Research article

Generating patient-specific virtual tumor populations with reaction-diffusion models and molecular imaging data

Nick Henscheid*

Department of Medical Imaging, University of Arizona, 1501 N. Campbell, Tucson, AZ 85724, USA

* **Correspondence:** Email: nph@email.arizona.edu.

Abstract: The use of mathematical tumor growth models coupled to noisy imaging data has been suggested as a possible component in the push towards precision medicine. We discuss the generation of population and patient-specific virtual populations in this context, providing in silico experiments to demonstrate how intra- and inter-patient heterogeneity can be estimated by applying rigorous statistical procedures to noisy molecular imaging data, and how the noise properties of such data can be analyzed to estimate uncertainties in predicted patient outcomes.

Keywords: mathematical onoclogy; precision medicine; molecular imaging; virtual populations; virtual clinical trials; emission computed tomography

1. Introduction

Broadly speaking, mathematical oncology seeks to develop a collection of in silico models that accurately predict the growth and progression of malignant tumors and their response to treatment [1]. While the basic scientific goal is to develop a rigorous quantitative science of cancer [2], a more readily attainable goal is to use existing models and data to inform clinical practice in some substantive way [3]. One approach to clinical mathematical oncology, which we call mathematical model-based precision medicine, is to use mathematical models to make patient-specific prognostic and therapeutic predictions which are subsequently employed to direct treatment choices and predict patient outcomes. In our view, such predictions must be conditioned on patient data in order for the strategy to qualify as precision medicine [4–6].

Two essential features of oncology are heterogeneity and uncertainty. Heterogeneity can manifest as interpatient (between-patient), intrapatient (within-patient) and intratumor (within-tumor); see [7] for a more complete discussion of the clinical impact of heterogeneity in cancer and [8] for a review of the mathematical strategies to address heterogeneity. From a modeling perspective, heterogeneity implies a statistical approach where model parameters are treated as spatially distributed random variables

with both population and patient-specific probability distributions [5]. Heterogeneity also highlights the issue of uncertainty: For a given patient, most if not all model parameters are unknown prior to the acquisition of patient-specific data, and even with such data, uncertainties may remain, since model parameters may only be partially identifiable and the data acquisition system may introduce additional irreducible noise. Image data provides a clear example of this issue, as imaging introduces information loss due to detector noise, low resolution and the possibility of null functions [9]. This uncertainty must be accounted for if model-based predictions are to be used reliably in the clinical context.

Once a relevant mathematical model has been selected, prediction takes place in silico, whereby a simulated patient is produced and Quantities-of-Interest (QoIs) are calculated. In this context, we say that the parameter values and resulting simulations correspond to a virtual patient. Owing to the heterogeneity and uncertainty issues discussed above, in both population and patient-specific applications of mathematical oncology it is desirable to generate suitably randomized Virtual Patient Populations (VPPs). If the VPP represents an untreated group, it is called a Virtual Control Population (VCP) [4, 10], while if the VPP represents a treatment group, it is called a Virtual Treatment Population (VTP). With appropriate VPPs, an in silico Virtual Clinical Trial (VCT) can be performed to assess treatment choices, again in either the population or patient-specific case. Population VPPs and VCTs can assist in making general prognosis and treatment recommendations or simulate a regulatory trial [11], while a patient-specific VPP and/or VCT can be used to assess uncertainty in a particular patient's disease progression or treatment response [6, 12]. While population VCTs have found usage in device design and regulatory contexts [13], patient-specific VPPs and VCTs have seen limited usage.

In previous work [5, 6], we have highlighted the application of spatiotemporal random field modeling to address heterogeneity and uncertainty. In this work, we discuss specific VPP-based strategies towards the practical implementation of the framework presented in [5, 6]. Specifically, we describe a method for generating VCPs by pairing a spatiotemporal reaction-diffusion equation (RDE) model for avascular, pre-metastatic tumor growth with spatially inhomogenous random field coefficients. We discuss the RDE model in sections 2.1 and 2.2 and the random field models in section 2.3. Together, these models provide a solution of the 'forward problem' for VPP generation. Two important inverse problems then arise. In the first, we consider in sections 4.1 and 4.2 the problem of estimating the RDE coefficient fields for a specific patient from noisy molecular imaging data; proof-of-concept simulation results are given in sections 5.1 and 5.2. For the second inverse problem, we must estimate the statistical parameters that define the random field coefficient models from a collection of many patients' data, hence solving the statistical calibration problem for virtual populations; we discuss this problem in section 4.3. Both maximum likelihood and Bayesian techniques are presented, and both require an accurate likelihood model connecting model parameters to available noisy data. To this end, in section 3 we discuss statistical models for molecular Emission Computed Tomography (ECT) data. We make this emphasis both because well-validated likelihood models for ECT are available [9], and because it is possible to model a direct connection between physiological processes such as tumor cell density and spatial growth rate, and certain emission imaging techniques and tracers.

Note that the results given in sections 5.1 and 5.2 are generated in a pure simulation setting, where both ground-truth tumors and image data are simulated. This allows for methodological comparisons with known ground truth and tight control of all nuisance parameters that influence the data. A

summary of future work that extends these results is provided in section 6.

2. Tumor growth modeling with heterogeneity and uncertainty

In this section, we introduce the tumor growth model that we employ for the remainder of the paper. We emphasize that the models we present are simplifications of real tumor biology, which is complex and not yet fully understood. These models are intended to be more 'clinically relevant' [3] than biologically accurate, due to their relative simplicity and amenability to patient-specific inference from spatiotemporal imaging data. In section 2.2 we discuss the parameter-to-solution map, in section 2.3 we discuss random field models for the tumor growth model coefficients, and in section 2.4 discuss the notions of virtual populations and virtual clinical trials.

2.1. Reaction-diffusion tumor growth models

A wide variety of Partial Differential Equation (PDE) models have been employed to model spatiotemporal tumor growth, and the derivation and mathematical analysis of such models is adequately described elsewhere [14–21]. Here, we are interested in computational techniques for generating virtual patient populations, and thus consider a relatively simple PDE model (2.1) that has gained traction in the context of GlioBlastoma Multiforme (GBM). The randomization and statistical calibration techniques we present are independent of the tumor growth model, however, so the choice of Eq (2.1) is largely for illustration purposes.

We denote a spatiotemporal density of tumor cells by $\boldsymbol{n} = n(\boldsymbol{x}, t)$, with $\boldsymbol{x} \in V$ being a $d-$dimensional position vector in the spatial domain V and $0 \le t \le T$, where T is the maximum number of days simulated. The units of $\boldsymbol{n}$ are cells per unit volume, and we take $(\boldsymbol{x}, t)$ to have units of cm and days, respectively. For in vivo tumors $d = 3$, though for simulation purposes and modeling of in vitro and certain in vivo experimental setups we also consider spatial domains in $d = 2$ as a reasonable approximation. The general reaction-diffusion equation for $\boldsymbol{n}$ is

$$\begin{cases} \frac{\partial \boldsymbol{n}}{\partial t} = \nabla \cdot (\boldsymbol{D}\nabla\boldsymbol{n}) + \mathcal{G}(\boldsymbol{n}) & (\boldsymbol{x}, t) \in V \times [0, T] \\ n(\boldsymbol{x}, 0) = n_0(\boldsymbol{x}) \\ \hat{\boldsymbol{v}} \cdot \nabla\boldsymbol{n} = 0 & \text{boundary conditions} \end{cases} \tag{2.1}$$

The spatially varying scalar diffusion coefficient $\boldsymbol{D} = D(\boldsymbol{x})$ describes the rate of apparent cell diffusion, while the growth function $\mathcal{G}(\boldsymbol{n})$ describes the growth and competition between the tumor and its environment. The boundary condition is taken to be the no-flux (Neumann) condition $\hat{\boldsymbol{v}} \cdot \nabla\boldsymbol{n} = 0$, where $\hat{\boldsymbol{v}}$ is the outward boundary normal of the domain V. Note that a treatment function $\mathcal{T}(\boldsymbol{n})$ can be added to Eq (2.1) to model an intervention such as chemotherapy, radiotherapy or surgical resection, but we postpone discussion of such models to a future work.

While in the scalar tumor growth context the Gompertz function $\mathcal{G}(\boldsymbol{n}) = -\rho\boldsymbol{n}\ln(\boldsymbol{n}/\kappa)$ has gained plausible mechanistic support via the notion of quiescence [22], to our knowledge the spatial RDE growth function is largely a phenomenological choice. The Fisher-KPP growth function is frequently selected in the GBM context, and is the one we will use in this paper:

$$\mathcal{G}(\boldsymbol{n}) = \rho(\boldsymbol{x})n(\boldsymbol{x}, t)\left(1 - \frac{n(\boldsymbol{x}, t)}{\kappa(\boldsymbol{x})}\right) \tag{2.2}$$

The growth $\rho(\boldsymbol{x})$ and carrying capacity $\kappa(\boldsymbol{x})$ functions, which have respective units of cells per day and cells, will be discussed further below.

2.2. Patient-specific parameters and the parameter-to-solution map

Note that in Eqs (2.1) and (2.2), several additional parameter functions are required to fully specify the solution to the PDE. These include the diffusion coefficient $D(\boldsymbol{x})$, the initial cell density $n_0(\boldsymbol{x})$, the growth $\rho(\boldsymbol{x})$ and carrying capacity $\kappa(\boldsymbol{x})$. These parameter functions are certainly patient-specific, since they correspond to local physiology and nutrient concentrations. For notational simplicity, we will use the symbol $\boldsymbol{\beta}$ to denote the collection of all parameters necessary to specify a solution to Eq (2.1). In the context of this paper, a virtual patient then corresponds to a particular choice of $\boldsymbol{\beta}$, which we will indicate by saying that patient j has parameter $\boldsymbol{\beta}_j$. We will assume for the remainder of this article the initial condition is given by a Gaussian

$$n_0(\boldsymbol{x}) = A \exp\left(\frac{1}{2\sigma^2} - \|\boldsymbol{x} - \boldsymbol{x}_0\|^2\right) \tag{2.3}$$

with $\boldsymbol{x}_0 = [0.5, 0.5]^T$, $A = 5/2\pi\sigma^2$ and $\sigma^2 = 1e-4$; this represents an initial collection of 5 tumor cells, well-localized around the center of our computational domain. The collection of all remaining unknown parameter fields is thus $\boldsymbol{\beta} = (\boldsymbol{D}, \boldsymbol{\rho}, \boldsymbol{\kappa})$. To avoid pathological situations, we also assume the technical condition that $\boldsymbol{\beta} \in \mathbb{B}$, where $\mathbb{B}$ is a set of parameter functions selected to guarantee that a positive, bounded solution $n(\boldsymbol{x}, t)$ of Eq (2.1) exists for all $\boldsymbol{\beta} \in \mathbb{B}$ and for $0 \le t \le T$. Selecting such a $\mathbb{B}$ requires careful technical analysis of Eq (2.1); see [23].

While we assume that each patient corresponds to a unique set of parameters $\boldsymbol{\beta} = \boldsymbol{\beta}_j$, it is likely that $\boldsymbol{\beta}_j$ is effectively unknown or uncertain for a given patient. To address this, we treat $\boldsymbol{\beta}$ as a function space-valued random vector in both the population and patient-specific cases, and provide corresponding random field models in section 2.3 below. In a simulation setting, we must choose a computational representation for $\boldsymbol{\beta}_j$, which we indicate with a superscript v for 'virtual'. So, $\boldsymbol{\beta}_j^{(v)}$ is a virtual (i.e., computational) sample corresponding to a sample for $\boldsymbol{\beta}$.

Assuming Eq (2.1) is classically well-posed for every (or almost every) $\boldsymbol{\beta} \in \mathbb{B}$, we can construct the deterministic parameter-to-solution map

$$\boldsymbol{n}_j^{(model)} = \mathcal{F}(\boldsymbol{\beta}_j) \tag{2.4}$$

where $\boldsymbol{n}_j^{(model)} = n_j^{(model)}(\boldsymbol{x}, t)$ solves Eq (2.1) with the parameter vector $\boldsymbol{\beta} = \boldsymbol{\beta}_j$, corresponding to patient j. In a simulation setting, the forward map (2.4) is implemented via a numerical differential equation solver. We have employed a finite difference method-of-lines approach with a five-point spatial stencil and Runge-Kutta time stepper, as discussed in [24]. The result is an approximate solution to Eq (2.1), where the virtual coefficient sample $\boldsymbol{\beta}_j^{(v)}$ is processed through computational algorithm $\mathcal{F}^{(v)}$ to produce the virtual tumor cell density $\boldsymbol{n}_j^{(v)}$:

$$\boldsymbol{n}_j^{(v)} = \mathcal{F}^{(v)}\left(\boldsymbol{\beta}_j^{(v)}\right) \tag{2.5}$$

The choice of discretization specifies the format for the virtual cell density $\boldsymbol{n}_j^{(v)}$; with our finite difference scheme, we can assume that $\boldsymbol{n}_j^{(v)}$ takes the form of an $N_x \times N_y \times N_t$ voxel array (for $d = 2$).

The specification of $\boldsymbol{\beta}_j^{(v)}$ is discussed below; we will assume a voxel-free form (Eq (2.8) below) that can be evaluated on an arbitrary grid when needed. Note that the virtual cell density will differ both from a patient's true cell density and from a measured cell density by discretzation, model and measurement errors; see [8] for a discussion of these errors. A demonstration of our scheme implementing (2.5) in $d = 2$ is shown in Figure 1.
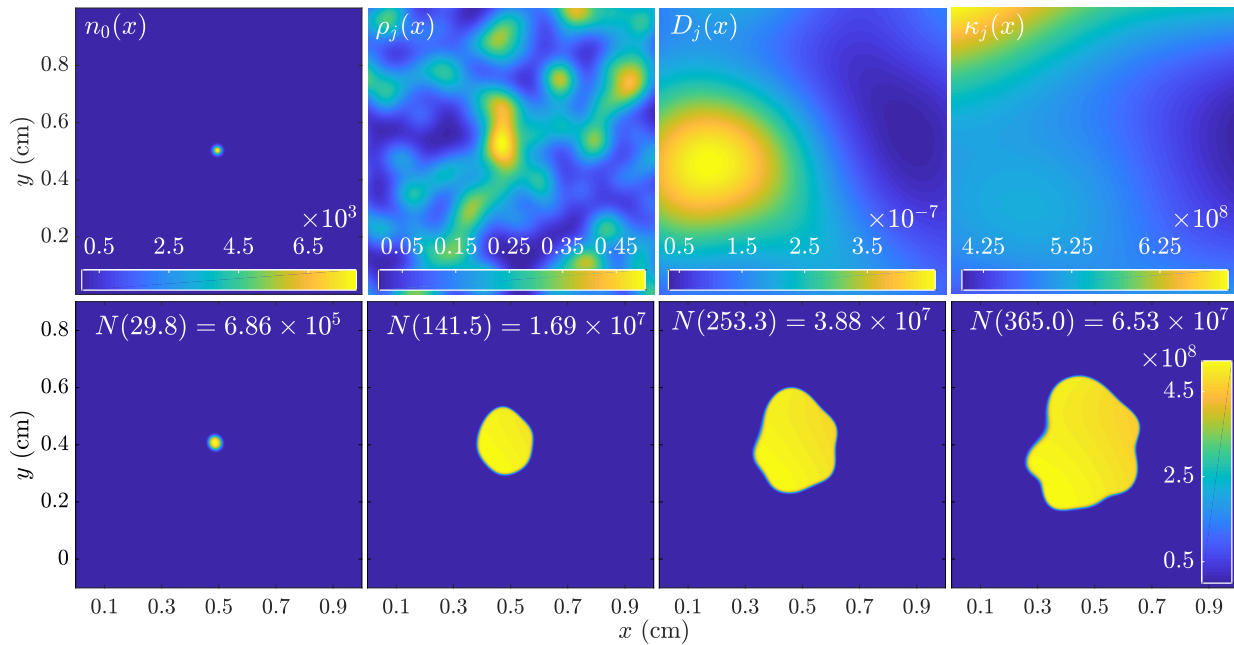
Figure 1. A realization of the tumor growth model (2.1) with Fisher-KPP growth term (2.2) (bottom row). The time-independent random field coefficients (top row, frames 2–4) are draws from $\boldsymbol{\rho} \sim LB(200, 0.1, 2e{-}3)$, $\boldsymbol{D} \sim LB(20, 1e{-}7, 0.04)$ and $\boldsymbol{\kappa} \sim LB(100, 5e7, 0.1)$, as defined in section 2.3. The tumor burden $N(t)$, defined in Eq (2.13), is shown for each time point. The initial condition $n_0(\boldsymbol{x})$ is shown top left.

In the next section, we will introduce random field models that allow Eqs (2.4) and (2.5) to be used to generate virtual tumor populations. Note briefly that we have made assumptions that allow us to analyze and simulate Eq (2.1) as a Random PDE (RPDE) model, that is as a classically well-posed PDE whose coefficients are randomized. Note that this is in contrast to a Stochastic PDE (SPDE) model, which requires more advanced analysis and simulation techniques. See e.g., [5] for discussion of a technique based on characteristic functionals (introduced briefly below), and [25–27] for further discussion of the distinction between SPDE and RPDE.

2.3. Random field models for RDE coefficients

As discussed in the introduction, any mathematical model that is employed in a clinical precision medicine context must account for the possibility of heterogeneity and uncertainty by assuming that model parameters are sampled from a probability distribution. Since the model parameters that specify Eqs (2.1) and (2.2) are spatial functions, we employ the theory of random fields. We will avoid technical discussions, sufficing to say that one can treat $\boldsymbol{\beta}$ as a function space-valued random

vector, written formally as $\boldsymbol{\beta} : \Omega \to \mathbb{B}$, where $(\Omega, \mathscr{F}, \mathbb{P})$ is a probability space and $\mathbb{B}$ is a set of functions mentioned previously, selected to guarantee the existence of the forward map (2.4). We assume that $\mathbb{B} \subset \mathbb{X}$, where $\mathbb{X}$ is a Hilbert space with inner product $(\boldsymbol{\beta}, \boldsymbol{\varphi})_{\mathbb{X}}$, taken to be the $L^2(V)$ inner product for vector-valued functions. In the most general setting, the complete statistics of a second-order random process are fully specified by its marginal distributions (through the Kolmogorov construction), its abstract probability distribution $\mathbb{P}_\beta$, or its characteristic functional: [5, 9, 28, 29]

$$\Psi_\beta(\boldsymbol{\varphi}) = \mathbb{E}\left[\exp\left(i(\boldsymbol{\beta}, \boldsymbol{\varphi})_{\mathbb{X}}\right)\right]. \tag{2.6}$$

Frequently, a process is completely specified by its first few moments e.g., its mean function $\bar{\beta}(\boldsymbol{x}) = \mathbb{E}[\boldsymbol{\beta}]$ and its covariance function $\boldsymbol{k}(\boldsymbol{x}, \boldsymbol{x}') = \mathbb{E}[(\boldsymbol{\beta}(\boldsymbol{x}) - \bar{\beta}(\boldsymbol{x}))(\boldsymbol{\beta}(\boldsymbol{x}') - \bar{\beta}(\boldsymbol{x}'))^{\mathsf{T}}]$. While its use requires functional calculus, the characteristic functional (2.6) is typically a compact and efficient way to describe the complete statistics of a random field, and its use in the biological context is a topic of current investigation [5]. Refer to the recent books [27] and [30] for a complete discussion of the theory of Hilbert space-valued random vectors.

In this work, we will simulate random fields by first assuming a convenient functional form for the realizations of the process, then randomizing the corresponding (finitely many) parameters to achieve desired statistics. This method allows us to guarantee that realizations of $\boldsymbol{\beta}$ satisfy conditions for Eq (2.1) to be well-posed, and allows for straightforward simulations using standard code packages. In particular, the random field models we use will take the form of a randomized synthesis map, which for a scalar-valued function $\varphi_j(\boldsymbol{x})$ takes the form:

$$\boldsymbol{\varphi}_j = \boldsymbol{\Phi}(\boldsymbol{\theta}_j) \tag{2.7}$$

where $\boldsymbol{\Phi}(\cdot) : \mathbb{R}^N \to \mathbb{X}$ is a synthesis map that maps a (finite-dimensional) patient-specific parameter $\boldsymbol{\theta}_j$ to the function $\varphi_j^{(v)}(\boldsymbol{x})$. By selecting a probability distribution $\mathbb{P}(\boldsymbol{\theta}_j | \boldsymbol{\theta}_p)$ that depends on the population parameter $\boldsymbol{\theta}_p \in \boldsymbol{\Theta} \subset \mathbb{R}^p$, the synthesis map (2.7) produces a random field. In sections 4.1 and 4.2, we discuss ways to estimate a particular $\boldsymbol{\theta}_j$ from image data, while in section 4.3 we discuss ways to estimate the population parameter $\boldsymbol{\theta}_p$ from a collection of images.

Because the coefficients in the RDE (2.1) must be non-negative and bounded to be physiologically realistic, it is advantageous to begin with a random field model that always produces non-negative, bounded realizations. For our simulations, we employ 'lumpy-type' random field models [31, 32], a realization of which has the functional form

$$\varphi(\boldsymbol{x}) = \sum_{l=1}^{L_{max}} b_l \ell(\boldsymbol{x} - \boldsymbol{x}_l; \boldsymbol{\gamma}). \tag{2.8}$$

If the lump function $\ell(\boldsymbol{x}; \boldsymbol{\gamma})$ and the lump amplitudes b_l are non-negative and bounded, and the number of lumps is finite, then Eq (2.8) is guaranteed to produce non-negative and bounded realizations. For simplicity, we use isotropic unnormalized Gaussian lumps:

$$\ell(\boldsymbol{x}; \sigma^2) = \exp\left(-\frac{1}{2\sigma^2}\|\boldsymbol{x}\|^2\right) \tag{2.9}$$

Random fields with this choice of lump function were proposed to model soft tissues with smoothly varying characteristics [31], while more sophisticated choices of $\ell(\boldsymbol{x})$, such as the clustered lump technique [32], can model more complex tissues such as breast or brain tissue. Ultimately, the choice of $\ell(\boldsymbol{x})$ and its shape parameter $\boldsymbol{\gamma}$ requires a statistical calibration step such as that discussed in section 4.3.

We randomize Eq (2.8) as follows. We first select a lump shape parameter $\boldsymbol{\gamma} = \sigma^2 = \sigma_0^2$, a constant lump amplitude b_0, and a mean number of lumps $\bar{L} \ll L_{max}$. We then draw $L \sim \text{Poi}(\bar{L})$, and set $b_\ell = b_0$ for $1 \le \ell \le L$, and $b_\ell = 0$ otherwise. Lastly, we draw $\boldsymbol{x}_1 \ldots, \boldsymbol{x}_L$, I.I.D. uniform in the domain V and form the sum (2.8). Note that the resulting random field realization can be written in terms of a nonlinear synthesis map (2.7) with $N = (d + 1)L_{max}$. We use the notation $LB(\bar{L}, b_0, \sigma_0^2)$ to indicate a lumpy-type field with mean number of lumps $\bar{L}$, amplitude b_0 and isotropic lump variance σ_0^2. A demonstration of the behavior of lumpy-type random fields of the form given in Eq (2.8) is shown in Figure 1, top row.

We reiterate that a single realization of the random field specified in Eq (2.8) is fully specified by the parameters σ_0^2, b_0 and the lump center list $[\boldsymbol{x}_l]$, while the complete statistics of the random field are fully specified by $\bar{L}, \sigma_0^2$ and b_0. This distinction is important when we discuss parameter estimation problems in sections 4.1–4.3.

2.4. Virtual populations and quantitative biomarkers

Suppose that $\boldsymbol{q} \in \mathbb{R}^n$ is a low-dimensional vector quantifying patient status or prognosis, for instance (informally) $\boldsymbol{q} = $ [metric of tumor burden, metric of normal tissue function]. In the model-based precision medicine context, such a vector is generally called a quantitative biomarker [6, 33, 34]. The purpose of a virtual patient population, and more generally a virtual clinical trial, is to predict the distribution of $\boldsymbol{q}$ for either a control or treatment group. We will assume for convenience that $\boldsymbol{q} = q \in \mathbb{R}$ is a scalar functional of $\boldsymbol{\beta}$ and $\boldsymbol{n}$:

$$q = \mathcal{M}(\boldsymbol{\beta}, \boldsymbol{n}) \tag{2.10}$$

Given a virtual population of simulated parameters and tumors, $\left(\boldsymbol{\beta}_1^{(v)}, \boldsymbol{n}_1^{(v)}\right), \ldots, \left(\boldsymbol{\beta}_J^{(v)}, \boldsymbol{n}_J^{(v)}\right)$, we estimate the biomarker q for each virtual patient using a method $\mathcal{M}^{(v)}$, resulting in a vector

$$\boldsymbol{Q}^{(v)} = \left[\mathcal{M}^{(v)}\left(\boldsymbol{\beta}_1^{(v)}, \boldsymbol{n}_1^{(v)}\right), \ldots, \mathcal{M}^{(v)}\left(\boldsymbol{\beta}_J^{(v)}, \boldsymbol{n}_J^{(v)}\right)\right] \in \mathbb{R}^J \tag{2.11}$$

If $\boldsymbol{Q}^{(v)}$ is generated from a population distribution, that is, the random vector $\boldsymbol{\beta}$ models inter-patient heterogeneity, then we say that $\boldsymbol{Q}^{(v)}$ is a *population* virtual biomarker sample. If $\boldsymbol{\beta}$ models patient-specific uncertainties, for instance if $\boldsymbol{\beta}$ is sampled from the conditional random vector $\boldsymbol{\beta}_j|\boldsymbol{g}_j$ where $\boldsymbol{g}_j$ is some patient-specific data, then we say that $\boldsymbol{Q}^{(v)}$ is a patient-specific virtual biomarker sample. Statistical properties of q such as its mean, standard deviation and distribution can be estimated from $\boldsymbol{Q}^{(v)}$. By generating virtual control and treatment groups, Eq (2.11) can be used to simulate a clinical trial. If the VPP employed in Eq (2.11) represents a treatment group, decisions regarding the intervention can be made based on these statistics. For instance, a rigorous decision-theoretic approach to patient-specific treatment selection would define an additional function $U(\boldsymbol{q})$ which measures the utility of achieving the biomarker $\boldsymbol{q}(\tau)$ via the intervention τ. Since $\boldsymbol{q}$ is a

random vector, the typical approach is to maximize the expected utility, that is, we seek the intervention τ^* which solves

$$\tau^* = \arg\max_{\tau} \mathbb{E}\left[U(\boldsymbol{q}(\tau))\right] \tag{2.12}$$

While numerous patient-specific therapy optimization schemes have been suggested [35], to our knowledge a scheme based on Eq (2.12) has not seen practical implementation in this context. The VPP methodologies presented here may prove applicable to a future application of Eq (2.12), since the expected value in Eq (2.12) can be approximated via Eq (2.11).

An example of a quantity-of-interest that we will use to illustrate our methods is the tumor burden, which is the total number of tumor cells as calculated from $n(\boldsymbol{x}, t)$:

$$N(t) = \int_V n(\boldsymbol{x}, t) \, d\boldsymbol{x} \tag{2.13}$$

If the goal is to predict $N(t)$ for some t in the future, a VPP allows one to assess the uncertainty in this value as a result of uncertainty in the patient-specific parameter $\boldsymbol{\beta}_j$ being 'propagated' through the models $\mathcal{F}$ and $\mathcal{M}$ [25, 30]. An illustration of uncertainty propagation is demonstrated in Figure 2, where several choices of the population parameter $\boldsymbol{\theta}_p$ are chosen and the resulting (estimated) probability distribution of $N(100)$ is shown. Further examples are given in section 5 below. Another possible QoI, the integrated log-kill, is similar to Eq (2.13) but with the integrand $n(\boldsymbol{x}, t)$ replaced by $\ln(n(\boldsymbol{x}, t)/n_0(\boldsymbol{x}))$. This metric is derived from a mass action drug effect and is applicable in chemotherapy evaluation [5]. Note that we do not discuss the important topic of exactly which scalar QoIs have real clinical impact in terms of patient outcomes; refer to e.g., [33, 34].
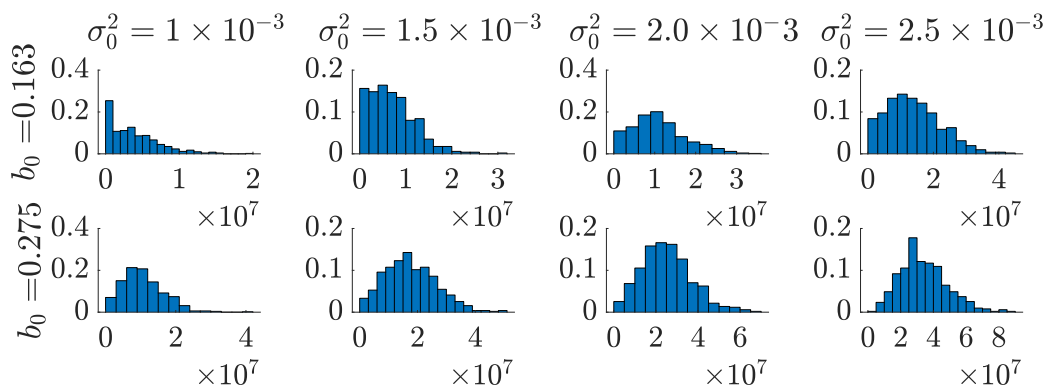
Figure 2. Demonstration of the influence of the random field parameters $\boldsymbol{\theta}_p$ on the distribution of the biomarker (2.13). The parameters (b_0, σ_0^2) controlling the statistics of $\rho(\boldsymbol{x}) \sim LB(200, b_0, \sigma_0^2)$ are changed, and for each pair, a virtual biomarker sample of size $J = 512$ is computed with $q = \mathcal{M}^{(v)}(\boldsymbol{n}^{(v)}) = N(100)$. Relative frequency histograms are displayed for each virtual biomarker sample.

3. Molecular emission imaging data

In a pure simulation setting, a virtual biomarker sample of the form shown in Eq (2.11) can be computed without reference to any particular patient-specific data, so long as the corresponding

probability distributions are well-calibrated. However, the model-based precision medicine context requires that patient-specific data be used to create individualized VPPs, hence making uncertainty-informed decisions about patient j possible. We will focus on the case of molecular Emission Computed Tomography (ECT) data, for reasons discussed below.

3.1. Preclinical and clinical imaging modalities

In the clinic, image data typically takes the form of reconstructed tomographic imaging studies such as MRI, CT, or [18]F-FDG Positron Emission Tomography (PET). Note that while it can be difficult to precisely relate these standard tomographic data to either n_j or β_j, it is nonetheless possible to postulate such relationships and thereby extract useful information from these images. An example of such an approach is the usage of Apparent Diffusion Coefficient (ADC) maps, generated using Diffusion-Weighted MRI (DWMRI), to estimate n_j through the application of a model which relates tissue cellularity to ADC [3,36,37]. The basic difficulty in such efforts is that malignancy is not directly measured by any of these standard techniques: For example, ADC maps estimate the apparent diffusion coefficient of water molecules and CT measures X-Ray attenuation, neither of which is uniquely altered by the presence of a malignant tumor. While increased [18]F-FDG uptake implies increased cellular metabolism along the glycolytic pathway, a byproduct of malignant growth, this effect is not unique to malignant cells, since benign tumors can also display an increased metabolic rate [38].

In this article we consider a wide class of imaging techniques known as Emission Computed Tomography (ECT), broadly defined as any imaging modality (such as PET) which detects a molecular tracer distribution via optical or nuclear detection techniques. We choose to focus on ECT data because the wide variety of available tracers allows for more direct modeling relating these data to tumor cell density n_j and the RDE coefficient fields β_j. As an example, genetically engineered cancer cells can be made to express Green Fluorescent Protein (GFP), which acts as a photon activity distribution when excited [39], and thus ECT, typically in the form of intravital microscopy (e.g., window chambers [40]) can be used to image n_j directly in vivo, at least in the small animal setting. Meanwhile, the growth factor ρ_j can potentially be measured using the PET tracer [18]F-FLT, which is a radioactive version of the molecule thymadine used in DNA synthesis [41, 42]. Ex vivo, the Ki-67 immunohistochemistry stain would typically be employed to measure proliferation, but this requires either a biopsy or sacrifice of the animal. Current progress in small animal imaging is working towards producing in vivo imaging agents for common immunohistochemistries, but to our knowledge an in vivo Ki-67 analog is not available. The other advantages of ECT are the possibility to perform polyscopic imaging (that is, measurement of several distinct physiological processes simultaneously) [43], and the wealth of literature on the statistics of raw and reconstructed ECT data [9].

To employ the statistical estimation strategies discussed in section 4, we require statistical models for ECT data. We emphasize the case of raw, un-reconstructed ECT data because the statistics are better understood: Most reconstructed images produced in the clinic are generated by a proprietary 'black box' method provided by the imaging system manufacturer. Our framework can still be applied to such data by carefully modifying the $\mathcal{H}$ operator defined in Eq (3.3) to account for the reconstruction step, however it can be difficult to calibrate an appropriate noise model for likelihood-based estimation. If the MLEM algorithm is known to have been employed for reconstruction, see [44] for a discussion of the resulting noise characteristics.

3.2. Statistics of ECT data

In an ECT system, a chemical tracer is engineered to emit some measurable particle, such as photons, charged particles or nuclear decay products [45]. A tracer corresponds to a particle activity distribution $f_j(\boldsymbol{x}, t)$, having units of particles emitted per unit volume per unit time. For the case of GFP imaging discussed above, a reasonable model is $f_j^{(GFP)}(\boldsymbol{x}, t) \propto n_j(\boldsymbol{x}, t)$, that is, the photon activity is proportional to the actual tumor cell density, with constant of proportionality related to the photon yield per cell (i.e., the number of photons detected, on average, per cell). Similarly, we can assume $f_j^{(^{18}\text{F-FLT})} \propto \boldsymbol{\rho}_j$, the growth factor. The energy produced by $f_j(\boldsymbol{x}, t)$ propagates through the tissue medium and is detected by specialized ECT imaging hardware, which can range from optical CCD and CMOS detectors to scintillation-type detectors for high-energy particles. The statistics of raw ECT data are discussed at length elsewhere [9, 46], so we summarize by noting that two data formats are commonly available, known as binned-mode data and particle-processed data. Because the presentation of particle-processing requires a more technical discussion of imaging detectors, we only consider binned-mode data in this work; see e.g., [45–47] for a discussion of particle processing. In a future work we will address the usage of particle processing data in mathematical oncology.

Binned-mode ECT data arises when detected particles are sorted into spatial bins corresponding to a discretization of the detector. These spatial bins need not correspond to physical detector pixels, which many ECT imaging detectors do not have. The physics-based statistics of particle counting nearly always lead to a Poisson data model for particle count data, which can be written in terms of a noisy linear functional of the activity f_j as follows. Suppose that the imaging system has M detector bins, with data for patient j labeled $g_{j1}, \ldots, g_{jM}$. The statistical model for binned-mode ECT data is that $\boldsymbol{g}_j$ is a Poisson random vector with mean

$$\bar{g}_{jm} = \int_0^T \int_V h_m(\boldsymbol{x}, t) f_j(\boldsymbol{x}, t) \, d\boldsymbol{x} dt \tag{3.1}$$

where $h_m(\boldsymbol{x}, t)$ is called the m-th sensitivity function of the imaging system (assumed to be fixed for all patients), $f_j(\boldsymbol{x}, t)$ is the particle activity distribution corresponding to the tracer administered to patient j, and T is the total imaging time. For tomographic systems, $h_m(\boldsymbol{x}, t)$ must account for both the propagation within the patient, typically following a Radiative Transport Equation (RTE), as well as the geometry and blur characteristics of the imaging detector [6, 9]. For this work, we will consider a simplified imaging geometry, intended to model intravital microscopy techniques such as mouse window chambers, where optical and/or gamma-ray images are collected with a planar detector above a pseudo-2D activity distribution. To a reasonable approximation, a Gaussian blur model

$$h_m(\boldsymbol{x}, t; A, \sigma_{blur}^2) = \frac{A}{2\pi\sigma_{blur}^2} \exp\left(-\|\boldsymbol{x} - \boldsymbol{x}_m\|^2 / 2\sigma_{blur}^2\right) \tag{3.2}$$

is then appropriate. This blur accounts for both imaging elements such as lenses and collimators as well as position estimation blur [9, 46]. More sophisticated imaging system models, for instance 3D tomographic systems, will be discussed in a future work (see section 6).

Expressing Eq (3.1) as a linear continuous-to-discrete operator $\mathcal{H} : \mathbb{X} \to \mathbb{R}^M$ allows us to write a binned-mode system model succinctly as

$$\boldsymbol{g}_j = \bar{\boldsymbol{g}}_j + \boldsymbol{\eta}_j = \mathcal{H} f_j + \boldsymbol{\eta}_j, \tag{3.3}$$

where $\bar{g}_j = [\bar{g}_{j1}, \ldots, \bar{g}_{jM}]^\top = \mathcal{H}f_j \in \mathbb{R}^M$ is defined via Eq (3.1). Considering f_j as a parameter, the probability distribution for the count vector g_j is thus

$$P(g_j|f_j) = \prod_{m=1}^{M} \frac{\exp(-\bar{g}_{jm})\bar{g}_{jm}^{g_{jm}}}{g_{jm}!} = \prod_{m=1}^{M} \frac{\exp\left(-(\mathcal{H}f_j)_m\right)(\mathcal{H}f_j)_m^{g_{jm}}}{g_{jm}!} \tag{3.4}$$

If f_j has been parameterized using a finite-dimensional parameter θ_j, for instance via a synthesis map of the form shown in Eq (2.8), we can define a nonlinear function $H : \mathbb{R}^N \to \mathbb{R}^M$, where for image reconstruction typically $N < M$, as follows:

$$H(\theta_j) = \mathcal{H}\Phi(\theta_j) \tag{3.5}$$

This function computes the mean image data $\bar{g}_j$ for the synthesized object $f_j = \Phi(\theta_j)$, and is implemented computationally via numerical quadrature methods applied to Eq (3.1). The component functions of H are denoted $H_m(\theta_j) : \mathbb{R}^N \to \mathbb{R}$. Note that if the synthesis method is nonlinear, as in Eq (2.8), then H is a nonlinear function of the vector θ_j, while if Φ is linear, then H is linear. In either case, we can write a probability distribution for $g_j|\theta_j$ as follows:

$$P(g_j|\theta_j) = \prod_{m=1}^{M} \frac{\exp\left(-H_m(\theta_j)\right)H_m(\theta_j)^{g_{jm}}}{g_{jm}!} \tag{3.6}$$

We will make use of Eqs (3.4)–(3.6) for likelihood-based estimation methods in section 4 below.

4. Statistical estimation and inversion techniques

Assume that for patient j, ECT data $g_j \in \mathbb{R}^M$ has been collected, where g_j can either be monoscopic (a single imaging study) or polyscopic (consisting of multiple imaging studies, corresponding to multiple tracers). In the previous section, we provided statistical models for g_j that relate it to the activity distribution f_j via Eq (3.3). As discussed there, ECT tracers allow f_j to be directly related to n_j and ρ_j. In this section, we discuss statistical techniques for estimating f_j (hence n_j and ρ_j) from g_j. In section 4.1, we discuss point estimation techniques for an individual patient, where a single estimate of f_j is produced, then in section 4.2 we discuss patient-specific Bayesian strategies where the solution consists of samples from a posterior distribution. In section 4.3, we discuss methods for estimating a population parameter θ_p from a database $G = [g_1, \ldots, g_J]$ of images collected for J patients.

4.1. Patient-specific maximum likelihood parameter estimation

Given data $g_j \in \mathbb{R}^M$ for patient j, we consider the estimation of the underlying activity distribution f_j using the imaging system model $g_j = \mathcal{H}f_j + \eta_j$, given in Eq (3.3). To approach this infinite-dimensional inverse problem, we select a finite-dimensional parameterization $f_j = \Phi(\theta_j)$, where $\theta_j \in \mathbb{R}^N$, and consider the problem of estimating θ_j from g_j, that is, solving the inverse problem $g_j = H(\theta_j) + \eta_j$. Note that moving from f_j to θ_j might introduce an approximation error, owing to the fact that f_j need not be expressible in same form as the synthesis used for the reconstruction; in ECT imaging, f_j is typically quite smooth, so this approximation error is typically insignificant when compared to

the influence of detector resolution and noise. While typical image reconstruction algorithms assume a voxel-based parameterization of f_j, we elect to use a lumpy-type synthesis of the form given in Eq (2.8). As discussed, the activity distribution can correspond to the cell density n_j, or the growth factor $\rho_j(x)$. We denote a generic estimate of θ_j produced from g_j by a hat, i.e., $\hat{\theta}_j$ is an estimate of θ_j produced from the data g_j. While a wide array of strategies exist to solve this inverse problem, we consider only likelihood-based methodologies.

Recall from section 3.2 that the probability distribution for binned-mode ECT data is Poisson. For an activity distribution f_j parameterized using the synthesis map Φ, i.e. $f_j = f_j(\theta_j)$, we thus have a likelihood function

$$\mathcal{L}(\theta_j|g_j) = P(g_j|\theta_j) = P(g_j|f_j(\theta_j)) = P(g_j|\Phi(\theta_j)) = \prod_{m=1}^{M} \frac{\exp(-H_m(\theta_j))H_m(\theta_j)^{g_{jm}}}{g_{jm}!} \tag{4.1}$$

where $P(g_j|f_j)$ is the Poisson distribution given in Eq (3.4), and $H_m(\theta_j)$ are the components of the function defined in Eq (3.5). Note the abuse of notation in Eq (4.1): The two functions $P(g_j|\theta_j)$ and $P(g_j|f_j)$ have different parameter types, but we take Eq (4.1) to define the former in terms of the latter. Taking its logarithm (and ignoring the θ_j-independent constant) results in the log-likelihood function

$$\ell(\theta_j|g_j) = \ln(\mathcal{L}(\theta_j|g_j)) = \sum_{m=1}^{M} \left[g_{jm} \ln(\bar{g}_{jm}) - \bar{g}_{jm} \right] = \sum_{m=1}^{M} \left[g_{jm} \ln(H_m(\theta_j)) - H_m(\theta_j) \right] \tag{4.2}$$

We now define the Maximum Likelihood Estimate (MLE) as the maximum of either Eq (4.1) or Eq (4.2) over the set Θ of allowed parameters:

$$\hat{\theta}_j^{(ML)} = \arg\max_{\theta\in\Theta} \mathcal{L}(\theta|g_j) = \arg\max_{\theta\in\Theta} \ell(\theta|g_j) = \arg\min_{\theta\in\Theta} \left[-\ell(\theta|g_j) \right] \tag{4.3}$$

A variety of optimization algorithms exist to compute $\hat{\theta}_j^{(ML)}$, including the Expectation Maximization (EM) algorithm and Newton-type methods [48, 49]. If the imaging operator $H(\theta_j)$ is linear, the EM algorithm applied to Poisson data leads to an iterative method known in the imaging community as the MLEM algorithm. Starting from $\hat{\theta}_j^{(0)}$, the MLEM algorithm performs the iterative update

$$\hat{\theta}_{jn}^{(k+1)} = \frac{\hat{\theta}_{jn}^{(k)}}{s_n} \sum_{m=1}^{M} H_{mn} \frac{g_m}{\left(H\hat{\theta}_j^{(k)}\right)_m}, \quad s_n = \sum_{m=1}^{M} H_{mn} \tag{4.4}$$

We will employ the method (4.4) to compute some MLEs for linear imaging models in section 5, while for nonlinear models we use the Matlab algorithm `fmincon`, which is a based on an interior point method with BFGS quasi-Newton steps [50]. A demonstration of algorithm (4.4) applied to a reconstruction of $n_j(x, t)$ from image data of the form (3.1) is shown in Figure 3. Note that the reconstruction displays visual artifacts due to the function $H(\theta)$ having a nontrivial null space. We have found that these artifacts have minimal impact on the task at hand, which is to predict the progression of $n(x, t)$ and $N(t)$. Future work will analyze in more detail the question of task-based image quality, as defined in [9, 51], in the mathematical oncology context.

We note briefly that under certain technical conditions, an MLE is asymptotically consistent and efficient, meaning for well-specified likelihood models, the MLE converges to the true parameter and
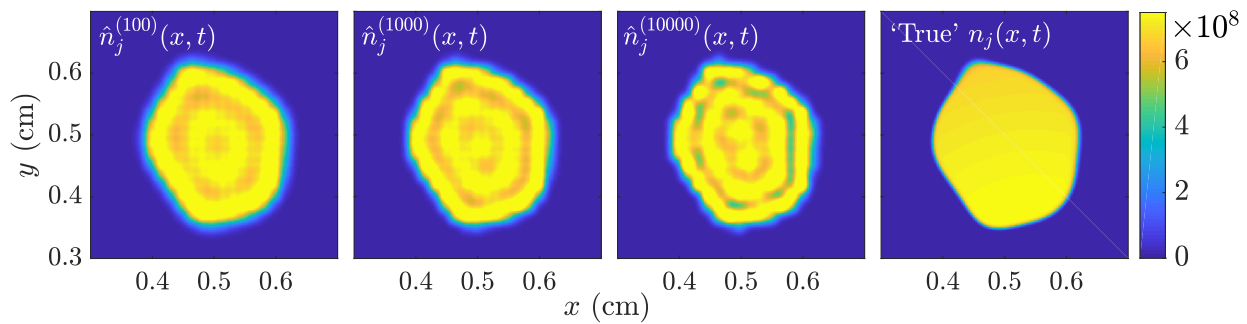
Figure 3. Illustration of the MLEM image reconstruction algorithm (4.4) for $k = 10^2, 10^3$ and 10^4 iterations. The true $n_j(x, t)$ shown on the right. The visible artifacts are due to the fact that $\mathcal{H}$ has a nontrivial null space which is not accounted for in the MLE (4.3). Regularization strategies can alleviate this, though we have found that for the task of predicting tumor burden $N(t)$, these artifacts have minimal impact.

attains the Cramér-Rao lower bound. Asymptotically, the distribution of the estimate $\hat{\boldsymbol{\theta}}_j^{(ML)}$ (as a random variable depending on the data $\boldsymbol{g}_j$) converges to a normal with covariance given by the inverse of the Fisher information [9]:

$$\sqrt{M}\left(\hat{\boldsymbol{\theta}}_j^{(ML)} - \boldsymbol{\theta}_j\right) \to \mathcal{N}(0, \boldsymbol{F}^{-1}), \quad F_{nn'} = -\mathbb{E}_{\boldsymbol{g}_j}\left(\frac{\partial^2 \ell(\boldsymbol{\theta}_j|g_{jm})}{\partial \theta_{jn} \partial \theta_{jn'}}\right) \tag{4.5}$$

The Fisher information matrix $\boldsymbol{F} = \boldsymbol{F}(\boldsymbol{\theta}_j)$ can be used to design more efficient measurement systems by 'tuning' the design parameters of the system to modify the likelihood in a desirable manner [52]. Note also that the Fisher information matrix can be evaluated 'off-line', by computing the expected values in Eq (4.5), which allows one to analyze the uncertainty in the MLE in a frequentist setting. While $\boldsymbol{F}$ depends on the true $\boldsymbol{\theta}_j$, for many ECT systems this dependence is minimal (i.e., the Fisher information matrix is nearly constant with respect to $\boldsymbol{\theta}_j$), which allows one to compute a single (system-dependent) $\boldsymbol{F}$ to be used for all MLEs derived from the system. One can also estimate confidence intervals for scalar quantities of interest derived from estimates of $\boldsymbol{\theta}_j$ using transformation rules for the Fisher information matrix [53].

To use the MLE procedure to generate a patient-specific virtual population, we assume that $\boldsymbol{g}_j = \mathcal{H}\boldsymbol{f}_j$, where either $\boldsymbol{f}_j = n_j(\boldsymbol{x}, t_0)$, $\boldsymbol{f}_j = \boldsymbol{\rho}_j$, or $\boldsymbol{f}_j = [n_j(\boldsymbol{x}, t_0), \rho_j(\boldsymbol{x})]^T$, that is, we image either the cell density at a single time point, or the growth factor, or both the cell density and the growth factor. From this data, we take the following steps:

(i) Use MLE to obtain $\hat{\boldsymbol{\theta}}_j^{(n)}$ (and $\hat{\boldsymbol{\theta}}_j^{(\rho)}$, if measured) from which we synthesize $\hat{\boldsymbol{n}}_j$ (and $\hat{\boldsymbol{\rho}}_j$, if measured).

(ii) Randomize the remaining parameters $\boldsymbol{D}$, $\boldsymbol{\kappa}$ (and $\boldsymbol{\rho}$, if not measured) using a well-calibrated random field model such as Eq (2.8), resulting in $\boldsymbol{D}_j^{(1)}, \ldots, \boldsymbol{D}_j^{(J)}, \boldsymbol{\kappa}_j^{(1)}, \ldots, \boldsymbol{\kappa}_j^{(J)}$ (and $\boldsymbol{\rho}_j^{(1)}, \ldots, \boldsymbol{\rho}_j^{(J)}$, if required).

(iii) Generate a patient-specific virtual population by using $\hat{\boldsymbol{n}}_j$ as the initial condition, either $\hat{\boldsymbol{\rho}}_j$ or $\boldsymbol{\rho}_j^{(1)}, \ldots, \boldsymbol{\rho}_j^{(J)}$ for the growth factor, and $\boldsymbol{D}_j^{(1)}, \ldots, \boldsymbol{D}_j^{(J)}, \boldsymbol{\kappa}_j^{(1)}, \ldots, \boldsymbol{\kappa}_j^{(J)}$ for the diffusion coefficient and carrying capacity, and solving the forward problem (2.5) for each sample.

(iv) With the resulting VPP, compute any desired quantitative biomarkers as in Eq (2.11).

Note that in the above procedure, we make no attempt to estimate D and κ by performing the nonlinear parameter estimation task implied by Eq (2.4), nor do we attempt to estimate $n(x, t)$ for any $t < t_0$. From single time point data such as this, an identifiability (i.e., uniqueness) issue arises preventing such estimation from succeeding. This procedure is used in sections 5.1 and 5.2 to produce virtual populations.

Note that one issue that is not addressed in the above technique is uncertainty inherent in the noisy data g_j, namely that g_j being random means the MLE is also random. While Fisher information analysis can be employed to analyze this uncertainty in an off-line frequentist setting, this analysis does not provide a clear method for producing alternates to $\hat{\theta}_j$, since one only has access to a single realization of g_j. The Bayesian methodology discussed below provides an avenue towards producing many plausible estimates from a single realization of g_j.

4.2. Patient-specific Bayesian solution

In a patient-specific Bayesian solution to the inverse problem $g_j = \mathcal{H}f_j + \eta_j$, we assume that f_j were sampled from a prior distribution $\mathbb{P}_0(f)$. Note that $\mathbb{P}_0(f)$ need not correspond to a well-calibrated model of the true population in order to proceed with the inference, but the method is typically much more reliable if it does. Then, since g_j is statistically coupled to f_j through Eq (3.3), we can consider the conditional random vector $g_j|f_j$, whose statistics are described by the likelihood function $\mathcal{L}(f_j|g_j) = \mathbb{P}(g_j|f_j)$ as discussed previously. A Bayesian inverse solution to the inference problem is to consider the posterior random vector $f_j|g_j$ instead of a single point estimate $\hat{f}_j$. Samples from the posterior can be used to generate a virtual population as described in the procedure above, where now $\hat{n}_j$ and $\hat{\rho}_j$ are replaced by samples from their corresponding posteriors to account for uncertainty in the data g_j.

Bayesian personalization of tumor growth has been considered in [21, 54], for example. The former considers an RDE similar to Eq (2.1) but requires two segmented MRI images to perform the personalization. The latter generates ensemble members (i.e., virtual patients) by using fixed RDE coefficients (β in our notation) and imposing an ad hoc additive Gaussian noise model on the system evolution, in order to leverage the ensemble Kalman filter technique.

The distribution of the posterior $f_j|g_j$ is given via Bayes rule, which states informally that $\mathbb{P}_{post}(f_j|g_j) \propto \mathbb{P}_0(f)\mathcal{L}(f_j|g_j)$. More rigorously [55], we would say that

$$\frac{d\mathbb{P}_{post}}{d\mathbb{P}_0} \propto \mathcal{L}(f_j|g_j) = \exp\left(-\ell(f_j|g_j)\right)$$

Practically speaking, in this paper we consider (via Eqs (2.7) and (2.8)) finite-dimensional parameterizations of f, so that when coupled with binned-mode ECT data, the result is a finite-dimensional Bayesian inverse problem [56]. Taking as before a lumpy-type parameterization $f_j = \Phi(\theta_j)$, and with the data described by the nonlinear function $\bar{g}_j = H(\theta_j) = \mathcal{H}\Phi(\theta_j)$ defined in section 3.2, we now assign a prior on the vector θ_j, which we assume has PDF $p_0(\theta_j)$. The Poisson likelihood $\mathcal{L}(\theta_j|g_j)$ is given in Eq (3.6), leading to a description of the posterior as

$$p(\theta_j|g_j) \propto p_0(\theta_j)\mathcal{L}(\theta_j|g_j) \tag{4.6}$$

If either the prior is infinite-dimensional or the ECT data is assumed to be of particle-processing form, a fully infinite-dimensional Bayesian approach must be considered [55].

To draw samples from the posterior (4.6), we apply a Markov Chain Monte Carlo (MCMC) algorithm based on the one presented in [57]. First, we assume for each tracer f_j an expansion of the form Eq (2.8), so that $f_j = \Phi(\theta_j)$ with $\theta_j = [b_1, \ldots, b_{L_{max}}, x_1, \ldots, x_{L_{max}}, \gamma_1, \ldots, \gamma_{L_{max}}]$. Note that in some cases, we reduce the computational burden by fixing the x_l and/or γ_l and only sample the remaining parameters. The MCMC procedure is then as follows:

(i) Given the current sample $\theta_j^{(n)}$, we propose a new sample $\theta'_j \sim \pi\left(\theta'_j | \theta_j^{(n)}\right)$ by randomly selecting L_{mcmc} lumps and perturbing their amplitudes b_l, location x_l and shape γ_l by a Gaussian with diagonal covariance Σ. This proposal is symmetric.

(ii) From θ'_j, the activity f'_j is synthesized using Eq (2.8) and the proposed mean image $\bar{g}' = \mathcal{H}f'_j$ is simulated using Eq (3.5).

(iii) The acceptance probability is computed using the Metropolis-Hastings rule [25]:

$$Q(\theta'_j, \theta_j^{(n)}) = \min\left\{1, \frac{p(\theta'_j | g_j)}{p(\theta_j^{(n)} | g_j)}\right\} = \min\left\{1, \frac{p_0\left(\theta'_j\right)\mathcal{L}\left(\theta'_j | g_j\right)}{p_0\left(\theta_j^{(n)}\right)\mathcal{L}\left(\theta_j^{(n)} | g_j\right)}\right\},$$

after which we accept $\theta_j^{(n+1)} = \theta'_j$ with probability $Q(\theta'_j, \theta_j^{(n)})$, and take $\theta_j^{(n+1)} = \theta_j^{(n)}$ with probability $1 - Q(\theta'_j, \theta_j^{(n)})$.

By construction of the Metropolis-Hastings algorithm, the Markov chain samples $\theta_j^{(1)}, \theta_j^{(2)}, \ldots$ have stationary distribution given by the posterior $p(\theta_j | g_j)$, and can thus be used to synthesize virtual patients for the calculation of any quantitative biomarkers as discussed in section 2.4. An example of this approach is discussed in section 5.1. Note that the standard Metropolis-Hastings algorithm performs best for θ with relatively small dimension; various modifications for high-dimensional problems have also been proposed [25, 30, 55].

4.3. Calibration of population distributions

As illustrated in Figure 2, the population parameter θ_p which specifies the random field models used to generate the virtual population has a substantial influence on the statistics of any biomarkers calculated using Eq (2.11). The question of how to estimate θ_p from data thus arises. Such calibration problems in the context of VCTs have been discussed in [58, 59], but their methodology accounts only for ODE models arising in systems pharmacology and does not directly apply to our context of PDE tumor growth models, random fields and imaging data. In this section, we discuss briefly a technique to estimate the population parameters which specify a lumpy-type model (2.8) from a database of ECT images $G = [g_1, \ldots, g_J]$.

Suppose we assume a lumpy-type model for each of the RDE coefficient fields, $\beta = [D, \rho, \kappa]$, say each is of the form given in Eq (2.8) with $\ell(x)$ given by a Gaussian (2.9). While the lump function $\ell(x)$ can itself be the target of calibration (see [60], for instance), we assume Eq (2.9) for simplicity. The statistics of the random field $LB(\bar{L}, b_0, \sigma^2)$ are thus fully specified by the three scalars $\bar{L}, b_0$ and σ^2, which we write as $\theta_p = [\bar{L}, b, \sigma^2] \in \Theta = (0, \infty)^3$. The problem of calibrating a population distribution such as this is to take a database of noisy image data $G = [g_1, \ldots, g_J]$ for a population of J patients and produce an estimate $\hat{\theta}_p = \left[\widehat{\bar{L}}, \hat{b}, \widehat{\sigma^2}\right]$. We can again apply either MLE or the Bayes approach as discussed in sections 4.1 and 4.2.

To develop a likelihood model for $\boldsymbol{\theta}_p$, we consider first the generic case where $\boldsymbol{g}_j$ consists of direct images of some activity $\boldsymbol{f}_j$, for instance $\boldsymbol{f}_j \propto \boldsymbol{\rho}_j$ as discussed in section 5.1. This case corresponds to that considered in [61]. Assuming inter-subject independence, we have

$$\mathcal{L}(\boldsymbol{\theta}_p|\boldsymbol{G}) = p(\boldsymbol{G}|\boldsymbol{\theta}_p) = \prod_{j=1}^{J} P(\boldsymbol{g}_j|\boldsymbol{\theta}_p) \tag{4.7}$$

Using the law of total probability and the fact that the distribution of the image data is independent of the object statistics, we can write

$$P(\boldsymbol{g}_j|\boldsymbol{\theta}_p) = \mathbb{E}_{\boldsymbol{f}_j|\boldsymbol{\theta}_p}\left[P(\boldsymbol{g}_j|\boldsymbol{f}_j)\right] \tag{4.8}$$

In (binned-mode) ECT imaging, the probability distribution $P(\boldsymbol{g}_j|\boldsymbol{f}_j)$ is given by the Poisson (3.4), which depends on the imaging operator $\mathcal{H}$. The expression (4.8) states that the probability of observing a given image $\boldsymbol{g}_j$ if the population parameter is $\boldsymbol{\theta}_p$ is the probability of observing $\boldsymbol{g}_j$ given the object $\boldsymbol{f}_j$, averaged over of all possible random field realizations $\boldsymbol{f}_j$ with statistics specified by the parameter $\boldsymbol{\theta}_p$. The expectation in Eq (4.8) is ostensibly infinite dimensional (since the $\boldsymbol{f}_j$ are functions), though it can be approximated using Monte Carlo:

$$P(\boldsymbol{g}_j|\boldsymbol{\theta}_p) \approx \frac{1}{J'} \sum_{j'=1}^{J'} P(\boldsymbol{g}_j|\boldsymbol{f}_{j'}), \quad \boldsymbol{f}_{j'} \sim \mathbb{P}(\boldsymbol{f}_j|\boldsymbol{\theta}_p) \tag{4.9}$$

Equation (4.9) specifies a computational method for approximating the likelihood (4.7): For fixed $\boldsymbol{\theta}_p$, compute Eq (4.9) for each image in the database, then compute the product (or its logarithm) in Eq (4.7). The MLE of $\boldsymbol{\theta}_p$ is then defined as before in Eq (4.3), with $\boldsymbol{\theta}_p \in \boldsymbol{\Theta}_p$, where $\boldsymbol{\Theta}_p$ is the set of admissible population parameters. Note that the estimate (4.9) can be computed in parallel using graphics processors, which drastically reduces the time required for each likelihood calculation, which would have previously been a bottleneck in the application of a traditional MLE. Despite this, the MLE optimization problem (4.3) applied to the likelihood (4.7) poses a challenge for standard numerical optimization routines due to scaling and non-convexity issues. See section 6 for a brief discussion of future work aimed at addressing these issues.

5. Results

In this section, we discuss two in silico experiments that illustrate the methodologies discussed above. For both, a common 'ground truth' cell density path $n(\boldsymbol{x}, t)$ is generated by sampling a random field for each coefficient and simulating until $t = T = 365$ days. In the first experiment, we use (simulated) ECT data for $f_j(\boldsymbol{x}) = n(\boldsymbol{x}, 100)$ to compute both an MLE and a collection of posterior samples for the tumor cell density at $t = t_0 = 100$ days. Then, a virtual population is generated by using the resulting $\hat{\boldsymbol{n}}_j$ (or posterior samples) as the initial condition in Eq (2.1), randomizing the remaining coefficients and simulating until $t = T$. The resulting populations are then compared to the 'true' $n(\boldsymbol{x}, t)$, using the tumor burden $N(t)$ defined in Eq (2.13) as the quantity-of-interest. In the second experiment, we use ECT data and MLE for both the growth factor $\rho(\boldsymbol{x})$ and the cell density $n(\boldsymbol{x}, t_0)$, again generating a VPP by using the estimated cell density as the initial condition $n_0(\boldsymbol{x})$ and $\boldsymbol{\rho}_j = \hat{\boldsymbol{\rho}}_j$, then randomizing both $\boldsymbol{D}$ and $\boldsymbol{\kappa}$.

For both experiments, we use the following random fields, both to generate the initial 'true' tumor path $n(\boldsymbol{x}, t)$ and to randomize unmeasured coefficients. We assume $\boldsymbol{D} \sim LB(20, 1\mathrm{e}{-7}, 0.04)$, $\boldsymbol{\rho} \sim LB(200, 0.25, 0.002)$, and $\boldsymbol{\kappa} \sim LB(100, 5\mathrm{e}7, 0.1)$. These parameters were selected to achieve an average tumor size of approximately 10^8 cells after the first year; obviously a calibration procedure such as the one outlined in section 4.3 would be required to estimate population parameters for a real population of tumors. The exact sample used is shown in Figure 4.
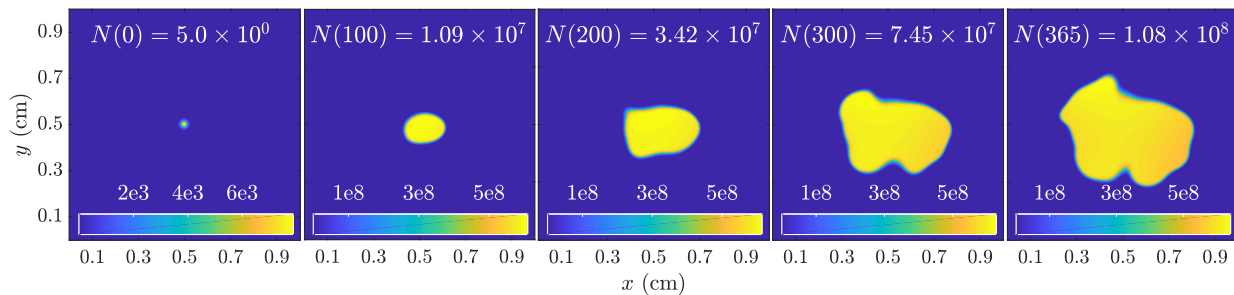

Figure 4. The 'true' tumor cell density, generated using Eq (2.1) with $\boldsymbol{\beta} = (\boldsymbol{D}, \boldsymbol{\rho}, \boldsymbol{\kappa})$ sampled from the random field models specified in the section 5 introduction above. The simulated imaging studies are performed at $t = 100$ days (second panel).

### 5.1. Experiment 1: Maximum likelihood estimation and posterior sampling of cell density $\boldsymbol{n}$

In this experiment, we suppose that $\boldsymbol{g}_j^{(n)} = \mathcal{H}^{(n)} \boldsymbol{n}_j + \boldsymbol{\eta}_j^{(n)}$ consists of ECT measurements of the cell density $n_j(\boldsymbol{x}, t_0)$, with $t_0 = 100$ days, and that no other measurements are available. We assume that the experimental setup is confined to a thin slab so that a 2D simulation described in section 3.1 is appropriate, and that the imaging system corresponds to Eq (3.1) with blur kernel given by Eq (3.2). The blur is taken to be $\sigma_{blur}^2 = 0.0212$, which corresponds to an imaging resolution (Full-Width-Half-Max, FWHM) of $500 \, \mu$m, and we assume that the photon yield is $A = 1\mathrm{e}{-3}$ (detected photons per cell). Once the mean image data $\bar{\boldsymbol{g}}_j = \mathcal{H} \boldsymbol{f}_j$ is computed using Eq (3.1), a randomized image is produced using `poissrnd`.

To perform the image reconstruction for $n_j(\boldsymbol{x}, t_0)$, we apply the linear MLEM algorithm (4.4) with a reconstruction of the type (2.8), with fixed lump centers and lump variance, so that $\boldsymbol{\theta} = [b_1, \ldots, b_{L_{max}}]$. Specifically, we fix the lump centers on a grid of size 25×25, uniformly spaced on $[0.35, 0.65]^2$, and set the lump variance to $\sigma^2 = 1\mathrm{e}{-2}/256$. The matrix $\boldsymbol{H}$ with columns given by $\bar{\boldsymbol{g}}_l = \mathcal{H}(\ell(\boldsymbol{x} - \boldsymbol{x}_l))$ is then computed using Eq (3.1), and the MLEM algorithm (4.4) applied for $K_{max} = 5000$ iterations, with initial guess $\hat{\boldsymbol{\theta}}_j^{(0)} = \boldsymbol{H}^T \boldsymbol{g}_j$. From the final $\hat{\boldsymbol{\theta}}_j^{(K_{max})}$, we synthesize the estimate $\hat{n}_j(\boldsymbol{x}, t_0)$ using Eq (2.8). The results of this reconstruction are shown in Figure 5.

To demonstrate the usage of this estimate to generate a virtual tumor population as discussed in section 4.1, we use the resulting $\hat{n}_j(\boldsymbol{x}, t_0)$ as the initial condition $n_0(\boldsymbol{x})$ in Eq (2.1), and randomize the remaining coefficient fields using the lumpy-type random field models described in the introduction to this section to produce a virtual population of size $J = 64$. In Figure 6, the quantity-of-interest $N(t)$ is shown for the resulting virtual population, while in Figure 7, a collection of 10 virtual tumors is shown for $t = 150$ (i.e., 50 days post-imaging).
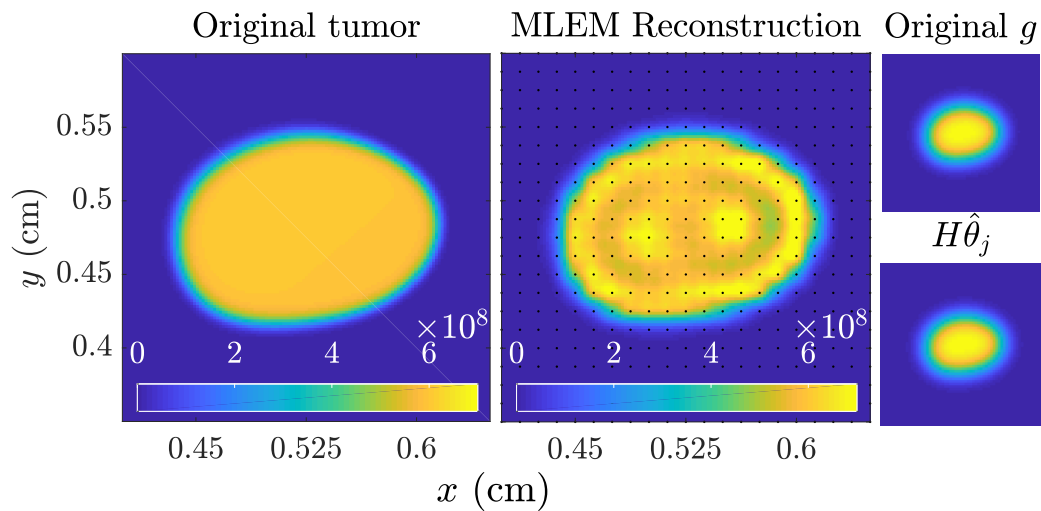
Figure 5. Comparison of the 'true' tumor shown in Figure 4 (left, imaged at $t_0 = 100$ days) to the maximum likelihood reconstruction $\hat{\boldsymbol{n}}_j^{(ML)}$ (middle). The original noisy image $\boldsymbol{g}_j = \mathcal{H}\boldsymbol{n}_j + \boldsymbol{\eta}_j$ and the mean image evaluated at the MLE $\bar{\boldsymbol{g}}_j = \mathcal{H}\hat{\boldsymbol{n}}_j^{(ML)} = \boldsymbol{H}\hat{\boldsymbol{\theta}}_j^{(ML)}$ are shown to the right. The MLEM algorithm (4.4) with 5000 iterations was employed. The reconstruction grid is shown on the middle plot (black dots). The computational time required for this reconstruction was about 30 seconds.
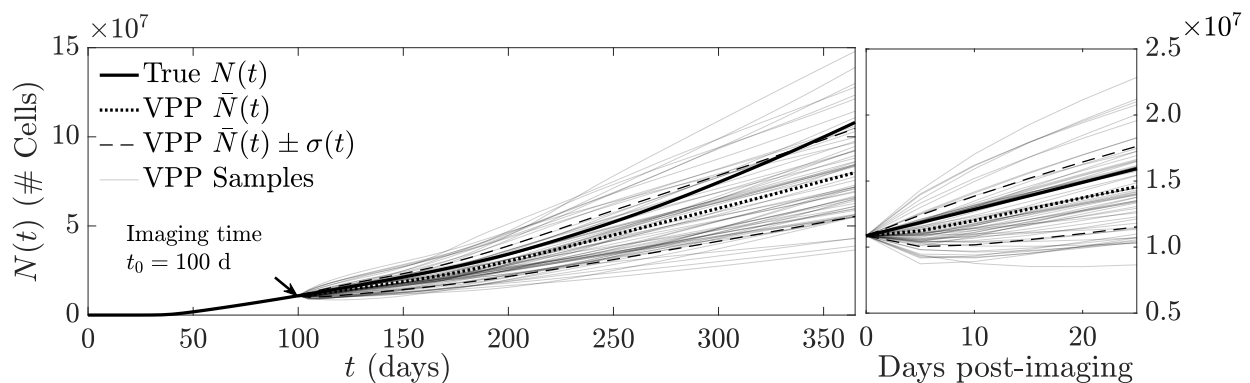
Figure 6. Illustration of a VPP-based prediction of tumor burden $N(t)$ when an MLE of $n(\boldsymbol{x}, t_0)$ is used as the initial condition and remaining coefficients are randomized ($J = 64$). The 'true' tumor path (bold), empirical average of the VPP (dotted) and a one standard deviation confidence band (dashed line) are shown. Individual VPP samples are shown in light grey. The left plot shows the time interval from 0 (tumor initiation) to 1 year, while the right shows the interval from imaging to 30 days post-imaging. The VPP appears biased towards slower trajectories than the actual.

To demonstrate the usage of the Bayesian methods discussed in section 4.2, we generated samples from the posterior $\boldsymbol{n}_j | \boldsymbol{g}_j$ to use as part of the virtual population generation procedure. For the MCMC procedure, we use the algorithm outlined in section 4.2, moving $L_{mcmc} = 3$ lumps for each proposal
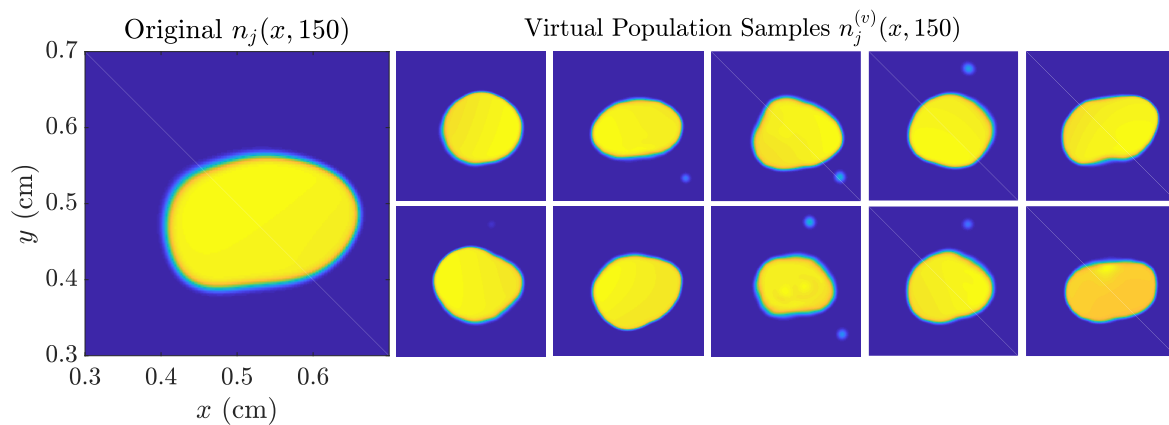
Figure 7. Example VPP Samples for $t = 150$, i.e., 50 days post-imaging, generated using the MLEM estimate of $n_j(x, 100)$ as the initial condition and randomizing the other coefficient fields as described in the main text. All plots are shown on $[0.3, 0.7]^2$ and with a common colorbar. It is apparent that the tumor morphology is largely consistent across the population, though some virtual tumors demonstrate small 'islands' of growth and the overall rate of growth is variable, as seen in Figure 6.

with a proposal distribution of $\mathcal{N}(0, \sigma_{mcmc}^2 \boldsymbol{I})$, taking $\sigma_{mcmc} = 1e4$. The prior for $\boldsymbol{\theta}_j = [b_1, \ldots, B_{L_{max}}]$ is chosen to be I.I.D. uniform on the interval $[0, 1e10]$. Starting from $\boldsymbol{\theta}_j^{(0)} = \hat{\boldsymbol{\theta}}_j^{(ML)}$, i.e., starting at the maximum likelihood estimate, we generate $J = 512$ samples, then synthesize $\boldsymbol{n}_j^{(1)}, \ldots, \boldsymbol{n}_j^{(J)}$. Note that the subscript remains j as these samples correspond to patient j, while the superscript indicates that these are posterior samples, generated using MCMC. Starting at the maximum likelihood estimate (which for our prior corresponds to the MAP estimate) reduces the need for an expensive burn-in period in which the first J_0 samples are thrown out to ensure that the chain is sampling from the posterior. The acceptance rate for our chain (number of proposals accepted divided by total number of samples) is 0.495. Using each posterior sample as an initial condition, we again generate a virtual population by solving Eq (2.1) with the remaining coefficients $\boldsymbol{D}, \rho, \kappa$ randomized according to the lumpy-type field models specified previously. The results of this virtual population are shown in Figure 8.

For both the maximum likelihood and Bayesian solutions, the resulting virtual populations underestimate, on average, the true tumor burden both over the short- and long-term, which indicates that additional data may be needed to estimate other model coefficients to gain a more accurate patient-specific prediction; we discuss on possibility in section 5.2 below. While the Bayesian technique accounts for uncertainty inherent in the image data $\boldsymbol{g}_j^{(n)}$, this uncertainty appears relatively small as it relates to the prediction of $N(t)$; the variance of $N(t)$ for the virtual population generated using the Bayesian technique is only slightly larger than the variance for the MLE population.

5.2. *Experiment 2: Maximum likelihood estimation of both $\boldsymbol{n}$ and ρ*

In Experiment 1, we saw that the virtual population under-estimated the true tumor path in both the MLE and Bayesian examples, suggesting that additional data may be required. In this experiment, we assume that the image data $\boldsymbol{g}_j$ consists of (simulated) polyscopic data of the following form. We assume that we have access first to cell density measurements of the form $\boldsymbol{g}_j^{(n)} = \mathcal{H}^{(n)} \boldsymbol{n}_j + \boldsymbol{\eta}_j^{(n)}$, where $\mathcal{H}^{(n)}$ is
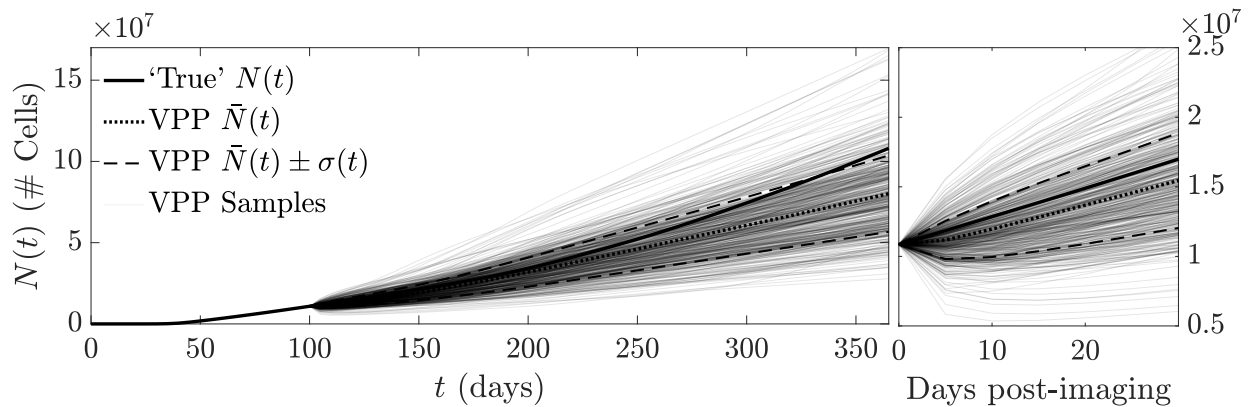
Figure 8. Illustration of the virtual population generated using Bayesian posterior sampling for the initial condition $\boldsymbol{n}_0(\boldsymbol{x})$. The empirical average path $\bar{N}(t)$ (dotted) and confidence band (dashed) are shown in comparison to the true tumor path (solid), as well as the $J = 512$ paths constituting the virtual population (light grey).

the same as in Experiment 1. The second dataset is of the form $\boldsymbol{g}_j^{(\rho)} = \mathcal{H}^{(\rho)}\boldsymbol{\rho}_j + \boldsymbol{\eta}_j^{(\rho)}$, where we now have $\sigma_{blur}^2 = 0.0425$ (1 mm FWHM). As in the first experiment, we assume that imaging is performed at $t = t_0 = 100$. The image reconstruction for $\boldsymbol{n}_j$ is performed using MLEM as in experiment 1 (Figure 5). For $\rho_j(\boldsymbol{x})$, we use MLE with a nonlinear synthesis map as follows. The lumpy-type expansion (2.8) with Gaussian lump (2.9) is assumed where now $\boldsymbol{\theta}_j = [\boldsymbol{x}_1, \ldots, \boldsymbol{x}_{L_{max}}, b, \sigma^2]$, that is, we fit the lump positions and a common lump amplitude and variance. The result is a reconstruction of the form

$$\hat{\boldsymbol{\rho}}_j = \hat{b} \sum_{l=1}^{L_{max}} \exp\left(-\frac{1}{2\hat{\sigma}^2}\|\boldsymbol{x} - \hat{\boldsymbol{x}}_l\|^2\right),$$

where we have chosen $L_{max} = 60$ for this experiment. We assume further the constraints that $\boldsymbol{x}_l \in [0,1]^2$, $b \in (0, \infty)$, and $\sigma^2 \in (0, 0.1]$, and use the Matlab algorithm `fmincon` to perform the MLE optimization (4.3). The results of this reconstruction are shown in Figure 9.

In Figure 10, we compare the tumor burden predicted using the virtual population to the 'true' tumor path. Over the first 30 day post-imaging period, the ensemble mean $\bar{N}(t)$ for the virtual population tracks the true tumor path closely, indicating that if both $\boldsymbol{n}$ and $\boldsymbol{\rho}$ are estimated, the result appears unbiased over this short time period. This is also demonstrated by looking at the individual virtual population samples in Figure 11: Visual inspection reveals a closer overall appearance to the 'true' tumor than was apparent in Figure 7. Over a longer time period, the ensemble still trends towards a lower overall growth rate than the 'true' tumor path, as seen previously in Figure 6.

6. Discussions

In this article, we have presented a rigorous, end-to-end statistical framework for addressing both inter- and intra-patient heterogeneity for the commonly used reaction-diffusion tumor growth model (2.1), emphasizing the usage of raw molecular emission imaging data to estimate unknown model parameters. We demonstrated these methods in two in silico experiments, using a simulated
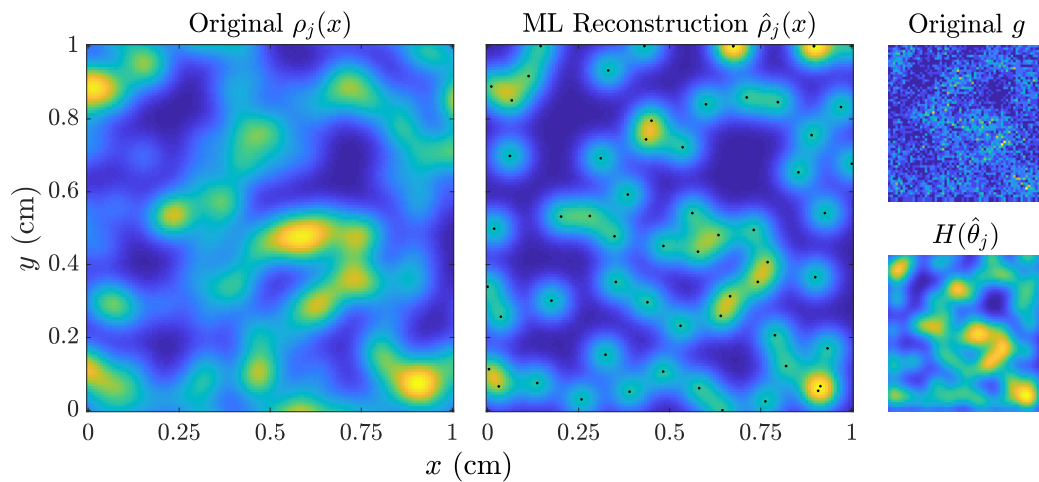
Figure 9. Maximum likelihood reconstruction of the growth factor $\boldsymbol{\rho}_j$ (left) from simulated imaging data (top right). The reconstruction (middle) with the resulting lump centers highlighted (black dots) are shown, as well as the mean image at the MLE $\bar{\boldsymbol{g}}_j^\rho = \mathcal{H}^{(\rho)}\hat{\boldsymbol{\rho}}_j$ (bottom right). The computational time for this reconstruction was several hours. While the reconstruction may appear visually poor, it is sufficient for the task of predicting $n(\boldsymbol{x}, t)$ and $N(t)$, as illustrated in Figures 10 and 11.
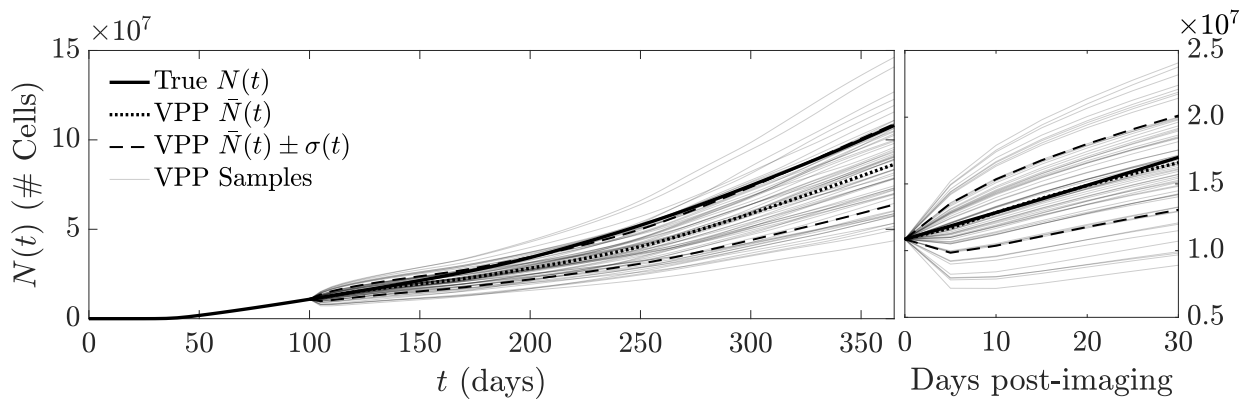
Figure 10. Prediction of tumor burden $N(t)$ when MLEs of $n(\boldsymbol{x}, t_0)$ and $\rho(\boldsymbol{x})$ are used for the initial condition and growth factor. $D(\boldsymbol{x})$ and $\kappa(\boldsymbol{x})$ are randomized to produce a VPP ($J = 64$). The original tumor path (bold) and empirical average of the VPP (dotted), as well as a one-σ confidence band (dashed) and individual VPP samples (light grey) are shown. The left plot shows $t = 0$ to $t = 365$, while the right plot shows $t = 100$ (imaging time) to $t = 130$. Over the 30 day post-imaging window, the VPP appears nearly unbiased (compare to Figure 6).

ground truth tumor path and simulated imaging data to illustrate patient-specific Virtual Patient Population (VPP) generation.

In section 2.3, we presented a family of random field models that are both easy to simulate and demonstrate a wide variety of spatial heterogeneity; in section 4.3, we discussed how the parameters defining the statistics of these random field models can be estimated from a collection of molecular
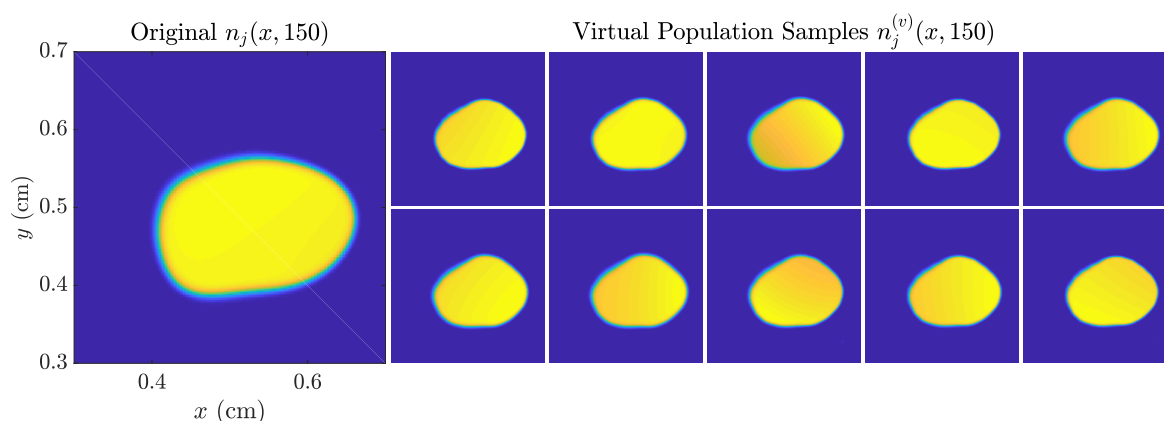
Figure 11. Example VCT Samples for $t = 150$, i.e. 50 days post-imaging, generated using the MLEM estimate of $n_j(\boldsymbol{x}, 100)$ as the initial condition, the MLE of $\rho_j(\boldsymbol{x})$ for the growth factor, and randomizing $D(\boldsymbol{x})$ and $\kappa(\boldsymbol{x})$ as described in the main text. All plots are shown on $[0.3, 0.7]^2$ and with a common colorbar. In contrast to Figure 7, the morphology apparent in the virtual population is nearly identical; heterogeneity in the overall growth is apparent in Figure 10.

images, and discussed some of the computational challenges inherent in this task. Future work will discuss algorithms for solving this MLE problem, as well as additional estimation strategies to calibrate more general random field models from noisy data.

In section 3, we discussed molecular Emission Computed Tomography (ECT) imaging systems and their associated statistical models. For simplicity, we emphasized the case of performing direct, possibly polyscopic, 2D imaging of a thin sample. While the 2D approximation is very limiting in general, it is applicable to the experimental setting of mouse window chambers, a topic of future work in our group. In another future work, we will address fully 3D tomographic ECT systems with particle-processing data and the associated computational and image quality questions that arise. Understanding image data quality in a task performance setting can help design better imaging systems [51], and to our knowledge, there is currently minimal effort formally evaluating the quality of image data in the mathematical oncology context.

In section 4, we presented general likelihood-based statistical procedures for estimating both patient-specific and population parameters from noisy molecular imaging data, and in section 5, we presented two experiments to illustrate these procedures. In the first experiment, measurements of $n_j(\boldsymbol{x}, 100)$ are used to construct a maximum likelihood estimate and posterior samples, which are subsequently used to generate a patient-specific VPP. In the second experiment, we add an additional imaging study, using only maximum likelihood to estimate both $n_j(\boldsymbol{x}, 100)$ and $\rho_j(\boldsymbol{x})$. The resulting virtual population demonstrated a more accurate prediction of the true $N_j(t)$, over a 30 day post-imaging prediction period. In a future work, we will investigate the question of optimizing imaging acquisition schemes for the task of predicting $N(t)$. We will also address the addition of treatment models and evaluation techniques discussed in [5] to our virtual population methodology, working towards providing a general strategy to solving the optimization problem (2.12).

Acknowledgments

Conflict of interest

The authors declare no conflict of interest for this paper.

References

1. R. C. Rockne, A. Hawkins-Daarud, K. R. Swanson, J. P. Sluka, J. A. Glazier, P. Macklin, et al., The 2019 mathematical oncology roadmap, *Phys. Biol.*, **16** (2019), 041005.

2. J. T. Oden, E. A. Lima, R. C. Almeida, Y. Feng, M. N. Rylander, D. Fuentes, et al., Toward predictive multiscale modeling of vascular tumor growth, *Arch. Comput. Methods Eng.*, **23** (2016), 735–779.

3. T. E. Yankeelov, N. Atuegwu, D. Hormuth, J. A. Weis, S. L. Barnes, M. I. Miga, et al., Clinically relevant modeling of tumor growth and treatment response, *Sci. Trans. Med.*, **5** (2013), 187ps9.

4. A. Baldock, R. Rockne, A. Boone, M. Neal, A. Hawkins-Daarud, D. Corwin, et al., From patient-specific mathematical neuro-oncology to precision medicine, *Front. Oncol.*, **3** (2013).

5. N. Henscheid, E. Clarkson, K. J. Myers, H. H. Barrett, Physiological random processes in precision cancer therapy, *PloS One*, **13** (2018), e0199823.

6. N. Henscheid, *Quantifying Uncertainties in Imaging-Based Precision Medicine*, Ph.D thesis, University of Arizona, 2018.

7. P. L. Bedard, A. R. Hansen, M. J. Ratain, L. L. Siu, Tumour heterogeneity in the clinic, *Nature*, **501** (2013), 355.

8. R. Everett, K. B. Flores, N. Henscheid, J. Lagergren, K. Larripa, D. Li, et al., A tutorial review of mathematical techniques for quantifying tumor heterogeneity, *Math. Biosci. Eng.*, **17** (2020), 3660.

9. H. Barrett, K. Myers, *Foundations of Image Science*, John Wiley & Sons, New York, 2004.

10. K. Swanson, R. Rostomily, E. Alvord Jr, A mathematical modelling tool for predicting survival of individual patients following resection of glioblastoma: a proof of principle, *Br. J. Cancer*, **98** (2008), 113.

11. A. Badano, C. G. Graff, A. Badal, D. Sharma, R. Zeng, F. W. Samuelson, et al., Evaluation of digital breast tomosynthesis as replacement of full-field digital mammography using an in silico imaging trial, *JAMA Network Open*, **1** (2018), e185474.

12. A. Karolak, D. A. Markov, L. J. McCawley, K. A. Rejniak, Towards personalized computational oncology: from spatial models of tumour spheroids, to organoids, to tissues, *J. R. Soc., Interface*, **15** (2018), 20170703.

13. E. Abadi, W. P. Segars, B. M. Tsui, P. E. Kinahan, N. Bottenus, A. F. Frangi, et al., Virtual clinical trials in medical imaging: a review, *J. Med. Imaging*, **7** (2020), 042805.

14. K. R. Swanson, C. Bridge, J. Murray, E. C. Alvord, Virtual and real brain tumors: using mathematical modeling to quantify glioma growth and invasion, *J. Neurol. Sci.*, **216** (2003), 1–10.

15. P. Tracqui, Biophysical models of tumour growth, *Rep. Prog. Phys.*, **72** (2009), 056701.

16. J. T. Oden, E. E. Prudencio, A. Hawkins-Daarud, Selection and assessment of phenomenological models of tumor growth, *Math. Models Methods Appl. Sci.*, **23** (2013), 1309–1338.

17. E. M. Rutter, T. L. Stepien, B. J. Anderies, J. D. Plasencia, E. C. Woolf, A. C. Scheck, et al., Mathematical analysis of glioma growth in a murine model, *Sci. Rep.*, **7** (2017), 2508.

18. J. A. Weis, M. I. Miga, L. R. Arlinghaus, X. Li, A. B. Chakravarthy, V. Abramson, et al., A mechanically coupled reaction–diffusion model for predicting the response of breast tumors to neoadjuvant chemotherapy, *Phys. Med. Biol.*, **58** (2013), 5851.

19. C. Hogea, C. Davatzikos, G. Biros, An image-driven parameter estimation problem for a reaction–diffusion glioma growth model with mass effects, *J. Math. Biol.*, **56** (2008), 793–825.

20. T. L. Stepien, E. M. Rutter, Y. Kuang, Traveling waves of a go-or-grow model of glioma growth, *SIAM J. Appl. Math.*, **78** (2018), 1778–1801.

21. E. J. Kostelich, Y. Kuang, J. M. McDaniel, N. Z. Moore, N. L. Martirosyan, M. C. Preul, Accurate state estimation from uncertain data and models: an application of data assimilation to mathematical models of human brain tumors, *Biol. Direct*, **6** (2011), 64.

22. M. Gyllenberg, G. F. Webb, Quiescence as an explanation of Gompertzian tumor growth, *Growth Dev. Aging.*, **53** (1989), 25–33.

23. J. D. Logan, *An Introduction to Nonlinear Partial Differential Equations*, John Wiley & Sons, 2008.

24. W. Hundsdorfer, J. G. Verwer, *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*, Springer Science & Business Media, 2013.

25. R. C. Smith, *Uncertainty Quantification: Theory, Implementation, and Applications*, SIAM, 2013.

26. X. Han, P. E. Kloeden, *Random Ordinary Differential Equations and their Numerical Solution*, Springer, 2017.

27. G. J. Lord, C. E. Powell, T. Shardlow, *Introduction to Computational Stochastic PDEs*, Cambridge, 2014.

28. E. Clarkson, H. H. Barrett, Characteristic functionals in imaging and image-quality assessment: tutorial, *JOSA A*, **33** (2016), 1464–1475.

29. I. Gel'fand, N. Vilenkin, *Generalized Functions, Vol. 4: Applications of Harmonic Analysis*, Academic Press, New York, 1964.

30. T. Sullivan, *Introduction to Uncertainty Quantification*, Springer, 2015.

31. J. Rolland, H. H. Barrett, Effect of random background inhomogeneity on observer detection performance, *JOSA A*, **9** (1992), 649–658.

32. F. Bochud, C. Abbey, M. Eckstein, Statistical texture synthesis of mammographic images with clustered lumpy backgrounds, *Opt. Express*, **4** (1999), 33–42.

33. L. G. Kessler, H. X. Barnhart, A. J. Buckler, K. R. Choudhury, M. V. Kondratovich, A. Toledano, et al., The emerging science of quantitative imaging biomarkers terminology and definitions for scientific studies and regulatory submissions, *Stat. Methods Med. Res.*, **24** (2015), 9–26.

34. J. P. O'connor, E. O. Aboagye, J. E. Adams, H. J. Aerts, S. F. Barrington, A. J. Beer, et al., Imaging biomarker roadmap for cancer studies, *Nat. Rev. Clin. Oncol.*, **14** (2017), 169.

35. A. M. Jarrett, D. Faghihi, D. A. H. II, E. A. Lima, J. Virostko, G. Biros, et al., Optimal control theory for personalized therapeutic regimens in oncology: Background, history, challenges, and opportunities, *J. Clin. Med.*, **9** (2020), 1314.

36. A. Anderson, J. Xie, J. Pizzonia, R. Bronen, D. Spencer, J. Gore, Effects of cell volume fraction changes on apparent diffusion in human cells, *Magn. Reson. Imaging*, **18** (2000), 689–695.

37. D. A. Hormuth II, J. A. Weis, S. L. Barnes, M. I. Miga, E. C. Rericha, V. Quaranta, et al., Predicting in vivo glioma growth with the reaction diffusion equation constrained by quantitative magnetic resonance imaging data, *Phys. Biol.*, **12** (2015), 046006.

38. J. Czernin, Oncological applications of FDG-PET, in *PET Molecular Imaging and its Biological Applications*, Springer, (2004), 321–388.

39. R. M. Hoffman, Application of GFP imaging in cancer, *Lab. Invest.*, **95** (2015), 432–452.

40. R. Schafer, H. M. Leung, A. F. Gmitro, Multi-modality imaging of a murine mammary window chamber for breast cancer research, *BioTechniques*, **57** (2014), 45–50.

41. H. Vesselle, J. Grierson, M. Muzi, J. M. Pugsley, R. A. Schmidt, P. Rabinowitz, et al., In vivo validation of 3 deoxy-3-[18F] fluorothymidine ([18F] FLT) as a proliferation imaging tracer in humans: correlation of [18F] FLT uptake by positron emission tomography with KI-67 immunohistochemistry and flow cytometry in human lung tumors, *Clin. Cancer Res.*, **8** (2002), 3315–3323.

42. E. T. McKinley, G. D. Ayers, R. A. Smith, S. A. Saleh, P. Zhao, M. K., et al., Limits of [18 F]-FLT PET as a biomarker of proliferation in oncology, *PloS One*, **8** (2013), e58938.

43. E. Clarkson, M. A. Kupinski, H. H. Barrett, L. Furenlid, A task-based approach to adaptive and multimodality imaging, *Proc. IEEE*, **96** (2008), 500–511.

44. H. H. Barrett, D. W. Wilson, B. M. Tsui, Noise properties of the EM algorithm. I. theory, *Phys. Med. Biol.*, **39** (1994), 833.

45. Y. Ding, L. Caucci, H. H. Barrett, Charged-particle emission tomography, *Med. Phys.*, **44** (2017), 2478–2489.

46. H. H. Barrett, W. C. Hunter, B. W. Miller, S. K. Moore, Y. Chen, L. R. Furenlid, Maximum-likelihood methods for processing signals from gamma-ray detectors, *IEEE Trans. Nuclear Sci.*, **56** (2009), 725–735.

47. A. K. Jha, H. H. Barrett, E. C. Frey, E. Clarkson, L. Caucci, M. A. Kupinski, Singular value decomposition for photon-processing nuclear imaging systems and applications for reconstruction and computing null functions, *Phys. Med. Biol.*, **60** (2015), 7359.

48. A. P. Dempster, N. M. Laird, D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. R. Stat. Soc. Ser. B (methodol.)*, **39** (1977), 1–22.

49. G. McLachlan, T. Krishnan, *The EM Algorithm and Extensions*, John Wiley & Sons, 2007.

50. R. H. Byrd, J. C. Gilbert, J. Nocedal, A trust region method based on interior point techniques for nonlinear programming, *Math. Program.*, **89** (2000), 149–185.

51. H. H. Barrett, K. J. Myers, C. Hoeschen, M. A. Kupinski, M. P. Little, Task-based measures of image quality and their relation to radiation dose and patient risk, *Phys. Med. Biol.*, **60** (2015), R1.

52. H. H. Barrett, J. Denny, R. F. Wagner, K. J. Myers, Objective assessment of image quality. II. Fisher information, Fourier crosstalk, and figures of merit for task performance, *JOSA A*, **12** (1995), 834–852.

53. Y. Pawitan, *In All Likelihood: Statistical Modelling and Inference Using Likelihood*, Oxford University Press, 2001.

54. M. Lê, H. Delingette, J. Kalpathy-Cramer, E. R. Gerstner, T. Batchelor, J. Unkelbach, et al., MRI based Bayesian personalization of a tumor growth model, *IEEE Trans. Med. Imaging*, **35** (2016), 2329–2339.

55. A. M. Stuart, Inverse problems: a Bayesian perspective, *Acta Numerica*, **19** (2010), 451–559.

56. J. Kaipio, E. Somersalo, *Statistical and Computational Inverse Problems*, Springer Science & Business Media, 2006.

57. M. A. Kupinski, J. W. Hoppin, E. Clarkson, H. H. Barrett, Ideal-observer computation in medical imaging with use of markov-chain monte carlo techniques, *JOSA A*, **20** (2003), 430–438.

58. R. Allen, T. R. Rieger, C. J. Musante, Efficient generation and selection of virtual populations in quantitative systems pharmacology models, *CPT: Pharmacometrics Syst. Pharmacol.*, **5** (2016), 140–146.

59. T. R. Rieger, R. J. Allen, L. Bystricky, Y. Chen, G. W. Colopy, Y. Cui, et al., Improving the generation and selection of virtual populations in quantitative systems pharmacology models, *Prog. Biophys. Mol. Biol.*, **139** (2018), 15–22.

60. B. Galerne, A. Leclaire, L. Moisan, Texton noise, in *Computer Graphics Forum*, Wiley Online Library, 2017.

61. M. A. Kupinski, E. Clarkson, J. W. Hoppin, L. Chen, H. H. Barrett, Experimental determination of object statistics from noisy images, *JOSA A*, **20** (2003), 421–429.