exacerbation</i>				
k	2.467	[2.395, 2.537]	2.564	[2.487, 2.634]
Intercept	- 3.705	[- 3.988, - 3.416]	- 3.568	[- 3.804, - 3.296]
θ_0 (random intercept)	-	-	- 0.0086	[- 0.0137, - 0.0036]
θ_1 (random slope)	-	-	- 0.3532	[- 0.4104, - 0.2979]
Baseline age	- 0.0021	[- 0.0187, 0.0147]	0.0023	[- 0.0131, 0.0176]
Baseline FEV1	- 0.0158	[- 0.0178, - 0.0138]	- 0.0198	[- 0.0218, - 0.0179]
Male	- 0.2973	[- 0.3618, - 0.2314]	- 0.2940	[- 0.3606, - 0.2267]
Birth cohort				
< 1981	- 0.3785	[- 0.7538, - 0.0028]	- 0.5348	[- 0.9056, - 0.1726]
1981–1988	- 0.1805	[- 0.3947, 0.0182]	- 0.2348	[- 0.4405, - 0.0339]
1989–1994	0.1431	[0.031, 0.2544]	0.1261	[0.0178, 0.2335]
> 1994 (reference)				
Lower SES	0.1173	[0.049, 0.1872]	0.1059	[0.0382, 0.1804]

* In the Bayesian sense, a 95% CI that excludes zero indicates statistical significance at the 0.05 level. Parameters b_1 – b_6 are regression coefficients of the cubic truncated power functions defined in "Appendix 1"

in an older birth cohort was associated with higher FEV₁. Although results were consistent between the two models in Table 1, estimates from the joint model tended to be lower for some of the covariate effects. This could be attributable to incorporating the effect of pulmonary exacerbation through the shared intercept and slope terms in the joint modeling. Both models had similar variance component estimates and indicated substantial heterogeneity in the FEV₁ response, in terms of measurement error, as well as between and within subjects.

We examined the smoothed posterior estimates of individual FEV₁ obtained from the longitudinal submodel. The estimated curves of FEV₁ for adolescents and young adults who were observed with pulmonary exacerbation showed more rapid decline than those who had not experienced a pulmonary exacerbation (Fig. 2).

Time-to-event submodel for pulmonary exacerbation

Both Weibull event models indicated that the hazard of pulmonary exacerbation significantly increased over age. Neither model suggested that age at entry was a significant factor (zero was an element of each 95% CI). The separate model indicates a negative association between age and risk of pulmonary exacerbation and the joint model suggests that the association was positive; however, these associations were not statistically significant. Both models imply that being male, having higher FEV₁ at entry, and belonging to one of the earlier birth cohorts were associated with decreased risk of pulmonary exacerbation; however, being born between 1989 and 1994 corresponded to a decreased risk of pulmonary exacerbation, compared to the youngest birth cohort (those born after 1994). Lower SES was associated with an increase in the hazard of pulmonary exacerbation.

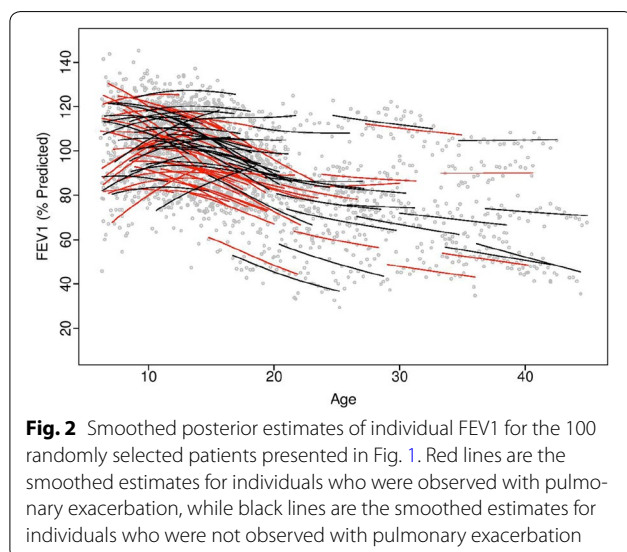


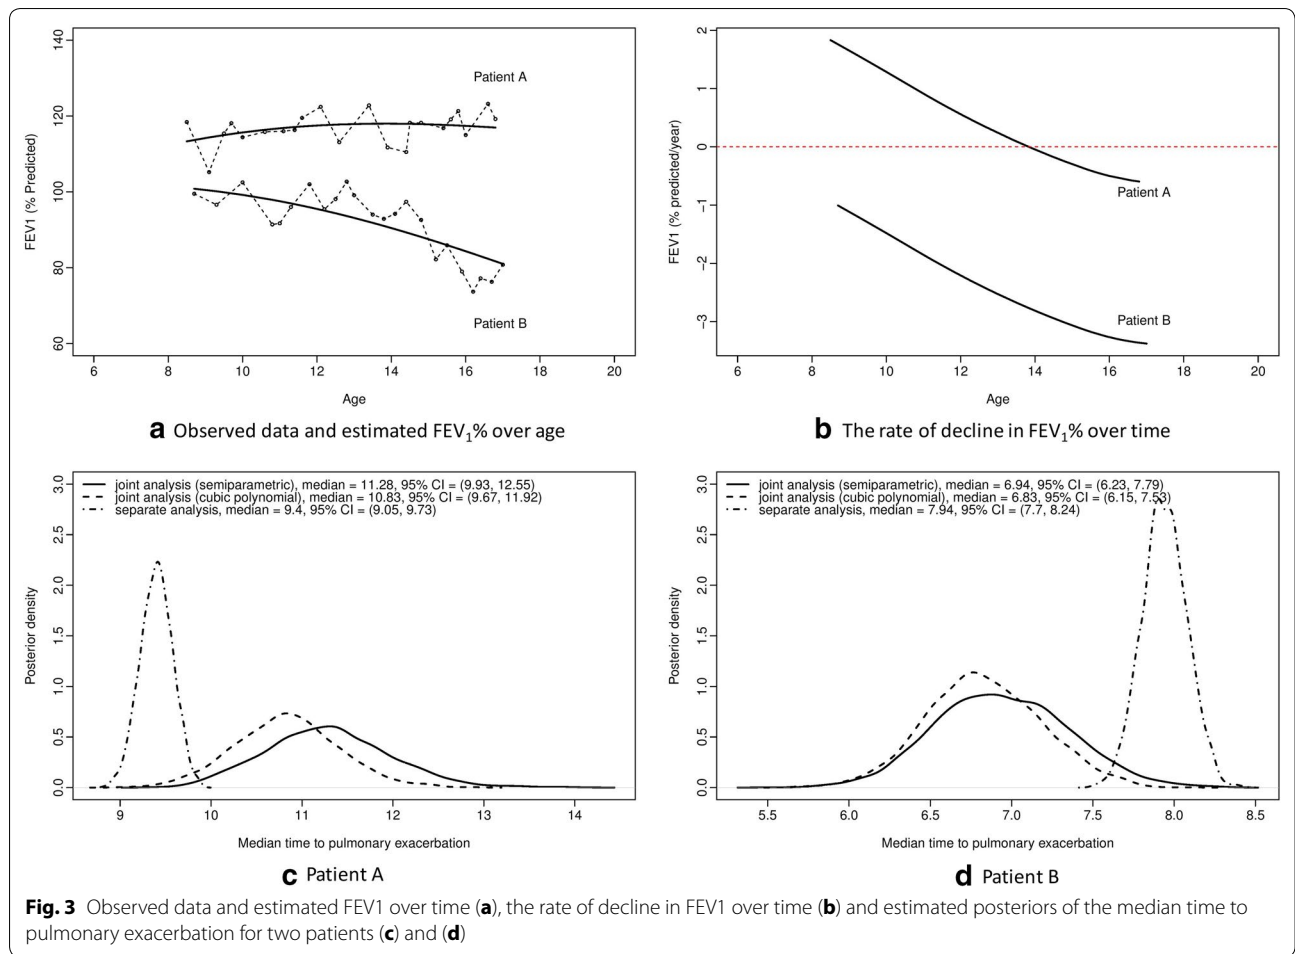
Fig. 2 Smoothed posterior estimates of individual FEV₁ for the 100 randomly selected patients presented in Fig. 1. Red lines are the smoothed estimates for individuals who were observed with pulmonary exacerbation, while black lines are the smoothed estimates for individuals who were not observed with pulmonary exacerbation

Submodel links

Parameter estimates from the joint model indicate that lung function trajectory is associated with risk of having a pulmonary exacerbation. Formally, the intercept and slope parameters corresponding to the shared random effects in the joint model (Table 1, θ_0 and θ_1 estimates, respectively) imply that higher values along the FEV₁ trajectory correspond to a decreased hazard of having a pulmonary exacerbation ($HR = \exp(\hat{\theta}_0) = 0.991$ for every 1% predicted increase in FEV₁, 95% CI 0.986–0.996; HR per one SD increase in FEV₁ = 0.856, 95% CI 0.781–0.937) and positive rates of change or improvements in the FEV₁ trajectory also correspond to decreased hazard ($HR = \exp(\hat{\theta}_1) = 0.702$ for every increase of one percentage point (1% predicted) in the rate of change per year in FEV₁, 95% CI 0.663–0.742; HR per one SD increase in FEV₁ rate of change = 0.566, 95% CI 0.516–0.619). This is clinically reasonable, since a higher level of FEV₁ represents better lung function; patients with FEV₁ measurements that are low or with more rapid decline would be expected to have a higher hazard of pulmonary exacerbation.

Model comparisons

Obvious differences between separate and joint modeling can be graphically observed. We investigated patients with observed FEV₁ trajectories but unknown pulmonary exacerbation event times. We selected two such patients, Patient A and Patient B and compared their estimated posterior median pulmonary exacerbation event time distributions using an established approach [20]. Although neither patient had an observed pulmonary exacerbation, they were censored at ages 16.8 and 17, respectively. Patient A's FEV₁ trajectory began relatively high and remained steady, but Patient B had a trajectory that started relatively low and declined over time (Fig. 3a, b). We compared the median time to pulmonary exacerbation for the two selected patients using posterior distributions from the separate pulmonary exacerbation event model and the two joint models. The predicted median time to pulmonary exacerbation for Patient A was younger under the separate model, compared to each joint model (Fig. 3c). The semiparametric joint model with cubic splines indicated that pulmonary exacerbation onset would occur at an older age than estimated for the joint model with only cubic polynomials and no splines; however, the distributions for the joint model parameters had substantial overlap. The two joint models were in closer agreement regarding time to pulmonary exacerbation for Patient B, but the estimate from the separate pulmonary exacerbation model projected a later event time (Fig. 3d). This is due to joint model's accounting of the correlation between the longitudinal and survival data.



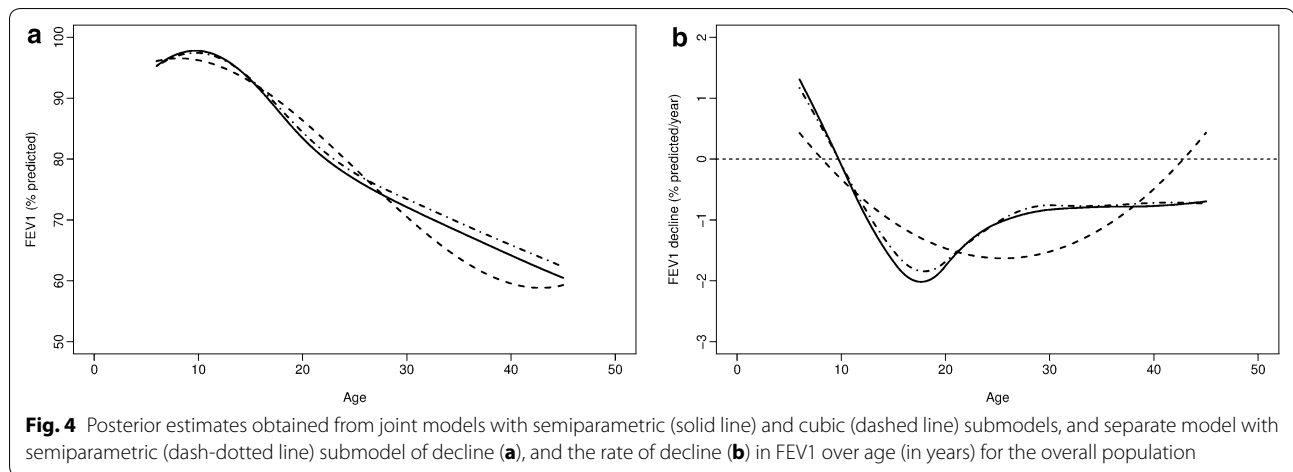
The semiparametric mixed effects joint model had the best fit of all models considered based on total DIC (Table 2), followed by the cubic polynomial model without splines, then the separate modeling of the longitudinal and event processes. The association between submodels through the common random intercepts and slopes in each joint model yielded a slight decrease in DIC₁ for the longitudinal submodel and a substantial decrease in DIC₂ for the survival submodel, compared to separate modeling.

In comparing the two joint models, which only differed based on inclusion of splines in the longitudinal submodel, posterior estimates of age-related FEV₁ progression are quite similar until adolescence (Fig. 4a). Based on rates of change, the semiparametric model estimates more rapid decline into early adulthood (Fig. 4b). Although the cubic regression closely resembles the spline-based estimates, it shows an upward trend in FEV₁ at the later range of age. Patients attained their most rapid decline at 17.6 (6.1) years; median (IQR) estimates

Table 2 Model comparison results

Model	Submodel 1	Submodel 2	DIC ₁	DIC ₂	\bar{D}	p_D	DIC _{total}
<i>Separate model</i>							
I	Semiparametric mixed effects	Weibull	952,566	21,110.9	962,619	11,057.7	973,676
<i>Joint model</i>							
II	Cubic mixed effects	Weibull	952,915	20,712.5	962,485	11,142.6	973,628
III	Semiparametric mixed effects	Weibull	952,409	20,605.4	961,923	11,091.4	973,015

Each of DIC₁ and DIC₂ is the sum of the posterior mean deviance (\bar{D}) and the effective number of parameters (p_D); DIC_{total} obtained by summing DIC₁ and DIC₂ from separate or submodels (see "Appendix 4" for details). Lower values indicate better model fit



are 17.5 (IQR 13.7–18) years. The rate of decline continues at a slower pace until patients approximately reach 30 years of age. Figure 4b indicates the cubic regression provided poorer fit and unrealistic estimates of the trend in FEV₁ decline at the later range of age. The improvement afforded by the semiparametric model was likely due to the localized splines, which captured more variability in FEV₁ than cubic polynomial terms alone, and in particular allowed for a more reasonable fit at later ages.

As a sensitivity analysis, we compared the joint model to a “two-stage” version of the joint model in “Appendix 5”. The two-stage model yielded similar estimates, and detected the same significant effects in the random intercept and slope as the joint model. Based on DIC, the joint model had a better fit to the data than the two-stage model, while the two-stage model was superior to the separate modeling. It demonstrates that the joint model not only allows uncertainty in the random effects to carry through to the event model, but also allows the estimation of random effects to depend on both FEV₁ decline and the occurrence of pulmonary exacerbation.

Discussion

In this paper, we described a flexible joint modeling technique aimed at analyzing long sequences of longitudinal and time-to-event data, and used it to simultaneously characterize the nonlinear progression of FEV₁ and assess risk of pulmonary exacerbation events for individual CF patients. We demonstrated how this approach could be used to inform patient management regarding rapid decline in lung function and assessment of pulmonary exacerbation risk over time. On the population level, we identified clinical and demographic risk factors associated with more rapid decline in FEV₁ and earlier onset of pulmonary exacerbation, which could be used to target subpopulations at increased risk of rapid decline or

pulmonary exacerbation. Translating these individualized results into clinical care is important because, once a pulmonary exacerbation occurs, it is possible that patients will not recover to their previous FEV₁ levels [5].

In addition to individualized estimates of pulmonary exacerbation risk that account for patient-specific longitudinal trajectories and sources of variation in FEV₁, this study examined the effect of established risk factors for lung function decline and onset of pulmonary exacerbation. We found significant effects of birth cohort, socioeconomic status and infections with *MRSA*, *Pa* and *B. cepacia* on FEV₁ decline, as well as significant effects of gender, socioeconomic status and birth cohort on pulmonary exacerbation risk. These findings corroborate what has been shown in separate modeling of FEV₁ and pulmonary exacerbation onset [2, 5, 6]; thus, findings from joint and separate modeling of pulmonary exacerbation appear consistent on the levels of population and subpopulations stratified by risk factors. In addition, this study is the first to our knowledge that has implemented a joint modeling approach to examine the relationship between FEV₁ decline and pulmonary exacerbation events. Through this approach, further insights are gained on an individual patient, the level at which joint and separate modeling results appear to diverge.

Our findings have a number of implications for the epidemiologic study of CF lung disease. Longitudinal modeling of FEV₁, joint or separate, shows that there is an inverse relationship between starting FEV₁ and rate of decline (Table 1), which corresponds to a correlation of -0.60 between initial FEV₁ at age 6 years (the age at which patients typically begin performing reliable pulmonary function testing) and slope estimates in our selected model. This association echoes the “ceiling effect” discovered in a previous study of pediatric CF registry data [6]. Our model setup allows us to estimate the population- and

subject-specific progression of lung function decline, with the latter being important for personalized medicine. We demonstrated through a specific example that a patient with a less severe FEV₁ trajectory is predicted to be free from a pulmonary exacerbation much longer than a patient with a more severe trajectory. It is likely that these differences among the separate and joint models regarding median pulmonary exacerbation time are attributable to the joint model's correctly accounting for the correlation between the longitudinal and survival data. In terms of model fit, including the time-to-event submodel for pulmonary exacerbation appeared to have a greater impact than including the splines to accommodate nonlinear FEV₁ progression, although both features improved fit (Table 2). Furthermore, having the localized differences fitted from the FEV₁ trajectory has been shown to offer more insights into rate of decline [2, 3], particularly in examining the first derivative estimates from our study (Fig. 4). Higher-order random effects could be included to estimate more complex changes over time. Adding a frailty term to Eq. (2) did not improve the total DIC (results not shown). Similar findings have been reported regarding the fits of models with and without the frailty term, all other portions of the model being the same [17].

This study has some limitations that should be considered, which are related to survivor and delayed entry biases, reference equations for lung function and lack of a clear definition of pulmonary exacerbation. Our joint modeling approach corresponds to a missing not at random mechanism as described on pg. 90 by Rizopoulos [21]; the pulmonary exacerbation event process, which is modeled through the random effects, corresponds to the attrition mechanism in a drop-out model. There is a clear birth cohort effect on results from both separate and joint models, which has been noted in other CF registry studies [2, 14, 22] and could be reflective of advancements in care that were largely unavailable to older patients or survivor bias. Our inclusion of birth cohort in the modeling is only a partial adjustment for left truncation bias, as the correlation between age at registry entry and birth cohort will dictate the extent to which this approach combats left truncation bias. Current WinBUGS implementation is not flexible enough to specify the likelihood for each patient as being conditional on his or her entry time into the registry, which would be necessary to model the delayed entry; we refer to Crowther and colleagues' Eq. (5) [23]. Although there are alternative R packages, such as 'rstan' [24], for flexible likelihood formulations to include delayed entry, estimation can be slow with a large number of patients. Faster computing algorithms would be needed to practically implement this model in large patient registry studies. FEV₁ trajectories, which are modeled based on % predicted values over age, may differ

according to the type of spirometry reference equation applied to the raw FEV₁ data, which is expressed in liters. In terms of fitting the FEV₁ trajectory, CFFPR data for patients aged 8–17 years taken from 2013 was analyzed in another study; Wang and Hankinson equations yielded higher median FEV₁ values, compared to values obtained using the Global Lung Initiative equations [25]. Monitoring changes in FEV₁, by contrast, appear to be less susceptible to this effect. There is no standard definition for what constitutes a pulmonary exacerbation; however, CFFPR studies typically infer its occurrence based on intravenous treatment with antibiotics documented in the registry [10]. This definition overlooks less severe exacerbations that impact lung function but do not warrant intravenous antibiotics. Sensitivity to this definition could be assessed with additional data from an individual CF center, which would likely have documentation on mild pulmonary exacerbation events. These events periodically occur throughout the course of CF; thus, extensions to the joint modeling approach presented here could be used to examine risk of recurrence (Additional files 1, 2).

In conclusion, we have utilized novel statistical modeling of data from a national patient registry to provide more realistic estimates of the FEV₁ trajectory and individualized assessment of pulmonary exacerbation risk in patients with CF. Through the Bayesian approach implemented here using existing software, posterior predictive distributions of this model could be used to aid clinicians in estimating risk of pulmonary exacerbation and rate of lung function decline for individual patients. This approach could be extended to a multivariate joint model [26], in which temporal associations of the evolutions of infections and other characteristics are assessed in conjunction with lung function and exacerbations in CF patients. Furthermore, this joint modeling approach could be used to characterize lung function decline in other diseases and disorders, and to identify subgroups of individuals who may benefit from novel therapeutics being tested in clinical trials.

Additional files

Additional file 1. R file for running WinBUGS from R.

Additional file 2. WinBUGS file for analysis implementation.

Abbreviations

B. cepacia: *Burkholderia cepacia*; CF: Cystic fibrosis; CFFPR: Cystic fibrosis foundation patient registry; CFRD: Cystic fibrosis related diabetes; MRSA: Methicillin-resistant staphylococcus aureus; Pa: *Pseudomonas aeruginosa*; SES: Socioeconomic status.

Author details

¹ Alzheimer's Therapeutic Research Institute, Keck School of Medicine, University of Southern California, 9860 Mesa Rim Rd, San Diego, CA 92121, USA.

² Department of Medical Statistics, London School of Hygiene and Tropical Medicine, Keppel Street, London WC1E 7HT, UK. ³ Division of Pulmonary Medicine, Cincinnati Children's Hospital Medical Center, MLC 2021, 3333 Burnet Ave, Cincinnati, OH 45229, USA. ⁴ Division of Biostatistics and Epidemiology, Cincinnati Children's Hospital Medical Center, MLC 5041, 3333 Burnet Ave, Cincinnati, OH 45229, USA.

Authors' contributions

DL, RDS: Developed the study concept, analysis plan, results summaries and drafted the initial manuscript. DL: Implemented all model analyses. RK: Critically reviewed the manuscript, provided statistical input on the evaluation of the models and interpretation of findings, and contributed to the writing in all sections. JPC: Critically reviewed the manuscript, provided interpretation of the statistical model results into the clinical context, and contributed to the writing in all sections. All authors read and approved the final manuscript.

Acknowledgements

The authors would like to thank the Cystic Fibrosis Foundation for the use of CF Foundation Patient Registry data to conduct this study. Additionally, we would like to thank the patients, care providers, and clinic coordinators at CF centers throughout the United States for their contributions to the CF Foundation Patient Registry.

Competing interests

The authors declare that they have no competing interests.

Availability of data and materials

The dataset analyzed for this study is not publicly available. Data is available upon request through the Cystic Fibrosis Foundation Patient Registry Comparative Effectiveness Research Committee. Additional information may be available by contacting datarequests@cff.org. Restrictions on access to data are to ensure patient privacy for all persons in the Cystic Fibrosis Foundation Patient Registry. Implementation materials in the form of program files are available as additional files accompanying this work.

Consent for publication

Not applicable.

Ethics approval and consent to participate

This retrospective study of the United States Cystic Fibrosis Foundation CF Patient Registry was approved by the Institutional Review Board of the Cincinnati Children's Hospital Medical Center (IRB # 2015-4515). The Cystic Fibrosis Foundation Patient Registry Committee approved the request for data access and dispensed the data (Request # PRR088).

Financial disclosure statement

The authors have no financial relationships relevant to this article to disclose.

Funding

This study was funded by grants K25 HL125954 from the National Heart, Lung and Blood Institute of the National Institutes of Health and R457-CR11 from the Cystic Fibrosis Foundation Research and Development Program.

Appendix 1: Regression splines used in the longitudinal submodel

The function $f(t)$ is defined as a cubic truncated power basis function for the semiparametric representation:

$$f(t) = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3 + \sum_{l=1}^L b_l (t - \kappa_l)_+^3,$$

where $\beta_0, \beta_1, \beta_2$ and β_3 are respectively the coefficients for intercept, linear, quadratic and cubic terms; the term $\sum_{l=1}^L b_l (t - \kappa_l)_+^3$ represents basis functions for the cubic splines evaluated at time s , and $(t - \kappa_l)_+ = \max(t - \kappa_l, 0)$. We selected the number and

location of knots ($\kappa_l, l = 1, \dots, L$) for the corresponding spline basis functions using the quantile-based approach described in Ngo and Wand [27]. Please see "Appendix 2" for more discussion on knot specification. Six knots used across the [6, 45] age (in years) interval were $\kappa = (10.78, 15.31, 19.89, 24.96, 30.49, 37.34)$.

The second formulation of the longitudinal submodel defines $f(t)$ without the basis functions, which we refer to as the cubic representation:

$$f(t) = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3$$

Appendix 2: Knot specification

The main idea of choosing knots is to choose enough knots to resolve the essential structure in the underlying regression function, but to keep the number of knots relatively low considering the computational advantages. A reasonable default is to choose the knots to ensure that there are a fixed number of unique observations (e.g., 4–5) between each knot. However, for large data sets, this may lead to an excessive number of knots, so a maximum number of knots (e.g., 20–40) is recommended. A simple default choice of L that always works well is:

$$L = \max \left\{ 5, \min \left(\frac{1}{4} \times \text{number of unique } t_i\text{'s}, 35 \right) \right\}.$$

A reasonable default rule for the knot location is:

$$\kappa_l = \{(l + 1)/(L + 2)\} \text{th sample quantile of the unique } t_i\text{'s,}$$

for $l = 1, \dots, L$. The default choice of knots can be generated in R using the function defined as below:

```
default.knots <- function(x, num.knots) {
  if (missing(num.knots))
    num.knots <- max(5, min(floor(length(unique(x))/4), 35))
  return(quantile(unique(x), seq(0, 1, length =
    (num.knots+2))[-c(1, (num.knots+2))]))
}
```

Appendix 3: Prior specification and sampling

We used a fully Bayesian version of the joint modeling approach implemented via Markov chain Monte Carlo (MCMC) method to fit the joint model. Although we focus on a particular form, the Bayesian framework allows for easier comparison among many forms of association between W_1 and W_2 . The data analyses are conducted using MCMC through WinBUGS software [28]. R and WinBUGS codes implementing the proposed method are available in Additional files 1 and 2. We selected vague prior distributions in the Bayesian analysis. For each of the regression coefficients of the cubic penalized spline and the subject-level covariate effects β, \mathbf{b} and α_1 , we used a multivariate normal prior with mean vector 0 and precision matrix $0.01 \cdot \mathbf{I}$, where \mathbf{I} is the identity matrix; and an inverse gamma prior¹⁹ $IG(0.01, 0.01)$ for the measurement error variance σ_ϵ^2 .

Similarly, we selected a vague normal prior in the survival submodel, $\alpha_1 \sim N(0, 0.01 \cdot \mathbf{I})$. For the random effects \mathbf{U}_i common to both submodels, we took a non-informative multivariate inverse-Wishart distribution with degrees of freedom equals to 2 (the number of random effect components). Finally, for the association parameters, we set $\theta_0 \sim N(0, 100)$ and $\theta_1 \sim N(0, 100)$.

Complete conditional distributions were used to approximate the posterior distributions for the parameters of interest. MCMC diagnostics to assess convergence were performed in accordance with standard techniques [29]. Early samples were discarded based on graphical inspection of burn-in across MCMC iterations. The 95% credible intervals (CIs) were obtained by taking the 2.5th and 97.5th percentiles of the posterior samples post burn-in. For each model, we conducted the MCMC simulation for 40,000 iterations. The first 20,000 draws were discarded as a burn-in phase. The effective iterations were thinned by storing every 10th iteration in order to reduce level of correlation between successive values of the sample. The posterior mean and 95% confidence interval (CI)

were calculated with the resulting 2000 values. The computing time for each model is shown in Table 3.

Appendix 4: DIC formulation

Each deviance information criterion (DIC) for assessing model fit is calculated as $DIC = D + p_D$, with the posterior mean of deviance D penalized by the effective number of parameters p_D under the Bayesian framework. Since small D indicates good fit while small p_D indicate a parsimonious model, smaller values indicate better model fit, compared to larger values. We respectively write the components of DIC for the longitudinal and survival submodels as DIC_1 and DIC_2 , which evaluate the contributions to the overall DIC, denoted as DIC_{total} in Table 2. All DIC values are output directly from WinBUGS.

Appendix 5: Comparison between two-stage model and joint model

In the joint modeling literature, various two-stage methods have been proposed. Here we compare the predictive power of a joint model with that of a simple (naïve) two-stage model. The naïve two-stage model consists of: (Stage I) longitudinal model for FEV_1 ; and (Stage II) Weibull model for pulmonary exacerbation with random effects from Stage I as covariates. To be consistent with the joint model (Model III), we use the same semiparametric mixed effects modeling for FEV_1 . The posterior means of subject-specific random effects (i.e., random intercept \tilde{U}_{0i} and slope \tilde{U}_{1i}) from Stage I are included as covariates in the Weibull model in Stage II. The parameters of the two-stage model are given the same priors as the joint model, and are estimated using the fully Bayesian-MCMC (WinBUGS) approach. The results of

Table 3 Computing time for each model in WinBUGS

Model	Submodel 1	Submodel 2	Computing time
<i>Separate model</i>			
I	Semiparametric mixed effects	Weibull	9 h 57 min
<i>Joint model</i>			
II	Cubic mixed effects	Weibull	8 h 45 min
III	Semiparametric mixed effects	Weibull	11 h 14 min

* The data set contains 7672 individuals, and the total number of longitudinal observations (FEV_1 measurements) is 36,051. For each model, we run 40,000 iterations, discard the first 20,000 as a warm-up phase and thin by 10, yielding a total of 2000 posterior samples

Table 4 Posterior estimates under two-stage model and joint model for lung function decline and pulmonary exacerbation onset

Parameter	Two-stage model		Joint model (Model III)	
	Posterior mean	95% HPD CI	Posterior mean	95% HPD CI
<i>Semiparametric longitudinal submodel for FEV₁</i>				
Curve				
β_0 , intercept	18.15	[15.54, 20.95]	16.89	[14.63, 20.39]
β_1 , age	2.86	[2.477, 3.189]	3.121	[2.74, 3.382]
β_2 , age ²	-0.1306	[-0.1418, -0.117]	-0.1367	[-0.1481, -0.124]
β_3 , age ³	-0.0011	[-0.0013, -0.0009]	-0.0016	[-0.0017, -0.0014]
b_1	0.0051	[0.0039, 0.0061]	0.0073	[0.0058, 0.0084]
b_2	0.0108	[0.0097, 0.0119]	0.0103	[0.0095, 0.0111]
b_3	-0.0177	[-0.0226, -0.0133]	-0.0206	[-0.0248, -0.0167]
b_4	-0.0001	[-0.0044, 0.0065]	0.0029	[-0.001, 0.0085]
b_5	0.0033	[-0.0024, 0.0083]	0.0015	[-0.0033, 0.007]
b_6	-0.0009	[-0.0114, 0.0095]	0.0008	[-0.0095, 0.0101]
Baseline FEV1	0.686	[0.6728, 0.6978]	0.6834	[0.6676, 0.6945]
Male	-0.15	[-0.582, 0.2736]	-0.1650	[-0.6054, 0.2971]
Birth cohort				
< 1981	11.43	[9.837, 13.29]	12.39	[10.89, 13.95]
1981–1988	4.0	[2.992, 4.978]	4.936	[3.995, 5.916]
1989–1994	-0.1924	[-0.7038, 0.3305]	0.0608	[-0.4059, 0.5908]
> 1994 (reference)				
Lower SES	-0.623	[-1.091, -0.175]	-0.5984	[-1.062, -0.1119]
CFRD	-0.4613	[-0.7895, -0.1506]	-0.3695	[-0.6731, -0.0633]
MRSA	-0.816	[-1.026, -0.6022]	-0.8114	[-1.016, -0.6134]
<i>B. cepacia</i>	-0.7964	[-1.465, -0.1291]	-0.7705	[-1.441, -0.1022]
<i>Pa</i>	-0.5643	[-0.6953, -0.4345]	-0.5673	[-0.6971, -0.4373]
Sources of variation				
σ^2 (measurement error)	59.29	[58.83, 59.78]	59.24	[58.77, 59.73]
σ_0^2 (between patients, intercept)	324.1	[307.6, 341.3]	326.8	[311, 343.8]
σ_1^2 (between patients, slope)	2.564	[2.443, 2.69]	2.602	[2.484, 2.727]
σ_{01} (covariance, intercept and slope)	-25.46	[-26.81, -24.18]	-25.76	[-27.11, -24.49]
<i>Weibull event submodel for pulmonary exacerbation</i>				
k	2.524	[2.456, 2.601]	2.564	[2.487, 2.634]
intercept	-3.509	[-3.802, -3.26]	-3.568	[-3.804, -3.296]
θ_0 (random intercept)	-0.0098	[-0.0149, -0.0048]	-0.0086	[-0.0137, -0.0036]
θ_1 (random slope)	-0.3273	[-0.3825, -0.2728]	-0.3532	[-0.4104, -0.2979]
Baseline age	0.0002	[-0.0136, 0.0161]	0.0023	[-0.0131, 0.0176]
Baseline FEV1	-0.0193	[-0.0211, -0.0175]	-0.0198	[-0.0218, -0.0179]
Male	-0.2919	[-0.36, -0.2194]	-0.2940	[-0.3606, -0.2267]
Birth cohort				
< 1981	-0.5046	[-0.8512, -0.1807]	-0.5348	[-0.9056, -0.1726]
1981–1988	-0.2297	[-0.4382, -0.0415]	-0.2348	[-0.4405, -0.0339]
1989–1994	0.1171	[0.0086, 0.2197]	0.1261	[0.0178, 0.2335]
>1994 (reference)				
Lower SES	0.1091	[0.0369, 0.1799]	0.1059	[0.0382, 0.1804]

Table 5 Model comparison results

Model	Submodel 1	Submodel 2	DIC ₁	DIC ₂	\bar{D}	p_D	DIC _{total}
Two-stage model	Semiparametric mixed effects	Weibull	952,553	20,740.6	951,177	11,063.7	973,294
Joint model (Model III)	Semiparametric mixed effects	Weibull	952,409	20,605.4	961,923	11,091.4	973,015

posterior estimates and model comparison for two-stage model and joint model (Model III) are given in Tables 4 and 5.

The two-stage model yields similar estimates, and detects the same significant effects in the random intercept and slope as the joint model. However, DIC in Table A3 suggests that the joint model (DIC_{total} = 973,015) gives a much better fit of the data than the two-stage model (DIC_{total} = 973,294).

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations. Received: 11 April 2017 Accepted: 3 November 2017

Published online: 14 November 2017

References

- Farrell PM. The prevalence of cystic fibrosis in the European Union. *J Cyst Fibros*. 2008;7(5):450–3.
- Szczesniak RD, et al. A semiparametric approach to estimate rapid lung function decline in cystic fibrosis. *Ann Epidemiol*. 2013;23(12):771–7.
- Moss A, et al. A comparison of change point models with application to longitudinal lung function measurements in children with cystic fibrosis. *Stat Med*. 2016;35(12):2058–73.
- Ferkol T, Rosenfeld M, Milla CE. Cystic fibrosis pulmonary exacerbations. *J Pediatr*. 2006;148(2):259–64.
- Sanders DB, et al. Failure to recover to baseline pulmonary function after cystic fibrosis pulmonary exacerbation. *Am J Respir Crit Care Med*. 2010;182(5):627–32.
- Konstan MW, et al. Risk factors for rate of decline in forced expiratory volume in one second in children and adolescents with cystic fibrosis. *J Pediatr*. 2007;151(2):134–9.
- Asar O, et al. Joint modelling of repeated measurement and time-to-event data: an introductory tutorial. *Int J Epidemiol*. 2015;44(1):334–44.
- Schluchter MD, Konstan MW, Davis PB. Jointly modelling the relationship between survival and pulmonary function in cystic fibrosis patients. *Stat Med*. 2002;21(9):1271–87.
- Piccorelli AV, Schluchter MD. Jointly modeling the relationship between longitudinal and survival data subject to left truncation with applications to cystic fibrosis. *Stat Med*. 2012;31(29):3931–45.
- Knapp, E.A., et al., *The Cystic Fibrosis Foundation Patient Registry. Design and Methods of a National Observational Disease Registry*. *Ann Am Thorac Soc*, 2016;13(7):1173–79.
- Wang X, et al. Pulmonary function between 6 and 18 years of age. *Pediatr Pulmonol*. 1993;15(2):75–88.
- Hankinson JL, Odencrantz JR, Fedan KB. Spirometric reference values from a sample of the general US population. *Am J Respir Crit Care Med*. 1999;159(1):179–87.
- Corey M, et al. Longitudinal analysis of pulmonary function decline in patients with cystic fibrosis. *J Pediatr*. 1997;131(6):809–14.
- Taylor-Robinson D, et al. Understanding the natural progression in %FEV1 decline in patients with cystic fibrosis: a longitudinal study. *Thorax*. 2012;67(10):860–6.
- Henderson R, Diggle P, Dobson A. Joint modelling of longitudinal measurements and event time data. *Biostatistics*. 2000;1(4):465–80.
- Tsatis AA, Davidian M. Joint modeling of longitudinal and time-to-event data: an overview. *Stat Sin*. 2004;14:809–34.
- Fitzmaurice G, Davidian M, Verbeke G, Molenberghs G. Longitudinal data analysis. In: Molenberghs G, Fitzmaurice G, editors. *Incomplete data: introduction and overview*. Boca Raton: Taylor & Francis Group; 2009. p. 398–400.
- Lunn DJT, Best N, Spiegelhalter D. WinBUGS: a Bayesian modelling framework—concepts, structure, and extensibility. *Stat Comput*. 2000;10:325–37.
- Spiegelhalter DJ, Best NGC, Van der Linde A. Bayesian measures of model complexity and fit (with discussion). *J R Stat Soc Ser B*. 2002;64(4):583–616.
- Guo XC, Carlin BP. Separate and joint modeling of longitudinal and event time data using standard computer packages. *Am Stat*. 2004;58(1):16–24.
- Rizopoulos D. Joint models for longitudinal and time-to-event data: with applications in R, vol. xiv, Chapman & Hall/CRC biostatistics series Boca Raton: CRC Press; 2012. p. 261.
- VanDevanter DR, Pasta DJ, Konstan MW. Improvements in lung function and height among cohorts of 6-year-olds with cystic fibrosis from 1994 to 2012. *J Pediatr*. 2014;165(6):1091e2–1097e2.
- Crowther MJ, et al. Joint modelling of longitudinal and survival data: incorporating delayed entry and an assessment of model misspecification. *Stat Med*. 2016;35(7):1193–209.
- Stan Development Team. Rstan: the R interface to Stan, version 2.16.2; 2017. <http://mc-stan.org>. Accessed 9 Nov 2017.
- Foundation CF. Cystic fibrosis foundation patient registry. In: 2013 annual report to the center directors. Cystic Fibrosis Foundation: Bethesda; 2014.
- Albert PS, Shih JH. An approach for jointly modeling multivariate longitudinal measurements and discrete time-to-event data. *Ann Appl Stat*. 2010;4(3):1517–32.
- Ngo L, Wand M. Smoothing with mixed model software. *J Stat Softw*. 2004;9:1–54.
- Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS—a Bayesian modelling framework: concepts, structure, and extensibility. *Stat Comput*. 2000;10:325–37.
- Spiegelhalter D, Thomas A, Best N, Gilks W. WinBUGS user manual. <http://www.mrc-bsu.cam.ac.uk/wp-content/uploads/manual14.pdf>.

Submit your next manuscript to BioMed Central and we will help you at every step:

- We accept pre-submission inquiries
- Our selector tool helps you to find the most relevant journal
- We provide round the clock customer support
- Convenient online submission
- Thorough peer review
- Inclusion in PubMed and all major indexing services
- Maximum visibility for your research

Submit your manuscript at
www.biomedcentral.com/submit

