

PAPER • OPEN ACCESS

Estimation of the cell membrane permeability for gas transport from surface pH measurements

To cite this article: A Bocchinfuso *et al* 2023 *Inverse Problems* **39** 094004

View the [article online](#) for updates and enhancements.

You may also like

- [Review—Silicon Based ISFETs: Architecture, Fabrication Process, Sensing Membrane, and Spatial Variation](#)
Anita Gupta, Santanu Sharma and Rupam Goswami
- [Metrological challenges for measurements of key climatological observables. Part 3: seawater pH](#)
A G Dickson, M F Camões, P Spitzer et al.
- [Establishment of Traceability Chain For pH Measurement in Indonesia](#)
Ayu Hindayani, Nuryatini Hamim, Christine Elishian et al.

Estimation of the cell membrane permeability for gas transport from surface pH measurements

A Bocchinfuso , D Calvetti*  and E Somersalo

Department of Mathematics, Applied Mathematics, and Statistics, Case Western Reserve University, 10900 Euclid Avenue, Cleveland, OH 44106, United States of America

E-mail: dxs57@case.edu

Received 28 October 2022; revised 23 May 2023

Accepted for publication 28 July 2023

Published 18 August 2023



CrossMark

Abstract

Bayesian particle filters (PFs) are a viable alternative to sampling methods such as Markov chain Monte Carlo methods to estimate model parameters and related uncertainties when the forward model is a dynamical system, and the data are time series that depend on the state vector. PF techniques are particularly attractive when the dimensionality of the state space is large and the numerical solution of the dynamical system over the time interval corresponding to the data is time consuming. Moreover, information contained in the PF solution can be used to infer on the sensitivity of the unknown parameters to different temporal segments of the data. This, in turn, can guide the design of more efficient and effective data collection procedures. In this article the PF method is applied to the problem of estimating cell membrane permeability to gases from pH measurements on or near the cell membrane. The forward model in this case comprises a spatially distributed system of coupled reaction–diffusion differential equations. The high dimensionality of the state space and the need to account for the micro-environment created by the pH electrode measurement device are additional challenges that are addressed by the solution method.

* Author to whom any correspondence should be addressed.



Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Keywords: particle filter, gas transport, membrane permeability

(Some figures may appear in colour only in the online journal)

1. Introduction

The estimate of the parameters of a dynamical system from partial and indirect noisy observations of the time evolving state vector arises in numerous applications in science and engineering, including weather forecasting, epidemiology and process tomography. A common characteristic in many of those applications is the non-linear dependency of the data on the parameters, often paired with a high computational complexity of the forward model. The nonlinear relation connecting the data and the parameters to be estimated may render the problem too challenging for simple optimization approaches: in many cases the objective functions turn out to be non-convex, with no guarantee of a unique minimizer. An alternative approach is to recast the inverse problem within a Bayesian framework, and to solve it using sampling techniques such as Markov Chain Monte Carlo (MCMC) methods. Among the advantages over single point estimates is the fact that sampling methods can deal better with non-uniqueness and in addition make it possible to quantify uncertainties in the estimates. On the other hand, the computational complexity of the forward model typically is a bottleneck for sampling approaches, because each proposal requires the computation of the full solution to be compared to the data. A reasonable compromise is provided by Bayesian filtering techniques such as particle filtering, a sequential sampling method that at each update requires the time propagation of the forward model over a short time interval.

Although particle filters (PFs) are a natural choice when the number of parameters to be estimated in a dynamical system is high, as is the case for distributed parameters, they may be very effective also when the number of unknown parameters is low, but the forward model poses significant computational challenges, as in the application of interest in this manuscript. When the parameter space is low dimensional, it may be tempting to reduce the forward model accordingly. It is not uncommon, however, that a reduced model is unable to reproduce the data, thus requiring the addition of details that in turn introduce additional nuisance parameters of little or no interest that cannot be neglected because of their effect on the output. Statistical marginalization of the nuisance parameters is possible, but, requires the use of sampling methods.

In some applications the low dimensionality of the parameters to be estimated may also be deceiving, as the complexity of the problem is determined by an underlying state space of high dimensionality: while the data may depend on tiny fraction of it, the entire state space variable must be computed to predict the output corresponding to the data. This is the case for the target application in this article, addressing a fundamental problem in biophysics, namely the estimation of a few parameters characterizing the properties of cell membranes, in particular permeability to gases, from pH measurements near the membrane. While the number of significant parameters is very low, the forward model comprises a spatially distributed reaction–diffusion model posing several computational challenges.

In this article we formulate the inverse problem of estimating the model parameters of cross-membrane gas transfer from the surface pH time series within the Bayesian framework, and we solve it using particle methods for the simultaneous estimation of the state vector and the parameters of interest. To the best of our knowledge, this is the first time that the problem of cell membrane permeability to gases has been formulated as an inverse problem.

A methodological contribution of this work is to show that the particle filtering approach is a computationally feasible alternative for MCMC methods when the forward model is time consuming. In addition, an analysis of the particle filtering results sheds some light on the time dependent sensitivity of the data to the unknown parameters, with information about the different parameters coming at different time instances. This suggests that the output of the PF can be used to design experiments that minimize the number of measurements needed.

The rest of the paper is organized as follows. In section 2, we present the biophysical background of the problem, and provide an overview of the modeling approaches of the forward problem proposed in the literature together with their pros and cons. We then concentrate on the forward model that we will use for the solution of our inverse problem, and describe some of the parameter transformations that we use in our computations. Next we formulate the inverse problem of membrane permeability to gases, discuss various possible solution methods to motivate our choice of the PF. Section 3 presents the details of the PF solution of the inverse problem of membrane permeability and includes a detailed description of the algorithm. The computed results are presented and analyzed in section 4. In section 5 we discuss our results and present our conclusions and outline future work.

2. Forward and inverse problem

In this section we provide a biophysical motivation of the problem considered in this paper, discuss the various modeling approaches proposed in the literature, and present the details of the forward model that we are going to use.

2.1. Biophysical background

The problem that motivated the work for this article is understanding the mechanism of gas transport across cell membranes, one of the fundamental processes for oxygen-based life forms, as can be easily understood by following the trajectory of oxygen molecules from the pulmonary alveoli to the cell mitochondria, or the reverse journey of carbon dioxide molecules. Of particular importance are the membrane properties of the red blood cells that play a central role in the oxygen delivery. The classical theory of gas transport through membrane, known as Overton's rule, dating back to studies from the nineteenth century [20], is based on the notion of gases dissolving in the lipid phase of the membrane and diffusing through the membrane driven by the concentration gradient. The discovery of membranes that are practically impermeable to gases like CO₂ [24] has been a great motivation for revisiting and possibly revising the theory: An alternative for the classical 'Overton's rule' is the hypothesis that gases permeate the cell membrane through specific gas channels associated with membrane bound proteins, such as rhesus proteins (Rh) and, more prominently, aquaporins (AQPs), the latter abundantly present in some of the membranes of key importance, erythrocyte and astrocytic endfeet membranes. A detailed account of the history of the discovery can be found in [4]. For a more complete recent review on the topic, see [15].

In addition to its role in cell respiration, the permeability of cell membranes to gases plays a central role in another important life-supporting process, the pH regulation of cells: The passage of carbon dioxide or ammonia through the membrane is effectively a way for the cell to virtually move protons between the extracellular and intracellular space. More precisely, in view of the limited permeability of the membrane to many solutes, including protons H⁺ and




Figure 1. pH control through gas transport by a shuttle mechanism for the virtual transport of a proton through cell membrane: The proton H^+ outside the cell (gray) associates with bicarbonate HCO_3^- to form carbonic acid H_2CO_3 , subsequently dissociating into CO_2 and water. CO_2 is driven by the concentration gradient through the membrane, and the reverse reaction lowers the pH inside the cell (blue).

bicarbonate HCO_3^- , one way to regulate the pH is through the shuttle mechanism in figure 1, comprising a sequence of chemical reactions. The observation that the pH balance depends indirectly on the membrane permeability to carbon dioxide has led to the idea that pH data could be used to estimate the membrane permeability [19, 22]. The formulation and solution of the inverse problem of estimating the membrane permeability from pH measurements within a Bayesian framework is the main contribution of this article.

The dependency of the surface pH on membrane properties and, in particular, on the presence of AQPs, has been well established experimentally and theoretically. The experimental data presented in [17] was collected using the oocyte of the African clawed frog (*Xenopus laevis*), a large cell that can be made to heterologously express membrane proteins such as AQPs by injection of foreign RNA [7, 17–19]. While these experimental results provided strong evidence for a role of AQPs in pH regulation, estimating the permeability of the membrane requires the solution of an inverse problem that to our knowledge has never been systematically attempted. The successful solution of the inverse problem in turn depends on the availability of an accurate and computationally feasible mathematical model mapping the unknown properties of the membrane to the observable surface pH data.

2.2. Review of earlier models

A number of mathematical models of cross-membrane gas transport and its connection to surface pH have been proposed in the literature, see, e.g. [3, 9, 10, 16, 23]. In [22], the forward gas transport problem including the diffusion and chemistry was solved by using a finite difference discretization of the spatial dependency in the diffusion equations. Taking advantage of the spherical symmetry of the oocyte, the diffusion model was reduced to a single (radial) space dimension, and the computational challenges are mostly related to the stiffness of the system arising from the diffusion model as well as from the widely different time scales in the reaction dynamics. While qualitatively correct, the model proposed in [22] was unable to reproduce the dynamical range of the observed surface pH. More specifically, the predicted pH increase on the membrane surface was less than one tenth of the observed one, strongly suggesting that the model was missing some significant key factor. The mismatch between the data and the model predictions raised the question whether the pH sensitive liquid electrode itself,

used for measuring the pH would create a microenvironment where the observed pH differs significantly from the free membrane pH addressed by the model. In [5], the microenvironment question was addressed by adding to the model a detailed description of the electrode tip. The pH measurement device is a thin cylindrical liquid electrode pushed against the cell membrane, and the pH is measured in the gap between the tip and the membrane. Numerical simulations showed that if the replenishment of substances by diffusion in the domain below the electrode tip is reduced by the electrode-membrane contact around the domain, the concentration profiles in the resulting microenvironment will differ significantly from the free membrane profiles, exhibiting a dynamical range similar to those experimentally observed. The detailed model in [5] is computationally very costly due to the need for a fine scale spatial discretization around the electrode tip, the multi-scale nature of the time dependency, and the introduction of additional model parameters describing the contact that affect the output, but are poorly known. These reasons made the model unsuitable for the inverse problem, as the simulation of a single pH trace requiring more than ten hours on a standard computer, and motivated the need for a computationally less demanding but quantitatively accurate model that would be suitable for repeated evaluations. The approximate reduced model proposed in [2] is a one-dimensional radial reaction-diffusion model, comprising a mini-compartment describing the microenvironment under the electrode which exchanges solutes through the cell membrane with the interior domain and with the exterior domain through a leak under the rim of the electrode tip. The sensitivity of the model to various model parameters was tested systematically in the cited article. The details of this model are presented in the following subsection.

2.3. Forward model

Consider a spherical oocyte of radius $R > 0$ centered at the origin of $\mathbb{R}^3$. The interior and exterior domains are filled with homogeneous diffusive environment with N solutes with concentrations $u_n^\pm(x, t)$, $1 \leq n \leq N$; the index ‘+’ referring to the exterior domain, $|x| > R$ and ‘−’ to the interior domain, $|x| < R$. The reaction–diffusion equation satisfied by each of these solutes is

$$\frac{\partial u_n^\pm}{\partial t} = \nabla \cdot \kappa_n^\pm \nabla u_n^\pm + \sum_{j=1}^M S_{nj} \varphi_j^\pm, \quad (1)$$

where $\kappa_n^\pm$ is the diffusion coefficient characteristic to the solute, $\varphi_{nj}^\pm = \varphi_{nj}^\pm(x, t)$ is the reaction flux of the j th chemical reaction and S_{nj} is the corresponding stoichiometric matrix entry. We assume that the initial concentration values at $t = 0$ are given, and at infinity $|x| \rightarrow \infty$ the concentrations converge to known constant asymptotic values, $u_n^+ \rightarrow u_n^\infty$ for all n . In practice, we assume that the asymptotic values are reached at some finite distance $r = R_\infty$ from the origin, see figure 2. Before describing the boundary conditions at the interface $|x| = R$, we fix the chemical system by specifying the solutes and the reactions included in the model. We consider six solutes, present both inside and outside the cell: Carbon dioxide, $[\text{CO}_2] = u_1$, carbonic acid, $[\text{H}_2\text{CO}_3] = u_2$, bicarbonate, $[\text{HCO}_3^-] = u_3$, proton, $[\text{H}^+] = u_4$, an unspecified buffer, e.g. HEPES, denoted by HA, $[\text{HA}] = u_5$, as well as the corresponding anion, $[\text{A}^-] = u_6$. Six chemical reactions are included in the model, and they take place both inside and outside the cell, namely






Figure 2. Description of the computational model (left panel): the oocyte is modeled as a sphere of radius R , and the concentrations are assumed to reach the asymptotic values at an artificial boundary at $r = R_\infty$. The cell membrane is characterized by Fick's law for CO_2 , while the membrane is assumed impermeable for other substances. The area in the red square contains the microenvironment under the electrode (right panel).



The reaction fluxes are modeled using the mass action formulas, that is,

$$\varphi_1(x, t) = k_1 u_1(x, t), \quad (5)$$

$$\varphi_2(x, t) = k_{-1} u_2(x, t), \quad (6)$$

$$\varphi_3(x, t) = k_2 u_2(x, t), \quad (7)$$

$$\varphi_4(x, t) = k_{-2} u_3(x, t) u_4(x, t), \quad (8)$$

$$\varphi_5(x, t) = k_3 u_5(x, t), \quad (9)$$

$$\varphi_6(x, t) = k_{-3} u_6(x, t) u_4(x, t), \quad (10)$$

where to simplify the notation we have suppressed the indication of exterior or interior domains. In reality, the picture is slightly more complex due to the presence of carbonic anhydrase (CA) inside and near the cell membrane outside the cell, catalyzing the reactions (2). The amount and distribution of CA are poorly known, and estimating them is part of the inverse

problem that we address in this paper. Here, to account for the action of CA we modify the reactions (5) and (6) as

$$\varphi_1^-(x, t) = Ak_1 u_1^-(x, t), \quad \varphi_1^+(x, t) = \begin{cases} Ak_1 u_1^+(x, t), & \text{for } R < |x| < R + \delta, \\ k_1 u_1(x, t), & \text{for } |x| \geq R + \delta, \end{cases}$$

where $\delta > 0$ is a small constant and $A > 1$ is the enhancement factor of the reaction rate. Similarly,

$$\varphi_2^-(x, t) = Ak_{-1} u_2^-(x, t), \quad \varphi_1^-(x, t) = \begin{cases} Ak_{-2} u_2^+(x, t), & \text{for } R < |x| < R + \delta, \\ k_{-2} u_1(x, t), & \text{for } |x| \geq R + \delta, \end{cases}$$

For simplicity, we assume that the enhancement factor A is fixed and known, and the thickness δ of the layer where CA is active outside the cell is small enough that the effect of CA in the computations can be limited to the immediate vicinity of the cell membrane.

Finally, to assign the boundary conditions at $|x| = r = R$, we assume that only carbon dioxide can cross the membrane, and that its flux is described by Fick's law,

$$\kappa_1^+ \frac{\partial u_1^+}{\partial r}(x, t) \Big|_{r=R} = \kappa_1^- \frac{\partial u_1^-}{\partial r}(x, t) \Big|_{r=R} = -\lambda (u_1^+(x, t) - u_1^-(x, t)) \Big|_{r=R}, \quad (11)$$

while the other solutes have vanishing normal derivatives at the boundary. Assuming a perfectly spherical cell, we have $u_n^\pm(x, t) = u_n^\pm(r, t)$, and the three-dimensional diffusion term reduces to a radial model of the form

$$\frac{\partial u_n^\pm}{\partial t} = \frac{1}{r^2} \frac{\partial}{\partial r} \left(\kappa_n^\pm r^2 \frac{\partial u_n^\pm}{\partial r} \right) + \sum_{j=1}^M S_{nj} \varphi_j^\pm, \quad (12)$$

implemented in [22] using a finite difference semi-discretization of the spatial part. In [5], the model was enhanced by placing a thin cylindrical electrode with circular cross section along the z -axis in the exterior domain, with the electrode tip almost touching the membrane. To account for the limited diffusion between the exterior domain and the small cylindrical domain between the electrode tip and the membrane, the diffusion coefficient is reduced by a multiplicative factor within the finite elements under the electrode rim simulating the partial clamping caused by the electrode pushing against the membrane. Although the presence of the electrode breaks the spherical symmetry, the electrode axial symmetry still allows a dimension reduction to two space dimensions. The geometric modeling of the tip region to sufficiently high accuracy requires a significant mesh refinement, increasing the computational complexity to the point of making the model impractical for solving the inverse problem. The detailed electrode tip model demonstrates the importance of accounting for the tip domain, supporting the hypothesis that the discrepancy between observations and model predictions could be explained by the micro-environment unaccounted for in previous models.

Recently, a simplified approximate model also accounting for the presence of the electrode was introduced in [2]. The main idea behind the reduced model is that the thin electrode disturbs minimally the radial reaction–diffusion model away from the electrode, and the domain under the tip can be modeled with sufficient accuracy by a homogenous mini-compartment communicating with the interior domain through the membrane transport of CO_2 , and with the exterior domain near the membrane by diffusion of all solutes with reduced diffusion. In

practice, the model augments the reaction–diffusion model (11) and (12) with the equations of the mini-compartment below the electrode tip

$$\frac{du_n^0}{dt}(t) = \delta_{n1} \frac{\lambda}{h} (u_n^0(t) - u_n^-(R, t)) - \frac{2\gamma}{P} (u_n^+(R, t) - u_n^0(t)) + \sum_{j=1}^m S_{mj} \varphi_j^0(t). \quad (13)$$

Here $u_n^0 = u_n^0(t)$ is the mean concentration of the n th solute in the mini-compartment, h is the distance of the electrode tip from the membrane, P is the radius of the circular electrode tip, and γ is a fictitious permeability describing the partial clamping of the mini-compartment from the exterior domain. The fact that the membrane is permeable only to carbon dioxide is the reason for the presence of the Kronecker delta δ_{n1} in the transport term. We model the reaction fluxes φ_j^0 using mass action formulas similar to (5)–(10) in terms of average concentrations. Since CA was assumed to be active in the immediate vicinity of the membrane, but not necessarily in the entire compartment, we write

$$\varphi_1^0(t) = A^0 k_1 u_1^0(t), \quad \varphi_2^0(t) = A^0 k_2 u_2^0(t),$$

where $1 \leq A^0 \leq A$. We refer to [2] for additional details pertaining the derivation of the reduced model.

2.4. Standard experiment and the inverse problem

In this section we describe the standard experiment carried out to measure the pH data. At $t = 0$, both the inside and the outside domain of the radial model are in equilibrium, and the concentrations in the exterior domain correspond to the asymptotic values at infinity:

$$u_n^+(x, 0) = u_n^{\text{ext}} = u_n^\infty, \quad u_n^-(x, 0) = u_n^{\text{int}}, \quad 1 \leq n \leq 6,$$

where the initial values satisfy the equilibrium conditions: If the fluxes at $t = 0$ are calculated by using the formulas (5)–(10), then

$$\sum_{j=1}^6 S_{nj} \varphi_j^\pm = 0, \quad 1 \leq n \leq 6.$$

We assume that initially,

$$0 = u_1^{\text{int}} < u_1^{\text{ext}},$$

which implies, in particular, that the pH in the exterior domain is lower than in the interior domain, and furthermore the concentrations in the electrode induced mini-compartment coincide with those in the exterior domain, $u^0(0) = u^\infty$. Before the experiment starts, the oocyte is not in contact with the ambient fluid, and no diffusion or gas transport processes are active. At $t = 0$, the oocyte is immersed in the ambient liquid, and the gas transport starts, guided by the carbon dioxide gradient across the membrane. Asymptotically, the carbon dioxide concentration equilibrates, and the pH returns to the initial value in the exterior domain, while in the interior domain, the asymptotic pH value is lower than the initial value. We refer to this experiment as standard experiment.

The standard experiment implicitly defines the following inverse problem. *Given the surface pH in the mini-compartment corresponding to the standard experiment,*

$$\text{pH}^0(t) = \log_{10} u_4^0(t), \quad 0 \leq t < \infty,$$

estimate the parameter vector $(\lambda, A^0, h, \gamma)$.

Numerical experiments reported in [2] show that the pH curve at the membrane depends only weakly on the parameter h . We believe that the size of the physical gap modeled here in terms of h and the known radius P of the electrode tip affect the data in two indirect ways: first, the closer the electrode dip is to the cell membrane, the stronger the clamping effect encoded in parameter γ is. Second, the smaller the mini-compartment is, the closer to the CA-laden membrane the reactions within the compartment take place, thus the average CA-enhancement A^0 becomes more pronounced as h decreases. For these reasons, we omit the parameter h from the list of unknowns, and treat it as a nuisance parameter that is kept at a fixed value in the computed examples.

Before going into further details, a discussion of the methodological choices available to solve the inverse problem is in order. Numerical tests with a simple least squares optimization yield estimates of the parameters that depends heavily on the initial values, indicating that the forward map $(\lambda, A^0, \gamma) \mapsto \text{pH}^0(t)$ is non-convex. On the other hand without a good starting point, an MCMC chain may require a long burn-in phase, which with the current model that requires a few seconds of computing time to produce one full pH trace makes standard Bayesian methods impractical. Particle filtering methods, on the other hand, require only the propagation of each particle over a short time interval $[t_{k-1}, t_k]$, where t_k is the k th observation time, making them an attractive alternative for an inverse problem where standard optimization based approaches are unfeasible or unsuccessful.

The biophysics of the problem dictates that the unknown parameters must satisfy certain conditions that can be either imposed as constraints, or implicitly enforced through suitable changes of variables. We adopt the latter approach. The next subsection describes the parameter transformations and the rationale behind them.

2.5. Transformations of the parameters

We express the unknowns of the inverse problem in terms of appropriately scaled dimensionless particles corresponding to the prior understanding of the model. First, consider the membrane permeability $\lambda > 0$. To obtain a reasonable upper bound, consider the value obtained by thinking the cell membrane as of a water layer of the thickness of the membrane, which is then pressed to a fictitious infinitely thin layer. Arguably, the non-resistant water layer is an ideal barrier, and realistic values that take into account different factors resisting the gas transport are smaller than the no-resistance value, denoted by λ_0 . We express these considerations by writing

$$\lambda = \xi_\lambda \lambda_0, \quad 0 \leq \xi_\lambda \leq 1.$$

Similarly, we account for the fact that CA cannot enhance the reactions in the mini-compartment more than it does at the cell membrane by letting

$$A^0 = \xi_A S, \quad 0 \leq \xi_A \leq 1.$$

To treat the fictitious transport parameter γ , we start by observing that, according to the numerical simulations in [2], assigning a value $\gamma \approx \gamma_0 = 1 \mu\text{m s}^{-1}$ will yield realistic increases in the surface pH during the standard experiment, while a tenfold increase of γ will result in unrealistically low dynamical range of the pH change. These observations justify the logarithmic bounds

$$-0.5 \leq \log_{10} \frac{\gamma}{\gamma_0} \leq 1,$$

leading to the natural parametrization,

$$\gamma = \gamma_0 10^{[-0.5\xi_\gamma + (1-\xi_\gamma)]}, \quad 0 \leq \xi_\gamma \leq 1.$$

In this manner, we have effectively reduced the range of all model parameters between 0 and 1, i.e. $\xi = (\xi_\lambda, \xi_A, \xi_\gamma) \in [0, 1]^3$. Since the presence of bound constraints would create technical problem with sampling, we use the logit-transform,

$$\text{logit} : (0, 1) \rightarrow \mathbb{R}, \quad \text{logit}(\xi) = \log \left(\frac{\xi}{1-\xi} \right),$$

with the inverse expit-transform,

$$\text{expit} : \mathbb{R} \rightarrow (0, 1), \quad \xi = \text{expit}(s) = \frac{\exp(s)}{1 + \exp(s)},$$

and formulate the inverse problem in terms of the unconstrained parameter vector

$$s = (s_1, s_2, s_3) = (\text{logit}(\xi_\lambda), \text{logit}(\xi_A), \text{logit}(\xi_\gamma)) \in \mathbb{R}^3.$$

2.6. Reparametrization of the state variable

The state variable of our model, after semidiscretization in the spatial direction, is the vector of all concentration values at the discretization nodes. Denote the concentration vector of a solute by $u(t)$, the dimension of the vector being equal to the number of the discretization points in the FEM mesh. Since the concentrations need to be non-negative but imposing a hard bound constraint is computationally inconvenient for random perturbations, we reparametrize the state vector as follows. Introduce the new variable,

$$x(t) = \log \left(\delta + \frac{u(t)}{\bar{u}} \right),$$

where $\bar{u}$ is a scalar representing a typical value of the concentration, and $\delta > 0$ is a safeguard parameter introduced to regularize in case $u(t) = 0$. The back-transform gives us

$$u(t) = \bar{u} \left(e^{x(t)} - \delta \right).$$

If $x(t)$ is perturbed by an additive term η , we have

$$x(t) + \eta = \log \left(\delta + \frac{u_\eta(t)}{\bar{u}} \right),$$

and solving for u_η we obtain

$$u_\eta(t) = \bar{u} \left(e^{x(t)+\eta} - \delta \right).$$

Observe that for $\delta > 0$ small, the above formula yields

$$u_\eta(t) = u(t)e^\eta + \delta\bar{u}(1 - e^\eta) \approx u(t)e^\eta \text{ for } \delta \ll 1,$$

that is, an additive perturbation of the transformed quantity x translates into a multiplicative perturbation of the concentration. However, in the numerical computations, we choose $\delta = 1$, leading in practice to well-behaving perturbations.

3. Parameter estimation via PF

The forward model of our problem consists of an initial value problem model

$$\frac{du}{dt}(t; s) = F(u(t; s), s), \quad u(0; s) = U_0, \quad (14)$$

where $u(t; s) \in \mathbb{R}^n$ is the spatially discretized time-dependent state vector comprising all of the concentrations at the nodal values of the FEM mesh and s is a vector of unknown parameters to be estimated, that we assume to be time invariant. In the reminder of the paper we will denote $u(t; s)$ by $u(t)$, thus omitting explicit reference to the parameter dependency. The measurement data

$$b_k = G(u(t_k)) + \varepsilon_k, \quad 1 \leq k \leq K,$$

are functions of the state vector at discrete times t_k , $0 < t_1 < \dots < t_K$, and ε_k denotes the observation noise vector at time t_k , whose statistics we assume to be approximately known.

After the logarithmic transformation discussed in the previous subsection and a suitable time discretization, the forward model can be expressed in terms of the variable $x = x(t)$. Numerical integration of the forward model (14) using, e.g. a standard numerical integrator from t_{k-1} to t_k yields a Markov model for the discrete transformed variable $x_k = x(t_k)$ of the form

$$x_k = M(x_{k-1}, s) + \eta_k, \quad (15)$$

where M is a non-linear propagation operator, and η_k accounts for the numerical approximation errors.

We write the observation model in terms of the approximate state vector x_k ,

$$b_k = H(x_k) + \varepsilon_k, \quad (16)$$

where H is the discretized version of the observation function G expressed in terms of x_k . For simplicity, the observation error is assumed to be the same as in the exact continuous model, neglecting errors that might have been introduced by the discretization process.

In the Bayesian framework, the unknowns states x_k , $1 \leq k \leq K$, and parameter vector s are modeled as random variables X_k and S_k , respectively. Observe that while s is assumed time invariant, the virtual time dependency of S_k is an indication of the temporal evolution of the knowledge, or belief, about the value of s . The PF estimates the probability density functions in a sequential manner, so that the probability density functions at the next time instance are determined from the probability density functions at the current time instance and the aggregate observations up to the current time. More precisely, assume that $\pi_{k-1}(x_{k-1}, s_{k-1})$ is the probability density of (X_{k-1}, S_{k-1}) at time t_{k-1} . In the rest of the paper, our particles consist of the concatenation of the state and parameter vectors.

The Bayesian filter estimation of $\pi_k(x_k, s_k)$ proceeds in two steps:

1. **Prediction Step:** Based on the Markov model (15) and the current distribution $\pi_{k-1}(x_{k-1}, s_{k-1})$, a predicted probability distribution $\pi_{k|k-1}(x_k, s_k)$ is determined.
2. **Correction:** As a new observation b_k becomes available, the predicted distribution is corrected based on the observation model (16), yielding the updated density $\pi_k(x_k, s_k)$.

The prediction step, which can be thought of as a push-forward of the current density using the propagation model, can be computed by using the Chapman–Kolmogorov formula, while the correction step, also referred to as analysis step, is based on an application of Bayes' formula.

In Bayesian particle filtering, the densities are approximated by Monte Carlo sampling, making the computation of the push-forward map particularly attractive as it entails only a propagation of sample points, referred to as particles. For details, we refer to [1, 6, 8, 11–14].

While the general ideas behind particle filtering are very straightforward, successful implementations typically require addressing several questions, including how to organize the computations efficiently when the forward model is computationally demanding, and how to choose the initial cloud of particles.

In the following, we discuss the computational details related to the implementation of the PF for our problem. We begin with addressing the problem of the selection of an initial particle cloud, which for our application turned out to be particularly challenging.

3.1. Initialization

In the particle filtering algorithm, the current information about the state of the system is encoded in the ensemble

$$S_k = \{(x_k^{(1)}, s_k^{(1)}, w_k^{(1)}), \dots, (x_k^{(N)}, s_k^{(N)}, w_k^{(N)})\},$$

where N is the ensemble size, and $w_k^{(j)}$ is the weight of the j th particle, specified later. In the initial ensemble of particles, all the state vector portions $x_0^{(j)}$ are set equal to the known initial conditions of the problem, listed in table 4. The initialization of the parameter portion $s_0^{(j)}$ of the particles is less obvious and requires some careful considerations, because of the nature of the problem. While, in principle, the values of the parameters could be assigned rather arbitrarily over a wide range, assuming that the uncertainty decreases as the data becomes available, a good coverage of a large interval leads easily to an ensemble size that for the current model is computationally unfeasible. Feasible, yet not too committal initial intervals for the parameters were determined partly on the basis of the biophysical information, partly on preliminary numerical experiments. We draw the initial parameter vector ensemble from a discrete uniform distribution over suitably selected support. More precisely, we start with a uniform grid of size $40 \times 40 \times 40$ over the domain

$$(\xi_\lambda, \xi_A, \xi_\gamma) \in [0.6, 1] \times [0.6, 1] \times [0, 1] = Q, \quad (17)$$

where the intervals have been chosen in a way that the curves generated with the corresponding λ, A^0 (with $A = 20$) and γ are compatible with the measurements reported in the literature. In particular, γ can substantially change the shape of the curve, and initial values within the selected interval yield responses whose shapes are in line with the experimental measurements reported in the literature. The motivation for the choice of initial values for the parameter A^0 is based on the consideration that the CA action in the sensor micro-environment is unlikely to be strongly different from that on the membrane and inside the cell. Finally, the permeability to CO_2 is assumed to take on any value less or equal to the theoretical no-resistance value λ_0 .

In light of the above considerations, the initial ensemble of parameter values are chosen as a discrete subset of an equidistant grid over the parallelepiped Q (17). More specifically, out of $m = (40)^3 = 64000$ parameter vectors regularly spaced over (17), we select N particles, $N < m$, so as to cover evenly the parameter space. The choice of N depends on computational resources, and will be specified when computed examples are discussed. Each particle in the initial cloud is the concatenation of an initial state vector and a parameter vector, and since *a priori* every particle is equally likely to be a realization, all particles are assigned the same weight, $w_0^{(j)} = 1/N$.

3.2. Prediction step

Assume that we have at our disposal an ensemble $\{(x_{k-1}^{(j)}, s_{k-1}^{(j)}, w_{k-1}^{(j)})\}_{j=1}^N$ of particles and corresponding weights such that $\sum_{j=1}^N w_{k-1}^{(j)} = 1$. The sample-based approximation of the probability distribution $\pi_k(x_{k-1}, s_{k-1})$ can then be expressed in terms of point masses as

$$\pi_{k-1}(x_{k-1}, s_{k-1}) \approx \sum_{j=1}^N w_{k-1}^{(j)} \delta_{x_{k-1}^{(j)}}(x_{k-1}) \delta_{s_{k-1}^{(j)}}(s_{k-1}),$$

the approximation being understood in the weak* sense for Radon measures. By propagating the particles using the evolution model, we obtain the approximate push-forward distribution,

$$\pi_{k|k-1}(\hat{x}_k, \hat{s}_k) \approx \sum_{j=1}^N w_{k-1}^{(j)} \delta_{\hat{x}_k^{(j)}}(\hat{x}_k) \delta_{\hat{s}_k^{(j)}}(\hat{s}_k), \quad (18)$$

where the point mass centers $\hat{x}_k^{(j)}$ are computed by using the propagation model,

$$\hat{x}_k^{(j)} = M(x_{k-1}^{(j)}, s_{k-1}^{(j)}), \quad \hat{s}_k^{(j)} = s_{k-1}^{(j)},$$

that is, the parameter particles $s_{k-1}^{(j)}$ are propagated as constants. In order to enrich the parameter ensemble through resampling, following [6, 13], we compute the sample mean and covariance of the parameter particles,

$$\begin{aligned} \bar{s} &= \sum_{j=1}^N w_{k-1}^{(j)} \hat{s}_{k-1}^{(j)}, \\ \mathbf{C}_s &= \sum_{j=1}^N w_{k-1}^{(j)} (s_{k-1}^{(j)} - \bar{s})(s_{k-1}^{(j)} - \bar{s})^T, \end{aligned}$$

and replace the point mass approximation of the parameter density by a Gaussian mixture model. Let h be a small positive constant, $0 < h \ll 1$, and let

$$\bar{s}_k^{(j)} = a \hat{s}_k^{(j)} - (1-a)\bar{s},$$

where

$$a = \sqrt{1-h^2}.$$

The points $\bar{s}_k^{(j)}$ are obtained by contracting the prediction points $\hat{s}_k^{(j)}$ slightly towards the ensemble mean. It has been shown [6, 25] that the Gaussian mixture,

$$\pi_{GM}(s, h) = \sum_{j=1}^N w_{k-1}^{(j)} \mathcal{N}(s | \bar{s}_k^{(j)}, h^2 \mathbf{C}_s),$$

has the same mean and covariance as the corresponding point mass approximation corresponding to the limit $h \rightarrow 0$. Based on this observation, we replace the approximation (18) by

$$\pi_{k|k-1}(\hat{x}_k, s_k) \approx \sum_{j=1}^N w_{k-1}^{(j)} \delta_{\hat{x}_k^{(j)}}(\hat{x}_k) \mathcal{N}(s_k | \bar{s}_k^{(j)}, h^2 \mathbf{C}_s). \quad (19)$$

This predicted density is further updated as the next observation arrives. Observe that this model does not take into account the approximation error (15), which we add later. In the following, the propagated variable is denoted by x_k .

3.3. Correction step

The correction step updates the probability density (19) determined in the prediction step in light of the new observation b_k . If we assume that the observation noise follows a Gaussian distribution with zero mean and symmetric positive definite covariance $\mathbf{S}$, the likelihood density is of the form

$$\pi_{B_k|X_k,S_k}(b_k | x_k, s_k) \propto \exp\left(-\frac{1}{2}(b_k - H(x_k, s_k))^T \mathbf{S}^{-1}(b_k - H(x_k, s_k))\right).$$

Considering the predicted density (19) as a prior, the posterior density can be written as

$$\begin{aligned} \pi_k(x_k, s_k) &\propto \sum_{j=1}^N w_{k-1}^{(j)} \pi_{B_k|X_k,S_k}(b_k | x_k, s_k) \delta_{\hat{x}_k^{(j)}}(x_k) \mathcal{N}(s_k | \bar{s}_k^{(j)}, h^2 \mathbf{C}_s) \\ &= \sum_{j=1}^N (w_{k-1}^{(j)} g_k^{(j)}) \left(\frac{\pi_{B_k|X_k,S_k}(b_k | x_k, s_k)}{g_k^{(j)}} \right) \delta_{\hat{x}_k^{(j)}}(x_k) \mathcal{N}(s_k | \bar{s}_k^{(j)}, h^2 \mathbf{C}_s), \end{aligned} \quad (20)$$

where we introduce the fitness weights,

$$g_k^{(j)} = \pi_{B|X,S}(b_k | \hat{x}_k^{(j)}, \hat{s}_k^{(j)}) \propto \exp\left(-\frac{1}{2}(b_k - \hat{b}_k^{(j)})^T \mathbf{S}^{-1}(b_k - \hat{b}_k^{(j)})\right),$$

with

$$\hat{b}_k^{(j)} = H(\hat{x}_k^{(j)}, \hat{s}_k^{(j)}).$$

These weights can be interpreted as a measure of how well the predicted particles are able to explain the new data. To draw the next generation of particles, we first define the importance weights

$$\hat{w}_k^{(j)} = \frac{w_{k-1}^{(j)} g_k^{(j)}}{\sum_{j=1}^m w_{k-1}^{(j)} g_k^{(j)}},$$

taking into account both the likelihood and the importance of the predicted particles. Next we generate a new sample representation of the density π_k by drawing N indices $j_1, j_2, \dots, j_N$ with replacement from the index set $\{1, 2, \dots, N\}$, using the weights $\hat{w}_k^{(j)}$ as probabilities, so that the resampled particles have equal weights.

Recalling the model (15), we introduce the approximation error and write

$$x_k = \hat{x}_k + \eta_k,$$

which is tantamount to replacing the point mass functions in (19) by a probability density centered around the noiseless propagated value. Assuming that the innovation η_k is a realization of a random variable E_k with probability density π_{E_k} , we draw the next generation of particles,

$$x_k^{(j)} = \hat{x}_k^{(j)} + \eta_k^{(j)}, \quad \eta_k^{(j)} \sim \pi_{E_k}, \quad s_j^{(k)} \sim \mathcal{N}(s_j | \bar{s}_k^{(j)}, h^2 \mathbf{C}_s).$$

Finally, we correct the weights based on (20), defining

$$f^{(j)} = \frac{\pi_{B_k|X_k,S_k}(b_k | x_k^{(j)}, s_k^{(j)})}{\pi_{B_k|X_k,S_k}(b_k | \hat{x}_k^{(j)}, \hat{s}_k^{(j)})},$$

which, after normalization, become the new weights

$$w_k^{(j)} = \frac{f^{(j)}}{\sum_{j=1}^N f^{(j)}}.$$

This completes the correction step. The sequence of steps described here is summarized below in the form of an algorithm.

Particle Filtering algorithm.

1. **Initialize:** Draw parameter values $s_0^{(j)} \in Q$, set $x_0^{(j)} = x_0$, $w_0^{(j)} = \frac{1}{N}$, $1 \leq j \leq N$, Set the counter $k = 1$.

2. **Repeat** for $k = 1 : T$

(a) **Prediction:** Propagate the particles:

$$\hat{x}_k^{(j)} = M(x_{k-1}^{(j)}, s_{k-1}^{(j)}), \quad \hat{s}_k^{(j)} = s_{k-1}^{(j)}, \quad \hat{w}_k^{(j)} = w_{k-1}^{(j)}.$$

(b) **Correction:**

i. Compute mean and covariance of the parameters:

$$\bar{s}_k = \sum_{j=1}^N \hat{w}_k^{(j)} \hat{s}_k^{(j)}, \quad \mathbf{C}_k = \frac{1}{N-1} \sum_{j=1}^N \hat{w}_k^{(j)} (\hat{s}_k^{(j)} - \bar{s}_k) (\hat{s}_k^{(j)} - \bar{s}_k)^\top.$$

ii. Shrinkage: With $a = \sqrt{1 - h^2}$, define

$$\bar{s}_k^{(j)} = a \hat{s}_k^{(j)} + (1 - a) \bar{s}_k.$$

iii. Compute the fitness weights,

$$g_k^{(j)} = \pi_{B|X,S}(b_k | \hat{x}_k^{(j)}, s_k^{(j)}),$$

and update the weights as

$$\bar{w}_k^{(j)} = \frac{\hat{w}_k^{(j)} g_k^{(j)}}{\sum_{j=1}^m \hat{w}_k^{(j)} g_k^{(j)}}.$$

iv. Survival of the fittest resampling: Draw N indices $j_1, \dots, j_N$ with replacement from $\{1, 2, \dots, N\}$ using the weights $\bar{w}_k^{(j)}$ as probabilities, and set

$$(\hat{x}_k^{(j)}, w_k^{(j)}) \leftarrow (\hat{x}_k^{(j_j)}, w_k^{(j_j)}), \quad 1 \leq j \leq N.$$

v. Update the particles,

$$x_k^{(j)} = \hat{x}_k^{(j)} + \eta_k^{(j)}, \quad \eta_k^{(j)} \sim \pi_{E_k}, \quad s_k^{(j)} \sim \mathcal{N}(\bar{s}_k^{(j)}, h^2 \mathbf{C}).$$

vi. Update the weights,

$$f^{(j)} = \frac{\pi_{B_k|X_k,S_k}(b_k | x_k^{(j)}, s_k^{(j)})}{\pi_{B_k|X_k,S_k}(b_k | \hat{x}_k^{(j)}, \hat{s}_k^{(j)})}, \quad w_k^{(j)} = \frac{f^{(j)}}{\sum_{j=1}^N f^{(j)}}.$$

3. **Advance the counter:** $k \leftarrow k + 1$.

In the following section, we use this algorithm to solve the membrane permeability inverse problem.

4. Computed examples

In this section we test the particle filtering algorithm for estimating the membrane parameters on synthetic data. The main goal is to demonstrate the viability of the approach, to understand whether the parameters can be learned from the data, and to test whether there are differences

Table 1. Parameters used in the generative model. CA strength A is assumed to be known, while the other parameters are to be estimated from the data.

λ	Membrane permeability	$30.78 \mu\text{m s}^{-1} = 0.9 \cdot 34.2 \mu\text{m s}^{-1}$
A	CA factor	20
A^0	CA factor (around sensor)	$16 = 0.8 \cdot A$
γ	quench factor	$6.31 \cdot 10^{-5} \mu\text{m s}^{-1} = 10^{-4.2} \mu\text{m s}^{-1}$

Table 2. Geometric parameters corresponding to non-resistant membrane hypothesis.

R	Oocyte radius	$650 \mu\text{m}$
R_∞	External radius	$800 \mu\text{m}$
w	Radius of the electrode tip	$10 \mu\text{m}$
h	Distance of the electrode tip	$10 \mu\text{m}$

Table 3. Diffusion coefficients κ of the substances.

Substance	Inside ($(\mu\text{m})^2 \text{s}^{-1}$)	Outside ($(\mu\text{m})^2 \text{s}^{-1}$)
CO_2	1.71×10^3	1.71×10^3
H_2CO_3	1.11×10^3	1.11×10^3
HCO_3^-	1.11×10^3	1.11×10^3
H^+	8.69×10^3	8.69×10^3
HA	1.56×10^3	1.56×10^3
A^-	1.56×10^3	1.56×10^3

in the learning rates of the parameters. Moreover, we test the robustness of the method with respect to the ensemble size, a detail of crucial importance for the computational time.

We generate the synthetic data using the parameter values listed in table 1. The other model parameters such as the geometry, reaction rates, or initial concentrations are given in tables 2, 3, 4, and 5 and are assumed to be known.

We discretize the radial variable using a mesh refinement near the cell membrane to catch the boundary effects, including the effect of the CA that is assumed to accelerate the reaction rates in the exterior domain only in the immediate vicinity of the membrane. More precisely, adjacent to the membrane the element size is $0.1 \mu\text{m}$, and when moving away from the membrane, the size is dilated by a factor of 1.05 at each step. This results in a state space discretization of 205 nodal points for each of the six substances. With the addition of the six concentrations in the microenvironment and the three model parameters, the dimensionality of the state space becomes 1239. By comparison, the surface pH data is a function of only one of the components but it is determined by the entire state vector, thus effectively making this is a high dimensional problem. We assume that the pH is measured with 1 Herz frequency, and we add scaled white noise with standard deviation $\sigma = 0.01$ to the noiseless computed data, so that the noise covariance matrix $\mathbf{S}$ is a scaled identity. The noiseless pH curve and the noisy data are shown in figure 3.

To generate the initial particle cloud, we define a regular grid of size $(40)^3 = 64000$ of the parameter domain Q defined in (17), and subsample the grid values retaining every tenth point, yielding a parameter ensemble of size $N = 6400$. The restriction of the sample size is needed to keep the computing time feasible. We run the particle filtering algorithm over a time interval of $T = 500$ s, estimating at each time step the particle mean as well as 75% and 90% quantiles of each parameter.

Table 4. Initial values of the concentrations in the standard experiment. The model assumes that at $t = 0$, the oocyte is added to the bath with the pH electrode already in place.

Substance	Inside (mM)	Outside (mM)
CO ₂	0	0.4720
H ₂ CO ₃	0	0.0013
HCO ₃ [−]	0	9.901
H ⁺ (pH)	6.310×10^{-5} (7.2)	3.162×10^{-5} (7.5)
HA	12.09	2.500
A [−]	15.22	2.500




Figure 3. The simulated noiseless membrane pH curve generated with the parameters listed in table 1, along with the noisy data collected with 1 Herz sampling frequency, indicated by the dots. The figure also shows the pH curve computed by using the posterior mean estimates of the parameter corresponding to the particle sample size $N = 6400$.

Figure 4 shows the results of the parameter estimation as quantile credible envelope plots. The generative values given in table 1 are indicated in the plots by a horizontal line.

The plots in figure 4 show that the credible intervals of different parameters converge at different times, indicating which part of the pH trace curve is sensitive to which parameter. We observe that in this run, γ is the first parameter to converge around $t = 150$ s, while λ and A^0 converge at around $t = 200$ s. Comparison of the results and the pH curve in figure 3 indicates that all the parameters are learned well only after the pH has passed its peak. The average over the particle values from $t = 200$ s until the end of the simulation ($t = 500$ s) yields the mean estimates for the parameters,

$$\begin{aligned}\bar{\lambda} &= 31.52 \mu\text{m s}^{-1}, \\ \bar{A}^0 &= 15.57, \\ \bar{\gamma} &= 6.34 \cdot 10^{-5} \mu\text{m s}^{-1} = 10^{-4.20} \mu\text{m s}^{-1}.\end{aligned}$$

Observe that, as far as inferring on the membrane properties is concerned, A^0 and γ are nuisance parameters and therefore of secondary interest, and only the membrane permeability λ




Figure 4. Particle filtering results of estimating the three model parameters. The figures show the 25% and 75% quantiles corresponding to the darker envelope, and 10% and 90% quantiles corresponding to the lighter envelope. The values used in the generative model are indicated by the horizontal dashed line. In the left column, the particle size was $N = 6400$, and the experiment was run for $T = 500$ s, and on the right, $N = 10000$.

is of interest. Running the forward model using the estimated parameter values yields a pH trace, plotted in figure 3, that reproduces well the noiseless data.

To test the robustness of the method with respect to the ensemble size, we perform an independent run generating a new data set with the same parameter values, but with different

Table 5. Reaction rates. Observe that we have two fast time scale parameters, ε and ε' , whose precise values are not well known. The reaction rates $k_{\pm 1}$ are enhanced by the acceleration factor A .

Reaction	k_ℓ	$k_{-\ell}$	$K = k_\ell/k_{-\ell}$
$\text{CO}_2 + \text{H}_2\text{O} \xrightleftharpoons[k_{-1}]{k_1} \text{H}_2\text{CO}_3$	$0.0302 \text{ (1 s}^{-1}\text{)}$	$10.9631 \text{ (1 s}^{-1}\text{)}$	2.7547×10^{-4}
$\text{H}_2\text{CO}_3 \xrightleftharpoons[k_{-2}]{k_2} \text{HCO}_3^- + \text{H}^+$	$\varepsilon = 10^{-9} \text{ (1 s}^{-1}\text{)}$	ε/K_2	$K_2 = 0.2407 \text{ (mM)}$
$\text{HA}_{\text{in}} \xrightleftharpoons[k_{-3}]{k_3} \text{A}_{\text{in}}^- + \text{H}_{\text{in}}^+$	$\varepsilon' = 10^{-6} \text{ (1 s}^{-1}\text{)}$	$\varepsilon'/K_{\text{HA}}$	$K_{\text{HA}} = 7.9433 \times 10^{-5} \text{ (mM)}$
$\text{HA}_{\text{out}} \xrightleftharpoons[k_{-3}]{k_3} \text{A}_{\text{out}}^- + \text{H}_{\text{out}}^+$	$\varepsilon' = 10^{-6} \text{ (1 s}^{-1}\text{)}$	$\varepsilon'/K_{\text{HA}}$	$K_{\text{HA}} = 3.1623 \times 10^{-5} \text{ (mM)}$

realization of the noise. The data are collected again for $T = 500$ s with known initial conditions at time $t = 0$ s, measurements and sampling step set like in the previous run, but with $N = 10000$ particles.

The estimates of the three parameters are not very different from those in the previous experiment where the number of particles was smaller, as shown in the right column of figure 4. In this particular run, the averaged values from 200 s to the end of the experiment are

$$\begin{aligned}\bar{\lambda} &= 30.86 \mu\text{s}^{-1}, \\ \bar{A}^0 &= 16.46, \\ \bar{\gamma} &= 6.79 \cdot 10^{-5} \mu\text{s}^{-1} = 10^{4.17} \mu\text{s}^{-1}.\end{aligned}$$

We observe that the value of λ is slightly closer to the one used in the generative model. The reproduced pH curve (not shown) is very similar to the one obtained in the previous run.

The sequential way in which the PF learns the parameters and how the parameters are correlated is summarized in figure 5, showing the scatter plots of the parameters at four different time instances. Notice how the cloud of parameter values tighten up around the value used in the generative model as time progresses.

As already pointed out, the fact that the number of unknown of interest is low and only one of the state space components is defining the data, makes the problem seem very simple, raising the question whether a simple optimization approach would yield reasonable estimates for the parameters. To test whether this is the case, we implemented a simple Fletcher-Reeves conjugate gradient algorithm, estimating the gradient of the standard least squares functional by using a finite difference scheme. It turns out that while the algorithm converges relatively fast, typically in less than 20 iterations, the outcome depends strongly on the initial point, pointing towards a non-convexity of the underlying objective function. Of more concern is that by starting at initial values relatively close to the generative values, the algorithm finds a minimizer yielding an output pH curve that has a good overall fit to data, although the parameter values are significantly different from the generative ones. A closer inspection reveals that the discrepancy between generative and estimated pH curves occurs at the very initial time steps that, according to the particle filtering analysis, are crucial for the identification of the parameters.




Figure 5. Illustration of how the learning of the parameters of the model by the particle filter progresses over time. The frames show the scatter plots of the parameters at six different times instances. Notice how the spread of the clouds tightens up as time passes. Top row: time instances $t = 1$ and $t = 10$; middle row: time instances $t = 50$ and $t = 100$; bottom row: time instances $t = 200$ and $t = 500$.

5. Conclusions

Bayesian particle filtering methods for solving inverse problem with an underlying forward model that has a structure of an evolution-observation model are a competitive alternative to traditional MCMC sampling methods even in the case when the data needs not to be interpreted as time series arriving in a sequential manner. When the state space dimension is high

and the forward model is computationally demanding as in the example considered in this article, the advantage of particle filtering algorithms is that the model needs to be propagated only a short time step at a time for each particle, rather than requiring to compute the entire time series to generate one proposal for the MCMC sampler. In the problem considered in this article, the dimensionality of the state space is of the order of 1200, comprising about two hundred nodal values for each of the six concentrations, and the computation time of producing one realization of length 500 s is about 2.5 s on a personal computer. MCMC algorithms, even with a reasonably high acceptance rate, may become prohibitively slow, in particular if there is no way to determine a good starting point for the chain, because a significantly long burn-in sequence may be needed before the sampler starts to explore the underlying posterior distribution. In the current application, the particle size was of the order $N = 10000$, which in light of the $1/\sqrt{N}$ convergence law yields a reasonable accuracy for estimating the few parameters of interest. Moreover, from the practical point of view, the particle filtering algorithm provides useful information by indicating when the algorithm starts learning the parameters: In the present problem, the results indicate that after about 200–250 s, the parameter estimates do not improve significantly, thus indicating that there is no need to measure the surface pH further beyond that point.

Another natural question is whether a particle smoothing algorithm such as backward-simulation particle smoothers [21] would be preferable. The problem with these methods in the case where the state space dimensionality is high as in the present problem is the large memory allocation needed, as the methods require that the particles are saved at each step.

This article is the first instance in which the fundamental question of how gasses cross the cell membrane is formulated as an inverse problem. The successful solution of the inverse problem using a forward model accounting for the mini-compartment created by the presence of the electrode pushing against the membrane is a demonstration of how sophisticated inverse problems can help answer fundamental questions in biophysics. Computed examples for synthetic realistic data show that the particle filter is able to estimate quite accurately the generative values of the parameters of interest, in particular the surface permeability for CO_2 in a simulated experiment. The outputs generated with the estimated parameters are very similar to the simulated noiseless data using the generative parameter values, thus suggesting that the filter is able to find combinations of λ , A^0 and γ yielding a response close to that produced by the generative model.

It is well known that particle filtering methods should be run with very large sample sizes. While the computational complexity of the reduced forward model proposed in [2] is significantly lower than that of detailed electrode tip model [5], the computing time of the forward solution remains the limiting factor for the ensemble size. While the present article is not addressing in depth the effect of the limited particle size on the results, however, the two test runs shown here demonstrate good consistency of the results. We conclude by pointing out that in the computed experiments reported in this article, the synthetic data was generated by the same forward model that is used in the inversion. Therefore, while the article provides primarily a proof of concept of the inverse solver, and demonstrates the identifiability of the parameters of interest in the model, it does not include any analysis of the effects of the discrepancy between the detailed and the reduced model or the model discrepancy between the reduced model and measured data. Furthermore, the model contains parameters whose values are assumed to be known, such as the CA activity within the cell and on the surface. It may be of interest to investigate how robust the method is to perturbations in their values, and how well the particle filter will be able to estimate the membrane permeability when the pH measurement are not generated with the same additive model used to solve the inverse problem. These are important questions that are left to be addressed in future work.

Data availability statement

No new data were created or analysed in this study.

Acknowledgments

This work made use of the High Performance Computing Resource in the Core Facility for Advanced Research Computing at Case Western Reserve University. The work of D C was supported in part by NSF Award DMS 1951446. The work of E S was supported in part by NSF Award DMS 2204618.

ORCID iDs

A Bocchinfuso  <https://orcid.org/0000-0002-6063-4131>

D Calvetti  <https://orcid.org/0000-0001-5696-718X>

References

- [1] Arnold A, Calvetti D and Somersalo E 2013 Linear multistep methods, particle filtering and sequential Monte Carlo *Inverse Problems* **29** 085007
- [2] Bocchinfuso A, Calvetti D and Somersalo E 2022 Modeling surface pH measurements of oocytes *Biomed. Phys. Eng. Express* **8** 045006
- [3] Boron W F and DeWeer P 1976 Intracellular pH transients in squid giant axons caused by CO₂, NH₃ and metabolic inhibitors *J. Gen. Physiol.* **67** 91–112
- [4] Boron W F 2010 Sharpey–Schafer lecture: gas channels *Exp. Physiol.* **95** 1107–30
- [5] Calvetti D, Prezioso J, Occhipinti R, Boron W F and Somersalo E 2020 Computational model for electron-induced micro-environmental effects on pH measurements near cell membrane *Multiscale Model. Simul.* **18** 1053–75
- [6] Calvetti D and Somersalo E 2023 *Bayesian Scientific Computing* (Springer)
- [7] Cooper G J and Boron W F 1997 The CO₂ permeability of the AQP1 water channel, expressed in *Xenopus* oocytes *J. Am. Soc. Nephrol.* **8** A0074
- [8] Doucet A and Johansen A M 2009 A tutorial on particle filtering and smoothing: fifteen years later *Handbook of Nonlinear Filtering* vol 12 (Oxford University Press) pp 656–704
- [9] Endeward V and Gros G 2005 Low carbon dioxide permeability of the apical epithelial membrane of guinea-pig colon *J. Physiol.* **567** 253–65
- [10] Endeward V and Gros G 2009 Extra- and intracellular unstirred layer effects in measurements of CO₂ diffusion across membranes—a novel approach applied to the mass spectrometric ¹⁸O technique for red blood cells *J. Physiol.* **587** 1153–67
- [11] Fearnhead P and Künsch H R 2018 Particle filters and data assimilation *Annu. Rev. Stat. Appl.* **5** 421–49
- [12] Kaipio J and Somersalo E 2004 *Computational and Statistical Inverse Problems* (Springer)
- [13] Liu J and West M 2001 Combined parameter and state estimation in simulation-based filtering *Sequential Monte Carlo Methods in Practice* (Springer) pp 197–223
- [14] Maskell S and Gordon N 2002 A tutorial on particle filters for on-line nonlinear/non-Gaussian Bayesian tracking *IEEE Target Tracking: Algorithms and Applications* (Ref. No. 2001/174)
- [15] Michenkova M et al 2021 Carbon dioxide transport across membranes *Interface Focus* **11** 20200090
- [16] Missner A, Kugler P, Saparov S M, Sommer K, Mathai J C, Zeidel M L and Pohl P 2008 Carbon dioxide transport through membranes *J. Biol. Chem.* **283** 25340–7
- [17] Musa-Aziz R, Chen L, Pelletier M F and Boron W F 2009 Relative CO₂/NH₃ selectivities of AQP1, AQP4, AQP5, AmtB and RhAG *Proc. Natl Acad. Sci. USA* **106** 5406–11
- [18] Musa-Aziz R, Occhipinti R and Boron W F 2014 Evidence from simultaneous intracellular- and surface-pH transients that carbonic anhydrase IV enhances CO₂ fluxes across *Xenopus* oocyte plasma membranes *Am. J. Physiol. Cell Physiol.* **307** C814–40

- [19] Nakhoul N L, Davis B A, Romero M F and Boron W F 1998 Effect of expressing the water channel aquaporin-1 on the CO₂ permeability of *Xenopus* oocytes *Am. J. Physiol.* **274** C543–8
- [20] Overton E 1897 Über die osmotischen Eigenschaften der Zelle in ihrer Bedeutung für die Toxicologie und Pharmacologie *Z. Phys. Chem.* **22** 189–209
- [21] Särkkä S and Svensson L 2023 *Bayesian Filtering and Smoothing* vol 17 (Cambridge University Press)
- [22] Somersalo E, Occhipinti R, Boron W F and Calvetti D 2012 A reaction-diffusion model of CO₂ influx into an oocyte *J. Theor. Biol.* **309** 185–203
- [23] Vaughan-Jones R D, Percy B E, Keener J P and Spitzer K W 2002 Intrinsic H⁺ ion mobility in the rabbit ventricular myocyte *J. Physiol.* **541** 139–58
- [24] Waisbren S J, Geibel J P, Modlin I M and Boron W F 1994 Unusual permeability properties of gastric gland cells *Nature* **368** 332–5
- [25] West M 1993 Approximating posterior distributions by mixtures *J. R. Stat. Soc. B* **55** 409–22