

Adding Constraints to Bayesian Inverse Problems

Jiacheng Wu

University of California
Berkeley, CA, 94706

Jian-Xun Wang

University of Notre Dame
Notre Dame, IN, 46556

Shawn C. Shadden

University of California
Berkeley, CA, 94706

Abstract

Using observation data to estimate unknown parameters in computational models is broadly important. This task is often challenging because solutions are non-unique due to the complexity of the model and limited observation data. However, the parameters or states of the model are often known to satisfy additional constraints beyond the model. Thus, we propose an approach to improve parameter estimation in such inverse problems by incorporating constraints in a Bayesian inference framework. Constraints are imposed by constructing a likelihood function based on fitness of the solution to the constraints. The posterior distribution of the parameters conditioned on (1) the observed data and (2) satisfaction of the constraints is obtained, and the estimate of the parameters is given by the maximum a posteriori estimation or posterior mean. Both equality and inequality constraints can be considered by this framework, and the strictness of the constraints can be controlled by constraint uncertainty denoting a confidence on its correctness. Furthermore, we extend this framework to an approximate Bayesian inference framework in terms of the ensemble Kalman filter method, where the constraint is imposed by re-weighting the ensemble members based on the likelihood function. A synthetic model is presented to demonstrate the effectiveness of the proposed method and in both the exact Bayesian inference and ensemble Kalman filter scenarios, numerical simulations show that imposing constraints using the method presented improves identification of the true parameter solution among multiple local minima.

1 Introduction

Computational models are pervasively used in wide-ranging engineering applications. Recent advances in computer platforms and numerical methods have enabled models to be increasingly sophisticated and comprehensive. With greater model complexity comes greater challenge to determine model parameters (including initial/boundary conditions), which are often unknown or uncertain. To address this challenge, one typically solves an *inverse problem* by using observational data to specify model parameters so that the model output matches the observational data. Many inversion techniques have been developed and used by different communities. These methods can be roughly cat-

egorized into two classes: variational and statistical approaches (Asch, Bocquet, and Nodet 2016). The variational approach aims to minimize a specific cost function based on classical optimization theory and calculus of variations (Smedstad and O'Brien 1991), while the statistical approach aims to evaluate or maximize posterior functions based on statistics and Bayesian theory (Cotter et al. 2009).

Because of its robustness and capability for uncertainty quantification, Bayesian inversion techniques are widely used for hidden state and parameter estimation for many physical systems (Iglesias, Lin, and Stuart 2014; Wang et al. 2017; Wang and Xiao 2016; Li et al. 2017). In the Bayesian framework, both the hidden state/parameters (prior) and observable quantities (likelihood) are described as random variables with statistical distributions. The Bayesian estimation aims to calculate the posterior distributions of the inferred quantities from the prior and likelihood based on Bayes theorem. Directly computing the posterior distribution based on the prior and likelihood functions is referred to as the exact Bayesian approach. In general, the posterior is obtained by sampling the prior and likelihood distributions based on efficient Monte Carlo sampling such as the Markov chain Monte Carlo (MCMC) method. However, since MCMC requires an enormous amount of samples, which is computationally infeasible when the likelihood calculation involves expensive model evaluations, many approximate Bayesian inversion approaches have been developed, such as the extended Kalman filter (EKF) (Haykin 2004), unscented Kalman filter (UKF) (Wan and Van Der Merwe 2000), ensemble Kalman filter (EnKF) (Evensen 2003), and sequential Monte Carlo (SMC) method (Del Moral, Doucet, and Jasra 2006).

A common challenge in solving inverse problems is identifiability. Measurement data is typically very limited and solutions for the hidden states and parameters are nonunique. Moreover, numerical stability of the inversion can be significantly reduced for ill-posed problems and uncontrolled inference may happen due to small random noise in the observation data. To address these issues, a general strategy is to incorporate additional information into the inversion process, either by including more observation data or imposing additional constraints. In most mission-critical applications, data are difficult to collect and limited in quality. In such cases, additional constraints can be significantly useful to

help regularize the inversion results to consistent ranges and relieve ill-posedness. Fortunately, for many physical systems, constraints on the state and parameters are available based on existed knowledge (He and Xiu 2016). Nonetheless, most existing Bayesian methods do not take constraints into account (Shao, Huang, and Lee 2010).

Initial progress has been made to incorporate constraints into certain Bayesian filters. For example, Simon et al. (Simon and Chia 2002) considered equality constraints in the standard Kalman filter by projecting the Kalman updated solution onto the state constraint surface. Shao et al. (Shao, Huang, and Lee 2010) developed a constrained sequential Monte Carlo algorithm based on acceptance/rejection and an optimization strategy. Most recently, Gardner et al. (Gardner et al. 2014) also considered inequality constraints in the context of Bayesian optimization. However, the existing approaches to incorporate constraints have been developed for each specific Bayesian filter, and most of them are based on a linearized form of the constraints, which is limiting when constraint functions are complicated and highly nonlinear. Moreover, in many complex systems, the constraints are approximations to reality and formulating the constraint in a deterministic way may neglect the uncertainties associated with constraint itself. In this work, we proposed a general approach to incorporate physics-based constraints into the Bayesian inversion framework, where uncertainty associated with the constraint itself can also be considered. Moreover, this idea is also extended to an approximate Bayesian approach—the ensemble Kalman filter.

2 Methodology

A mathematical model of system defines a *forward problem* that can be formulated as

$$\mathbf{x} = F(\theta), \quad (1)$$

where $\theta \in \mathbb{R}^{d_\theta}$ are model parameters and $\mathbf{x} \in \mathbb{R}^{d_x}$ are the states of the system. The forward operator F is nominally assumed to describe a physical system, whereby F typically represents a suite of algebraic and/or differential equations. In most cases, the model parameters θ are uncertain or unknown, and the state variables \mathbf{x} are largely unobservable. Therefore, the unknown parameters and hidden states need to be inferred from observations $\mathbf{y} \in \mathbb{R}^{d_y}$. These observations indirectly and incompletely describe the state of the system, which can be formulated mathematically as

$$\mathbf{y} = H\mathbf{x} + \epsilon, \quad (2)$$

where H is a projection operator projecting the full state to the observed space and ϵ represents measurement error. The standard inverse problem deals with estimating the unknown parameters θ (or the hidden states \mathbf{x}) based on the observations \mathbf{y} . In practice, approximate Bayesian inversion frameworks, such as Kalman filtering and Sequential Monte Carlo, are used for computational efficiency.

Constraints in Exact Bayesian Inference

Inverse problems are typically ill-posed because the observational data is not sufficient to uniquely determine the unknown parameters. Thus, specification of additional constraints can be useful to regularize the inverse problem.

Equality constraints can be defined with respect to the state variables x as,

$$G(x) = [g_1(x), g_2(x), \dots, g_{d_g}(x)]^T = \mathbf{0}, \quad (3)$$

where $g_i(x)$, $i = 1, 2, \dots, d_g$ represent different equality constraints. In many application, the constraint only approximates reality. Thus, instead of directly imposing a hard constraint, we assume that each constraint satisfies a zero-mean Gaussian distribution, expressed as

$$G(x) \sim \mathcal{N}(\mathbf{0}, \Sigma_c). \quad (4)$$

where the Σ_c is a covariance matrix used to control the strictness of each constraint. Since x is intrinsically a function of θ , the constraints can alternatively be expressed in terms of the parameters

$$G(x) = G(F(\theta)) \sim \mathcal{N}(\mathbf{0}, \Sigma_c). \quad (5)$$

As such, these constraints on the parameters θ can be more naturally considered within the Bayesian framework by imposing additional likelihood functions introduced by these nondeterministic constraints.

Without loss of generality, both the prior and likelihood are assumed Gaussian. Namely, the prior of the parameters θ is defined by

$$p(\theta) = \frac{1}{\sqrt{(2\pi)^{d_\theta} |\Sigma_\theta|}} \exp\left(-\frac{1}{2}(\theta - \hat{\theta})^T \Sigma_\theta^{-1}(\theta - \hat{\theta})\right), \quad (6)$$

where $\hat{\theta}$ and Σ_θ are the prior mean and covariance, which are based on existing knowledge or preliminary estimation. The observation data errors are also assumed to follow a zero-mean Gaussian distribution, i.e., $\epsilon \in \mathcal{N}(0, \Sigma_l)$, thus the likelihood of the observed data set on \mathbf{y} is,

$$p(\mathcal{D}|\theta) = \frac{1}{\sqrt{(2\pi)^{d_y} |\Sigma_l|}} \exp\left(-\frac{1}{2}(\mathbf{y} - H F(\theta))^T \Sigma_l^{-1}(\mathbf{y} - H F(\theta))\right). \quad (7)$$

The covariance matrix Σ_l is obtained by estimating the sample variance of the observed data sets \mathcal{D} . The constraints are imposed by considering the following likelihood function,

$$\begin{aligned} p(G(x) = \mathbf{0} | \theta) \\ = \frac{1}{\sqrt{(2\pi)^{d_g} |\Sigma_c|}} \exp\left(-\frac{1}{2}G(F(\theta))^T \Sigma_c^{-1}G(F(\theta))\right). \end{aligned} \quad (8)$$

The likelihood of the constraints defines a fitness of a specific value of θ based on the satisfaction of the constraints. By introducing this Gaussian-type likelihood function, we enable a “soft” enforcement of the constraints. The strictness of the constraint can be controlled by the diagonal variance matrix Σ_c ,

$$\Sigma_c = \text{diag}\{\sigma_{c,1}^2, \sigma_{c,2}^2, \dots, \sigma_{c,d_g}^2\}. \quad (9)$$

where the variance σ_i represent a confidence on the accuracy of the constraint. Smaller $\sigma_{c,i}$ corresponds to a stricter constraint.

Inequality constraints can be converted to equivalent equality constraints. For example, a scalar inequality constraint $g(x) \leq 0$ can be expressed as

$$\max(0, g(x)) = 0, \quad (10)$$

and thus the corresponding likelihood can be expressed as

$$p(g(x) \leq 0 \mid \theta) = \frac{1}{\sqrt{2\pi\sigma_c^2}} \exp\left(-\frac{1}{2\sigma_c^2} [\max(0, g(x))]^2\right). \quad (11)$$

Imposing constraints through a likelihood function can also be extended to disjunctive constraints. For example, consider a constraint of the form $g_1(x) = 0 \vee g_2(x) = 0$. By the union rule of probability

$$\begin{aligned} p(g_1(x) = 0 \vee g_2(x) = 0 \mid \theta) \\ &= p(g_1(x) = 0 \mid \theta) + p(g_2(x) = 0 \mid \theta) - p(g_1(x) = 0 \wedge g_2(x) = 0 \mid \theta) \\ &= \frac{1}{\sqrt{2\pi\sigma_{c,1}^2}} \exp\left(-\frac{g_1(F(\theta))^2}{2\sigma_{c,1}^2}\right) + \frac{1}{\sqrt{2\pi\sigma_{c,2}^2}} \exp\left(-\frac{g_2(F(\theta))^2}{2\sigma_{c,2}^2}\right) \\ &\quad - \frac{1}{\sqrt{(2\pi)^2 |\Sigma_c|}} \exp\left(-\frac{1}{2} \begin{bmatrix} g_1(F(\theta)) \\ g_2(F(\theta)) \end{bmatrix}^T \Sigma_c^{-1} \begin{bmatrix} g_1(F(\theta)) \\ g_2(F(\theta)) \end{bmatrix}\right), \quad (12) \end{aligned}$$

where Σ_c again defines the covariance matrix of constraints.

With the prior distribution, likelihood of the data, and likelihood of the constraints now defined, the posterior probability distribution conditioned on the observed data \mathcal{D} and the constraints $G(x)$ can be defined as

$$\begin{aligned} p(\theta \mid \mathcal{D}, G(x) = \mathbf{0}) &= \frac{p(\mathcal{D} \mid \theta) p(G(x) = \mathbf{0} \mid \theta) p(\theta)}{p(\mathcal{D}, G(x) = \mathbf{0})} \\ &\propto p(\mathcal{D} \mid \theta) p(G(x) = \mathbf{0} \mid \theta) p(\theta). \quad (13) \end{aligned}$$

Since the posterior distribution cannot be solved analytically in general, it is commonly evaluated based on MCMC sampling.

Constraints in Approximate Bayesian Inference

The direct Bayesian inference based on MCMC sampling is usually intractable when the likelihood calculation involves a computationally expensive model; instead approximate Bayesian approaches are commonly used to provide a more computationally tractable solution. The EnKF is one such method, which is a variant of the standard Kalman filter where the covariance matrix is replaced by Monte Carlo samples.

For EnKF, we combine the original hidden states x and the unknown parameters θ into a new augmented state

$$z = [\theta^T, x^T]^T, \quad (14)$$

which will be updated during the filtering process according to the observed data \mathcal{D} . The initial ensemble is first obtained by sampling the prior distribution $p(\theta)$ and evaluating the model at each ensemble member

$$\{\hat{z}^{(j)}\}_{j=1}^J = \left\{ \begin{bmatrix} \hat{\theta}^{(j)} \\ \hat{x}^{(j)} \end{bmatrix} \right\}_{j=1}^J = \left\{ \begin{bmatrix} \hat{\theta}^{(j)} \\ F(\hat{\theta}^{(j)}) \end{bmatrix} \right\}_{j=1}^J, \quad (15)$$

where J is the number of ensemble members. The probability associated with each ensemble member is initially set to be uniform

$$w_j \triangleq p(z = \hat{z}^{(j)}) = \frac{1}{J}, \quad j = 1, 2, \dots, J. \quad (16)$$

Then the expectation and covariance matrix of the state variables are estimated from the ensemble as

$$\mathbb{E}(\hat{z}) = \sum_{j=1}^J w_j \hat{z}^{(j)}, \quad (17)$$

$$\mathbb{C}(\hat{z}) = \sum_{j=1}^J w_j (\hat{z}^{(j)} - \mathbb{E}(\hat{z}))(\hat{z}^{(j)} - \mathbb{E}(\hat{z}))^T. \quad (18)$$

If the observed data follows a normal distribution $N(\bar{y}, \Sigma_l)$, the prior ensemble can be updated by the observed data \mathcal{D} according to the Kalman update

$$\begin{aligned} z^{(j)} &= \hat{z}^{(j)} + \mathbb{C}(\hat{z}) H^T (H \mathbb{C}(\hat{z}) H^T + \Sigma_l)^{-1} (\bar{y} - H \hat{z}^{(j)}), \\ &j = 1, 2, \dots, J. \quad (19) \end{aligned}$$

The posterior ensemble $\{z^{(j)}\}_{j=1}^J$ represents a sampling for the posterior probability distribution $p(z \mid \mathcal{D})$, with the probability associated with each ensemble member equal to

$$p(z^{(j)} \mid \mathcal{D}) = w_j, \quad \forall j = 1, 2, \dots, J. \quad (20)$$

Now we consider inclusion of constraints. The likelihood of the constraint $G(x) = \mathbf{0}$ to be satisfied conditioned on each member of the posterior ensemble can be computed as

$$\begin{aligned} Lg^{(j)} &\triangleq p(G(x) = \mathbf{0} \mid z^{(j)}) \\ &= \frac{1}{\sqrt{(2\pi)^d |\Sigma_c|}} \exp\left(-\frac{1}{2} G(x^{(j)})^T \Sigma_c^{-1} G(x^{(j)})\right). \quad (21) \end{aligned}$$

By Bayes theorem, the posterior probability density of each ensemble member conditioned on the observed data \mathcal{D} and constraints $G(x) = \mathbf{0}$ is given by

$$\begin{aligned} p(z^{(j)} \mid \mathcal{D}, G(x) = \mathbf{0}) &= \frac{1}{Z} p(G(x) = \mathbf{0}, z^{(j)} \mid \mathcal{D}) \\ &= \frac{1}{Z} p(G(x) = \mathbf{0} \mid z^{(j)}) p(z^{(j)} \mid \mathcal{D}) = \frac{1}{Z} w_j Lg^{(j)} \quad (22) \end{aligned}$$

where Z is the normalization constant defined as

$$Z = \sum_{j=1}^J p(G(x) = \mathbf{0} \mid x^{(j)}) p(x^{(j)} \mid \mathcal{D}) = \sum_{j=1}^J w_j Lg^{(j)}. \quad (23)$$

The empirical distribution for $p(z \mid \mathcal{D}, G(x) = \mathbf{0})$ can be described by the posterior ensemble $\{z^{(j)}\}_{j=1}^J$, and the associated probability mass

$$p(z = z^{(j)} \mid \mathcal{D}, G(x) = \mathbf{0}) = \frac{w_j Lg^{(j)}}{\sum_{p=1}^J w_j Lg^{(p)}}, \quad \forall j = 1, 2, \dots, J, \quad (24)$$

for each ensemble member. We here re-define the new weights for each ensemble members as

$$w'_j \triangleq \frac{w_j Lg^{(j)}}{\sum_{p=1}^J w_j Lg^{(p)}}. \quad (25)$$

The state estimation for the current iteration step is thus computed as the expectation of the empirical posterior distribution conditioned on the data \mathcal{D} and the prior knowledge

that $G(x) = \mathbf{0}$ according to

$$\begin{aligned}\bar{z} &= \mathbb{E}(z|\mathcal{D}, G(x) = \mathbf{0}) = \sum_{j=1}^J p\left(z = z^{(j)}|\mathcal{D}, G(x) = \mathbf{0}\right) z^{(j)} \\ &= \sum_{j=1}^J w'_j z^{(j)}.\end{aligned}\quad (26)$$

Then the estimation of the unknown parameters $\bar{\theta}$ can be extracted from the estimation of the full augmented state. Also, the covariance of the parameter θ with respect to $p(z|\mathcal{D}, G(x) = \mathbf{0})$ can be computed as

$$\Sigma_\theta = \sum_{i=1}^J \left(\theta^{(j)} - \bar{\theta}\right) \text{diag}\{w'_j\}_{j=1}^J \left(\theta^{(j)} - \bar{\theta}\right)^T. \quad (27)$$

The new prior ensemble for the next iteration step is obtained by sampling the following normal distribution

$$\hat{\theta}^{(j)} \sim N(\bar{\theta}, \Sigma_\theta), \quad j = 1, 2, \dots, J, \quad (28)$$

to maximize the next step prior entropy (Penfield Jr 2010) while keeping the mean and covariance the same as the previous posterior distribution. The iterative process continues until a stopping criterion is satisfied or the maximum iteration number is reached.

3 Results and Discussion

Model Test Problem

To verify the effectiveness of the constrained Bayesian inference framework described above, a simple test case is presented here. The forward model mapping from the parameter space $\Theta \subset \mathbb{R}^2$ to the state space $X \subset \mathbb{R}^2$ is defined as

$$\begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = F(\theta) = \begin{bmatrix} \exp(-(\theta_1 + 1)^2 - (\theta_2 + 1)^2) \\ \exp(-(\theta_1 - 1)^2 - (\theta_2 - 1)^2) \end{bmatrix}. \quad (29)$$

The projection matrix mapping from state space to output is given by

$$H = [-1.5, -1.0], \quad (30)$$

and thus the reconstructed output is

$$\begin{aligned}HF(\theta) &= -1.5 \exp(-(\theta_1 + 1)^2 - (\theta_2 + 1)^2) \\ &\quad - 1.0 \exp(-(\theta_1 - 1)^2 - (\theta_2 - 1)^2),\end{aligned}\quad (31)$$

where $HF(\theta) \in \mathbb{R}^1$. We consider the following constraint:

$$G(x) = -0.25 \log x_1 + 0.25 \log x_2 - 2 = 0, \quad (32)$$

which can be equivalently written in terms of θ ,

$$G(F(\theta)) = \theta_1 + \theta_2 - 2 = 0. \quad (33)$$

We assume the observed data follow the normal distribution $N(\bar{y}, \Sigma_l)$ where the mean $\bar{y} = -1.0$ and the covariance matrix Σ_l is chosen based on the uncertainty associated with data.

This model is chosen to create a simple scenario with multiple local minimums. Namely, regardless of the prior information and constraints, we seek the model parameters that minimize the difference between the observed output and reconstructed output, quantified by the cost function

$$\begin{aligned}I(\theta) &= \|\bar{y} - HF(\theta)\|^2 \\ &= \left(1.5 \exp(-(\theta_1 + 1)^2 - (\theta_2 + 1)^2) \right. \\ &\quad \left. + 1.0 \exp(-(\theta_1 - 1)^2 - (\theta_2 - 1)^2) - 1.0\right)^2.\end{aligned}\quad (34)$$




Figure 1: Contour plot of the cost function $I(\theta)$ with respect to parameters $\theta = [\theta_1, \theta_2]$. The red “+” denote the local minima of the cost function.




Figure 2: (a) Sampling of the prior distribution; (b) Sample + data likelihood $p(\mathcal{D}|\theta)$; (c) Sample + constraint likelihood $p(G(x) = \mathbf{0}|\theta)$; (d) Sample + posterior distribution $p(\theta|\mathcal{D}, G(x) = \mathbf{0})$

which has minimums at (a) $\theta^* = (1, 1)$; and (b) θ on the circle defined by

$$(\theta_1 + 1)^2 + (\theta_2 + 1)^2 = \log 1.5. \quad (35)$$

The contour plot of the cost function and the local minimums are visualized in Figure 1. Here we assume $\theta^* = (1, 1)$ is the true value of the parameter θ , and the constraint (33) will help to eliminate convergence to other local minima.

The situation of multiple local minima is a common challenge in solving inverse problems, because the observed output information is usually not enough to uniquely determine the unknown parameters. We show below that imposing constraints can help the solution to converge to the true value.

Exact Bayesian Inference

The prior distribution (6) is first sampled with $J = 5000$ samples. The mean and the covariance matrix are set to

$$\hat{\theta} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad \Sigma_\theta = \begin{bmatrix} 3 & 0 \\ 0 & 3 \end{bmatrix}. \quad (36)$$

The distribution of the samples is visualized in Figure 2(a). The trivial zero mean represents non-informative prior

	θ_1	θ_2	x_1	x_2	y
True values	1	1	0	1	-1
No constraint	-0.0448	-0.0722	0.1698	0.1063	-0.3610
With constraint: EXP	0.9926	0.9449	0.0004	0.9969	-0.9976
With constraint: MAP	0.9845	0.9698	0.0004	0.9988	-0.9994

Table 1: Parameters θ , states x and output y estimated using sample-based Bayesian inference with no constraint imposed and with constraint imposed.

knowledge of θ , and the large variance defined by Σ_θ above denotes large uncertainty about the prior. Ideally if better prior knowledge exists, we can specify a better prior here, with more accurate mean and less uncertainty. After sampling the prior, the likelihood function of data $p(\mathcal{D}|\theta)$ is evaluated at each individual sample point $\{\theta^{(j)}\}_{j=1}^J$. The likelihood of the data is plotted with respect to each sample in Figure 2(b). The value of the likelihood is indicated by the brightness of the sample. It can be clearly observed that the brightest regions coincide with the local minimums in Figure 1, which shows the region of the highest likelihood of data. Similarly, we evaluate the likelihood function of the constraint $p(G(x) = 0|\theta)$ at each sample, and the likelihood is visualized in Figure 2(c). The region with the highest likelihood represents the form of the constraint in (θ_1, θ_2) space, which is $\theta_1 + \theta_2 - 2 = 0$. The variance for the constraint is set to be $\Sigma_c = 0.5$, which controls how strict the constraint is enforced. Lastly, the posterior $p(\theta|\mathcal{D}, G(x) = 0)$ is evaluated at each samples, and the distribution of the sample along with the posterior weights are plotted in Figure 2(d). It is clearly observed that the location with the highest posterior density correspond to the intersection between the regions with high likelihood of data \mathcal{D} and high likelihood of the constraint satisfaction. This intersection region picks out the true value of the parameter θ . Computing the weighted sum of the parameter samples $\{\theta^{(j)}\}_{j=1}^J$ with respect to the posterior weights yields the final estimation of the unknown parameter $\theta_{\text{Exp}}^* = \sum_{j=1}^J p(\theta^{(j)}|\mathcal{D}, G(x) = 0) \theta^{(j)} = \sum_{j=1}^J w_j \theta^{(j)}$. Or simply taking the sample $\theta^{(j)}$ that maximize the posterior $p(\theta^{(j)}|\mathcal{D}, G(x) = 0)$ yields the maximum a posteriori estimation (MAP) of the unknown parameter, i.e. θ_{MAP}^* . Once the parameter is estimated, the estimated value of the state variables x and output y can be computed by evaluating the forward model $F(\theta^*)$. These estimated values are listed in Table 1 for the case of including and not including constraint. It can be seen from this table that imposing the constraint significantly increase the estimation accuracy in the case where multiple local minimums exist.

Approximate Bayesian Inference

No constraint imposed. Iterative ensemble Kalman filter estimates the unknown model parameters θ in a iterative manner. Since the cost function $I(\theta)$ has multiple local minimums, different initial guesses of θ will converge to different




Figure 3: Different initial guesses of the unknown parameter θ (marked with red triangles) and their corresponding converged values after 1000 EnKF iteration steps (marked with blue circles), with no constraint imposed.

local minimums. We here define

$$\text{Group I} \triangleq \{\theta^* \in \mathbb{R}^2 | \theta^* = (1, 1)\}, \quad (37)$$

$$\text{Group II} \triangleq \{\theta^* \in \mathbb{R}^2 | (\theta_1^* + 1)^2 + (\theta_2^* + 1)^2 = \log 1.5\}, \quad (38)$$

which represent two different local minimum regions. θ^* represents the converged value of the parameter after ensemble Kalman filter iterations.

Here we simulated three different cases with different initial guesses: (a) $\theta^0 = (-2, -2)$; (b) $\theta^0 = (0, 0)$; (c) $\theta^0 = (2, 2)$. The covariance matrix of the prior and the covariance of the data likelihood are given as $\Sigma_\theta = [1, 0; 0, 1]$ and $\Sigma_l = 0.01$. The results for the three different simulations are visualized in the parameter space of θ in Figure 3. It can be seen that the upper right initial guess at $(2, 2)$ converges to local minimum Group I, and the other two initial guesses both converge to local minimum Group II. The convergence processes of the parameter θ for the three different initial guesses are plotted in Figure 4. It can be seen that all converge to the corresponding local minimum groups within about 400 iterations. The main difference is that while Case (c) converges to the local minimum $(1, 1)$, i.e., Group I, directly after a few iterations, Case (a) and (b) converge to $\theta = (-1, -1)$ first, which is the center of the local minimum circle of Group II, and then shift to a local minimum on the circle of Group II at around the 200th iteration (indicated by the “jump”). The reason behind this is that we use the mean of the ensemble of each step as the estimated parameter value. When the ensemble converges to the local neighborhood of Group II, the mean of the ensemble will generally be the center of the local minimum circle because the high likely ensemble members are roughly symmetrically distributed around the center $(-1, -1)$. The mean of the ensemble will gradually shift to certain points on the circle based on the distribution of the ensemble members. The evolution of the ensemble for different initial guesses is shown in Figure 5. It can be seen that the variance of the ensemble gradually decrease until all ensemble members collapse to the corresponding local minimum.

The convergence results for the three different initial




Figure 4: Results for the convergence of the parameter θ with no constraint imposed. (a) Initial guess $\theta^0 = (-2, -2)$; (b) initial guess $\theta^0 = (0, 0)$; (c) initial guess $\theta^0 = (2, 2)$.




Figure 5: Evolution of the ensemble of the parameter θ with no constraint imposed. The points on the red circle centered at $(-1, -1)$ and the red point at $(1, 1)$ denote the local minimums of the cost function $I(\theta)$. (a) Initial guess $\theta^0 = (-2, -2)$; (b) initial guess $\theta^0 = (0, 0)$; (c) initial guess $\theta^0 = (2, 2)$.

guesses are summarized here,

$$\begin{aligned}\theta^0 = (-2, -2) &\rightarrow \theta^* = (-0.4080, -0.7598) \in \text{Group II}, \\ \theta^0 = (0, 0) &\rightarrow \theta^* = (-0.3654, -1.0421) \in \text{Group II}, \\ \theta^0 = (2, 2) &\rightarrow \theta^* = (0.9998, 0.9812) \in \text{Group I}.\end{aligned}$$

The reconstructed outputs $HF(\theta^*)$ for the above cases all converge to the target value $\bar{y} = -1$ within 1000 EnKf iterations. However, with no constraint is imposed, the estimate of the parameter θ will converge to the closer local minimum group based on where the initial guesses are. The initial guess in the middle $(0, 0)$ converges to the local minimum Group II because Group II contains more local minimums than Group I, and therefore the the solution is more likely to converge to Group II when initial guess is in the middle. More broadly, there is no guarantee that the estimate of the parameter will converge to the the true parameter value $(1, 1)$.

With constraint imposed For the local minimums in Group I and Group II, only the true parameter value $(1, 1)$ satisfies the constraint. We test here whether imposing the constraint can help the convergence of the parameter estimation to the true value.

Three cases of different initial guesses are simulated with constraint imposed by re-weighting individual ensembles based on their likelihood of satisfying the constraint (see (26)). The covariances are $\Sigma_\theta = [1, 0; 0, 1]$ and $\Sigma_l = 0.01$, which are kept the same as previous simulations. The covariance of the constraint used here is $\Sigma_c = 2.0$, which defines a certainty about the constraint. The simulation results are shown in Table 2 and visualized in Figure 6 (left). It can be seen that the solution converges to the true value $(1, 1)$ when starting from

$(0, 0)$ and $(2, 2)$, and the solution converges to the local minimum Group II when starting from $(-2, -2)$. It is interesting to note that the middle initial point $(0, 0)$, which originally converges to Group II, now is able to converge to the true value $(1, 1)$.

The reason that the solution starting from the lower left initial guess $(-2, -2)$ cannot converges to the true value $(1, 1)$ is because it is too far way from the true value and the variance of the prior is not large enough to sample the parameter space near the true value $(1, 1)$. Therefore, even though the constraint has been imposed, the solution cannot converge to the true value. To verify this, we simulated the three different starting locations with a larger prior variance $\Sigma_\theta = [3, 0; 0, 3]$, while the covariance of the data likelihood $\Sigma_l = 0.01$ and the covariance of the the constraint $\Sigma_c = 2.0$ are kept the same. The simulation results are shown in Table 3 and visualized in Figure 6 (middle), demonstrating that all the three initial guesses lead to the true parameter value $(1, 1)$.

As a further test, we decreased the variance of the constraint Σ_c to 1.0 to see how this influences the parameter estimation. The results in Table 4 and Figure 6 (right) demonstrate that (contrary to original conditions in the left panel) all three initial guesses converge to the true parameter value $\theta^* = (1, 1)$. Thus decreasing the variance of the constraint can also improve convergence to the true parameter value in cases where the constraint is more certain.

Relation between Bayesian inference and optimization

Using Bayesian inference framework to estimate the unknown model parameters is intrinsically related to solving




Figure 6: Parameter convergence results with constraint imposed. (Left) When $\Sigma_\theta = [1, 0; 0, 1]$ and $\Sigma_c = 2.0$, the initial guesses at the (2, 2) and (0, 0) converge to the true local minimum (1, 1); (Middle) When $\Sigma_\theta = [3, 0; 0, 3]$ and $\Sigma_c = 2.0$, all initial guesses converge to the true local minimum (1, 1); (Right) When $\Sigma_\theta = [1, 0; 0, 1]$ and $\Sigma_c = 1.0$, all initial guesses converge to the true local minimum (1, 1).

θ^0	(-2, -2)	(0, 0)	(2, 2)
θ^*	(-0.5685, -0.5173)	(0.9839, 1.0013)	(0.9837, 1.0015)
$HF(\theta^*)$	-0.9948	-0.9958	-0.9958
Group	II	I	I

Table 2: Simulation results with the constraint imposed for different initial guesses with $\Sigma_\theta = [1, 0; 0, 1]$, $\Sigma_l = 0.01$ and $\Sigma_c = 2.0$.

θ^0	(-2, -2)	(0, 0)	(2, 2)
θ^*	(0.9947, 0.9962)	(0.9958, 0.9943)	(0.9955, 0.9942)
$HF(\theta^*)$	-0.9859	-0.9897	-0.9908
Group	I	I	I

Table 3: Simulation results with the constraint imposed for different initial guesses with $\Sigma_\theta = [3, 0; 0, 3]$, $\Sigma_l = 0.01$ and $\Sigma_c = 2.0$.

a corresponding optimization problem (Aravkin, Burke, and Pillionetto 2014). To see this, we write out the posterior probability distribution of the unknown parameter θ conditioned on the observed data \mathcal{D} and the fact that the constraint needs to be satisfied,

$$\begin{aligned}
 p(\theta|\mathcal{D}, G(x) = \mathbf{0}) &\propto p(\mathcal{D}|\theta)p(G(x) = \mathbf{0}|\theta)p(\theta) \\
 &= \frac{1}{Z'} \exp \left(-\frac{1}{2} \left\| \Sigma_\theta^{-\frac{1}{2}} (\theta - \hat{\theta}) \right\|^2 - \frac{1}{2} \left\| \Sigma_l^{-\frac{1}{2}} (y - HF(\theta)) \right\|^2 \right. \\
 &\quad \left. - \frac{1}{2} \left\| \Sigma_c^{-\frac{1}{2}} G(F(\theta)) \right\|^2 \right), \quad (39)
 \end{aligned}$$

where Z' is a normalization constant. As mentioned before, the final estimation of the parameter θ can be taken as the posterior expectation $\mathbb{E}(\theta|\mathcal{D}, G(x) = \mathbf{0})$, or as the value that maximizes the posterior probability (MAP)

$$\theta^* = \arg \max_{\theta} p(\theta|\mathcal{D}, G(x) = \mathbf{0}). \quad (40)$$

θ^0	(-2, -2)	(0, 0)	(2, 2)
θ^*	(0.9976, 0.9984)	(0.9888, 1.0063)	(0.9863, 1.0087)
$HF(\theta^*)$	-0.9888	-0.9957	-0.9962
Group	I	I	I

Table 4: Simulation results with the constraint imposed for different initial guesses with $\Sigma_\theta = [1, 0; 0, 1]$, $\Sigma_l = 0.01$ and $\Sigma_c = 1.0$.

Based on (39), solving the MAP estimation of θ is equivalent to the following optimization problem:

$$\min_{\theta} \left\| \Sigma_\theta^{-\frac{1}{2}} (\theta - \hat{\theta}) \right\|^2 + \left\| \Sigma_l^{-\frac{1}{2}} (y - HF(\theta)) \right\|^2 + \left\| \Sigma_c^{-\frac{1}{2}} G(F(\theta)) \right\|^2, \quad (41)$$

where the three terms in the cost function from left to right represent the contributions from the prior, the data and the constraint. Therefore using Bayesian inference to estimate model parameters is equivalently solving an optimization problem of minimizing the miss-match between the observed output and the reconstructed output, while penalizing based on the prior and satisfaction of the constraints. It is easier to see the relations between different terms if we assume all quantities are scalar:

$$\min_{\theta} \frac{1}{\sigma_\theta^2} \|\theta - \hat{\theta}\|^2 + \frac{1}{\sigma_l^2} \|y - HF(\theta)\|^2 + \frac{1}{\sigma_c^2} \|G(F(\theta))\|^2. \quad (42)$$

It can be seen that the variances of the prior, the data and the constraints define the relative importance of each individual terms. The smaller the variance, the more important the corresponding term is in the cost function. This is reasonable because the information source with smaller belief uncertainty should naturally get more weight. Increasing the variance of the prior σ_θ not only samples a broader region but also places more relative weight on satisfying the constraint, which is why increasing the variance of the prior led to convergence of all three initial guesses to the true value (see Table 3 and Figure 6 (middle)). Similarly, decreasing the variance of the constraint can also put more relative weight on the constraint, which is verified in the results shown in Table 4 and Figure 6 (right). In the cases of multiple constraints, the variance for each constraint can be used to tune the relative importance of the constraint.

Extensions

The constrained Bayesian inference approach here was developed to address non-uniqueness of the solutions for inverse problems. However, it can also be extended as a way to solve more general constrained optimization problems. An advantage of this approach, compared to traditional gradient-based optimization, is that it is derivative-free and does not require construction of the cost function gradient. Gradient information is implicitly represented by the ensemble. This approach can also provide a potential framework to incorporate domain knowledge in learning models to accelerate convergence and improve accuracy.

Although we assumed Gaussian distributions, this approach can be extended to non-Gaussian distributions. This ultimately leads to a different “weighting” on the ensemble members (see Eq. (25)). Applying different distributions for the constraints, prior or data can be useful. For example, assuming a Laplace likelihood for the constraint results in L_1 -regularization instead of L_2 -regularization in (41), and a skew likelihood for the constraint will result in different strictness on either side of the constraint surface in parameter space.

4 Conclusion

To address the non-uniqueness of the feasible solutions for ill-posed inverse problems due to model complexity and lack of observation dimension, we here proposed a general method to constrain the inverse problem in a Bayesian inference framework. The constraint is imposed by constructing a likelihood function denoting the fitness of a solution. Then the posterior distribution for the unknown parameter conditioned on both the observation data and the constraint is obtained, and the final parameter estimation is given by the MAP estimation or the posterior mean. This method was also extended to an approximate Bayesian inference framework in terms of the ensemble Kalman filter, which was shown to lead to a re-weighting of ensemble members based on their fitness to the constraint. Numerical simulations were carried out to demonstrate the effectiveness of this approach for basic proof-of-concept.

Acknowledgements

This work was supported in part by the American Heart Association, Award No. 18EIA33900046.

References

Aravkin, A. Y.; Burke, J. V.; and Pillonetto, G. 2014. Optimization viewpoint on Kalman smoothing with applications to robust and sparse estimation. In *Compressed Sensing and Sparse Filtering*. Springer. 237–280.

Asch, M.; Bocquet, M.; and Nodet, M. 2016. *Data assimilation: methods, algorithms, and applications*, volume 11. SIAM.

Cotter, S. L.; Dashti, M.; Robinson, J. C.; and Stuart, A. M. 2009. Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse problems* 25(11):115008.

Del Moral, P.; Doucet, A.; and Jasra, A. 2006. Sequential monte carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 68(3):411–436.

Evensen, G. 2003. The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean dynamics* 53(4):343–367.

Gardner, J. R.; Kusner, M. J.; Xu, Z. E.; Weinberger, K. Q.; and Cunningham, J. P. 2014. Bayesian optimization with inequality constraints. In *ICML*, 937–945.

Haykin, S. 2004. *Kalman filtering and neural networks*, volume 47. John Wiley & Sons.

He, Y., and Xiu, D. 2016. Numerical strategy for model correction using physical constraints. *Journal of Computational Physics* 313:617–634.

Iglesias, M. A.; Lin, K.; and Stuart, A. M. 2014. Well-posed Bayesian geometric inverse problems arising in subsurface flow. *inverse problems* 30(11):114001.

Li, Z.; Zhang, H.; Bailey, S. C.; Hoagg, J. B.; and Martin, A. 2017. A data-driven adaptive Reynolds-averaged Navier–Stokes $k-\omega$ model for turbulent flow. *Journal of Computational Physics* 345:111–131.

Penfield Jr, P. 2010. Principle of maximum entropy. *Information, Entropy and Computation* 104–12.

Shao, X.; Huang, B.; and Lee, J. M. 2010. Constrained Bayesian state estimation—a comparative study and a new particle filter based approach. *Journal of Process Control* 20(2):143–157.

Simon, D., and Chia, T. L. 2002. Kalman filtering with state equality constraints. *IEEE transactions on Aerospace and Electronic Systems* 38(1):128–136.

Smedstad, O. M., and O’Brien, J. J. 1991. Variational data assimilation and parameter estimation in an equatorial pacific ocean model. *Progress in Oceanography* 26(2):179–241.

Wan, E. A., and Van Der Merwe, R. 2000. The unscented Kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, 153–158. Ieee.

Wang, J.-X., and Xiao, H. 2016. Data-driven CFD modeling of turbulent flows through complex structures. *International Journal of Heat and Fluid Flow* 62:138–149.

Wang, J.-X.; Tang, H.; Xiao, H.; and Weiss, R. 2017. Inferring tsunami flow depth and flow speed from sediment deposits based on ensemble Kalman filtering. *Geophysical Journal International* 212(1):646–658.