Bayesian inference of substrate properties from film behavior

Bayesian inference of substrate properties from film behavior

R. Aggarwal[1], M.J. Demkowicz[2] and Y.M. Marzouk[3]

[1]Dept. of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[2]Dept. of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[3]Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

E-mail: raghav_a@mit.edu

Abstract.

We demonstrate that, by observing the behavior of a film deposited on a substrate, certain features of the substrate may be inferred with quantified uncertainty using Bayesian methods. We carry out this demonstration on an illustrative film/substrate model, where the substrate is a Gaussian random field and the film is a two-component mixture that obeys the Cahn-Hilliard equation. We construct a stochastic reduced order model to describe the film/substrate interaction and use it to infer substrate properties from film behavior. This quantitative inference strategy may be adapted to other film/substrate systems.

Keywords: bayesian inference, substrate, film, uncertainty quantification, reduced order models

1. Introduction

Materials scientists are often interested in measuring spatially varying properties of surfaces. When these properties cannot be observed directly, they may sometimes nevertheless be inferred from the behavior of a film deposited on the surface. The surface of interest then becomes a substrate that influences film behavior while the behavior of the film is an indicator of substrate properties. For example, Lewandowski and Greer inferred high temperature shear bands in metallic glasses through the localized melting of a film of tin deposited on the sample [1]. Aizenberg *et al.* inferred the distribution of local disorder in self-assembled alkanethiolate monolayers by observing water vapor condensation figures [2]. Bowden *et al.* found that buckling patterns of metal films on elastomer substrates indicate underlying relief structures of the substrate [3].

In all the examples above, a property of a substrate (e.g., temperature, local disorder, topography) may be inferred from film behavior (e.g., melting, condensation, buckling). We investigate how such intuitive inferences may be made quantitative

using Bayesian statistical methods [4; 5]. These methods enable the inference of model parameters from indirect and incomplete data. We will demonstrate two key benefits of Bayesian inference, namely quantification of uncertainty in the inferred parameters and the ability to improve these estimates sequentially as additional data become available.

We carry out our analysis on a model film/substrate system. The substrate property is modeled as a Gaussian random field with a characteristic length scale. The film is a two-component mixture described by an order parameter that obeys the Cahn-Hilliard equation. In this model system, the length scale of the substrate field is inferred from the phase separation behavior of the film. To assess the quality of the inference, we show how the Bayesian posterior concentrates as more data are employed, and we compare the posterior with the known correct substrate length scale. Although our model system does not correspond to a specific experiment, it nevertheless demonstrates certain general features of how Bayesian methods may be used to deduce substrate properties from film behavior. It also sheds light on the usefulness of developing stochastic reduced order models (ROMs) for the purpose of inference.

In Section 2, we introduce a numerical model of the substrate and the film. In Section 3, we formulate a simple ROM for the stochastic film-substrate interaction, based on simulations of the numerical model and on the Buckingham Pi theorem. In Section 4, we demonstrate Bayesian inference of the substrate length scale from film behavior using the ROM developed in Section 3. Section 5 discusses features of the inference strategy described here that may be adapted to other film/substrate systems.

2. Model problem

In our model problem, the substrate is described by a spatially varying scalar field $\mathcal{T}(x, y)$. The film is a two-component mixture that phase separates or remains uniform, depending on the value of $\mathcal{T}(x, y)$. It is described by order parameter $c(x, y, t)$, which is a function of both position and time. We assume that the substrate cannot be directly observed. Instead, some of its aggregate properties may be inferred from the behavior of the film, which may be directly observed.

2.1. Modeling the substrate

We model the substrate as a Gaussian random field [6] $\mathcal{T}(x, y)$ on a square domain Ω with side L_D, with zero mean and covariance function specified as follows:

$$\begin{aligned} \mathbb{C}\mathrm{ov}\left(\mathcal{T}(x_1, y_1), \mathcal{T}(x_2, y_2)\right) &= K_{1\mathrm{D}}\left(x_1, x_2\right) K_{1\mathrm{D}}\left(y_1, y_2\right) \\ (x_i, y_i) &\in \Omega = [0, L_D] \times [0, L_D] \end{aligned} \tag{1}$$

where

$$K_{1\mathrm{D}}\left(s_1, s_2\right) = \beta^2 \exp\left(-\frac{2\sin^2\left(|s_1 - s_2|\pi/L_D\right)}{\left(\ell_s/L_D\right)^2}\right). \tag{2}$$

This covariance kernel ensures that realizations of $\mathcal{T}(x, y)$ are smooth (infinitely differentiable) and periodic on the square with dimensionless edge length L_D. Note

that $\mathcal{T}(x, y)$ is a stochastic quantity; every realization of the Gaussian random field is a distinct substrate field $\mathcal{T}(x, y)$. At any fixed spatial location (x^*, y^*), $T(x^*, y^*)$ is a Gaussian random variable with mean zero and variance β^2. The covariance kernel above also encodes smoothness or structure by imposing nonzero covariance between $\mathcal{T}(x_1, y_1)$ and $\mathcal{T}(x_2, y_2)$. Thus it is different from uncorrelated Gaussian noise.

As the distance between locations increases, the covariance in (2) decays at a rate controlled by parameter ℓ_s. This distance-dependent covariance creates correlated variations in the substrate field with characteristic dimensionless length scale ℓ_s. Large values of ℓ_s result in long-wavelength variations in $\mathcal{T}(x, y)$ and, conversely, small values of ℓ_s give rise to short-wavelength $\mathcal{T}(x, y)$ variations.

Due to the periodic nature of the covariance kernel, peaks in the covariance kernel recur with a period of L_D. Therefore, the above described decay in covariance is to be understood to occur locally. To ensure that periodic images of the covariance peaks do not interact, ℓ_s must be much smaller than L_D. In this study, ℓ_s is never greater than $L_D/5$. We find that this condition is sufficient to prevent periodic image artifacts.

Figures 1(e) and 1(i) show two realizations of $\mathcal{T}(x, y)$ with ℓ_s values of 0.77 and 0.13, respectively, and $L_D = 5$. As expected, variations in $\mathcal{T}(x, y)$ occur over larger length scales in 1(e) than in 1(i). For contrast, Fig. 1(a) shows a uniform substrate field $\mathcal{T}(x, y) = -1$.

In our inference problem, the substrate field property of interest is ℓ_s. That is, we will not infer the particular realization of the substrate field giving rise to a single instance of film behavior. Rather, we are interested in learning the length scale parameter in the stochastic description of the substrate field. Since the substrate field cannot be observed directly, ℓ_s is inferred from the behavior of a film evolving on the substrate, described next.

2.2. Modeling the film behavior

The film deposited on top of the substrate is a two-component mixture represented by an order parameter field $c(x, y, t)$. The order parameter takes values in the range $[-1, 1]$, where $c = -1$ and $c = 1$ represent single-component phases and $c = 0$ represents a uniformly mixed phase. The behavior of the film is encoded in the substrate-dependent potential energy density

$$g(c, \mathcal{T}(x, y)) = \frac{c^4}{4} + \mathcal{T}(x, y)\frac{c^2}{2}. \tag{3}$$

For negative values of $\mathcal{T}(x, y)$, $g(c, \mathcal{T}(x, y))$ is a double-well potential and the film tends to separate into two distinct phases with $c = \pm\sqrt{-\mathcal{T}(x, y)}$. By contrast, when $\mathcal{T}(x, y)$ is positive, $g(c, \mathcal{T}(x, y))$ has a single well at $c = 0$ and the film remains uniform. The substrate property $\mathcal{T}$ may therefore be thought of as a temperature difference from some hypothetical critical temperature. It promotes mixing in regions where $\mathcal{T}(x, y) > 0$ and phase separation in regions where $\mathcal{T}(x, y) < 0$.

The total energy of the film may be expressed as

$$E = \int_{\Omega} \left(g(c, \mathcal{T}(x,y)) + \frac{\epsilon^2}{2} |\nabla c|^2 \right) dxdy, \tag{4}$$

where Ω is the spatial domain. In addition to the potential energy density $g(c, \mathcal{T}(x,y))$, the energy functional E contains a gradient energy penalty $\epsilon^2|\nabla c|^2/2$. The larger the value of ϵ, the larger the penalty for sharp gradients in $c(x,y,t)$. Thus, larger ϵ values lead to smoother $c(x,y,t)$ solutions. ϵ may be thought of as a characteristic length scale for composition fluctuations in the film. It is a material property of the film.

The temporal evolution of $c(x,y,t)$ is modeled by the Cahn-Hilliard equation [7]

$$\frac{\partial c}{\partial t} = \Delta \left(\frac{\partial g}{\partial c} - \epsilon^2 \Delta c \right), \tag{5}$$

which minimizes the energy functional in Equation (4). We solve this equation on a square domain Ω with edge length L_D and periodic boundary conditions using an unconditionally stable, linear splitting scheme [8]. Spectral methods are used for spatial discretization [9] and exponential time differencing is used for time integration [10]. In all solutions, the initial condition of the Cahn-Hilliard equation is taken to be a realization of a spatially uncorrelated Gaussian random field with variance 0.01. This initial condition represents a film that has been slightly perturbed from a purely uniform state. It is another independent source of stochasticity in the problem, in addition to the substrate field: different realizations of the initial condition give rise to different temporal evolution trajectories of the film.

The substrate variance parameter β^2 controls the amplitude of $\mathcal{T}(x,y)$, but does not affect the length scale over which $\mathcal{T}(x,y)$ varies. The choice of β^2 does not affect our results as long as there is sufficient contrast between the phase-separating and phase-mixing regions of the substrate. This requirement is met if $\epsilon^2 < \beta^2$. We choose to fix β^2 to a value of 0.09, resulting in $\mathcal{T}(x,y)$ lying largely within the range of $[-1,1]$. All ϵ^2 values we consider are smaller than 0.01, so there is always good contrast between the phase-separating and mixing regions in our simulations.

2.3. Understanding the film behavior

To illustrate the interaction of the film with the substrate, Figure 1 shows the evolution of the film for three different cases: a film with $\epsilon = 0.05$ on a uniform substrate with $\ell_s = \infty$, a film with $\epsilon = 0.02$ on a coarsely patterned substrate with $\ell_s = 0.77$, and a film with $\epsilon = 0.025$ on a finely patterned substrate with $\ell_s = 0.13$.

To describe the behavior of the film quantitatively, we calculate the characteristic length scale of composition fluctutations using the normalized autocorrelation function $\mathcal{C}(x,y,t)$ of the order-parameter field $c(x,y,t)$, calculated as

$$\mathcal{C}\,(x,y,t) = \frac{\int_{\Omega} c(x'+x, y'+y, t) c(x', y', t) dx'dy'}{\int_{\Omega} c(x', y', t)^2 dx'dy'}. \tag{6}$$

The length scale $\Lambda(t)$ is calculated as the radius of the central peak in the autocorrelation function $\mathcal{C}(x, y, t)$. The radius is calculated by least squares fitting of a covariance kernel

$$K(x, y, \lambda) = \exp\left(-\frac{2\left(\sin^2\left(x\pi/L_D\right) + \sin^2\left(y\pi/L_D\right)\right)}{\left(\lambda/L_D\right)^2}\right), \tag{7}$$

with an adjustable length scale parameter λ, to the autocorrelation function $\mathcal{C}(x, y, t)$:

$$\Lambda(t) = \arg\min_{\lambda\in\mathbb{R}^+}\left(\int_\Omega \left(\mathcal{C}(x, y, t) - K(x, y, \lambda)\right)^2 dxdy\right). \tag{8}$$

The uniform substrate with $\mathcal{T}(x, y) = -1$ shown in Figure 1(a) favors phase separation at all locations. Upon solving the Cahn-Hilliard equation (5) with $\epsilon = 0.05$, we observe classic spinodal decomposition [11], snapshots of which are shown in Figure 1(b)-(d). The two phases separate, coarsen, and eventually converge to a stationary, boundary length-minimizing shape. The phase separated structure of the film converges because the computational domain is finite. In principle, if the film were simulated over a hypothetical infinite computational domain ($L_D \to \infty$), the film would coarsen indefinitely, albeit at an ever-decreasing rate. In our finite-domain simulations, Figure 2(a) shows that $\Lambda(t)$ increases with time t and converges to a value Λ_∞ close to $L_D = 5$. Λ_∞ is calculated as the length scale of the film in its converged state:

$$\Lambda_\infty = \lim_{t\to\infty} \Lambda(t). \tag{9}$$

In practice, the film behavior is considered converged if the maximum change in the film between two subsequent time steps is below a threshold δ:

$$\max_{(x,y)\in\Omega} \left|c(x, y, t^{n+1}) - c(x, y, t^n)\right| < \delta. \tag{10}$$

For this study, $\delta = 10^{-3}$.

The substrate shown in Figure 1(e) is coarsely patterned ($\ell_s = 0.77$). We simulate the behavior of a film with $\epsilon = 0.02$ on this substrate and find that it phase separates in some regions, but not in others, as shown in Figure 1(f)-(h). Phase separation occurs where $\mathcal{T}(x, y) < 0$. The order parameter field in Figure 1(h) has a structure finer than the scale of the computational domain. Figure 2(a) plots $\Lambda(t)$ for this film/substrate system and shows that the converged length scale Λ_∞ is smaller than the computational domain size $L_D = 5$. Thus, unlike the case of the uniform substrate, the convergence of the film over this patterned substrate is due to the substrate itself, not the finite size of the computational domain.

To further illustrate this distinction, we plot Λ_∞ as a function of domain size L_D for both the uniform substrate and the patterned substrate in Figure 2(b). As expected, Λ_∞ increases with L_D for the uniform substrate while for the patterned substrate, Λ_∞ is independent of L_D. This confirms that the length scale of compositional fluctuations in the film on a patterned substrate converges because of the patterned substrate and not due to the finite computational domain size.

The third substrate, shown in Figure 1(i), is finely patterned ($\ell_s = 0.13$). A film with $\epsilon = 0.025$ is deposited on top of it. Figures 1(j)-(l) show the phase separation of the film. We do not find any correlation between the phase-separated regions of the film and the parts of the substrate where $\mathcal{T}(x, y) < 0$. This behavior contrasts with that of a film with $\epsilon = 0.02$ deposited on a substrate with $\ell_s = 0.77$. This comparison suggests that the behavior of the film/substrate system depends on both ϵ and ℓ_s. The nature of this dependence is explored in the next section.

3. Model simplification

The previous section describes a model film/substrate system where the substrate has characteristic length scale ℓ_s and the film has two characteristic length scales: the gradient energy penalty factor ϵ and the converged length scale Λ_∞. We propose to infer ℓ_s given ϵ and Λ_∞. Note that the relationship between these parameters, as described by the phase field model given in the previous section, is intrinsically stochastic due to randomness both in the initial conditions and in the substrate field itself (for a given ℓ_s). Direct evaluation of the resulting likelihood function in a Bayesian approach (to be described in Section 4) is thus intractable. Approximate Bayesian computation methods [12; 13] could be used instead, but these would require many thousands or millions of model evaluations.

Instead, we adopt a different approach to infer ℓ_s given ϵ and Λ_∞: we construct a simplified model that relates ℓ_s, ϵ, and Λ_∞, in a way that will give rise to a tractable likelihood function to be used in Section 4. This relation may be viewed as a "reduced order model" (ROM), because it ignores the detailed structure of the substrate, replacing it with ℓ_s, and because it also ignores the exact temporal evolution of the film, instead describing it using only ϵ and Λ_∞. This relation may be expressed as

$$\Lambda_\infty = \overbrace{\underbrace{f\left(\epsilon, \ell_s\right)}_{\text{deterministic term}} + \underbrace{\gamma\left(\epsilon, \ell_s\right)}_{\text{random term}}}^{\text{reduced order model}}. \tag{11}$$

The deterministic term captures the average response of the film/substrate system. The random term $\gamma(\epsilon, \ell_s)$ can account for both the inherent stochasticity of the film evolution and any systematic error in the deterministic term. We will construct both $f(\epsilon, \ell_s)$ and $\gamma(\epsilon, \ell_s)$ in turn.

First, we use the Buckingham Pi theorem [14] to simplify the putative relation $\Lambda_\infty = f(\epsilon, \ell_s)$ between ℓ_s, ϵ, and Λ_∞. Because ϵ, ℓ_s, and Λ_∞ all denote characteristic lengths, we may use them to write two independent parameter groups Λ_∞/ℓ_s and ℓ_s/ϵ. The Buckingham Pi theorem then allows us to simplify the relation to

$$\frac{\Lambda_\infty}{\ell_s} = F\left(\frac{\ell_s}{\epsilon}\right). \tag{12}$$

To get the actual form of function F, we perform runs of the full model over a range of ϵ and ℓ_s. ϵ was varied over ten logarithmically-spaced values between [0.01,

0.1]. Similarly, ℓ_s was systematically varied in the range [0.1, 1]. For each combination of ϵ and ℓ_s, ten simulations of the full model were performed, each with an independent realization of the substrate parameter field and the stochastic initial condition. The resulting values of Λ_∞ were recorded. The spatial discretization length h was always at least a factor of 2 smaller than ℓ_s, ϵ, or Λ_∞ and the computational domain size L_D was always at least a factor of 5 larger than ℓ_s, ϵ, or Λ_∞.

Figure 3(a) plots Λ_∞/ℓ_s as a function of ℓ_s/ϵ. The individual data points collapse onto a single curve, as expected. The curve has two asymptotes. When $\ell_s/\epsilon \gg 1$, $\Lambda_\infty/\ell_s \approx 1$. This limit is illustrated in Figure 1(f)-(h). In it, the film order parameter $c(x, y, t)$, upon convergence, exhibits variations over a distance of approximately ℓ_s. This is because the phase-separating regions of the substrate $((x, y) \mid \mathcal{T}(x, y) < 0)$ are on average separated by a distance ℓ_s. Since ℓ_s is much larger than ϵ, neighboring phase separating regions do not interact. The film phase separates to the full extent of the phase separating regions, but not beyond. Since the phase separating regions are on average of size ℓ_s, Λ_∞ is close to ℓ_s.

The second asymptote occurs at $\ell_s/\epsilon \approx 1$, where Λ_∞/ℓ_s appears to diverge. This is the behavior observed in Figure 1(j)-(l), where the film phase separates to a length scale larger than that of the substrate. Variations in the converged film occur over a distances larger than ℓ_s. This is because the characteristic spacing between phase separating regions ℓ_s, is close to ϵ. As a result, the film is no longer constrained to follow individual phase separating regions, and phase separates beyond length scale ℓ_s. This causes Λ_∞ to be larger than ℓ_s.

Given these observations of the simulated film behavior, we prescribe a functional form for F that is consistent with both asymptotes:

$$\frac{\Lambda_\infty}{\ell_s} = F\left(\frac{\ell_s}{\epsilon}\right) = a + \frac{b}{(\ell_s/\epsilon - 1)^c}. \tag{13}$$

We obtain point estimates of a, b, and c by least squares fitting to all model runs where $\Lambda_\infty < 1$. This is done to ensure that $\Lambda_\infty < L_D/5$ at all times. The estimates are:

$$a = 1.046, \quad b = 79.51, \quad c = 1.537. \tag{14}$$

Figure 3(a) shows the scatter of data around this parametric fit, which reflects the stochasticity of the full model. To simulate the likelihood function in Bayesian inference, we must capture this stochasticity. We do so by adding a zero-mean Gaussian random variable $\Gamma(\ell_s/\epsilon)$ to $F(\ell_s/\epsilon)$:

$$\frac{\Lambda_\infty}{\ell_s} = F\left(\frac{\ell_s}{\epsilon}\right) + \Gamma\left(\frac{\ell_s}{\epsilon}\right), \quad \Gamma\left(\frac{\ell_s}{\epsilon}\right) \sim N\left(0, \sigma^2\left(\frac{\ell_s}{\epsilon}\right)\right). \tag{15}$$

Any systematic offset between the deterministic model F and the mean of the scattered data in Figure 3(a) should be incorporated in the mean of the random variable Γ. However, this offset is negligible in the current results, and thus we choose Γ to have mean zero. The variance of the random term is chosen to be a function of ℓ_s/ϵ, due

to the inherent non-stationarity of the scatter in Figure 3(a). This non-stationarity is an intrinsic feature of the film evolution and its response to the stochastic inputs. We fit this dependence to the observed scatter using Gaussian process regression [6]. This approach captures the variance $\sigma^2(\ell_s/\epsilon)$ nonparametrically.

The variance of the residual $(\Lambda_\infty/\ell_s - F(\ell_s/\epsilon))$ at each (ℓ_s/ϵ) value is estimated directly from the corresponding simulation data. We then regress the logarithm of the variance on the logarithm of ℓ_s/ϵ. GP regression is applied to the logarithm of the variance, rather than the variance itself, to ensure that the latter remains positive. Our GP regression scheme uses a squared exponential covariance kernel; hyperparameters of this kernel are tuned by maximizing their marginal likelihood, as described in [6]. The mean function so obtained is used as a point estimate of σ^2. The variance function obtained from the GP is ignored. In Figure 3(b), we have plotted the variance $\sigma^2(\ell_s/\epsilon)$ as obtained from the procedure described above. The non-stationarity of the variance can be easily observed.

4. Bayesian inference

We use Bayesian methods and the simplified model described in Section 3 for the inference problem of interest here. In Bayesian inference, the parameter to be inferred, θ, is treated as a random variable. Thus, instead of producing a single estimate for the parameter, Bayesian inference yields the probability density function (PDF) of the parameter conditioned on the available data, $\mathscr{D}_1 : P(\theta|\mathscr{D}_1, \eta_1)$. In general, the parameter θ may also be conditioned on some known experimental parameter, η_1. This *posterior PDF* contains complete information about the uncertainty of the inference, as opposed to only a point estimate (or a point estimate and a Gaussian approximation). Particular measures of uncertainty, such as variances, credible intervals, or tail probabilities, may be extracted from the posterior PDF.

The posterior PDF may be calculated using Bayes formula:

$$\overbrace{P(\theta|\mathscr{D}_1, \eta_1)}^{\text{posterior}} = \frac{\overbrace{P(\mathscr{D}_1|\theta, \eta_1)}^{\text{likelihood}}\overbrace{P(\theta)}^{\text{prior}}}{\underbrace{\int_\theta P(\mathscr{D}_1|\theta, \eta_1)P(\theta)d\theta}_{\text{evidence}}}. \tag{16}$$

$P(\theta)$ is the prior PDF and represents knowledge about possible values of θ before any data is collected. The prior may reflect expert opinion, previous experiments, or physical constraints. In case of no prior information, one may select a maximum entropy prior [15; 16] or another kind of non-informative prior [17; 18]. The likelihood is the probability density of the data, conditioned on the parameter of interest. It encodes the relationship between the parameters and the observations. In our case, the likelihood is the simplified model from Section 3. The evidence is a normalizing term needed to ensure that the left hand side of (16) is a valid PDF.

An important feature of Bayesian inference is that it allows for sequential improvements the posterior as additional data become available. The previous posterior is simply assigned to be the new prior and an improved posterior, incorporating the new data, is obtained using

$$P(\theta|\mathscr{D}_{1:n}, \eta_{1:n}) \propto P(\mathscr{D}_n|\theta, \eta_n) P(\theta|\mathscr{D}_{1:n-1}, \eta_{1:n-1}). \tag{17}$$

Here $\mathscr{D}_{1:n}$ refers to the first n data points collected from experiments with known parameters $\eta_{1:n}$

4.1. Inferring ℓ_s

We consider a series of experiments where ϵ is a known experimental parameter, Λ_∞ is directly observed in the experiment, and ℓ_s is to be inferred. The quality of the inference is to be improved as more experiments are performed, i.e., as additional $(\Lambda_\infty, \epsilon)$ pairs become available. To model experiments that generate $(\Lambda_\infty, \epsilon)$ pairs, values of ϵ are picked at random from a uniform distribution in the log-space with range [0.01,0.1]. For each value of ϵ, a substrate is randomly generated with $\ell_s = 0.4$ and used to run the Cahn-Hilliard model described in Section 2. Λ_∞ are obtained from the converged order-parameter field of the film.

To infer ℓ_s, both sides of equation (15) are multiplied by ℓ_s/ϵ, giving

$$\frac{\Lambda_\infty}{\epsilon} = F'\left(\frac{\ell_s}{\epsilon}\right) + \Gamma'\left(\frac{\ell_s}{\epsilon}\right), \quad \Gamma'\left(\frac{\ell_s}{\epsilon}\right) \sim N\left(0, \sigma'^2\left(\frac{\ell_s}{\epsilon}\right)\right), \tag{18}$$

$$F'\left(\ell_s/\epsilon\right) = \left(\ell_s/\epsilon\right) F\left(\ell_s/\epsilon\right), \tag{19}$$

$$\sigma'\left(\ell_s/\epsilon\right) = \left(\ell_s/\epsilon\right) \sigma\left(\ell_s/\epsilon\right). \tag{20}$$

The likelihood of any candidate ℓ_s is obtained from the difference between the experimentally-obtained value of Λ_∞/ϵ and the value of Λ_∞/ϵ obtained from equation (18), given by

$$\overbrace{\frac{\Lambda_\infty}{\epsilon}}^{\text{Observed}} - \underbrace{\frac{\ell_s}{\epsilon} F\left(\frac{\ell_s}{\epsilon}\right)}_{\text{Predicted}} \sim N\left(0, \left(\frac{\ell_s}{\epsilon}\right)^2 \sigma^2\left(\frac{\ell_s}{\epsilon}\right)\right). \tag{21}$$

We may use equation (21) to construct the likelihood function, which incorporates both parts of the ROM: the parametric deterministic term and the non-parameteric stochastic term,

$$P\left(\Lambda_\infty|\ell_s, \epsilon\right) = \frac{1}{\sqrt{2\pi\sigma'(\ell_s/\epsilon)}} \exp\left(\frac{-\left(\Lambda_\infty/\epsilon - F'(\ell_s/\epsilon)\right)^2}{2\sigma'^2(\ell_s/\epsilon)}\right). \tag{22}$$

Figure 4(a) shows the likelihood for $\Lambda_\infty/\epsilon = 70$. Note that the likelihood is bimodal because the function F' in equation (18) is not monotonic. Thus, a unique value of ℓ_s cannot be inferred from a single $(\Lambda_\infty, \epsilon)$.

We use a truncated Jeffreys prior [17] for ℓ_s, $P(\ell_s) \propto 1/\ell_s$. The prior is truncated outside the range [0.1, 1], which includes the true value $\ell_s = 0.4$. This restriction is imposed for computational convenience and may be relaxed, if needed. The result of an iterative inference process that incorporates successive $(\Lambda_\infty, \epsilon)$ pairs is shown in Figure 4(b). The posterior PDF with zero data points is the prior. The posterior PDF with one data point is bimodal, as expected from the construction of the likelihood. However, the bimodality of the posterior vanishes with two or more data points. When multiple $(\Lambda_\infty, \epsilon)$ pairs are introduced, only the modes corresponding to the most probable ℓ_s will coincide, reinforcing each other. The modes corresponding to the unlikely values of ℓ_s will not coincide and will decay away after the inclusion of a few data points.

The bimodal nature of the posterior PDF accurately represents the non-monotonicity of the ROM given in equation (18), and the possibility of two correct values of ℓ_s for a given $(\Lambda_\infty, \epsilon)$ pair. A single estimate of ℓ_s and any single measure of uncertainty in this estimate (e.g., a variance or a standard error) would not have accurately represented the ambiguity in ℓ_s resulting from a single measurement. Inclusion of additional data points lifts this ambiguity, which is captured by the change of the posterior PDF from bimodal to unimodal.

As additional $(\Lambda_\infty, \epsilon)$ pairs are introduced, the peak in the posterior moves towards the true value of ℓ_s. Any number of point estimates of ℓ_s may be calculated from the posterior, such as the mean, median, or mode. The PDF also gives a full characterization of the uncertainty in ℓ_s. As an example, we have plotted in Figure 4(c) both the posterior variance (a measure of uncertainty) and the absolute difference between the posterior mean and the true value of ℓ_s (a measure of error), for different numbers of data points. As more data are used in the inference problem, both the posterior variance and the error in the posterior mean decrease.

It is interesting to consider whether the posterior mean of ℓ_s will converge to the true value of ℓ_s, i.e., the value used to generate the data, as the number of data pairs $(\Lambda_\infty, \epsilon)$ approaches infinity. It is also interesting to consider what happens to the posterior variance in the same limit. These are questions of posterior consistency, i.e., about the asymptotic (and frequentist) properties of the posterior distribution. If the data were being generated directly from the likelihood, then in problems with fixed and finite parameter dimension such as this one, standard results establish conditions under which the posterior mean converges to the truth and the posterior becomes asymptotically normal with decreasing variance [19; 20]. But the problem analyzed here is different, in that the data pairs $(\Lambda_\infty, \epsilon)$ are generated via solution of the Cahn-Hilliard equation while the ROM is used to define the likelihood function for inference. Thus there exists model error or misspecification, and posterior consistency is far less clear. Theoretically, there may be no particular reason to expect consistency in the present setting. But the empirical results in Figure 4(c) suggest that the ROM is sufficiently accurate for the posterior mean to provide a good representation of the true underlying parameter. Had this not been the case, we could have introduced a model discrepancy term to capture the difference between the ROM mean and the true mean response of the Cahn-Hilliard

system. This could be accomplished using the framework provided by Kennedy and O'Hagan [21], in which the model discrepancy is represented as a Gaussian process and conditioned on data.

4.2. Inferring ϵ

In addition to inferring ℓ_s from ϵ and Λ_∞, the modeling framework we have described also allows us to solve another inference problem, namely determining ϵ from (Λ_∞, ℓ_s) pairs. Thus, we infer a property of the film, ϵ, given that the length scale ℓ_s of the substrate is known and the converged length scale of the film, Λ_∞, may be observed.

To generate data in the form of (Λ_∞, ℓ_s) pairs, values of ℓ_s are picked at random from a uniform distribution in the range [0.1,1]. For each value of ℓ_s, a substrate field is randomly generated. For each substrate, the full Cahn-Hilliard model is run using $\epsilon = 0.04$. Values of Λ_∞ thereby obtained are used to construct (Λ_∞, ℓ_s) pairs and the likelihood is calculated from equation (15) in a manner similar to that in Section 4.1. The process for obtaining the likelihood is shown in Figure 5(a). The mode of the likelihood corresponds to the intersection of the horizontal line defined by the ratio of Λ_∞/ℓ_s with equation (15). The spread in the likelihood is controlled by the random term, Γ. In this case, the likelihood is unimodal because the function F in equation (15) is monotonic.

As before, we use a truncated Jeffreys prior [17] for ϵ, supported on the interval [0.01, 0.1]. Figure 5(b) shows the effect of incorporating successive data points on the posterior. Both the accuracy and confidence in the estimate of ϵ improve as the number of (Λ_∞, ℓ_s) pairs increases. Figure 5(c) plots the error in the posterior mean and the posterior variance for different numbers of data points used in the inference. As the number of data points increases, both the posterior variance and the error in the posterior mean decrease.

5. Discussion

We demonstrated the inference—with quantified uncertainty—of a substrate property from film behavior in an illustrative, model film/substrate system. Two features of Bayesian inference are demonstrated in this context. First is the ability to quantitatively describe uncertainty in the inferred parameters via a posterior probability distribution. This distribution describes the state of knowledge about the inferred parameters, after observing the data.

Second, Bayesian inference allows for sequential improvement of parameter estimates by using the posterior from a previous inference as the prior for inference with new data. We used the error in the posterior mean and variance of the posterior distribution as estimates of the error and uncertainty in the inference, respectively. It was shown that both the error and the uncertainty reduce as an increasing number of data points used for inference.

An important component of our inference procedure is the construction of a reduced order model (ROM) that directly connects the known quantities to those that are to be inferred. To construct our ROM, we built a simplified, parametric model using the Buckingham Pi theorem. The parametric model incorporates physical understanding of the asymptotic behavior of the film. The values of the parameters in the model were obtained by least squares fitting with simulation data. In addition to the parametric model, the ROM includes an additive Gaussian random term to account for the stochastic nature of the full model. The sources of stochasticity are the random initial condition and the randomness of the underlying substrate. The variance of the additive noise is a function of ℓ_s/ϵ. The exact form of the random term is obtained non-parametrically using a Gaussian process fit.

Similar approaches have been adopted in previous applications of Bayesian inference in materials science-related problems. For example, Davis and Guttierrez inferred the melting point of Ti_2GaN starting with what is effectively a ROM for the energy-temperature relation in isochoric molecular dynamics simulations of Ti_2GaN [22]. Yoo *et al.* modeled the creep rupture life of Ni-base superalloys as a function of their composition using Bayesian neural networks [23]. These neural networks are trained on data collected from existing literature and serve a purpose analogous to our ROM.

Other methods for constructing ROMs besides the one we used are available. For instance, a direct GP regression of $\log(\Lambda_\infty/\ell_s)$ on $\log(\ell_s/\epsilon)$ might be used as a ROM. However, this choice would not have allowed us to take advantage of known structure of the problem (i.e., the asymptotic behavior described in Section 3). Moreover, capturing the non-stationary stochastic response of the film/substrate system would require using a non-stationary prior covariance kernel — a non-trivial task.

Also, we could have used a more fully Bayesian procedure to *construct* the ROM. This approach would involve constructing posterior distributions for a, b, c and $\sigma^2(\ell_s/\epsilon)$, conditioned on the training data, instead of using point estimates. However, this additional complexity seemed unnecessary in our case, since we had a large and comprehensive training set from which to construct the ROM. If we had had a limited training set, as in [24], a Bayesian construction of the ROM would have been warranted.

Although the film/substrate model studied here does not correspond to a specific experiment, it nevertheless illustrates general features of how Bayesian inference may be applied to inferring substrate properties from film behavior. In particular, it demonstrates the utility of constructing a ROM that connects the quantities of interest in a way that may be inverted. Although it ignores much of the details of the film/substrate model under consideration, this approach nevertheless gives rise to accurate inferences with quantified uncertainty. Similar approaches may be adapted to particular film/substrate systems as well as to other situations in materials science where difficult-to-observe quantities are to be inferred from easily observable ones.

Acknowledgement

MJD acknowledges fruitful discussions with the organizers and attendees of the Uncertainty Quantification in Materials Modeling workshop organized in 2013 at the Institute for Mathematics and its Applications, University of Minnesota, Twin Cities. This project was funded by the US Department of Energy, Office of Basic Energy Sciences under award no. DE-SC0008926.

References

[1] J. Lewandowski and A. Greer, "Temperature rise at shear bands in metallic glasses," *Nature Materials*, vol. 5, pp. 15–18, 2006.

[2] J. Aizenberg, A. Black, and G. Whitesides, "Controlling local disorder in self-assembled monolayers by patterning the topography of their metallic supports," *Nature*, vol. 394, pp. 868–871, 1998.

[3] N. Bowden, S. Brittain, A. Evans, J. Hutchinson, and G. Whitesides, "Spontaneous formation of ordered structures in thin films of metals supported on an elastomeric polymer," *Nature*, vol. 393, pp. 146–149, 1998.

[4] D. Sivia and J. Skilling, *Data analysis: A Bayesian tutorial*. Oxford science publications, 2 ed., 2006.

[5] A. Gelman, J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin, *Bayesian data analysis*. CRC Press, 3 ed., 2013.

[6] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[7] J. Cahn and J. Hilliard, "Free energy of a nonuniform system. i. interfacial free energy," *J. Chem. Phys*, vol. 28, pp. 258–267, 1958.

[8] D. Eyre, "An unconditionally stable one-step scheme for gradient systems." Unpublished manuscript, University of Utah, Salk Lake City, UT, June 1998.

[9] L. Trefethen, *Spectral Methods in MATLAB*. SIAM, 2000.

[10] S. Cox and P. Matthews, "Exponential time differencing for stiff systems," *Journal of Computational Physics*, vol. 176, pp. 430–455, 2002.

[11] J. Cahn, "On spinodal decomposition," *Acta Met.*, vol. 9, pp. 795–801, 1961.

[12] M. Beaumont, W. Zhang, and D. Balding, "Approximate bayesian computation in population genetics," *Genetics*, vol. 162, pp. 2025–2035, 2002.

[13] P. Marjoram, J. Molitor, V. Plagnol, and S. Tavare, "Markov chain monte carlo without likelihoods," *Proceedings of the National Academy of Sciences*, vol. 100, no. 26, pp. 15324–15328, 2003.

[14] L. Yarin, *The Pi-Theorem: Applications to Fluid Mechanics and Heat and Mass Transfer*, vol. 1. Springer, 2012.

[15] E. Jaynes, "Information theory and statistical mechanics," *Physical Review*, vol. 106, no. 4, pp. 620–630, 1957.

[16] E. Jaynes, "Information theory and statistical mechanics. ii," *Physical Review*, vol. 108, no. 4, pp. 171–190, 1957.

[17] H. Jeffreys, "An invariant form for the prior probability in estimation problems," *Proceedings of the Royal Society*, no. 453-461, 1946.

[18] J. Berger and J. Bernardo, "On the development of the reference prior method," Tech. Rep. 91-15C, Department of Statistics, Purdue University, 1991.

[19] J. Doob, "Application of the theory of martingales," in *Le Calcul des Probabilites et ses Applications*, pp. 23–27, Colloques Internationaux du Centre National de la Recherche Scientifique, 1949.

[20] A. van der Vaart, *Asymptotic Statistics*. Cambridge University Press, 2000.

[21] M. Kennedy and A. O'Hagan, "Bayesian calibration of computer models," *Journal of the Royal Statistical Society B*, vol. 63, pp. 425–464, 2001.

[22] S. Davis and G. Gutierrez, "Bayesian inference as a tool for analysis of first-principles calculations of complex materials: an application to the melting point of Ti_2GaN," *Modelling and Simulation in Materials Science and Engineering*, 2013.

[23] Y. Yoo, C. Jo, and C. Jones, "Compositional prediction of creep rupture life of single crystal ni base superalloy by bayesian neural network," *Materials Science and Engineering*, pp. 22–29, 2001.

[24] S. Sarkar, D. Kosson, S. Mahadevan, J. Meeussen, H. van der Sloot, J. Arnold, and K. Brown, "Bayesian calibration of thermodynamic parameters for geochemical speciation modeling of cementitious materials," *Cement and Concrete Research*, vol. 42, pp. 889–902, 2012.

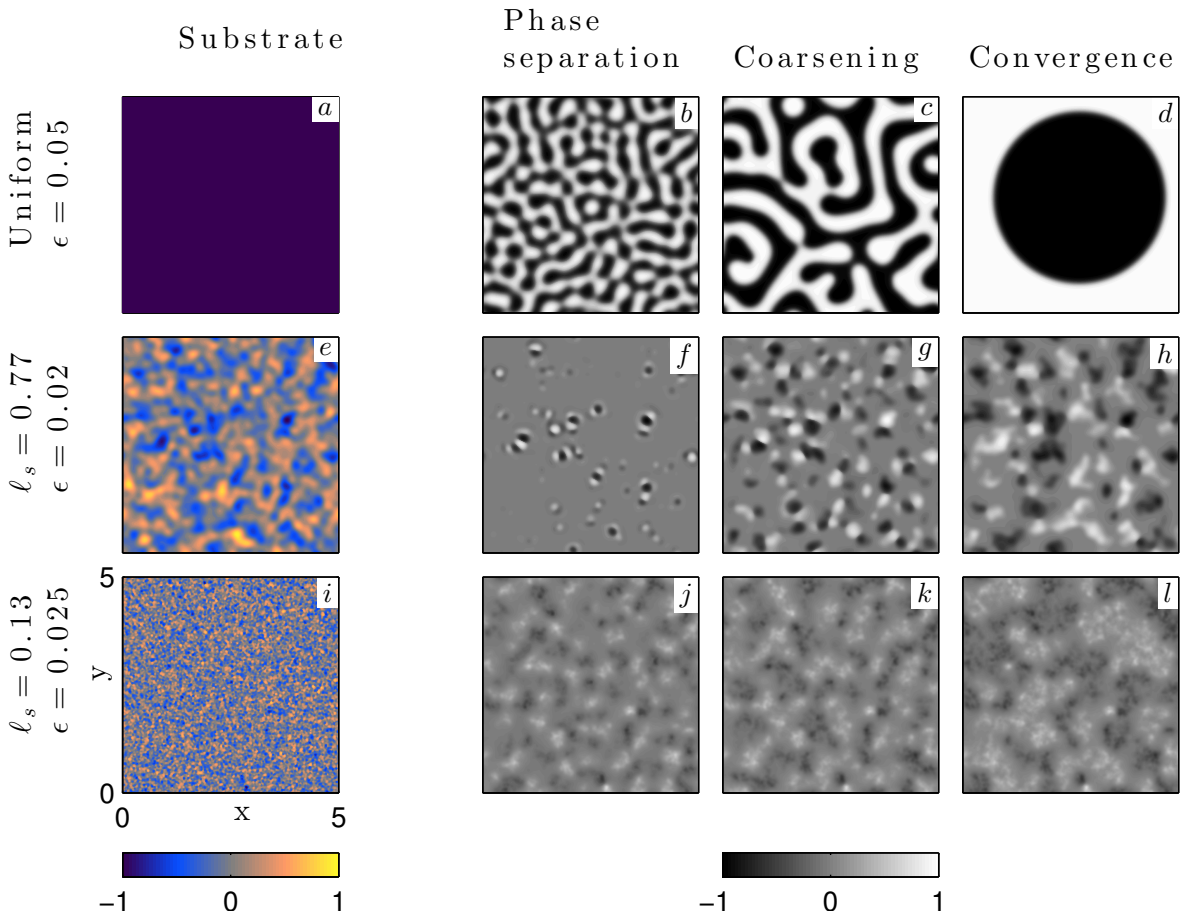

Figure 1. (a) Uniform substate ($\mathcal{T}(x, y) = -1$), (e) patterned substrate with $\ell_s = 0.77$ and (i) patterned substrate with $\ell_s = 0.13$. The rows (b)-(d),(f)-(h), and (j)-(l) show the time evolution of initially nearly uniform films with $\epsilon = [0.05, 0.02, 0.05]$ respectively, on the three substrates

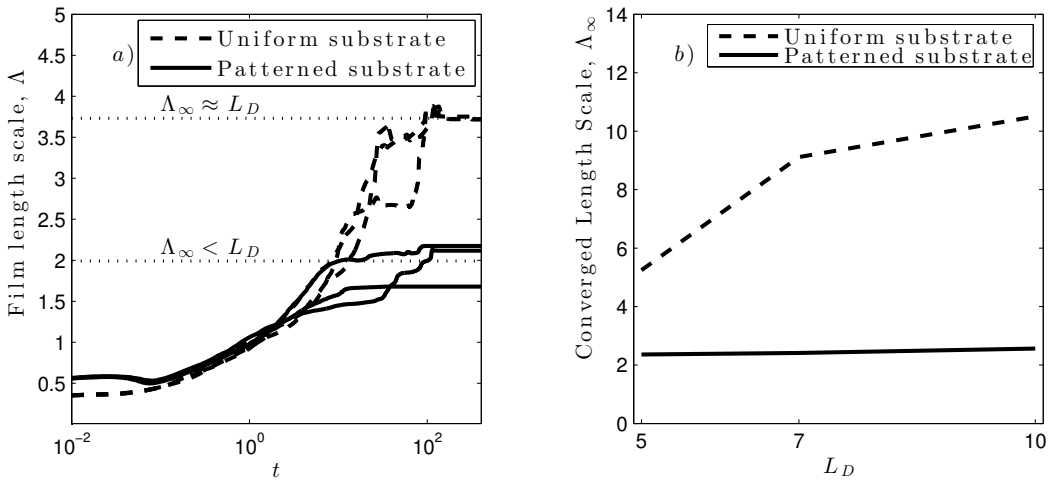

(a) (b)

Figure 2. (a) Three trajectories of the film length scale $\Lambda(t)$ for the uniform substrate in Figure 1(a) and the patterned substrate in Figure 1(e). (b) Variation of the converged film length scale Λ_∞ as a function of computational domain size L_D for the uniform substrate in Figure 1(a) and the patterned substrate in Figure 1(e).

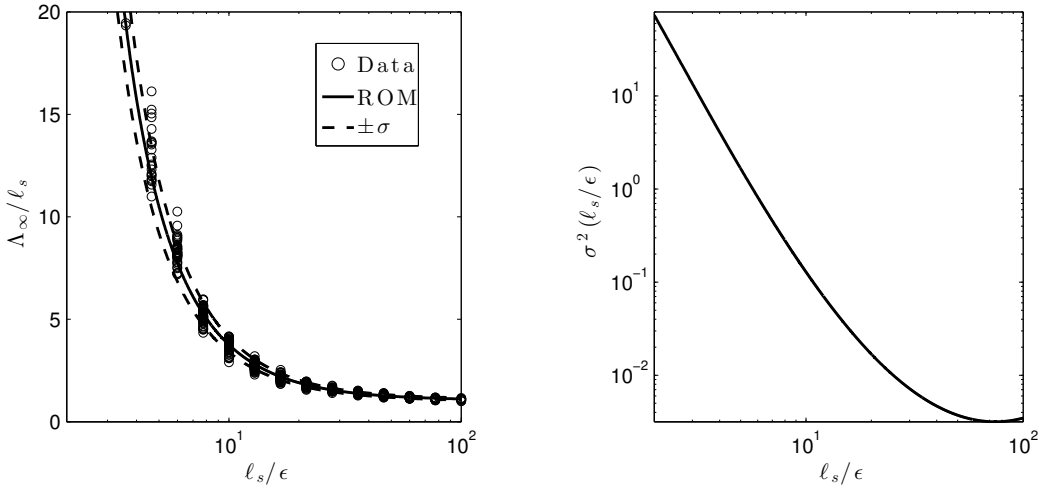

(a) (b)

Figure 3. (a) A plot of the reduced order model (ROM) given in equation (15). (b) A plot of the non-stationary variance of the random term $\Gamma(\ell_s/\epsilon)$ given in equation (15)

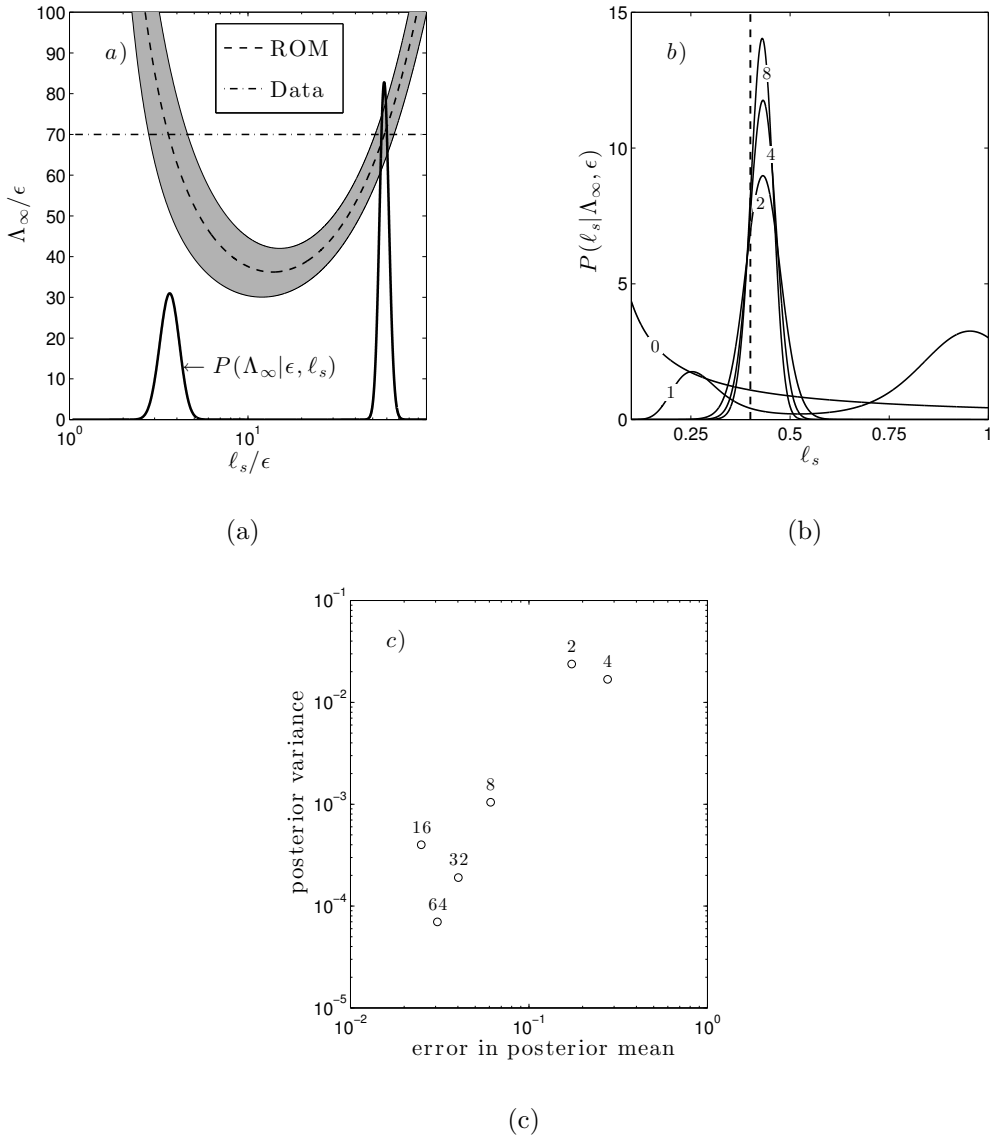

(a) (b)

(c)

Figure 4. (a) Likelihood $P(\Lambda_\infty|\ell_s,\epsilon)$ plotted as a function of ℓ_s/ϵ for fixed Λ_∞ and ϵ. Each $(\Lambda_\infty,\epsilon)$ pair represents a point on the y-axis. The discrepancy is given by the distance between the horizontal line at the given value of (Λ_∞/ℓ_s) and the ROM. The random term in the model, given in grey, is used to calculate the likelihood. (b) Posterior PDFs for different numbers of $(\Lambda_\infty,\epsilon)$ pairs. With the inclusion of ever more data, PDF centers closer to the true value of ℓ_s with a tighter uncertainty bound. (c) Posterior variance and error in posterior mean for different numbers of $(\Lambda_\infty,\epsilon)$ pairs. Both error and variance are reducing with increasing data points

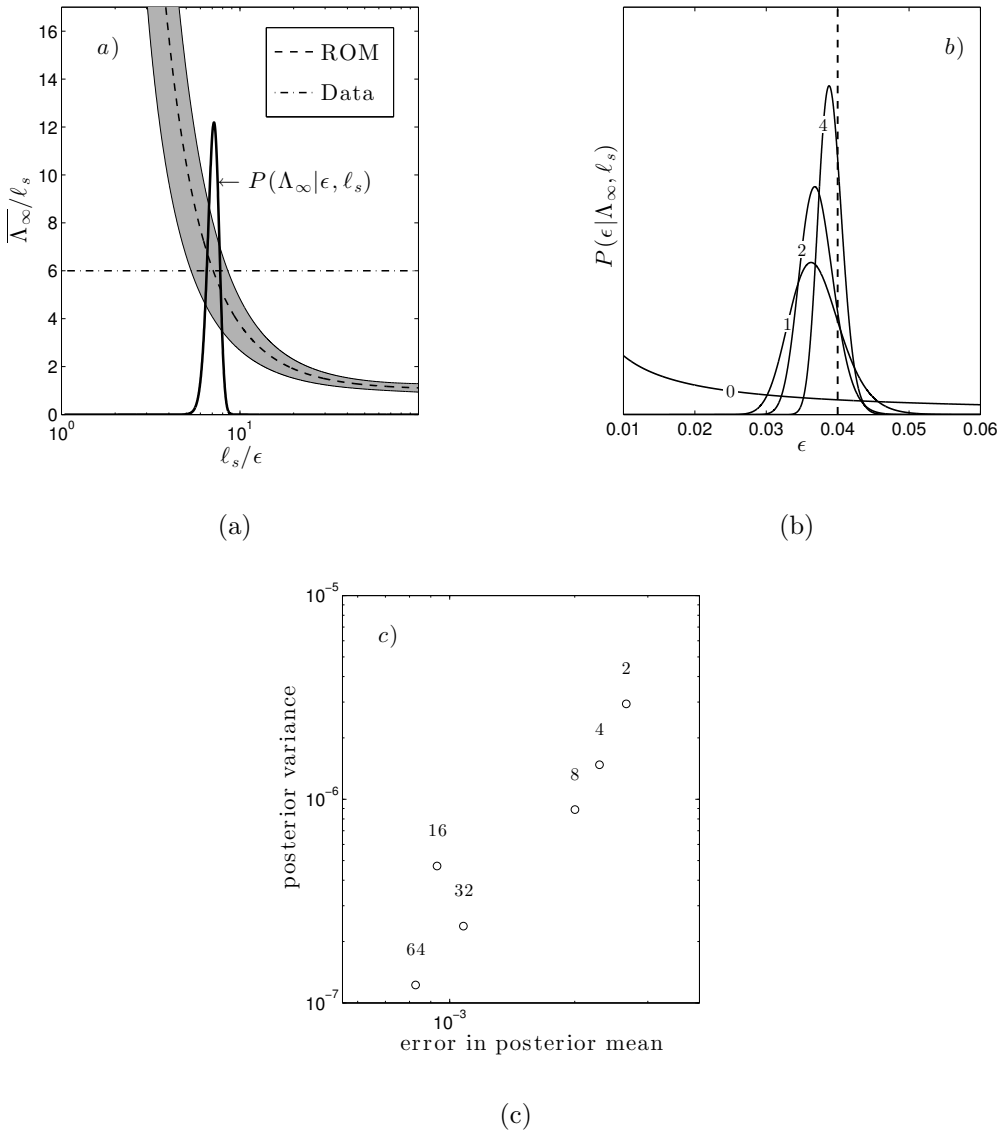

(a) (b)

(c)

Figure 5. (a) Likelihood $P(\Lambda_\infty|\ell_s,\epsilon)$ plotted as a function of ℓ_s/ϵ for fixed Λ_∞ and ℓ_s. Each (Λ_∞,ℓ_s) pair represents a point on the y-axis. The discrepancy is given by the distance between the horizontal line at the given value of (Λ_∞/ℓ_s) and the ROM. The random term in the model, given in grey, is used to calculate the likelihood. (b) Posterior PDFs for different numbers of (Λ_∞,ℓ_s) pairs. With the inclusion of ever more data, PDF centers closer to the true value of ls with a tighter uncertainty bound. (c) Posterior variance and error in posterior mean for different numbers of (Λ_∞,ℓ_s) pairs. Both error and variance are reducing with increasing data points