Bayesian change-point analysis for atomic force microscopy and soft material indentation

Daniel Rudoy, Shelten G. Yuen, Robert D. Howe and Patrick J. Wolfe

Harvard University, Cambridge, USA

Summary. Material indentation studies, in which a probe is brought into controlled physical contact with an experimental sample, have long been a primary means by which scientists characterize the mechanical properties of materials. More recently, the advent of atomic force microscopy, which operates on the same fundamental principle, has in turn revolutionized the nanoscale analysis of soft biomaterials such as cells and tissues. The paper addresses the inferential problems that are associated with material indentation and atomic force microscopy, through a framework for the change-point analysis of pre-contact and post-contact data that is applicable to experiments across a variety of physical scales. A hierarchical Bayesian model is proposed to account for experimentally observed change-point smoothness constraints and measurement error variability, with efficient Monte Carlo methods developed and employed to realize inference via posterior sampling for parameters such as Young's modulus, which is a key quantifier of material stiffness. These results are the first to provide the materials science community with rigorous inference procedures and quantification of uncertainty, via optimized and fully automated high throughput algorithms, implemented as the publicly available software package BayesCP. To demonstrate the consistent accuracy and wide applicability of this approach, results are shown for a variety of data sets from both macromaterials and micromaterials experiments—including silicone, neurons and red blood cells—conducted by the authors and others.

Keywords: Change-point detection; Constrained switching regressions; Hierarchical Bayesian models; Indentation testing; Markov chain Monte Carlo methods; Materials science; Young's modulus

1. Introduction

This paper develops a hierarchical Bayesian approach for contact point determination in material indentation studies and atomic force microscopy (AFM). Contemporary applications in materials science and biomechanics range from analysing the response of novel nanomaterials to deformation (Wong *et al.*, 1997) to characterizing disease through mechanical properties of cells, tissues and organs (Costa, 2004). Experimental procedures and analyses, however, remain broadly similar across these different types of material and physical scales (Gouldstone *et al.*, 2007), with the scientific aim in all cases being to characterize how a given material sample deforms in response to the application of an external force.

As illustrated in Fig. 1, indentation experiments employ a probe (or cantilever arm, in the case of AFM) to apply a controlled force to the material sample. This indenting probe displaces the sample while concurrently measuring resistive force, with the resultant force–position data used to infer material properties such as stiffness (by analogy with compressing a spring to determine

Address for correspondence: Patrick J. Wolfe, Statistics and Information Sciences Laboratory, Harvard University, Oxford Street, Cambridge, MA 02138-2901, USA.
E-mail: wolfe@stat.harvard.edu
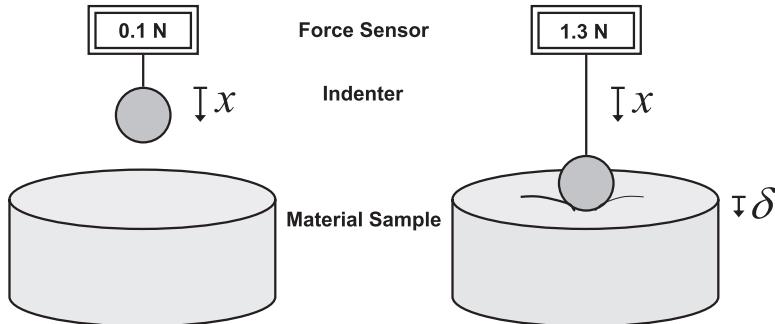
Fig. 1. Diagram of a macroscale indentation experiment in which a spherical probe, attached to a force sensor, indents a material sample and deforms it by a distance δ: in this hypothetical example, a net change of 1.2 N in resistive force is consequently observed

its spring constant experimentally). Before subsequent data analysis, a key technical problem is to determine precisely the moment at which the probe comes into contact with the material. Sample preparation techniques and sizes frequently preclude the direct measurement of this contact point, and hence its inference from indenter force–position data forms the subject of this paper.

At present, practitioners across fields lack an agreed standard for contact point determination; a variety of *ad hoc* data preprocessing methods are used, including even simple visual inspection (Lin *et al.*, 2007a). Nevertheless, it is well recognized that precise contact point determination is necessary to infer material properties accurately in AFM indentation experiments (Crick and Yin, 2007). For example, Dimitriadis *et al.* (2002) showed that, for small displacements of thin films, estimation errors of the order of 5 nm for a 2.7-μm sample can cause an increase of nearly 200% in the estimated Young's modulus—the principal quantifier of material stiffness. When soft materials such as cells are studied at microscopic scales, e.g. to determine biomechanical disease markers (Costa, 2004), the need for robust and repeatable AFM analyses becomes even greater (Lin *et al.*, 2007a).

In this paper, we present the first formulation of the contact point determination task as a *statistical change-point* problem, and we subsequently employ a switching regressions model to infer Young's modulus. Section 2 summarizes the basic principles of material indentation, showing that the resultant force–displacement curves are often well described by low order polynomials. Section 3 introduces a corresponding family of Bayesian models designed to address a wide range of experimental conditions, with specialized Markov chain Monte Carlo samplers for inference developed in Section 4. Following validation of the proposed inference procedures in Section 5, they are employed in Section 6 to infer material properties of mouse neurons and human red blood cells from AFM force–position data. The paper concludes in Section 7 with a discussion of promising methodological and practical extensions.

2. Material indentation

2.1. Indentation experiments and data

Indentation experiments proceed by carefully moving a probe from an initial non-contact position into a material sample, as shown in Fig. 1, while measuring the resistive force at some prescribed temporal sampling rate. After a small deformation has been made, the probe retracts to its initial position; during this stage the resistive force decreases with every subsequent
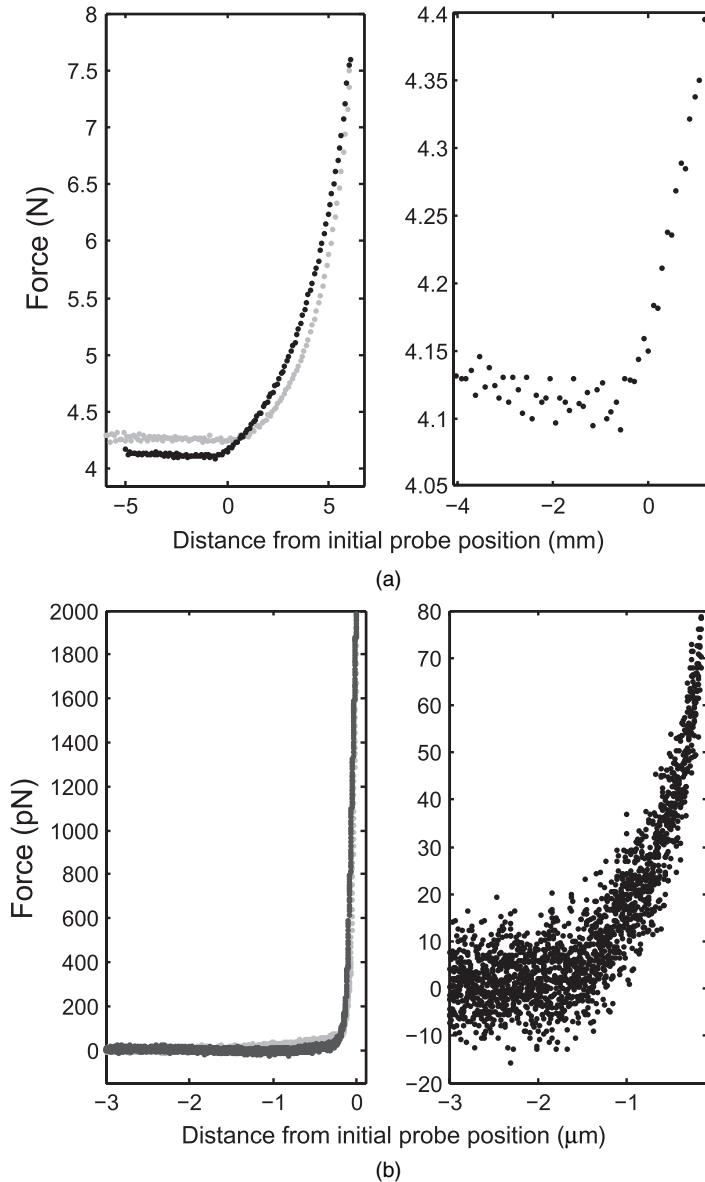
Fig. 2. Example of force–position curves for (a) the silicone and (b) the red blood cell indentation data, which from material stiffness properties are to be inferred, with force measurements during indentation and subsequent retraction shown in black and grey respectively and subsets of the indentation data shown near the presumed contact points: experiments were performed by using a mechanical arm for the soft silicone sample and an atomic force microscope for the red blood cell sample; note the differences in physical scale and level of noise

measurement. At the conclusion of each such experiment, two force–position curves are produced, examples of which are shown in Fig. 2. In this paper we consider only the forward indentation curves, as is standard practice (Lin *et al.*, 2007a), though the methods that we present are extendable to retraction data whenever suitable models are available.

Despite significant differences in physical scale and noise level, the curves in Fig. 2 (and,

indeed, most indentation data sets) feature several common characteristics. In the precontact region, the force response appears linear in the position of the indenter; drift due to experimental conditions is often present, yielding a non-zero slope as in Fig. 2(a). The post-contact data appear well modelled by low order polynomial functions of the corresponding displacement and, indeed, the key assumption of such experiments is that, conditioned on knowledge of the geometry of the indenting probe, the relationship between the degree of material deformation and the measured resistive force depends *in a known way* on the stiffness of the material. This is analogous to the case in which an ideal spring is compressed a specified distance δ by a force of known magnitude; a measure of the spring's stiffness is given by the spring constant k, which may be calculated by using Hooke's law: $F = -k\delta$.

2.2. Contact mechanics and the Hertzian model

Indentation data such as those shown in Fig. 2 are typically acquired for *linear elastic* materials. This implies, not only that the material instantaneously returns to its original shape following the cessation of an external force, but also that the relationship between the applied stress (the force per unit area) and the resultant strain (the deformation per unit length) is linear. This ratio is known as *Young's modulus*—the primary quantifier of material stiffness introduced above—and is reported in units of pascals.

In small deformation experiments to infer Young's modulus E for linear elastic materials, both the indenter geometry and the measured resistive force come into play, by way of the so-called *Hertzian model* (Lin *et al.*, 2007a). Specifically, the relationship between the sample deformation depth δ and the measured resistive force F takes the form

$$F \propto E\delta^{\beta}, \tag{1}$$

where the constant of proportionality and the fixed parameter β depend on the geometry of the indenter tip in a known way. Examples include indentation by a sharp pyramid with tip angle 2ϕ or a sphere of radius R, whereupon expression (1) respectively takes the following forms:

$$
\begin{aligned}
F &= \frac{1.5\tan(\phi)}{2(1-\nu^2)}E\delta^2, \\
F &= \frac{4R^{1/2}}{3(1-\nu^2)}E\delta^{3/2},
\end{aligned}
\tag{2}
$$

with ν a known dimensionless quantity termed Poisson's ratio. A subsequent fitting of the Hertzian model (1) to experimental data thus allows us to obtain an estimate for Young's modulus E once the post-contact region has been identified. (Below we retain the standard practitioner notation $(E, F, R, \delta, \nu, \phi)$, as distinct from other variables to follow.)

2.3. Precontact and post-contact data regimes

Observe that the Hertzian model (1) posits a relationship between force and indentation *depth*, whereas the data of Fig. 2 are seen to be a function of the *position* of the indenter. The Hertzian model thus describes the underlying physics of the *post*-contact stage of a typical indentation experiment, whereas the measured data also comprise a *precontact* stage. As we detail below, the union of these two regimes is well described by a switching regressions scenario, with force measurements before contact typically linear in the position of the probe. The corresponding

intercept represents the equilibrium tare that is required to achieve zero net force, whereas the slope captures constant velocity drift in force measurements that can arise in a variety of experimental settings.

In both the precontact and the post-contact data regimes, it is standard to assume independence of measurement errors—a reasonable assumption for all practically achievable force sampling rates. Errors before contact arise because of force sensor vibration in the experimental medium, thermal variations and other effects—whereas after contact they also depend on interactions between the probe and the sample such as frictional forces. Consequently, although error variances in these regimes can be expected to differ in practice (analysis of AFM data from a red blood cell, which is shown in Fig. 6 and described later, reveals one such example), their relative magnitudes are not known *a priori*. As a final consideration, uncertainty in the reported position of the indenting probe is typically several orders of magnitude smaller than the distance between consecutive force sampling points and can safely be disregarded.

As noted by Crick and Yin (2007), all the sources of variability that were mentioned above can lead to large relative errors in resistive force measurements near the contact point; Fig. 2(b) illustrates a typical scenario. This makes manual identification of the contact point difficult in many cases and motivates the model-based approach that we now describe.

3. Bayesian change-point model

Having outlined the basic principles of material indentation, we now formulate a Bayesian model for force–position data that encompasses both precontact and post-contact regimes. In light of the discussion above, the corresponding task of contact point determination is recognizable as a change-point estimation problem in the context of switching regressions. In turn, the Hertzian model (1) implies that an estimate of Young's modulus can be obtained as a *linear* function of the leading post-contact regression coefficient. Control over experimental conditions implies that each indentation data set contains precisely one change-point, thus obviating any need to estimate the presence or number of contact points.

During an indentation experiment, the indenter moves continuously through a sequence of n equispaced positions $\mathbf{x} = (x_1, x_2, \ldots, x_n)'$ and records a force measurement y_i at each position x_i, resulting in a sequence of force measurements $\mathbf{y} = (y_1, y_2, \ldots, y_n)'$. As we shall later treat models for soft material indentation, in which the precontact and post-contact curves are constrained to be continuous at the regression change-point, we begin by introducing a continuous parameter $\gamma \in (1, n)$ denoting the *contact point index*, with the corresponding *contact point* at which the indenter first contacts the sample denoted by $x_\gamma \in (x_1, x_n)$.

3.1. Data likelihood for indentation experiments

We adopt a classical switching regressions scenario for our model, in which y is assumed to be a polynomial function of known degree d_1 in position x before contact, and of known degree d_2 in deformation depth $\delta = x - x_\gamma$ after contact with the sample is made. This formulation encompasses the Hertzian model (1) if fractional powers are allowed; however, for clarity of presentation we consider d_2 to be an integer unless otherwise noted. Letting $p = d_1 + d_2 + 2$ denote the number of regression coefficients in our model, and with n the number of data points, the corresponding design matrix is hence of dimension $n \times p$. We subsequently employ the subscript γ to denote any quantity that depends on γ, and index via subscripts 1 and 2 the precontact and post-contact regression regimes.

We denote the regression coefficients by $\beta_1 \in \mathbb{R}^{d_1+1}$ and $\beta_2 \in \mathbb{R}^{d_2+1}$, with design matrices $\mathbf{X}_{1,\gamma}$ and $\mathbf{X}_{2,\gamma}$ defined as follows, for $\lfloor\gamma\rfloor$ the largest integer less than or equal to γ:

$$\mathbf{X}_{1,\gamma} = \begin{pmatrix} 1 & x_1 & \cdots & x_1^{d_1} \\ 1 & x_2 & \ldots & x_2^{d_1} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_{\lfloor\gamma\rfloor} & \cdots & x_{\lfloor\gamma\rfloor}^{d_1} \end{pmatrix}, \qquad \mathbf{X}_{2,\gamma} = \begin{pmatrix} 1 & x_{\lfloor\gamma\rfloor+1} - x_\gamma & \cdots & (x_{\lfloor\gamma\rfloor+1} - x_\gamma)^{d_2} \\ 1 & x_{\lfloor\gamma\rfloor+2} - x_\gamma & \ldots & (x_{\lfloor\gamma\rfloor+2} - x_\gamma)^{d_2} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & x_n - x_\gamma & \ldots & (x_n - x_\gamma)^{d_2} \end{pmatrix}. \quad (3)$$

The observed data $\mathbf{y}$ may likewise be partitioned into precontact and post-contact vectors

$$\mathbf{y}_{1,\gamma} = (y_1, y_2, \ldots, y_{\lfloor\gamma\rfloor})'$$

and

$$\mathbf{y}_{2,\gamma} = (y_{\lfloor\gamma\rfloor+1}, y_{\lfloor\gamma\rfloor+2}, \ldots, y_n)',$$

and, following our discussion in Section 2.3 regarding the noise characteristics that are typical of indentation experiments, we assume independent and normally distributed additive errors, with unknown variances σ_1^2 and σ_2^2. Thus, for $1 \leqslant i \leqslant n$, we have that y_i is distributed as follows:

$$y_i \sim \begin{cases} \mathcal{N}(\mathbf{X}_{1,\gamma}\beta_1, \sigma_1^2) & \text{if } 1 \leqslant i \leqslant \lfloor\gamma\rfloor, \\ \mathcal{N}(\mathbf{X}_{2,\gamma}\beta_2, \sigma_2^2) & \text{if } \lfloor\gamma\rfloor+1 \leqslant i \leqslant n. \end{cases} \quad (4)$$

The statistical model (4) is consistent with the Hertzian mechanics model (1), resulting in a post-contact force–response curve that is a power of the *displacement* $\delta = x - x_\gamma$ of the material sample, rather than the *position* x of the indenter. However, when d_2 is an integer, the coefficients of these two polynomials are related by a simple linear transformation. Consider, for instance, a quadratic curve in δ given by $f(\delta) = a_0 + a_1\delta + a_2\delta^2$. We may rewrite $f(\delta)$ as a quadratic polynomial in x as follows:

$$a_0 + a_1(x - x_\gamma) + a_2(x - x_\gamma)^2 = b_0 + b_1 x + b_2 x^2, \quad (5)$$

where $b_2 = a_2$, $b_1 = a_1 - a_2 x_\gamma$ and $b_0 = a_0 - a_1 x_\gamma + a_2 x_\gamma^2$. This transformation enables $\mathbf{X}_{2,\gamma}$ to be reformulated directly in terms of indenter position x, such that it no longer depends *continuously* on x_γ, in contrast with expression (3). Transformations akin to equation (5) do not apply, however, when the d_2 is a fraction, as in the case of expression (2) for a spherical indenter, or when a continuity constraint is enforced at the change-point; we detail such cases below.

3.2. General parametric Bayesian model for material indentation

The likelihood of expression (4), together with the presence of genuine prior information dictated by the underlying physics of material indentation experiments, suggests a natural hierarchical Bayesian model. In contrast with the semiconjugate approach that was taken by Carlin *et al.* (1992), we detail below a fully conjugate model, as this allows for analytical simplifications that we have observed to be important in practice. Integrating out nuisance parameters improves not only the mixing of the chains underlying the resultant algorithms and inferential procedures, but also their computational tractability when data sizes grow large.

We specify prior distributions for all model parameters, including the contact point index $\gamma \in (1, n)$, the precontact and post-contact regression coefficients β_1 and β_2, and the error variances σ_1^2 and σ_2^2. For $i \in \{1, 2\}$, we then assume that $\beta_i \sim \mathcal{N}(\mu_i, \sigma_i^2 \Lambda_i^{-1})$, with Λ_i a $(d_i+1) \times (d_i+1)$ diagonal positive definite matrix and $\mu_i \in \mathbb{R}^{d_i+1}$. A standard inverse gamma conjugate prior

$\mathcal{IG}(a_0, b_0)$ is adopted for both variances σ_1^2 and σ_2^2. Finally, recalling that a single regression change-point is always assumed within the set of equispaced force–position measurements, we employ a uniform prior distribution on the interval $(1, n)$ for the contact point index γ. In certain experimental settings, whereupon the initial position of the indenter is known to be at least a certain distance from the sample, an informative prior distribution may be available.

Because of sensitivity to prior parameters, we follow standard practice and adopt hyper-priors for increased model robustness (see Section 5). A gamma prior is assumed on b_0 so that $b_0 \sim \mathcal{G}(\kappa, \eta)$; however, we determined from simulations that the posterior estimators considered were not sensitive to the prior parameters a_0, $\boldsymbol{\mu}_i$ and $\boldsymbol{\Lambda}_i$, and so did not employ an additional level of hyperprior hierarchy for the regression coefficients. For notational convenience, we let $\mathbf{1}_k$ denote the 1-vector of dimension k whose entries are all equal to $\mathbf{1}$ and $\mathbf{0}$ the zero matrix of appropriate dimension, and define the variables $\mathbf{y}$, $\boldsymbol{\beta}$, $\boldsymbol{\mu}$, $\mathbf{X}_\gamma$, $\boldsymbol{\Lambda}$, $\boldsymbol{\Sigma}_\gamma$ and $\boldsymbol{\Sigma}$ as follows:

$$\mathbf{y} = \begin{pmatrix} \mathbf{y}_{1,\gamma} \\ \mathbf{y}_{2,\gamma} \end{pmatrix} \in \mathbb{R}^{n \times 1},$$

$$\boldsymbol{\beta} = \begin{pmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{pmatrix} \in \mathbb{R}^{p \times 1},$$

$$\boldsymbol{\mu} = \begin{pmatrix} \boldsymbol{\mu}_1 \\ \boldsymbol{\mu}_2 \end{pmatrix} \in \mathbb{R}^{p \times 1},$$

$$\mathbf{X}_\gamma = \begin{pmatrix} \mathbf{X}_{1,\gamma} & \mathbf{0} \\ \mathbf{0} & \mathbf{X}_{2,\gamma} \end{pmatrix} \in \mathbb{R}^{n \times p},$$

$$\boldsymbol{\Lambda} = \begin{pmatrix} \boldsymbol{\Lambda}_1 & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Lambda}_2 \end{pmatrix} \in \mathbb{R}^{p \times p},$$

$$\boldsymbol{\Sigma}_\gamma = \mathrm{diag}(\sigma_1^2 \mathbf{1}_{\lfloor \gamma \rfloor}, \sigma_2^2 \mathbf{1}_{n - \lfloor \gamma \rfloor}) \in \mathbb{R}^{n \times n},$$
$$\boldsymbol{\Sigma} = \mathrm{diag}(\sigma_1^2 \mathbf{1}_{d_1+1}, \sigma_2^2 \mathbf{1}_{d_2+1}) \in \mathbb{R}^{p \times p}.$$

The posterior probability distribution of the model parameters $(\gamma, \boldsymbol{\beta}, \sigma_1^2, \sigma_2^2, b_0)$, conditioned on the observations $\mathbf{y}$ and the fixed model parameters $\psi \overset{\Delta}{=} (\boldsymbol{\mu}, \boldsymbol{\Lambda}, a_0, \kappa, \eta)$, is then

$$\begin{aligned}
p(\gamma, \boldsymbol{\beta}, \sigma_1^2, \sigma_2^2, b_0 | \mathbf{y}; \psi) \propto\; & p(\mathbf{y} | \boldsymbol{\beta}, \sigma_1^2, \sigma_2^2, \gamma)\, p(\boldsymbol{\beta} | \sigma_1^2, \sigma_2^2; \boldsymbol{\mu}, \boldsymbol{\Lambda})\, p(\sigma_1^2 | b_0; a_0)\, p(\sigma_2^2 | b_0; a_0) \\
& \times p(b_0; \kappa, \eta)\, p(\gamma) \\
\propto\; & \sigma_1^{-2(a_0-1)} \exp(-b_0/\sigma_1^2) \sigma_2^{-2(a_0-1)} \exp(-b_0/\sigma_2^2) b_0^{\kappa-1} \exp(-b_0/\eta) \\
& \times (|\boldsymbol{\Sigma}_\gamma||\boldsymbol{\Sigma}|)^{-1/2} \exp[-\tfrac{1}{2}\{(\mathbf{y} - \mathbf{X}_\gamma \boldsymbol{\beta})' \boldsymbol{\Sigma}_\gamma^{-1} (\mathbf{y} - \mathbf{X}_\gamma \boldsymbol{\beta}) \\
& + (\boldsymbol{\beta} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} \boldsymbol{\Lambda} (\boldsymbol{\beta} - \boldsymbol{\mu})\}].
\end{aligned} \tag{6}$$

To confirm robustness, we also studied the effect of replacing the diagonal prior covariance $\boldsymbol{\Lambda}$ for the precontact and post-contact regression coefficients $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ with an appropriately adapted g-prior (Zellner, 1986) such that

$$\boldsymbol{\beta}_i | \rho_i, \gamma \sim \mathcal{N}\{\boldsymbol{\mu}_i, \sigma_i^2 \rho_i^2 (\mathbf{X}'_{i,\gamma} \mathbf{X}_{i,\gamma})^{-1}\},$$

with ρ_i^2 a scale parameter to which we ascribed a diffuse inverse gamma hyperprior. We observed no measurable effect of this change in priors on the resulting inference—further confirming the insensitivity of the adopted model to the prior covariance of the regression coefficients. Moreover, efficient sampling from the conditional distribution of ρ_i^2 is precluded by its dependence on the contact point index γ, reducing the overall efficacy of this approach in the Markov chain Monte Carlo approaches to posterior sampling that are described in Section 4.

3.3. Smoothness constraints at the change-point

Often force–position curves are continuous at the contact point x_γ. Especially for soft materials such as the red blood cells that we consider in Section 6, it is to be expected that the change in the force measurement is smooth, and a continuity constraint can serve to regularize the solution in cases where many different fits will have high likelihood. Imposing smoothness constraints dates back to at least Hudson (1966), who considered this constraint in deriving maximum likelihood estimators for switching regressions. More recently, Stephens (1994) used it in a hierarchical Bayesian setting.

In our setting, according to the likelihood of expression (4), a continuity constraint on the precontact and post-contact force–position curves at $x = x_\gamma$ implies that

$$\beta_{10} + \beta_{11} x_\gamma + \ldots + \beta_{1d_1} x_\gamma^{d_1} = \beta_{20}, \tag{7}$$

where $\boldsymbol{\beta}_1 = (\beta_{10}, \beta_{11}, \ldots, \beta_{1d_1})'$ and $\boldsymbol{\beta}_2 = (\beta_{20}, \beta_{21}, \ldots, \beta_{2d_2})'$ denote the vectors of precontact and post-contact regression coefficients respectively. Higher order smoothness can also be imposed: we say that the force–position curve is s times continuously differentiable at x_γ if the sth-order derivatives of the precontact and post-contact curves meet at x_γ, with equation (7) corresponding to the case $s = 0$. However, if $X_{2,\gamma}$ were a function of the position x, rather than the displacement $x - x_\gamma$, then the continuity constraint would become

$$\sum_{i=0}^{d_1} \beta_{1i} x_\gamma^i = \sum_{j=0}^{d_2} \beta_{2j} x_\gamma^j. \tag{8}$$

Either continuity constraint implies that the likelihood function is *non-linear* in the contact point x_γ; enforcing more degrees of smoothness at the change-point serves to exacerbate the non-linearity and makes the design of efficient inference methods increasingly difficult. Stephens (1994) imposed equation (8) in a Bayesian switching regressions setting and proposed a rejection sampling step within a Gibbs sampler to address the resultant non-linearity. Later, in Section 4, we describe a more efficient approach that can be applied when either equation (7) or (8) (or higher order analogues) are enforced.

3.4. Change-point estimation and contact point determination in the literature

As demonstrated above, inference for material indentation data is well matched to classical statistical frameworks for change-point estimation. Independently of the specifics of our contact point problem, the last half-century has seen a vast body of work in this area. Sequential and fixed sample size varieties have been considered from both classical and Bayesian viewpoints, with numerous parametric and non-parametric approaches proposed. We refer the interested reader to several excellent surveys, including those by Hinkley *et al.* (1980), Zacks (1983), Wolfe and Schechtman (1984), Carlin *et al.* (1992) and Lai (1995). Some of the earliest work on maximum likelihood estimation of a single change-point between two polynomial regimes was done by Quandt (1958) and Robison (1964).

Historically, Chernoff and Zacks (1964) were among the first to consider a parametric Bayesian approach to change-point estimation. Change-points arising specifically in linear models have been treated by many researchers, including Bacon and Watts (1971), Ferreira (1975), Smith (1975, 1980), Choy and Broemling (1980), Smith and Cook (1980) and Menzefricke (1981). The introduction of Markov chain Monte Carlo methods has led to more sophisticated hierarchical Bayesian models for change-point problems, beginning with the semiconjugate approach that was taken by Carlin *et al.* (1992), in which the prior variance of the regression coefficients is left unscaled by the noise variance. Advances in transdimensional simulation methods have rekindled

interest in multiple-change-point problems, as discussed by Stephens (1994), Punskaya *et al.* (2002) and Fearnhead (2006), among others.

In the context of material indentation, however, existing approaches to contact point determination do not make use of change-point estimation methodology. In fact, current methods are error prone and labour intensive—even consisting of visual inspection and manual thresholding (Lin *et al.*, 2007a). However, as described earlier, the many sources of variability in indentation data imply that one cannot always simply proceed 'by eye'. Moreover, in the context of AFM, most experiments aiming to characterize cell stiffness, for example, employ multiple *repeated* indentations at different spatial locations. These requirements have motivated a more recent desire for effective, high throughput *automated* techniques, as detailed in Lin *et al.* (2007a).

Interpreted in a statistical context, the procedures that have thus far been adopted by practitioners fall under the general category of likelihood fitting. Rotsch *et al.* (1999) suggested simply to take two points in the post-contact data and to solve for E and γ by using the appropriate Hertzian model; however, the resultant estimate of Young's modulus depends strongly on the depth of indentation of the points selected (Costa, 2004). Costa *et al.* (2006) proposed to minimize the mean-squared error of a linear precontact and quadratic post-contact fit to the indentation data, though under the assumption of equal precontact and post-contact variances.

None of the existing approaches that have been adopted by practitioners, however, provides any means of quantifying uncertainty in change-point estimation—an important consideration in practice, since measurement errors can be large relative to the reaction force of the material probed—and consequently may result in poor point estimates (Crick and Yin, 2007). Moreover, such approaches fail to capture necessary physical constraints of the material indentation problem, such as the smoothness constraints that were described in Section 3.3. Such shortcomings provide strong motivation for the hierarchical model that was developed above, as well as the robust and automated fitting procedures that we describe next.

4. Posterior inference via Markov chain Monte Carlo sampling

The hierarchical Bayesian modelling framework that was introduced above features a large number of unknowns, with constraints on certain parameters precluding closed form expressions for the marginal posteriors of interest. These considerations suggest a simulation-based approach to inference; indeed, it is by now standard to use Markov chain Monte Carlo methods to draw samples from the posterior in such cases. Though widely available software packages for Gibbs sampling are adequate for inference in certain hierarchical Bayesian settings, the complexity of the conditional distributions that we obtain here (after imposing constraints and integrating out parameters whenever possible) necessitates explicit algorithmic derivations case by case. For this, we build on the approaches of Carlin *et al.* (1992) and Stephens (1994) and employ Metropolis-within-Gibbs techniques to draw samples from the posterior of distribution (6) as well as under the smoothness constraints of Section 3.3.

4.1. Metropolized Gibbs samplers and variance reduction
The selection of conjugate priors in our model allows nuisance parameters to be integrated out, to reduce the variance of the resultant estimators. Following standard manipulations, we marginalize over the precontact and post-contact regression coefficients β_1 and β_2 respectively. This yields the following marginal posterior probability distribution:

$$p(\gamma, \sigma_1^2, \sigma_2^2, b_0 | \mathbf{y}; \boldsymbol{\psi}) \propto \sigma_1^{-2(a_0-1)} \exp(-b_0/\sigma_1^2) \sigma_2^{-2(a_0-1)} \exp(-b_0/\sigma_2^2) b_0^{\kappa-1} \exp(-b_0/\eta)$$
$$\times (|\boldsymbol{\Sigma}_\gamma| \, |\boldsymbol{\Sigma}| \, |\mathbf{A}_\gamma|)^{-1/2} \exp\{-\tfrac{1}{2}(\mathbf{y}'\boldsymbol{\Sigma}_\gamma^{-1}\mathbf{y} + \boldsymbol{\mu}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} - \mathbf{b}_\gamma'\mathbf{A}_\gamma^{-1}\mathbf{b}_\gamma)\}, \quad (9)$$

where

$$\mathbf{A}_\gamma \triangleq \mathbf{X}_\gamma'\boldsymbol{\Sigma}_\gamma^{-1}\mathbf{X}_\gamma + \boldsymbol{\Sigma}^{-1}\boldsymbol{\Lambda} \in \mathbb{R}^{p\times p}$$

is block diagonal and

$$\mathbf{b}_\gamma \triangleq \boldsymbol{\Lambda}\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu} + \mathbf{X}_\gamma'\boldsymbol{\Sigma}_\gamma^{-1}\mathbf{y} \in \mathbb{R}^{p\times 1}.$$

The marginal posterior of expression (9) factors into a gamma density in b_0, and inverse gamma densities in σ_1^2 and σ_2^2 by way of the following partitions of $\mathbf{A}_\gamma$ and $\mathbf{b}_\gamma$:

$$
\begin{aligned}
\mathbf{A}_\gamma &= \begin{pmatrix} \mathbf{A}_{1,\gamma} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_{2,\gamma} \end{pmatrix}, & \begin{array}{l} \mathbf{A}_{1,\gamma} \in \mathbb{R}^{(d_1+1)\times\lfloor\gamma\rfloor}, \\ \mathbf{A}_{2,\gamma} \in \mathbb{R}^{(d_2+1)\times(n-\lfloor\gamma\rfloor)}; \end{array} \\[2mm]
\mathbf{b}_\gamma &= \begin{pmatrix} \mathbf{b}_{1,\gamma} \\ \mathbf{b}_{2,\gamma} \end{pmatrix}, & \begin{array}{l} \mathbf{b}_{1,\gamma} \in \mathbb{R}^{\lfloor\gamma\rfloor\times 1}, \\ \mathbf{b}_{2,\gamma} \in \mathbb{R}^{(n-\lfloor\gamma\rfloor)\times 1}. \end{array}
\end{aligned}
\quad (10)
$$

Expressions (9) and (10) lead to the following Gibbs sampler for change-point estimation (algorithm 1).

(a) Draw $\gamma \sim p(\gamma|\sigma_1^2, \sigma_2^2, b_0, \mathbf{y}; \boldsymbol{\psi})$ according to distribution (9).
(b) Draw $\sigma_1^2 \sim \mathcal{IG}\{a_0 + \tfrac{1}{2}\lfloor\gamma\rfloor, b_0 + \tfrac{1}{2}(\mathbf{y}'_{1,\gamma}\mathbf{y}_{1,\gamma} + \boldsymbol{\mu}'_1\boldsymbol{\mu}_1 - \mathbf{b}'_{1,\gamma}\mathbf{A}_{1,\gamma}^{-1}\mathbf{b}_{1,\gamma})\}$.
(c) Draw $\sigma_2^2 \sim \mathcal{IG}\{a_0 + \tfrac{1}{2}(n - \lfloor\gamma\rfloor), b_0 + \tfrac{1}{2}(\mathbf{y}'_{2,\gamma}\mathbf{y}_{2,\gamma} + \boldsymbol{\mu}'_2\boldsymbol{\mu}_2 - \mathbf{b}'_{2,\gamma}\mathbf{A}_{2,\gamma}^{-1}\mathbf{b}_{2,\gamma})\}$.
(d) Draw $b_0 \sim \mathcal{G}(\kappa, \eta^{-1} + \sigma_1^{-2} + \sigma_2^{-2})$.

To simulate from the conditional distribution of γ, we employ as a Metropolis-within-Gibbs step a mixture of a local random-walk move with an independent Metropolis step in which the proposal density is a pointwise evaluation of distribution (9) on the grid $1, 2, \ldots, n$ of indenter location indices. It is also possible to integrate out both noise variances (or the hyperparameter b_0). In this case, additional Metropolis steps are required, as the resulting conditional density of b_0 is non-standard. Our simulation studies confirm that these variants exhibit less Monte Carlo variation than a Gibbs sampler that is based on the full posterior of distribution (6).

4.2. *Posterior inference in the presence of smoothness constraints*

Bearing in mind the underlying physics of soft materials, a sufficiently high sampling rate of the force sensor relative to the speed of the indenter may yield data that are consistent with a smoothness assumption of a given order s. In this setting, our inferential procedure may be modified accordingly to take this into account. Stephens (1994) considered continuity-constrained switching linear regressions and used rejection sampling to draw from the conditional distribution of the change-point γ given the remaining model parameters. A more effective procedure, however, is to transform the data such that all except one of the regression coefficients may be integrated out; this variance reduction yields a Gibbs sampler that is analogous to algorithm 1 which we detail below. In fact, it is possible to derive such an algorithm for any value of $s \geqslant 0$, though for ease of presentation we first describe the case $s = 0$, whereupon only continuity is enforced.

This approach to variance reduction in the presence of smoothness constraints proceeds as follows: define $\tilde{\beta}_1 = \beta_1$ and, without loss of generality, let $\tilde{\beta}_2$ contain the last d_2 elements of

β_2, so that $\tilde{\beta} = (\tilde{\beta}_1', \tilde{\beta}_2')'$ contains all $p = d_1 + d_2 + 1$ independent coefficients. Via the continuity constraint of expression (7), define a *linear* transformation $\mathbf{T}_\gamma$ of $\tilde{\beta}$ to β as follows:

$$\beta = \mathbf{T}_\gamma \tilde{\beta}, \qquad \mathbf{T}_\gamma \triangleq \begin{pmatrix} \mathbf{I}_{(d_1+1)\times(d_1+1)} & \mathbf{0}_{(d_1+1)\times d_2} \\ \mathbf{c}_{1\times(d_1+1)} & \mathbf{0}_{1\times d_2} \\ \mathbf{0}_{d_2\times(d_1+1)} & \mathbf{I}_{d_2\times d_2} \end{pmatrix}, \tag{11}$$

where $\mathbf{I}_{m\times m}$ is the $m \times m$ identity matrix, $\mathbf{0}_{m\times n}$ is the $m \times n$ matrix of 0s and $\mathbf{c}_{1\times(d_1+1)} \overset{\triangle}{=} (1, x_\gamma, x_\gamma^2, \ldots, x_\gamma^{d_1})$. The choice of which of the $d_1 + 2$ regression coefficients to select as the dependent variable is made without loss of generality, since a transformation similar to expression (11) can be defined for every such choice, as well as for the continuity constraint of expression (8).

Since one of the regression coefficients is now a *deterministic* function of the others and the change-point, we place a prior directly on $\tilde{\beta}$ rather than on β. We assume that $\tilde{\beta}_1 \sim \mathcal{N}(\mu_1, \sigma_1^2 \Lambda_1^{-1})$ and $\tilde{\beta}_2 \sim \mathcal{N}(\tilde{\mu}_2, \sigma_2^2 \tilde{\Lambda}_2^{-1})$, with $\tilde{\Lambda}_2$ a $d_2 \times d_2$ diagonal positive definite matrix. Using the transformation $\mathbf{T}_\gamma$ of expression (11), we obtain by analogy with expression (9) the full posterior

$$\begin{aligned} p(\gamma, \tilde{\beta}, \sigma_1^2, \sigma_2^2, b_0|\mathbf{y}; \psi) \propto &\, \sigma_1^{-2(a_0-1)} \exp(-b_0/\sigma_1^2) \sigma_2^{-2(a_0-1)} \exp(-b_0/\sigma_2^2) b_0^{\kappa-1} \exp(-b_0/\eta) \\ &\times (|\Sigma_\gamma| |\tilde{\Sigma}|)^{-1/2} \exp[-\tfrac{1}{2}\{(\mathbf{y} - \tilde{\mathbf{X}}_\gamma \tilde{\beta})' \Sigma_\gamma^{-1} (\mathbf{y} - \tilde{\mathbf{X}}_\gamma \tilde{\beta}) \\ &+ (\tilde{\beta} - \tilde{\mu})' \tilde{\Sigma}^{-1} \tilde{\Lambda} (\tilde{\beta} - \tilde{\mu})\}], \end{aligned} \tag{12}$$

where $\tilde{\mathbf{X}}_\gamma = \mathbf{X}_\gamma \mathbf{T}_\gamma \in \mathbb{R}^{n\times p}$ and, by analogy with the quantities (Σ, μ, Λ), we define

$$\tilde{\Sigma} \overset{\triangle}{=} \operatorname{diag}(\sigma_1^2 \mathbf{1}_{d_1+1}, \sigma_2^2 \mathbf{1}_{d_2}) \in \mathbb{R}^{p\times p},$$

$\tilde{\mu} = (\mu_1', \tilde{\mu}_2')' \in \mathbb{R}^{p\times 1}$ and $\tilde{\Lambda} = \operatorname{diag}(\Lambda_1, \tilde{\Lambda}_2) \in \mathbb{R}^{p\times p}$.

The transformation $\mathbf{T}_\gamma$ makes it possible to integrate out the regression coefficients $\tilde{\beta}$ by using standard manipulations. Indeed, introducing the terms

$$\tilde{\mathbf{A}}_\gamma \overset{\triangle}{=} \tilde{\mathbf{X}}_\gamma' \Sigma_\gamma^{-1} \tilde{\mathbf{X}}_\gamma + \tilde{\Sigma}^{-1} \tilde{\Lambda} \in \mathbb{R}^{p\times p}$$

and

$$\tilde{\mathbf{b}}_\gamma \overset{\triangle}{=} \tilde{\Lambda} \tilde{\Sigma}^{-1} \tilde{\mu} + \tilde{\mathbf{X}}_\gamma' \Sigma_\gamma^{-1} \mathbf{y} \in \mathbb{R}^{p\times 1}$$

as before, we obtain the marginal posterior

$$\begin{aligned} p(\gamma, \sigma_1^2, \sigma_2^2, b_0|\mathbf{y}; \psi) \propto &\, \sigma_1^{-2(a_0-1)} \exp(-b_0/\sigma_1^2) \sigma_2^{-2(a_0-1)} \exp(-b_0/\sigma_2^2) b_0^{\kappa-1} \exp(-b_0/\eta) \\ &\times (|\Sigma_\gamma| |\tilde{\Sigma}| |\mathbf{A}_\gamma|)^{-1/2} \exp\{-\tfrac{1}{2}(\mathbf{y}' \Sigma_\gamma^{-1} \mathbf{y} + \tilde{\mu}' \tilde{\Sigma}^{-1} \tilde{\mu} - \tilde{\mathbf{b}}_\gamma' \tilde{\mathbf{A}}_\gamma^{-1} \tilde{\mathbf{b}}_\gamma)\}. \end{aligned} \tag{13}$$

It is straightforward to generalize this notion to any $s \in \{-1, 0, \ldots, d_1 + d_2\}$; a prior is put on $d_1 + d_2 - s + 1$ regression coefficients and a transformation $\mathbf{T}_\gamma$ analogous to expression (11) is defined.

As noted previously, the smoothness constraint of expression (7) introduces dependence between the pre- and post-change-point regression coefficients. In contrast with the marginal posterior distribution of expression (9) that was derived earlier for the unconstrained case, enforcement of constraint (7) precludes integrating out the associated noise variances σ_1^2 and σ_2^2. In the former case, the block diagonal structure of $\mathbf{X}_\gamma$ (and therefore of $\mathbf{A}_\gamma$) implies that the induced joint distribution on the variances is separable. However, in the latter case of expression (13), $\tilde{\mathbf{X}}_\gamma$ is *not* block diagonal—owing to the action of $\mathbf{T}_\gamma$—and hence nor is $\tilde{\mathbf{A}}_\gamma$. Therefore, the

variances σ_1^2 and σ_2^2 are no longer conditionally independent, and their joint distribution does not take the form of known generalizations of the univariate gamma distribution to the bivariate case (Yue *et al.*, 2001). Consequently, only the conditional distribution of the hyperparameter b_0 is in standard form, and simulation from posterior (13) proceeds with all other variables drawn by using Metropolis–Hastings steps, as shown in the following smoothness-constrained Gibbs sampler algorithm for change-point estimation (algorithm 2).

(a) Draw $\gamma \sim p(\gamma|\sigma_1^2, \sigma_2^2, b_0, \mathbf{y}; \psi)$ according to posterior (13) by using a Metropolis-within-Gibbs step.
(b) Draw $\sigma_1^2 \sim p(\sigma_1^2|\gamma, \sigma_2^2, b_0, \mathbf{y}; \psi)$ likewise.
(c) Draw $\sigma_2^2 \sim p(\sigma_2^2|\gamma, \sigma_1^2, b_0, \mathbf{y}; \psi)$ likewise.
(d) Draw $b_0 \sim p(b_0|\gamma, \sigma_1^2, \sigma_2^2, \mathbf{y}; \psi) = \mathcal{G}(\kappa, \eta^{-1} + \sigma_1^{-2} + \sigma_2^{-2})$.

In contrast with the case of algorithm 1, where a mixture kernel was employed purely for computational efficiency, we emphasize here that such a move is in fact *required* to sample from the full support $(1, n)$ of the change-point index; otherwise mixing of the underlying chain is poor. As before, the mixture kernel consisted of a local random-walk move and an independent global move drawing from a discrete distribution derived as a pointwise evaluation of posterior (13) on the integers $1, 2, \ldots, n$. The coupling of noise variances suggests a joint Metropolis–Hastings move, but this requires the specification of a proposal covariance; in simulations we found that separate normal random-walk moves for $\ln(\sigma_1^2)$ and $\ln(\sigma_2^2)$ were adequate.

We re-emphasize that, in our experience, integrating out the regression coefficients is essential to obtain a Gibbs sampler with favourable mixing properties. In particular, a sampler drawing for each parameter of expression (12) in the presence of smoothness constraints is severely restricted in its ability to explore the state space if one co-ordinate at a time updates are employed. When s continuous derivatives are enforced at the change-point, the likelihood function becomes more non-linear in change-point x_γ as s increases. Thus, as the induced constraint set becomes more non-linear, local moves on the scale of the regression coefficients themselves will lie far from it—leading to small acceptance probabilities. To overcome these problems, one may design a high dimensional Metropolis–Hastings move to update all the smoothness-constrained regression coefficients jointly with the change-point; however, a unique move must be designed for each s to be considered. In contrast, integrating out the regression coefficients obviates this need by handling such constraints for all s simultaneously.

5. Experimental validation

To validate algorithms 1 and 2 experimentally before their application in the setting of AFM, we designed and performed two sets of special macroscale indentation experiments in which precise contact point identification was made possible by the use of an impedance-measuring electrode mounted on the indenter. Although this direct measurement procedure is precluded in the vast majority of biomaterials experiments involving cells and tissues, as such samples are submerged in an aqueous solution, we could measure the true contact point for experiments involving cantilever bending and silicone indentation, as described below.

In addition, we conducted some simulation studies to characterize uncertainty in change-point estimation as a function of various model parameters. On the basis of these simulation studies and subsequent experimental validation, we found that both algorithms were insensitive to the exact choice of hyperparameters ψ, and hence we retained the following settings throughout: $\boldsymbol{\mu} = \mathbf{0}$, $\boldsymbol{\Lambda} = 10^{-5}\mathbf{I}_{p \times p}$, $a_0 = 2$, $\kappa = 1$ and $\eta = 10^{-2}$, with p being the total number of regression coefficients. We set the precontact polynomial degree $d_1 = 1$ throughout, on the basis

of the discussion in Section 2 supporting the assumption of a linear precontact regime. When smoothness constraints were used, we retained identical parameter settings and decremented the value of p appropriately. All posterior distributions were obtained by running the appropriate Gibbs samplers for 50 000 iterations and discarding the first 5000 samples. Convergence was assessed by using standard methods, confirming that increasing the number of Gibbs iterations did not appreciably change the resulting inference.

5.1. *Validation of change-point inference via cantilever bending*

We first performed several trials of an experiment whereby a steel cantilever was bent through application of a downward uniaxial force. Here, the measured force F is expected to change linearly with displacement $\delta = x - x_\gamma$ according to Hooke's law, with $\gamma \in (1, n)$ representing the contact point index. Representative data that are shown in Fig. 3 were obtained by using a TestBench$^{\text{TM}}$ Series system with a high fidelity linear actuator (Bose Corporation EnduraTEC Systems Group, Minnetonka, Minnesota, USA), which moved a cylindrical indenter into a cantilevered piece of FSS-05/8-12 spring steel measuring approximately 1.27 cm $\times$ 2.54 cm (carbon content 0.9–1.05%; Small Parts, Inc., Miramar, Florida, USA) at a speed of 10 mm s^{-1}. According to the impedance measurement technique that was described above, the contact point index γ was determined to correspond to position index 48 of the indenter, corresponding to a contact point $x_\gamma \in [-1.358, -1.262]$ mm. Because of the hardness of steel and the speed of the indenter, we did not make a smoothness assumption and thus used the Gibbs sampler of algorithm 1 to evaluate the efficacy of our approach, with $p = 4$ based on the linear post-contact regime implied by Hooke's law.

A full 100% of posterior values for γ after a 10% burn-in portion took the value 48, indicating correct detection of the change-point. Fig. 3 shows the results of the curve fitting procedure, with
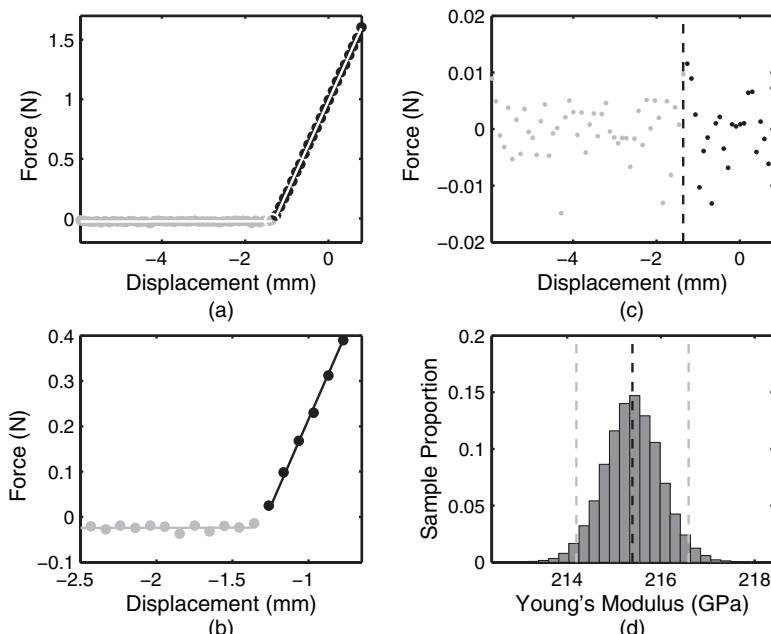
Fig. 3. Inference for the cantilever experiment of Section 5.1: (a) force–displacement data and posterior mean reconstructions obtained via algorithm 1 with (b) a close-up view near the change-point and (c) the associated residuals, and (d) marginal posterior of Young's modulus with its mean ( ) and 95% interval ( )

similar results obtained for varying indenter speeds. The minimum mean-square error (MMSE) estimate of Young's modulus was determined to be 215.3 GPa, with an associated 95% posterior interval of [214.0, 216.6]. By comparison, the range of values of Young's modulus for steel with similar carbon content is reported in the literature to be 210 ± 12.6 GPa (Kala and Kala, 2005). Despite its simple design, this experiment is not far removed from practice; a nearly identical procedure was employed by Wong *et al.* (1997) to probe the mechanical properties of silicon carbide nanorods, with each nanorod pinned on a substrate and subjected to a bending force along the unpinned axis.

5.2. Analysis of silicone indentation data

Whereas changes in slope at the contact point tend to be more pronounced in harder materials such as steel, the change in measured force is typically smoother in softer materials such as cells and tissues, making the contact point more difficult to detect. In earlier work (Yuen *et al.*, 2007) we detailed an indentation experiment using a soft silicone sample (Aquaflex; Parker Laboratories, Fairfield, New Jersey, USA), which was chosen both for its similarity to human tissues used in material indentation studies (Chen *et al.*, 1996), and because it enables direct contact point determination via an impedance-measuring electrode.

10 trials of this experiment were conducted, using a sample roughly 20 mm in depth, with a maximum indentation of approximately 8 mm; a typical force–displacement curve was shown earlier in Fig. 2. A hemispherical metal indenter of radius $R = 87.5$ mm compressed the sample at 10 mm s^{-1}, and the resulting forces were measured approximately every 10 μm to yield about 960 force–displacement data points. Both the unconstrained ($s = -1$) and continuity-constrained ($s = 0$) models were fitted in this setting, by way of algorithms 1 and 2 respectively. For a spherically tipped indenter, the Hertzian model (2) indicates that force is proportional to $(x - x_\gamma)^{3/2}$, and hence we employed a post-contact design matrix with only the fractional power $\frac{3}{2}$. In this regime, we have that $p = 4$ for algorithm 1 and $p = 3$ for algorithm 2. Since the initial distance from the indenter to the sample was approximately known, a uniform prior on $\gamma \in [125, 250]$ was assumed.

For each of the 10 data sets that were collected, the first $n = 450$ data points were taken to represent a conservative estimate of operation within the small deformation regime, and the Gibbs samplers of algorithms 1 and 2 were each run on these data. Results for both the cases are summarized in Table 1, along with the experimentally determined contact point, which varied from trial to trial owing to viscoelastic effects. Indeed, given that the spacing between data points is 0.01 mm on average, it may be deduced from Table 1 that the average error of 0.8–1% across trials corresponds to 8–10 sampled data points.

Marginal contact point posterior distributions for both the unconstrained and the continuity-constrained cases are summarized, pairwise by experiment, in the box plots of Fig. 4. These are seen to be notably more diffuse in the former case (left-hand side) than the latter (right-hand side); this is consistent with the softness of the silicone sample under study, which makes it difficult to reject *a priori* the possibility of continuity at the change-point. In the absence of this assumption, the experimentally determined contact point lies within the 95% posterior interval for each of the 10 trials. Once a continuity constraint has been imposed, the marginal posterior distributions of the change-points tighten noticeably; this is consistent with our expectation that constraining the model reduces the number of high likelihood fits.

A more subtle point can also be deduced from the slight yet consistent left-hand shift across 95% posterior intervals under the continuity-constrained regime relative to the unconstrained case. Observe the fourth row of Table 1, which reports the values that were obtained on extrapolating precontact and post-contact MMSE curve fits to their meeting point in this case. The

Table 1. Contact point estimates and associated errors for the 10 silicone trials of Section 5.2, based on the unconstrained and continuity-constrained models, with force–response data sampled every 0.01 mm on average†

| | *Results for the following trials:* | | | | | | | | | | *Average* |
	1	*2*	*3*	*4*	*5*	*6*	*7*	*8*	*9*	*10*	
Truth (x_γ)	5.52	5.49	5.47	5.44	5.42	5.38	5.39	5.48	5.42	5.38	—
MMSE (unconstrained)	5.42	5.44	5.42	5.43	5.47	5.39	5.46	5.35	5.41	5.39	—
MMSE (continuity constrained)	5.40	5.43	5.41	5.39	5.46	5.37	5.41	5.34	5.34	5.39	—
Extrapolated (unconstrained)	5.39	5.42	5.40	5.37	5.42	5.36	5.40	5.32	5.34	5.38	—
% error (unconstrained)	−1.78	−0.79	−0.83	−0.08	0.88	0.19	1.40	−2.23	−0.10	0.11	0.83
% error (continuity constrained)	−2.06	−1.02	−1.13	−0.95	0.31	−0.25	0.47	−2.44	−1.41	0.03	1.00
2.5 percentile	16.5	16.5	16.4	16.4	16.5	16.1	16.6	16.1	16.2	16.4	—
$\hat{E}$ (kPa)	17.2	17.3	17.2	17.0	17.2	16.8	17.4	16.8	17.0	16.9	17.1
97.5 percentile	17.9	18.4	17.9	17.6	18.0	17.6	18.1	17.7	17.8	17.9	—

†Young's modulus estimates $\hat{E}$ and 95% posterior intervals for the unconstrained case are also shown, along with averages over all 10 trials where appropriate.
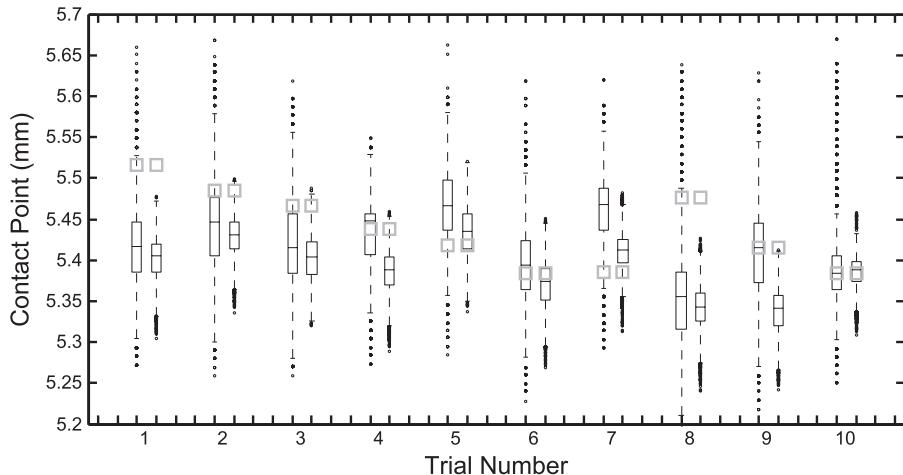
Fig. 4. Box plots of contact point marginal posterior distributions in each of 10 silicone indentation trials, shown side by side for the unconstrained (left) and continuity-constrained (right) models, with the centres of the grey squares indicating true contact point values: note the consistent decrease in posterior variance and slight downward shift of the posterior under the assumption of continuity

increase in force after contact implies that these values will always lie below the directly inferred contact point, as indeed they do. Nevertheless, enforcing continuity results in only a small increase in overall contact point estimation error and produces what practitioners might judge to be more physically feasible curve fits. Although independent verification of Young's modulus is not available for the silicone sample that was used in our study, the quality of our contact point estimates relative to a known ground truth leads to high confidence in the inferred values of Young's modulus.

Overall, the ability to obtain inferential results and accompanying uncertainty quantification, as exemplified by the cantilever bending and silicone indentation experiments, represents a significant improvement on current methods; a more detailed comparison with the method of Costa *et al.* (2006) is provided in our earlier work (Yuen *et al.*, 2007). In particular, the latter experiment demonstrates that reliable estimates of soft material properties can be obtained even in the presence of measurement error—a regime applicable to many AFM studies of cells and tissues, as we now describe.

6. Inference in the setting of atomic force microscopy

Having validated our algorithms on a macroscopic scale, we turn to analysing cellular biomaterials data collected by using AFM techniques. In contrast with our earlier examples, no direct experimental verification is available in this case, though we note that our resultant estimates of Young's modulus are considered plausible by experimentalists (Socrate and Suresh Laboratories, personal communications, 2008).

It is widely believed that cell biomechanics can shed light on various diseases of import, and hence a key research objective following the advent of AFM technology has been to analyse quantitatively the mechanical properties of various cell types. Indeed, numerous references over the past decade have linked stiffness and other related mechanical properties to cell malfunction and death (Costa, 2004); for example, the ability of cardiac myocytes in heart muscle tissue to contract is intimately linked to their cytoskeletal structure and its influence on cellular mechanical response (Lieber *et al.*, 2004). Here, we are similarly motivated to understand the stiffness properties of neuronal and red blood cells—currently a topic of intensive research in the biomechanics and bioengineering communities.

6.1. Indentation study of embryonic mouse cortical neurons

We first analysed *ex vivo* live mouse neurons, submerged in cell culture medium and repeatedly indented by an atomic force microscope (Asylum Research, Santa Barbara, California, USA) equipped with a spherically tipped probe. The indentation of each neuron (Socrate Laboratory, Massachusetts Institute of Technology) yielded approximately 700 force measurements, of which the first 500 were used in subsequent analysis to stay within the small deformation regime. Such data are of wide interest to neuroscientists and engineers, as traumatic damage to neurons is hypothesized to be related to their mechanical properties.

As a spherical probe tip was used for indenting each cell, we employed a post-contact design matrix with only the fractional power $\frac{3}{2}$, as in the case of our earlier silicone example, and the continuity-constrained sampler of algorithm 2. The results of a typical trial are shown in Fig. 5; the precontact and post-contact residuals were observed to be white. The primary inferential quantity of interest in such cases is Young's modulus, which is the principal characterization of cell material stiffness that was introduced in Section 2.2. According to the Hertzian model (2) for a spherical indenter, the regression coefficient corresponding to the $(x - x_\gamma)^{3/2}$-term of the fitted model is proportional to Young's modulus, with the constant of proportionality a function of the given radius $R = 10$ μm and Poisson's ratio $\nu = 0.5$. Thus, we can obtain the posterior distribution of Young's modulus by appropriately scaling the distribution of this regression coefficient, as shown in Fig. 5(c). We report the MMSE estimate of Young's modulus as $\hat{E} = 530.4$ Pa, and the corresponding 95% posterior interval as [518.8, 544.1] Pa. This estimate is in reasonable agreement with those previously reported for similar neurons (Lu *et al.*, 2006; Elkin *et al.*, 2007), with variability due primarily to differences in indentation speeds (Socrate Laboratory, personal communication, 2008).
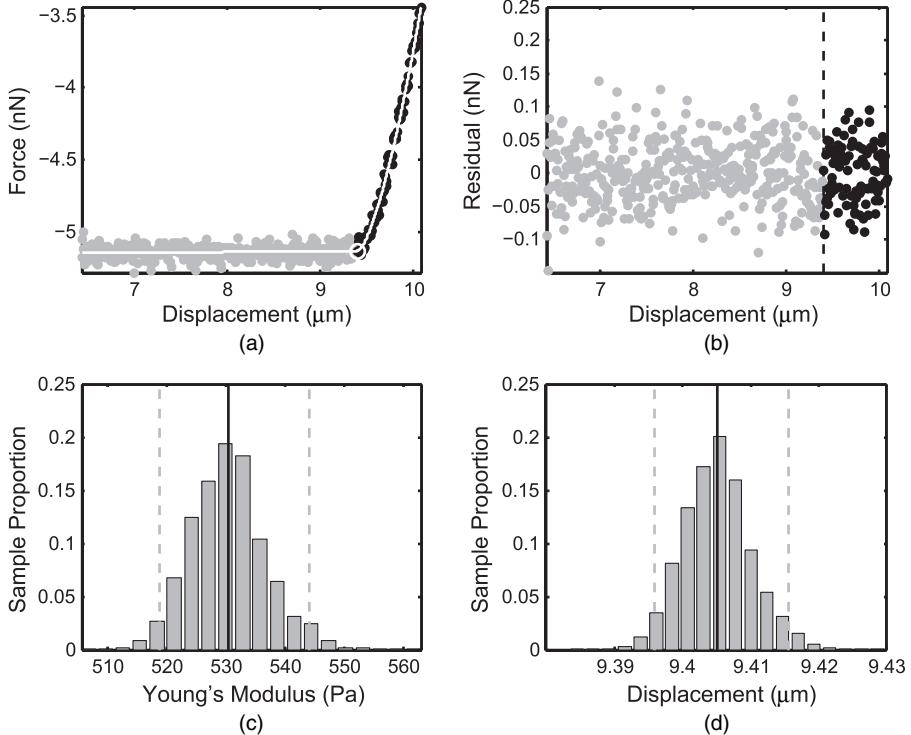
Fig. 5. Data collected during an AFM indentation of a mouse neuron together with (a) the posterior mean estimate of the underlying regressions and (b) the induced residuals: also shown are marginal posterior distributions of (c) Young's modulus and (d) the contact point, each overlaid with the posterior mean (|) and 95% posterior interval (⦙)

6.2. Indentation study of red blood cells

The mechanics of red blood cells have also been extensively studied by using AFM (Radmacher *et al.*, 1996). In this vein, we next analysed data from *ex vivo* live human erythrocytes (red blood cells, Suresh Laboratory, Massachusetts Institute of Technology) submerged in a cell culture medium, indented by an atomic force microscope (Asylum Research) equipped with a pyramidally tipped probe. The final 800 of approximately 8500 data points were discarded before the analysis, as they clearly lay outside the small deformation regime. Relative to the neuron AFM data that were considered in Section 6.1, the sampling rate of resistive force here is high, and consequently it is feasible to enforce continuity ($s = 0$) at the change-point. Algorithm 2 was therefore employed, with results from a typical trial shown in Fig. 6.

In the case at hand, the regression coefficient β_{12} corresponding to $(x - x_\gamma)^2$ is proportional to Young's modulus as per model (2), with the constant of proportionality depending on the indenter tip angle $2\phi = 70°$ and Poisson ratio $\nu = 0.5$, and, as we discuss below, no linear post-contact regression term was included. The inferred distribution of Young's modulus may be obtained by appropriately transforming the marginal posterior of β_{12}, as detailed in Section 3.1, and is shown in Fig. 6(c). The resultant MMSE estimate of $\hat{E} = 25.3$ Pa and corresponding posterior interval of [16.0, 34.9] Pa were confirmed to be consistent with various experimental assumptions (Suresh Laboratory, personal communication, 2008). Further, the precontact and post-contact error variances are determined to be unequal, as shown in Fig. 6(d).
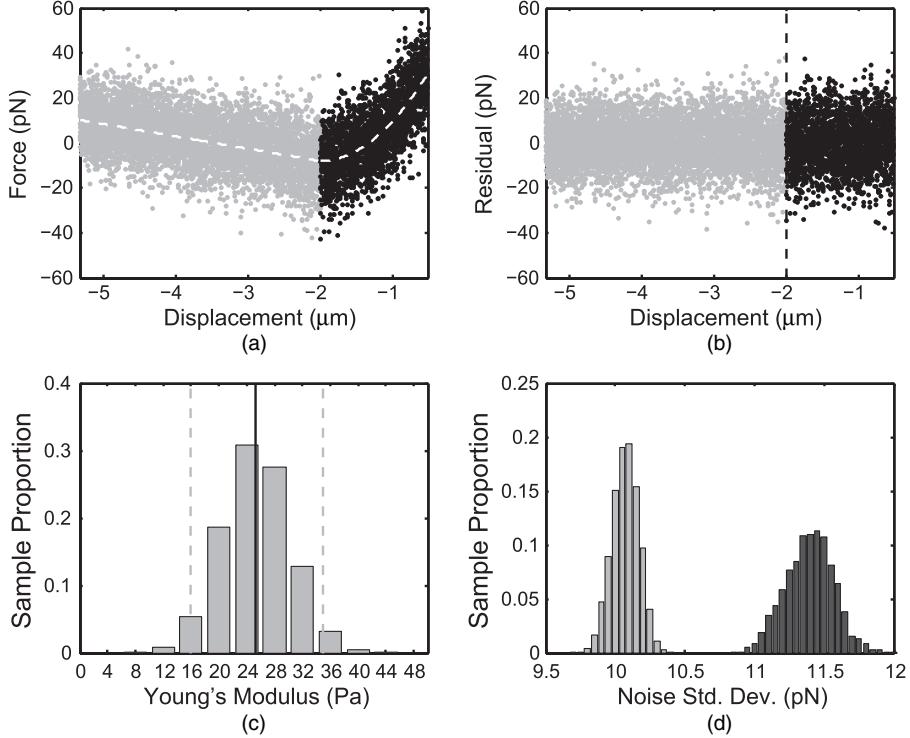
Fig. 6. Data collected during an AFM indentation of a human red blood cell together with (a) the posterior mean estimate of the underlying regressions and (b) the induced residuals: also shown are marginal posterior distributions of (c) Young's modulus overlaid with the posterior mean (│) and 95% posterior interval (┊) and (d) σ_1^2 (▨) and σ_2^2 (▰)

In the absence of a post-contact drift, enforcing continuous differentiability at the change-point ($s = 1$) constrains the precontact linear fit to have zero slope and is inconsistent with the precontact drift that is clearly visible in Fig. 6(a). However, if the post-contact polynomial were to include a linear drift term, then enforcing continuous differentiability would imply that the precontact and post-contact drifts are identical. Though this is appealing from a modelling viewpoint, as it eliminates an additional free parameter, practitioners lack evidence for such an equality. Moreover, in our experiments its inclusion had no appreciable effect on the inference of Young's modulus, and so we did not include a post-contact drift term in our final analysis of these data.

As in our earlier experiments, we compared our approach with the likelihood method of Costa *et al.* (2006), which yielded an estimate of x_γ that was shifted to the right by more than 1000 data points relative to that shown in Fig. 6. The corresponding estimate of Young's modulus in turn was found to be 34.8 Pa—an increase of 37.5% relative to the MMSE point estimate of 25.3 Pa, and close to the upper boundary of our estimated posterior interval. Although in this experimental setting one cannot conclude that either estimate is superior to the other, we note that the difference between them can be attributed in part to our model's incorporation of differing precontact and post-contact error variances, and smoothness constraints. Thus, we may view our inferential procedures as both a formalization and an extension of earlier likelihood-based approaches developed by practitioners, enabling both robust, automated fitting procedures and explicit uncertainty quantification.

7. Discussion

In this paper we have posed the first rigorous formulation of—and solution to—the key inferential problems arising in a wide variety of material indentation systems and studies. In particular, practitioners in the materials science community to date have lacked accurate, robust and automated tools for the estimation of mechanical properties of soft materials at either macroscales or microscales (Lin *et al.*, 2007a; Crick and Yin, 2007). A principal strength of our approach is its applicability to the analysis of biomaterials data that are obtained by indenting cells and tissues by using AFM; contact point determination is even more difficult in this setting, owing to the gradual change in measured resistive force that is a hallmark of soft materials. The Bayesian switching polynomial regression model and associated inferential procedures that we have proposed provide a means both to determine the point at which the indenting probe comes into contact with the sample and to estimate the corresponding material properties such as Young's modulus. In turn, its careful characterization holds open the eventual possibility of new biomechanical testing procedures for disease (Costa, 2004).

Our parametric approach is strongly motivated by—and exploits to full advantage—the Hertzian models governing the physical behaviour of linear elastic materials undergoing small deformations. The Bayesian paradigm not only enables quantification of uncertainty, which is crucial in applications, but also allows for the natural incorporation of physically motivated smoothness constraints at the change-point. Inference is realized through application of carefully designed Markov chain Monte Carlo methods together with classical variance reduction techniques. The resultant algorithms have been shown here to be both statistically and computationally efficient as well as robust to the choice of hyperparameters over a wide range of examples, and they are available on line for use by practitioners. Indeed, the direct applicability of our methods precludes any need for data preprocessing before analysis.

Outside the linear elastic materials that we considered here, it is of interest to apply the methodology to viscoelastic materials (e.g. biopolymers) which return to their precontact state slowly over time (Lin *et al.*, 2007b). In such cases, the amount of induced deformation depends not only on the indenter geometry but also on the rate of indentation. Another extension is to incorporate multiple spatially distributed change-points into our model, which is a key construct when atomic force microscopes are used to indent repeatedly a sample in order to characterize cell stiffness as a function of surface location (Geisse *et al.*, 2009). Finally, a sequential estimation scheme could be of great use in surgical robotics applications, where contact point determination plays a key role in enabling tactile sensing—a subject of current study by the authors.

Acknowledgements

References

Bacon, D. W. and Watts, D. G. (1971) Estimating the transition between two intersecting lines. *Biometrika*, **58**, 525–534.

Carlin, B. P., Gelfand, A. E. and Smith, A. F. M. (1992) Hierarchical Bayesian analysis of changepoint problems. *Appl. Statist.*, **41**, 389–405.

Chen, E. J., Novakofski, J., Jenkins, W. K. and Brien, J. W. D. (1996) Young's modulus measurements of soft tissues with application to elasticity imaging. *IEEE Trans. Ultrason. Ferrelect. Freq. Control*, **43**, 191–194.

Chernoff, H. and Zacks, S. (1964) Estimating the current mean of a Normal distribution which is subjected to changes in time. *Ann. Math. Statist.*, **35**, 999–1018.

Choy, J. H. C. and Broemling, L. D. (1980) Some Bayesian inferences for a changing linear model. *Technometrics*, **22**, 71–78.

Costa, K. D. (2004) Single-cell elastography: probing for disease with the atomic force microscope. *Dis. Mark.*, **19**, 139–154.

Costa, K. D., Sim, A. J. and Yin, F. C.-P. (2006) Non-Hertzian approach to analyzing mechanical properties of endothelial cells probed by atomic force microscopy. *ASME J. Biomech. Engng*, **128**, 176–184.

Crick, S. L. and Yin, F. C.-P. (2007) Assessing micromechanical properties of cells with atomic force microscopy: importance of the contact point. *Biomech. Modl Mechbiol.*, **6**, 199–210.

Dimitriadis, E. K., Horkay, F., Maresca, J., Kachar, B. and Chadwick, R. S. (2002) Determination of elastic moduli of thin layers of soft material using the atomic force microscope. *Biophys. J.*, **82**, 2798–2810.

Elkin, B. S., Azeloglu, E. U., Costa, K. D. and Morrison III, B. M. (2007) Mechanical heterogeneity of the rat hippocampus measured by atomic force microscope indentation. *J. Neurtraum.*, **24**, 812–822.

Fearnhead, P. (2006) Exact and efficient Bayesian inference for multiple changepoint problems. *Statist. Comput.*, **16**, 203–213.

Ferreira, P. E. (1975) A Bayesian analysis of a switching regression model: a known number of regimes. *J. Am. Statist. Ass.*, **70**, 370–374.

Geisse, N. A., Sheehy, S. P. and Parker, K. K. (2009) Control of myocyte remodeling in vitro with engineered substrates. *In Vitro Cell. Dev. Biol. Anim.*, **45**, 343–350.

Gouldstone, A., Chollacoop, N., Dao, M., Li, J., Minor, A. M. and Shen, Y.-L. (2007) Indentation across size scales and disciplines: recent developments in experimentation and modeling. *Acta Mater.*, **55**, 4015–4039.

Hinkley, D., Chapman, P. and Runger, G. (1980) Change point problems. *Technical Report*. University of Minnesota, Minneapolis.

Hudson, D. J. (1966) Fitting segmented curves whose join points have to be estimated. *J. Am. Statist. Ass.*, **61**, 1097–1129.

Kala, J. and Kala, Z. (2005) Influence of yield strength variability over cross-section to steel beam load-carrying capacity. *Nonlin. Anal. Modl Control*, **10**, 151–160.

Lai, T. L. (1995) Sequential changepoint detection in quality control and dynamical systems (with discussion). *J. R. Statist. Soc.* B, **57**, 613–658.

Lieber, S., Aubry, N., Pain, J., Diaz, G., Kim, S.-J. and Vatner, S. (2004) Aging increases stiffness of cardiac myocytes measured by atomic force microscopy nanoindentation. *Am. J. Physiol. Hrt Circuln Physiol.*, **287**, 645–651.

Lin, D. C., Dimitriadis, E. K. and Horkay, F. (2007a) Robust strategies for automated AFM force curve analysis—I: non-adhesive indentation of soft inhomogeneous materials. *J. Biomed. Engng*, **129**, 430–440.

Lin, D. C., Dimitriadis, E. K. and Horkay, F. (2007b) Robust strategies for automated AFM force curve analysis—II: adhesion-influenced indentation of soft, elastic materials. *J. Biomed. Engng*, **129**, 904–912.

Lu, Y.-B., Franze, K., Seifert, G., Steinhauser, C., Kirchoff, F., Wolburg, H., Guck, J., Janmey, P., Wei, E.-Q., Kas, J. and Reichenbach, A. (2006) Viscoelastic properties of individual glial cells and neurons in the CNS. *Proc. Natn. Acad. Sci. USA*, **103**, 17759–17764.

Menzefricke, U. (1981) A Bayesian analysis of a change in the precision of a sequence of independent normal random variables at an unknown time point. *Appl. Statist.*, **30**, 141–146.

Punskaya, E., Andrieu, C., Doucet, A. and Fitzgerald, W. J. (2002) Bayesian curve fitting using MCMC with applications to signal segmentation. *IEEE Trans. Signal Process.*, **50**, 747–758.

Quandt, R. (1958) The estimation of the parameters of a linear regression system obeying two separate regimes. *J. Am. Statist. Ass.*, **53**, 873–880.

Radmacher, M., Fritz, M., Kacher, C., Cleveland, J. and Hansma, P. (1996) Measuring the viscoelastic properties of human platelets with the atomic force microscope. *Biophys. J.*, **70**, 556–567.

Robison, D. E. (1964) Estimates for the points of intersection of two polynomial regressions. *J. Am. Statist. Ass.*, **59**, 214–224.

Rotsch, C., Jacobson, K. and Radmacher, M. (1999) Dimensional and mechanical dynamics of active and stable edges in motile fibroblasts investigated by using atomic force microscopy. *Proc. Natn. Acad. Sci. USA*, **96**, 921–926.

Smith, A. F. M. (1975) A Bayesian approach to inference about a change-point in a sequence of random variables. *Biometrika*, **62**, 407–416.

Smith, A. F. M. and Cook, D. G. (1980) Straight lines with a change-point: a Bayesian analysis of some renal transplant data. *Appl. Statist.*, **29**, 180–189.

Stephens, D. A. (1994) Bayesian retrospective multiple-changepoint identification. *Appl. Statist.*, **43**, 159–178.

Wolfe, D. A. and Schechtman, E. (1984) Nonparametric statistical procedures for the change point problem. *J. Statist. Planng Inf.*, **9**, 386–396.

Wong, E. W., Sheehan, P. E. and Liebert, C. M. (1997) Nanobeam mechanics: elasticity, strength, and toughness of nanorods and nanotubes. *Science*, **277**, 1971–1975.

Yue, S., Ouarda, T. B. M. J. and Bobeeb, B. (2001) A review of bivariate Gamma distributions for hydrological application. *J. Hydrol.*, **246**, 1–18.

Yuen, S. G., Rudoy, D. G., Howe, R. D. and Wolfe, P. J. (2007) Bayesian changepoint detection for switching regressions in material indentation experiments. In *Proc. 14th Wrkshp Statistics and Signal Processing*, pp. 104–108. New York: Institute of Electrical and Electronics Engineers.

Zacks, S. (1983) Survey of classical and Bayesian approaches to the change point problem: fixed sample and sequential procedures of testing and estimation. In *Recent Advances in Statistics: Herman Chernoff Festschrift*, pp. 245–269. New York: Academic Press.

Zellner, A. (1986) On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In *Bayesian Inference and Decision Techniques*, pp. 233–243. New York: Elsevier.