

Modeling recovery curves with application to prostatectomy

FULTON WANG*

*Department of EECS, Massachusetts Institute of Technology,
77 Massachusetts Ave, Cambridge, MA 02139, USA
fultonw@mit.edu*

CYNTHIA RUDIN†

*Department of Computer Science, Duke University,
LSRC D342, Research Drive, Durham, NC 27708, USA*

TYLER H. MCCORMICK†

*Department of Statistics and Department of Sociology, University of Washington
Box 354322, Seattle, WA 98195-4322, USA*

JOHN L. GORE

Department of Urology, University of Washington, 959 NE Pacific St, Seattle, WA 98195, USA

SUMMARY

In many clinical settings, a patient outcome takes the form of a scalar time series with a recovery curve shape, which is characterized by a sharp drop due to a disruptive event (e.g., surgery) and subsequent monotonic smooth rise towards an asymptotic level not exceeding the pre-event value. We propose a Bayesian model that predicts recovery curves based on information available before the disruptive event. A recovery curve of interest is the quantified sexual function of prostate cancer patients after prostatectomy surgery. We illustrate the utility of our model as a pre-treatment medical decision aid, producing personalized predictions that are both interpretable and accurate. We uncover covariate relationships that agree with and supplement that in existing medical literature.

Keywords: Bayesian methods; Interpretable modeling; Prostate cancer; Recovery curves.

1. INTRODUCTION

In the medical community, there is a pressing need for personalized predictions of how a disruptive event, such as a treatment or disease, will impact particular bodily function levels. Of particular interest is the extent to which the function is initially perturbed by the event and the ensuing pattern of recovery. In many contexts, such as mental acuity following a stroke or sexual function following prostatectomy, the

*To whom correspondence should be addressed.

†These two authors contributed equally.

post-event trajectory generally exhibits what we call a recovery curve shape, characterized by an initial instantaneous drop followed by a monotonic rise towards an asymptotic level not exceeding the original function level. Here, we propose a Bayesian model that can be used to predict a patient's expected recovery curve, given information about the patient that is available before the event.

This article presents a decision aid for patients considering a medical treatment who want to know what adverse side effect the treatment would have on a particular bodily function. In particular, our model will be used to display to the patient a distribution over post-treatment function trajectories, conveying the uncertainty in predictions that should be considered in decision-making. We assume that the function level lies in some closed interval, the pre-treatment function level is known, and the adverse effect of the event on the function is *a priori* known by the medical community to be immediate but wearing off over time.

If a model is to be widely adopted as a medical decision aid, it is not enough for it to merely produce predictions that are accurate; it must also be *interpretable*: it must give predictions that a health care provider or patient can readily understand. This is crucial in a clinical setting not only because of time constraints, but also because each additional point of confusion regarding the predictions decreases the flow of information to the patient and thus their trust in it. Any model, including ours, should be used only if it fits the data, and standard model checking should be performed. However, aggregate model fit and predictive accuracy are not sufficient; when a tool causes more questions than it answers, it will be less likely to be used.

In the applications we consider, we will predict time series that are expected to be recovery curves. These recovery curves occur in practice in situations where a medical procedure may cause a temporary inhibition or detriment in a patient, but will not improve the patient's condition over a baseline value. In our case, we examine sexual function in patients with prostate cancer who are considering a prostatectomy. Though a prostatectomy will likely decrease sexual function in some patients, it will not *improve* a patient's sexual function above the baseline.

We restrict the space of possible outputs to our model, therefore, to *only* include predictions that are recovery curves. In other words, our model outputs a distribution of post-event trajectories each of which is guaranteed to be a recovery curve, so that the function level drops instantaneously downwards at event time, and rises smoothly to approach an asymptotic level lying between the pre-treatment value and value immediately following treatment (we will use the words "level" and "value" interchangeably). Furthermore, our model encourages the posterior predictive distribution over trajectories to have a well-defined mode, so that the distribution of curves, when plotted, can be visually easily interpreted as a single maximum *a posteriori* prediction along with the uncertainty in that single prediction. Thus, not only do the predictions match prior expectations but they are also simple, defined by a small set of salient features (i.e. initial drop, recovery rate, and range of uncertainty), not containing extraneous artifacts that slightly increase accuracy but greatly reduce interpretability.

We use our method to create a decision aid for prostate cancer patients who are considering a prostatectomy, predicting their sexual function trajectory, should they undergo a prostatectomy. We fit our model using data from a study that tracked the quantified sexual function level, expressed as a number between 0 and 1, of 237 patients both before radical prostatectomy surgery and at a common set of time points in the 4 years immediately following surgery. These numerical measures of sexual function were obtained by administering the Prostate Cancer Index, which is a multiple choice questionnaire that first evaluates patient function and bother following prostate cancer treatment, and then converts the answers to a numerical score. Note that while our target population is patients merely considering a prostatectomy (and satisfies two properties listed shortly), our data set only contains data from patients who actually did undergo a prostatectomy.

Prostate cancer will affect 1 of 6 men, and low-stage patients usually have several viable treatment options, each with different side effects. Radical prostatectomy is known to adversely affect sexual




Fig. 1. Sexual function trajectory (fxn value) following prostatectomy. (a) The average function value and scaled function value over the prostatectomy data set exhibit a recovery curve shape. (b) Raw data of 12 randomly chosen patients who passed the filters as described in Section 6.1.

function. Thus, for patients considering prostatectomy, it is important to forecast the pattern of sexual function level should they undergo one. Past studies (Potosky and others, 2004) and our data set indicate that sexual function trajectories after prostatectomy follow a recovery curve (at least up to 5 years post-treatment), suggesting our model may fit such trajectories well.

To illustrate, in Figure 1a we show the data set-wide averages, for each time point, of sexual function level (fxn level), as well as data set-wide averages, for each time point, of the patients' sexual function values scaled by their respective value immediately before prostatectomy. We also show in Figure 1b the unscaled sexual function level time series of 12 randomly selected patients, which includes their unscaled sexual function level both post-treatment and right before treatment. We hypothesize that the function levels reported by individual patients are noisy versions of a latent smooth “true” function level. In this context, we use our method to study whether there are patient covariates that correlate with post-surgery sexual function trajectory.

The remainder of the article is organized as follows. In Section 2, we contextualize our approach by describing related work. Then, we formally define the recovery curve shape in Section 3 and present the specifics of our model in Section 4. We demonstrate that the model performs well in simulation studies in Section 5 and then using the aforementioned prostatectomy data in Section 6.

2. RELATED WORK

Past work on personalized prediction of sexual function following prostate cancer treatment has attempted to predict a post-operative binary outcome, typically whether one is able to achieve an erection sufficient for sex at some single time point following treatment (Descaseaux and others, 2006; Ayyathurai and others, 2008; Eastham and others, 2008; Regan and others, 2011), or the change in Sanda and others (2008) or absolute level (Talcott and others, 2003) of some continuous measure of sexual function such as the IIEF-5 score. Such models incorporated patient covariates in linear regression models for continuous outcomes and logistic regression models for binary outcomes even though sexual function is not a binary outcome (Briganti and others, 2011). Another deficiency of logistic and linear regression is that they are not suitable for modeling longitudinal outcomes, whereas a patient would want to know their entire post-treatment function trajectory. The only longitudinal model in the literature uses

linear regression to relate the change in function level between two fixed time points post-treatment (Potosky and others, 2004).

Functional data analysis (Denison and others, 1998; Ramsay and Silverman, 2002) and growth curve modeling (Jung and Wickrama, 2008) are rich areas of past study, but existing models from those fields do not guarantee that the predicted time series possesses a recovery curve shape, that it drops following the event and then monotonically approaches an asymptotic level no higher than the pre-event value. Parametric functions mentioned in Rogosa and Willett (1985) resemble the functional form we assume of recovery curves. However, they are modeling growth, not recovery after some disruptive event, and assume the initial level of the growth curve to be known, whereas we are trying to predict the entire post-treatment trajectory, which includes the initial post-treatment value. Furthermore, they do not place an upper bound on the asymptotic level of the predicted function. Isotonic regression models (Mammen, 1991; Neelon and Dunson, 2004; Cai and Dunson, 2007; Shively and others, 2009), enforce the predicted functions to be monotonic, but do not naturally output recovery curves as predictions.

Other statistical models have also been applied in contexts where the predicted time series is expected to exhibit a recovery curve shape. For example, in a medical context, growth curve techniques have been used to model recovery of a bodily function following a disruptive event. Warschausky and others (2001) model recovery of FIM-measured function following spinal surgery as being the sum of a linear and a plateauing function. Although the FIM-score must lie within a bounded interval, they do not guarantee that the predicted scores lie within that interval. Rolfe and others (2011) model verbal function following chemotherapy using a Bayesian latent basis model, but the model lacks incorporation of patient-correlates, and is instead specialized to infer average recovery at only two fixed time points. Tilling and others (2001) model a measure of quality of life—the Barthel Index—following stroke using a multilevel model where both patient-specific and time-specific contributions are modeled as a linear combination of a fractional polynomial basis. In all these models, the predicted time series is expected to possess a recovery curve shape, but there is no explicit constraint built into the models to ensure that the predicted trajectory actually does possess a recovery curve shape.

Two past works suggest that a model of sexual function following prostatectomy needs a significant amount of interpretability-promoting features if it is to be used in practice. One work involved eliciting and incorporating the preferences of patients, providers, and design experts via a three-step human-centered design process to design such dashboards (Hartzler and others, 2015). The second work measured patients' abilities to understand three different visual displays communicating the same information—bar charts, line graph, and table, and examined how a patient's understanding of those three displays related to their demographics (i.e., education level, race) and graphical and numerical literacy as measured through the REALM and SNS questionnaires, two standard medical instruments (Nayak and others, 2016). One takeaway from these works is just how much detail, care, and user feedback goes into designing visual dashboards and studying the impact of seemingly small changes to them. One example of a small design decision significant enough to warrant study was, in a pictogram used to communicate well-being, whether a sunny weather/cloudy weather icon should be used in place of a smiling/frowning face to represent well-being. These meticulous studies reflect the sensitivity of patient comprehension to dashboard features and suggest each additional feature improving the interpretability of our model can greatly improve patient comprehension and thus clinical applicability. A second takeaway from these works is that some patients prefer extremely simple dashboards. For example, some patients felt that putting confidence intervals on personalized predictions was too confusing, preferring a point prediction instead. This preference bolsters the case for our unimodality requirement—given that some patients do not even want to see uncertainty in predictions, were we to display them, we should do so in the simplest manner possible. In any case,

this second takeaway suggests many interpretability-promoting features may be necessary for any clinical applicability at all.

3. RECOVERY CURVES

3.1. Recovery curve definition

A *recovery curve* is a function $f(t)$ defined on $\mathbb{R}^+$. We will always define $f(t)$ piecewise as

$$f(t) = \begin{cases} S & \text{for } t = 0, \\ f^+(t) & \text{for } t > 0, \end{cases} \quad (3.1)$$

which is interpreted as the disruptive event occurring right after time 0, so that S is the (known) pre-treatment function value, and $f^+(t)$ is the post-treatment trajectory. A recovery curve will satisfy the following:

$$(f^+)' > 0 \text{ for } t > 0, \quad (3.2)$$

$$f^+(t) \leq S \text{ for } t > 0, \quad (3.3)$$

$$f^+(t) \geq 0 \text{ for } t > 0, \quad (3.4)$$

$$S \in [0, 1]. \quad (3.5)$$

3.2. Parameterizing recovery curves

We parameterize $f(t)$ scaled to the pre-treatment function value, instead of $f(t)$ itself, assuming:

$$f^+(t; S, \theta) = Sg(t; \theta). \quad (3.6)$$

Thus, the actual post-event trajectory is the shape of the post-event trajectory, $g(t; \theta)$, scaled to the pre-event function level. We choose this parameterization because in order to satisfy requirements of 3.3 and 3.4, we just need to ensure that $g(t; \theta) \in (0, 1)$ for $t > 0$. In this work, we will refer to the scaled post-event trajectory, denoted $g(t; \theta)$, as a *recovery shape* and use the term *scaled function value* to refer to a patient's function value normalized by their pre-treatment value; a patient's recovery shape is a time series over their scaled function values.

We parameterize recovery shape $g(t; \theta)$ with three parameters: A , the asymptotic drop in scaled function level after surgery, B , the initial drop in scaled function value in excess of the asymptotic drop, and C , the rate of recovery of the scaled function value.

$$f(t; S, A, B, C) = Sg(t; A, B, C), \text{ where} \quad (3.7)$$

$$g(t; A, B, C) = S \left(1 - A - B(1 - A) \exp\left(-\frac{t}{C}\right) \right), \quad (3.8)$$

$$A \in [0, 1], B \in [0, 1], C \geq 0. \quad (3.9)$$

$f(t; S, A, B, C)$ is a recovery curve if the constraints in 3.2 to 3.5 are satisfied, and we respect those constraints in our model.




Fig. 2. Two unrealistic predictive distributions. (a) This predictive distribution is unrealistic because some of the time series are not recovery curves, as their post-treatment function value exceeds that pre-treatment. (b) This predictive distribution is unrealistic because the distribution is not unimodal. One can see this because there are two dark sets of curves, one for each mode of the posterior curve distribution.

4. MODEL

The previous section formalized the definition of a recovery curve. To fit the curves to data and make meaningful predictions for patients, however, we need both the shape of the curve and a framework for inference. In this section, we describe a Bayesian approach to fitting recovery curves. We first describe the structure of the statistical model, then we note the properties of the model that make it well-suited for situations where we expect recovery curve trajectories.

Throughout the section, we refer to the i th patient's covariate vector as $X_i \in \mathbb{R}^K$ where K is the number of features per patient, and their observed function value at time t by $y_i(t)$. Here, $y_i(t)$ is considered a noisy measurement of their “underlying” function value at time t , $f(t; S_i, A_i, B_i, C_i)$, where $f(\cdot; S_i, A_i, B_i, C_i)$ is the parameterization of post-treatment function value described in Section 3.2. The noise in $y_i(t)$ could arise from a number of sources, including short-term fluctuations in patients' experiences or difficulty in recalling function between time periods. The “underlying” function value, $f(t; S_i, A_i, B_i, C_i)$, is a function of the patient's pre-treatment function value S_i and patient-specific random parameters A_i, B_i, C_i . To simplify notation, we abbreviate the latent function value as $f_i(t)$. When making a prediction for a new patient, we assume S_i is known based on the patient's experience before the procedure. In [supplementary material](#) available at *Biostatistics* online, we study the robustness of our model to measurement error in S_i .

4.1. Model components

As previously mentioned, we perform inference on recovery curves using a hierarchical Bayesian model. The Bayesian paradigm facilitates sharing information across similar but not identical patients. This information sharing is critical in our context as data on outcomes after radical prostatectomy are very difficult to collect and rare. The remainder of this subsection provides a detailed description of the model.

Recall that our model is designed to be *interpretable*. In particular, a patient's posterior over recovery curves firstly should only have support over the space of recovery curves, so that a posterior as in Figure 2a is not acceptable, containing support over trajectories whose asymptotic level exceeds that pre-treatment. Secondly, the posterior should be unimodal, so that a posterior as in Figure 2b is not acceptable. We encourage the first desiderata by appropriately constraining the support of relevant conditional distributions of the model, and the second by guaranteeing unimodality of those conditional distributions by using specialized distribution parameterizations.

First, recall the observed data $y_i(t)$ is a patient's *reported* function level at time t . We assume the reported function level comes from a likelihood that is a mixture distribution of the form:

$$y_i(t)|f_i(t), \theta, p, \phi_M \sim \theta \text{Bernoulli}(p) + (1 - \theta) \text{beta}_{m,\phi}(f_i(t), \phi_M), \text{ with } p, \theta, \phi_M \in (0, 1), X_i \in \mathbb{R}^K.$$

We use a mixture distribution because patients can and do report values of 0 and 1. In the data presented in Section 6, approximately 5% of patient responses are on the boundary of the unit interval. The mixture distribution, therefore, places finite mass on the 0 and 1 responses, but also allows responses between 0 and 1 to be modeled using a recovery curve.

For values other than 0 and 1, we propose a beta distribution that depends on the patient's (latent) recovery curve value at time t , $f_i(t)$. Even notwithstanding the potential for values on the boundary, parameterizing the unit interval in a way that is interpretable is challenging. To encourage unimodality of the $y_i(t)$, we would like this beta distribution to always be unimodal. For the typical parameterization of the beta distribution, $\text{beta}_{\alpha,\beta}(\alpha', \beta')$, the mean $(\alpha'/\alpha' + \beta')$ and mode $(\alpha' - 1/\alpha' + \beta' - 1)$ both depend on both parameters. Further, the distribution is only unimodal if α and β are both greater than one. As our goal is to develop a method that is easy to explain to clinicians and patients, we chose to reparameterize the beta distribution in terms of mode m and spread parameter ϕ . This $\text{beta}_{m,\phi}$ parameterization relates to the typical beta as:

$$\text{beta}_{m,\phi}(m', \phi') = \text{beta}_{\alpha,\beta}\left(1 + \left(\frac{1}{\phi'} - 1\right)m', 1 + \left(\frac{1}{\phi'} - 1\right)(1 - m')\right). \quad (4.1)$$

Critically, our $\text{beta}_{m,\phi}(m', \phi')$ distribution has mode m' and for *all* m' , is unimodal if and only if spread parameter $\phi' \in (0, 1)$. Examples of such distributions and a full description of the steps to reparameterize are in Figure S3 and Section S1 of the [supplementary material](#) available at *Biostatistics* online.

For values not on the boundary, each respondents' reported function value comes from a beta distribution centered on their true (latent) function value, $f_i(t)$, and with spread around that function value determined by parameter ϕ_M . The patient's latent function value takes the form of a recovery curve described in Section 3. Following the Bayesian paradigm, we specify prior distributions for the parameters of the recovery curve. We expect that patients that are observably similar will have similar recovery trajectories, so we model the parameters of each patient's recovery curve as a function of observable covariates. Recall that the recovery curve depends on individual specific parameters A_i , B_i , and C_i controlling the asymptotic decrease in function post-treatment, initial drop post-treatment, and rate of recovery, respectively. Since A_i and B_i are scaled to be consistent across patients, they have support on $(0, 1)$. The C_i parameter is a rate and thus has support on $\mathbb{R}^+$. We model each with a generalized linear model of the form:

$$A_i|b_A, \phi_A; z_A, X_i \sim \text{beta}_{m,\phi}(\text{logistic}(z_A + b_A^T X_i), \phi_A) \quad (4.2)$$

$$B_i|b_B, \phi_B; z_B, X_i \sim \text{beta}_{m,\phi}(\text{logistic}(z_B + b_B^T X_i), \phi_B) \quad (4.3)$$

$$C_i|b_C, \phi_C; z_C, X_i \sim \text{gamma}_{m,\phi}(\exp(z_C + b_C^T X_i), \phi_C), \text{ with} \quad (4.4)$$

$$b_A, b_B, b_C \in \mathbb{R}^K, \phi_A, \phi_B, \phi_C \in (0, 1), X_i \in \mathbb{R}^K. \quad (4.5)$$

Note that by assuming the recovery curve parameterization of 3.8, we are effectively modeling a patient's function values *scaled* by their known pre-treatment value. Justification for this approach is given in Section S5 of the [supplementary material](#) available at *Biostatistics* online.

To promote interpretability by encouraging unimodality of conditional distributions, we again use the alternative parameterization of the beta distribution described above for the likelihood. Each patient's initial drop in function value (B_i) is centered at a mode given by the expected drop based on patients

with similar observable characteristics but with spread ϕ_B . A similar interpretation applies to the eventual drop, A_i . To ensure that C_i is also interpretable, we perform a similar reparameterization for the gamma distribution. A $\text{gamma}_{m,\phi}(m', \phi')$ distribution has mode m' and for *all* m' , is unimodal if and only if spread parameter $\phi' \in (0, 1)$. Examples of such distributions and details of the reparameterization are in [Figure S9](#) and [Section S1](#) of the [supplementary material](#) available at *Biostatistics* online. Under this reparameterization, the interpretation of the model for C_i matches A_i and B_i . The patient's rate of recovery is modeled using a gamma distribution centered at the modal rate of observably similar patients, with spread around that mode given by ϕ_C .

The necessity for these specialized parameterizations of the beta and gamma distributions becomes clear if we consider a model that does not use them. Consider the more traditional $\text{beta}_{\mu,\beta}(\mu', \beta')$ parameterization, where a $\text{beta}_{\mu,\beta}(\mu', \beta')$ distribution has mean μ' , and β' is a spread parameter. Suppose we had let $A_i|b_A, \phi_A; z_A, X_i \sim \text{beta}_{\mu,\beta}(\text{logistic}(z_A + b_A X_i), \phi_A)$. A $\text{beta}_{\mu,\beta}(\mu', \beta')$ distribution is unimodal if and only if $\beta' > 1$ and $\mu' > \frac{1}{1+\beta'}$. Given β' , there is some μ' for which a $\text{beta}_{\mu,\beta}(\mu', \beta')$ distribution is not unimodal. Thus given b_A and ϕ_A , there would exist some X_i for which $A_i|b_A, \phi_A; z_A, X_i$ would not be unimodal, which violates Property 3. Similar reasoning applies to the gamma parameterization.

At this point, only the prior distributions for the hyperparameters must be specified to complete the model description. We encourage regularization on the regression coefficients by letting:

$$\phi_A; \lambda_A \sim \exp(\lambda_A, 1), \phi_B; \lambda_B \sim \exp(\lambda_B, 1), \phi_C; \lambda_C \sim \exp(\lambda_C, 1), \quad (4.6)$$

where $\exp(\lambda, 1)$ denotes an exponential distribution with rate parameter λ truncated on the right at 1 and $\lambda_A, \lambda_B, \lambda_C$ are hyperparameters.

Further, we assume there is some “average” recovery shape $g(\cdot; \mu_A, \mu_B, \mu_C)$ such that the prior expected recovery curve of a “average” patient (one whose value of each feature is equal to the mean of that feature in the data set) is centered about $S_i g(\cdot; \mu_A, \mu_B, \mu_C)$ (see 3.7). That is, for the “average” patient, we want the conditional prior distributions of A_i, B_i , and C_i to be centered at μ_A, μ_B , and μ_C , respectively. We will normalize all features to have mean 0 and unit standard deviation, so that the “average” patient has a feature vector consisting of all 0's. Thus in light of 4.2, 4.3, and 4.4, we let

$$z_A \sim \text{normal}(\text{logit}(\mu_A), s_A), z_B \sim \text{normal}(\text{logit}(\mu_B), s_B), z_C \sim \text{normal}(\exp(\mu_C), s_C), \text{ and} \quad (4.7)$$

$$b_A \sim \text{multi_normal}(\vec{0}, s_A I), b_B \sim \text{multi_normal}(\vec{0}, s_B I), b_C \sim \text{multi_normal}(\vec{0}, s_C I), \quad (4.8)$$

where $\mu_A, \mu_B, \mu_C \in \mathbb{R}$, $s_A, s_B, s_C \in \mathbb{R}^+$ are hyperparameters, $\vec{0}$ is the K -dimensional 0 vector, and I is the K -dimensional identity matrix. Note that the intercept for the regressions of 4.2, 4.3, and 4.4 is given a prior and not fixed.

Finally, without any prior belief about the parameters p, θ governing the likelihood, we let:

$$p \sim \text{unif}(0, 1), \theta \sim \text{unif}(0, 1). \quad (4.9)$$

4.2. Recap of model features

We recap below the desired properties of our model, and how our model satisfies those properties.

1. *Property:* Observed within-patient function values should be dependent, and post-treatment values for patients with similar covariates should be shrunk towards each other.

Solution: We adopted a hierarchical Bayesian model. Shrinkage was accomplished by letting A_i be drawn from a single covariate dependent distribution. In particular, A_i was modeled using (a variant of) a generalized linear model. An analogous approach models B_i and C_i .




Fig. 3. For each parameter, the mean signed error over the simulated data sets decreases with the size of the simulated data sets. Dotted lines denote 1 standard deviation.

2. *Property:* For the sake of interpretability, for each patient, their distribution over the underlying post-treatment function value should be a recovery curve—those functions satisfying requirements 3.2–3.5. A predictive distribution like that in Figure 2a is not acceptable.
Solution: We respect the constraints of 3.9 in modeling A_i, B_i, C_i , letting their generalized linear models have beta, beta, and gamma response distributions, respectively, as these are canonical distributions with the desired support.
3. *Property:* For the sake of interpretability, we want the posterior of $f_i(t)$ to be unimodal. For example, we do not want the predictive distribution to be bimodal, like that in Figure 2b.
Solution: The conditional distribution $A_i|b_A, \phi_A; X_i$ was constrained to be unimodal, for all X_i, b_A , and ϕ_A . An analogous approach and constraint were used to model B_i and C_i . Ensuring this unimodality required special parameterizations of the beta and gamma distributions.
4. *Property:* $y_i(t)$ should have support on the *closed* unit interval, because we observed that roughly 5% of the time, patients recorded a 0 or 1 response.
Solution: $y_i(t)$ comes from a mixture of a beta centered at $f_i(t)$ and a Bernoulli distribution.
5. *Property:* In the prior, a patient's distribution over recovery shapes should be centered about some “average” shape, given by curve parameters μ_A, μ_B, μ_C .
Solution: The GLM modeling A_i depends on hyperparameter bias term z_A . z_A was chosen so that in the prior, $A_i|B_A, \phi_A; X_i$ is centered at μ_A . Analogous approaches model B_i and C_i .

5. SIMULATION STUDIES

Here, we examine the ability of our model to recover the model parameters as the amount of data simulated using those parameters grew. We chose a single set of shared model parameters and hyperparameters μ_A, μ_B, μ_C . Then, we performed the following for several values of N , the number of patients in a simulated data set: We simulated 100 data sets, where for each data set we used that set of chosen parameters to simulate observed function values $y_i(t)$ for N patients at times $t \in \{1, 2, 4, 8, 12, 18, 24, 30, 36, 42, 48\}$, the same times at which data were observed in the prostate cancer data set. For each data set, we obtained for each parameter a point prediction as its posterior median and calculated two quantities: the signed error and unsigned error. Figure 3 shows the mean and standard deviation of the signed error across the 100 data sets for each parameter, for various values of N ($N \in \{50, 100, 250, 500, 1000, 2500, 5000\}$). Please see Figure S11 of the [supplementary material](#) available at *Biostatistics* online for the analogous information, for the unsigned error. Note that Figure 3 thus shows the bias and variance of the point estimates of the parameters, with the estimator being their respective posterior medians.

The set of parameters we used was simply one that was not pathological. We used $b_A = 1, b_B = 2, b_C = 3, \theta = 0.1, p = 0.3, \phi_A = \phi_B = \phi_C = \phi_M = 0.01$, and $\mu_A = 0.4, \mu_B = 0.7, \mu_C = 5$. We assume only one feature, which for each sample is generated from a unit normal distribution. For inference, we set

$s_A = s_B = s_C = 1$, $\lambda_A = \lambda_B = \lambda_C = 10$, and $\lambda_M = 10$. To obtain posterior samples, we used Stan (Hoffman and Gelman, 2014), obtaining 2500 samples from each of 4 chains with no thinning, using 2500 burn-in steps. We assessed convergence both by using the Gelman statistic (Gelman and Rubin, 1992) and visual examination of the traces for each parameter. We checked that in fitting the model to each simulated data set, the maximum Gelman statistic over parameters was less than 1.2. The meaning of errors in the regression parameter is provided by 4.2–4.4, and the meaning of errors in the spread parameters $\phi_A, \phi_B, \phi_C, \phi_M$ is provided by the plots of beta and gamma distributions in Figures S3 and S9 of the [supplementary material](#) available at *Biostatistics* online.

6. ANALYSIS OF PROSTATE CANCER DATA SET

6.1. Data set description

Our data come from a study (Gore and others, 2009, 2010) that prospectively tracked the sexual function as measured using the UCLA Prostate Cancer Index (Litwin and others, 1998) of 304 patients who underwent radical prostatectomy to treat clinically localized prostate cancer. After applying data set filters as detailed in Section S3 of the [supplementary material](#) available at *Biostatistics* online, data from 237 patients are retained. Their sexual function levels were collected right before treatment and over a 48-month post-treatment study period via mailed surveys at 1, 2, 4, 8, 12, 18, 24, 30, 36, 42, and 48 months after their respective treatments, and missing data was due to lack of survey response. The Prostate Cancer Index, derived from answers to a series of multiple choice questions, is a numerical measure of a patient's level of sexual function that lies between 0 and 100, which we scale to the unit interval. Various patient covariates were collected at time of treatment, including age, cancer grade/stage, physical/mental condition, urinary/bowel function, and comorbidity count.

Prostate cancer patients' post-prostatectomy sexual function outcomes can be modulated by non-mandatory post-prostatectomy treatments such as the use of an erectile aid. As such additional treatments are non-standard, our goal in this particular analysis is to model the sexual function outcomes for patients who would not receive them. Furthermore, we are not interested in modeling the post-prostatectomy sexual function of patients whose sexual function prior to the potential prostatectomy is already close to 0, as such patients' post-treatment sexual function would be expected to remain constant afterwards, following a different model that is uninteresting to analyze. Thus, we define the target population of this model to be patients considering a prostatectomy, who satisfy the following two properties: firstly, they would not use any additional remedial treatments post-prostatectomy, such as an erectile aid, and secondly, they would have a non-negligible level of sexual function prior to receiving the potential prostatectomy. The data set filters we applied retain members of this target population.

6.2. Choosing features

To identify potential correlates of recovery curve shapes, for every patient, we used curve fitting to find the A, B, C parameters corresponding to their post-event recovery shapes. We made scatter plots of each of those parameters against all available covariates to identify ones that correlated with curve parameters, and identified the pre-treatment sexual function level (referred to as “init” in all figures) and patient age (at treatment time) to be the two covariates most strongly correlated with curve parameters. From the scatter plots (in [Figure S14](#) of the [supplementary material](#) available at *Biostatistics* online), we saw the relationship between those two covariates and curve parameters is likely nonlinear. Thus, we created binned categorical features based on age and pre-treatment function level. The bins for the categorical features we used were as follows:

- age (in years): 0 to 55, 55 to 65, 65+
- pre-treatment function level: 0 to 41, 41 to 60, 60 to 80, 80 to 100.

We note that such subdivisions matches with the urologist co-author's clinical experience regarding how urologists categorize age and pre-treatment function level.

Thus, in our model, patients belong to 1 of 12 classes, depending on into which of the three age groups they fall into, and which of the four intervals their pre-treatment sexual function level lies. These features were normalized. To visualize the effect of these two covariates on recovery shape from another view, we stratified the patients by age category and pre-treatment sexual function level category, and plotted (see Figure S17 of the [supplementary material](#) available at *Biostatistics* online) the average shape of the patients in each category.

6.3. Fitting our model

Now, we describe how we chose hyperparameters and the fitting of the model. To choose μ_A, μ_B, μ_C , which describe the average recovery shape for the “average” patient in the target population, we fit a recovery shape using our parametric form to the training fold-wide average scaled function value shown in Figure 1a (labeled “average shape”). The values we used for the remaining hyperparameters were $s_A = s_B = s_C = 1$ and $l_A = l_B = l_C = l_M = 10$. We show in Section 6.1 of the [supplementary material](#) available at *Biostatistics* online that out-of-sample performance (as described in Section 6.4) is not sensitive to the particular choice of those hyperparameters.

To fit the model, we used Stan ([Hoffman and Gelman, 2014](#)), for each of 4 chains, running 2500 steps with 2500 burn-in steps and no thinning, and assessed convergence using the Gelman statistic ([Gelman and Rubin, 1992](#)) (The maximum value of the Gelman statistic over all parameters was 1.11). Please see Section 6.3 of the [supplementary material](#) available at *Biostatistics* online for posterior predictive checks.

6.4. Out-of-sample performance

We measure the performance of our model by its ability to predict $y_i(t)$, the observed function values. We obtain a point prediction of $y_i(t)$, denoted $\hat{y}_i(t)$, via the median of the posterior distribution of $f_i(t)$, the “underlying” function value. The loss function we use to measure performance was absolute prediction error: the absolute difference between $y_i(t)$ and $\hat{y}_i(t)$. To measure out-of-sample performance, we performed 5-fold cross-validation, obtaining, for each test sample, point predictions from our model, and examined the average, over the test folds, of loss at a given time as measured by absolute prediction error. (The average loss at time t for a test fold consisting of the patient index set I for which function values were recorded at time t is $\frac{1}{|I|} \sum_{i \in I} |\hat{y}_i(t) - y_i(t)|$, where $\hat{y}_i(t)$ is the point prediction of the function level of patient i at time t and $|I|$ is the size of index set I .) In particular, the entire time series for patients in the testing folds are predicted given the entire time series of the patients in the training fold. Data from the early part of one patient's time series are *not* used to predict the same patient's future values. We plot, over time, the out-of-sample performance of our model, as well as that of two baseline models in Figure 4a. Note that all comparison models, like our model, first predict the patient's *scaled* function values, and then multiply it by their pre-treatment value to obtain a prediction of *absolute* function value. To compare the improvement of our method to the *status quo*, in which a doctor merely tells a patient the population-wide average shape, we plotted the performance of simply predicting a patient to have the average recovery shape, labeled “mean”. We compared the performance of our model to a timewise scaled regression that, at each of the 11 common time points, uses a separate generalized linear regression model to relate the scaled function value at the time point to patient features. This model, labeled “scaled regression”, uses a logistic inverse link function and assumes a normal response distribution. Finally, because of the high variance in our data, we show for comparison the *in-sample* performance of a model that is prone to overfitting. This model, labeled “median”, in order to make a prediction for a patient at a given time, looks at which of the 12 patient classes the patient belongs to and then calculates the median scaled function




Fig. 4. Our model identifies the relationship between pre-treatment level, age, and a patient's latent recovery shape. (a) The out-of-sample performance of our model is comparable to that of timewise scaled regression, and approaches the *in-sample* performance of “median,” a method prone to overfitting. (b) The posterior predictive distribution over recovery curves (black) and timewise medians of it (red) convey more plausible predictions than that of timewise scaled regression (blue), whose prediction is not guaranteed to be a recovery curve. (c) age trend. (d) pre-treatment level trend.

value, at the given time, of patients who belong to that patient class, over the *entire* data set. As can be seen in Figure 4a, the out-of-sample performance of our model is roughly equivalent to that of the scaled regression model, though our model is more interpretable. Error bars show the variance in estimates of the expected loss at each time.

6.5. Interpretability of model

Our model achieves out-of-sample performance comparable to that of the timewise scaled regression described in Section 6.4. However, our model produces much more easily understood predictions, outputting a distribution over time series consisting solely of recovery curves, so that they are smoothly increasing monotonically towards an asymptote, and do not exceed the pre-treatment value. In contrast, the timewise scaled regression model produces a time series that is not guaranteed to be smooth or monotonically increasing. In addition to matching prior expectations, our model's predictions are more quickly processed by the patient, which, as our references to interpretability indicated, are crucial in a clinical setting. To illustrate, in Figure 4b for several of the 12 classes of patients, we plot the scaled function

values produced by scaled timewise regression, the distribution over f_i from our model, and the timewise median of that distribution. A patient, expecting to see a recovery curve, can with a quick glance of the red curves (timewise median of our predictions), pick up what the initial drop, asymptotic drop, and recovery rate of their predicted recovery curve are. On the other hand, with the scaled timewise regression, the patient tries to extrapolate what those same quantities are from the jagged predictions, finds it hard to do so, wondering whether the fluctuations are a real trend or just noise.

Furthermore, we have designed our model so that prediction *uncertainty* is easily interpretable when a patient's posterior distribution of curves is plotted. Because we encourage the a patient's recovery curve parameter distribution to be unimodal in the posterior, we expect the pointwise distribution of curve values, namely that of f_i , to be unimodal. This is why in Figure 4b, the distribution of curves appears clustered about the red curve. It is important that one can visually extract from a plot of posterior distribution of curves a *single* most likely curve. Then, such a plot can be interpreted as giving a single curve prediction, along with the uncertainty in that prediction. On the other hand, if the posterior distribution of curves were clustered around, say, two curves, there would be no such clear interpretation.

6.6. Dependence of recovery shape on covariates and comparison with literature

Our analysis teases apart the dependence of recovery shape on age and pre-treatment value. In Figure 4c, we examine the effect of patient age on recovery curve shape by stratifying those curves by pre-treatment level. We find when pre-treatment level is controlled for, patients younger than 55 years of age have a smaller asymptotic drop in sexual function level, proportional to their pre-treatment level. (We performed a one-sided z-test that the scaled function value at 48 months for patients younger than 55 years of age was larger than those not; p-value = 0.005.) This effect is diminished for patients with pre-treatment level higher than 0.80. When pre-treatment level is not controlled for, the asymptotic proportional drop in function level for younger patients is lower. In both cases, the proportional initial drop in function level does not depend on age. In Figure 4d, we examine the effect of pre-treatment sexual function level on recovery shape by stratifying those curves by age. We find that when age is controlled for, patients with pre-treatment level higher than 0.80 have a smaller asymptotic drop in function level, proportional to their pre-treatment level. (One sided z-test p-value = 4.17×10^{-7} .) However, this effect is diminished for patients younger than 55. The proportional initial drop in function depends mildly on pre-treatment level.

Unlike past methods, which have mostly focused on modeling a continuous or binary measure of sexual function at a single fixed time, our model makes predictions of the entire post-treatment function trajectory. Regardless, we can still compare our findings to them. Past work that modeled a continuous measure of sexual function found that lower age and higher pre-treatment sexual function level are statistically linked to higher *absolute* levels of that measure (Talcott and others, 2003), and that lower age is linked to a smaller *change* in that measure of function level (Sanda and others, 2008). Likewise, when a binary indicator of satisfactory sexual function has been logistically regressed against patient covariates, lower age (Ayyathurai and others, 2008; Regan and others, 2011) and higher pre-treatment function level (Regan and others) have been found to lead to a higher probability of having satisfactory sexual function. One can conclude from these past statistical analyzes, as well as model-free data analyzes (Rabbani and others, 2000; Michl and others, 2006), that lower age and higher pre-treatment sexual function level, by any measure, are linked to higher post-treatment sexual function level, agreeing with our findings. Though, we stress that unlike any previous analysis, we model the link between patient features and longitudinal sexual function levels *proportional* to the pre-treatment level.

7. CONCLUSION

We presented a Bayesian model that can be used to predict recovery curves, which arise in many medical contexts. Our overarching goal is to facilitate the flow of information from the data to the user, who may

not be statistically inclined. Towards this end, we impart interpretability to both the model and its output, a model that is easily explained and produces believable outputs is more clinically applicable. In particular, our model predicts quantities that are of natural interest, and guarantees that its output is in fact the recovery curve that we assume a domain expert to expect of a prediction. Furthermore, our model is designed for easy visualization of predictions and the associated uncertainty, as we encourage the posterior distribution over recovery curves to have a clear mode. We used our model to analyze the impact of prostatectomy on a patient's post-treatment sexual function trajectory, and characterized the extent of that impact on patient age and pre-treatment sexual function level, producing conclusions that agree with and supplement past findings. We believe our model can provide insights in other medical domains as well.

SUPPLEMENTARY MATERIAL

The [supplementary material](http://biostatistics.oxfordjournals.org) is available at <http://biostatistics.oxfordjournals.org>.

ACKNOWLEDGMENTS

We thank Dr. Jim Michaelson of Massachusetts General Hospital for helpful discussions. *Conflict of Interest*: None declared.

FUNDING

National Science Foundation CAREER grant (IIS-1053407 to C.R.) and National Science Foundation grant (DMS-1737673 to T.H.M).

REFERENCES

- AYYATHURAI, R., MANOHARAN, M., NIEDER, A. M., KAVA, B. AND SOLOWAY, M. S. (2008). Factors affecting erectile function after radical retropubic prostatectomy: results from 1620 consecutive patients. *BJU International* **101**, 833–836.
- BRIGANTI, A., GALLINA, A., SUARDI, N., CAPITANIO, U., TUTOLO, M., BIANCHI, M., SALONIA, A., COLOMBO, R., DI GIROLAMO, V., MARTINEZ-SALAMANCA, J. I. and others (2011). What is the definition of a satisfactory erectile function after bilateral nerve sparing radical prostatectomy? *The Journal of Sexual Medicine* **8**, 1210–1217.
- CAI, B. AND DUNSON, D. B. (2007). Bayesian multivariate isotonic regression splines. *Journal of the American Statistical Association* **102**, 1158–1171.
- DENISON, D. G. T., MALICK, B. K. AND SMITH, A. F. M. (1998). Automatic Bayesian curve fitting. *Journal of Royal Statistical Society Series B* **60**, 333–350.
- DESCAZEUD, A., DEBRÉ, B. AND FLAM, T. A. (2006). Age difference between patient and partner is a predictive factor of potency rate following radical prostatectomy. *The Journal of Urology* **176**, 2594–2598.
- EASTHAM, J. A., SCARDINO, P. T. AND KATTAN, M. W. (2008). Predicting an optimal outcome after radical prostatectomy: the trifecta nomogram. *The Journal of Urology* **179**, 2207–2211.
- GELMAN, A. AND RUBIN, D. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science* **7**, 457–511.
- GORE, J. L., GOLLAPUDI, K., BERGMAN, J., KWAN, L., KRUPSKI, T. L. AND LITWIN, M. S. (2010). Correlates of bother following treatment for clinically localized prostate cancer. *The Journal of Urology* **184**, 1309–1315.
- GORE, J. L., KWAN, L., LEE, S. P., REITER, R. E. AND LITWIN, M. S. (2009). Survivorship beyond convalescence: 48-month quality-of-life outcomes after treatment for localized prostate cancer. *Journal of the National Cancer Institute* **101**, 888–892.

- HARTZLER, A. L., IZARD, J. P., DALKIN, B. L., MIKLES, S. P. AND GORE, J. L. (2015). Design and feasibility of integrating personalized PRO dashboards into prostate cancer care. *Journal of the American Medical Informatics Association* **23**, 38–47.
- HOFFMAN, M. D. AND GELMAN, A. (2014). The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, **15**, 1593–1623.
- JUNG, T. AND WICKRAMA, K. (2008). An introduction to latent class growth analysis and growth mixture modeling. *Social and Personality Psychology Compass* **2**, 302–317.
- LITWIN, M. S., HAYS, R. D., FINK, A., GANZ, P. A., LEAKE, B. AND BROOK, R. H. (1998). The UCLA Prostate Cancer Index: development, reliability, and validity of a health-related quality of life measure. *Medical Care* **36**, 1002–1012.
- MAMMEN, E. (1991). Estimating a smooth monotone regression function. *The Annals of Statistics* **19**, 724–740.
- MICHL, U. H. G., FRIEDRICH, M. G., GRAEFEN, M., HAESE, A., HEINZER, H. AND HULAND, H. (2006). Prediction of postoperative sexual function after nerve sparing radical retropubic prostatectomy. *The Journal of Urology* **176**, 227–231.
- NAYAK, J. G., HARTZLER, A. L., MACLEOD, L. C., IZARD, J. P., DALKIN, B. M. AND GORE, J. L. (2016). Relevance of graph literacy in the development of patient-centered communication tools. *Patient Education and Counseling* **99**, 448–454.
- NEELON, B. AND DUNSON, D. B. (2004). Bayesian isotonic regression and trend analysis. *Biometrics* **60**, 398–406.
- POTOSKY, A. L., DAVIS, W. W., HOFFMAN, R. M., STANFORD, J. L., STEPHENSON, R., PENSON, D. F. AND HARLAN, L. C. (2004). Five-year outcomes after prostatectomy or radiotherapy for prostate cancer: the Prostate Cancer Outcomes Study. *Journal of the National Cancer Institute* **96**, 1358–1367.
- RABBANI, F., STAPLETON, A. M., KATTAN, M. W., WHEELER, T. M. AND SCARDINO, P. T. (2000). Factors predicting recovery of erections after radical prostatectomy. *The Journal of Urology* **164**, 1929–1934.
- RAMSAY, J. O. AND SILVERMAN, B. W. (editors). (2002). *Applied Functional Data Analysis: Methods and Case Studies*. New York, NY: Springer Series in Statistics.
- REGAN, M. M., COOPERBERG, M. R., WEI, J. T., MICHALSKI, J. M., SANDLER, H. M., LITWIN, M. S., KLEIN, E., KIBEL, A. S., HAMSTRA, D. A., PISTERS, L. L. and others (2011). Prediction of erectile function following treatment for prostate cancer. *Journal of the American Medical Association* **306**, 1205–1214.
- ROGOSA, D. R. AND WILLETT, J. B. (1985). Understanding correlates of change by modeling individual differences in growth. *Psychometrika* **50**, 203–228.
- ROLFE, M. I., Mengersen, K. L., VEARNCOMBE, K. J., ANDREW, B. AND BEADLE, G. F. (2011). Bayesian estimation of extent of recovery for aspects of verbal memory in women undergoing adjuvant chemotherapy treatment for breast cancer. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **60**, 655–674.
- SANDA, M. G., DUNN, R. L., MICHALSKI, J., SANDLER, H. M., NORTHOUSE, L., HEMBROFF, L., LIN, X., GREENFIELD, T. K., LITWIN, M. S., SAIGAL, C. S. and others (2008). Quality of life and satisfaction with outcome among prostate-cancer survivors. *The New England Journal of Medicine* **358**, 1250–1261.
- SHIVELY, T. S., SAGER, T. W. AND WALKER, S. G. (2009). A Bayesian approach to non-parametric monotone function estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **71**, 159–175.
- TALCOTT, J. A., MANOLA, J., CLARK, J. A., KAPLAN, I., BEARD, C. J., MITCHELL, S. P., CHEN, R. C., O'LEARY, M. P., KANTOFF, P. W. AND D'AMICO, A. V. (2003). Time course and predictors of symptoms after primary prostate cancer therapy. *Journal of Clinical Oncology* **21**, 3979–3986.

- TILLING, K., STERNE, J. A. C. AND WOLFE, C. D. A. (2001). Multilevel growth curve models with covariate effects: application to recovery after stroke. *Statistics in Medicine* **20**, 685–704.
- WARSCHAUSKY, S., KAY, J. B. AND KEWMAN, D. G. (2001). Hierarchical linear modeling of FIM instrument growth curve characteristics after spinal cord injury. *Archives of Physical Medicine and Rehabilitation* **82**, 329–334.

[Received September 27, 2016; revised November 3, 2017; accepted for publication November 22, 2017]