XX IMEKO World Congress
Metrology for Green Growth
September 9-14, 2012, Busan, Republic of Korea

FAST ESTIMATION AND UNCERTAINTY QUANTIFICATION IN ELECTRICAL CAPACITANCE TOMOGRAPHY USING SURROGATE TECHNIQUES

M. Neumayer, D. Watzenig, and G. Steiner

Institute of Electrical Measurement and Measurement Signal Processing,
Kopernikusgasse 24/4, A-8010 Graz, Austria, Graz University of Technology
Email: neumayer@tugraz.at

Abstract: The need for uncertainty quantification (UQ) in metrology has seen serious research efforts and is of ever-growing interest in order to quantify the quality of measurement results. This especially holds for indirect measurement problems. Bayesian methods allow a natural and universal access for UQ but become computationally expensive for scenarios with a complex interaction between the unknown quantity x and the raw measurements $\tilde{d}$. In this paper we present a surrogate approach for fast estimation and UQ of material distributions in the inverse problem of electrical capacitance tomography using a recursive Bayesian estimator.

Keywords: uncertainty quantification, electrical capacitance tomography, surrogate model, recursive Bayesian methods.

1. INTRODUCTION

Indirect measurement techniques such as parameter estimation tasks, remote sensing applications, or inverse problems in general, have seen rigorous growth in the past decades due to the availability of computational processing power for the numerical solution of them. This type of measurement can formally be written as $P : x \mapsto \tilde{d}$, where $x \in \mathbb{R}^N$ are the unknown parameters, $\tilde{d} \in \mathbb{R}^M$ are the (noisy) measurements and P describes the measurement process. P relates the N dimensional image space with the M dimensional data space. In this work we assume an additive noise model meaning that $\tilde{d}$ is given by $\tilde{d} = d + v$, $v \in \mathbb{R}^M$, where d should denote the noise free measurements which are not accessible. For the case of complex interactions between x and $\tilde{d}$ the determination of x given $\tilde{d}$ can become an intractable problem. This is typically the case for problems were P contains an underlying partial differential equation (PDE). Problems out of this class are nonlinear inverse problems like the industrial problem of electrical capacitance tomography (ECT) [5], which will also be the subject under investigation in this work. Due to their inherent ill-posed nature a reasonable solution cannot even be obtained. In this case prior knowledge has to be incorporated to obtain a meaningful solution. Like in all other fields of metrology the need for uncertainty quantification (UQ) exists for this class of indirect measurement problems, too. However, due to their complex nature the propagation of uncertainty (like the measurement noise v) in the data space $\mathbb{R}^M$ onto the image space $\mathbb{R}^N$ can typically not be treated by standard methods of UQ. Bayesian methods provide a powerful framework for UQ as it gives the possibility to compute any statistic about x from the posterior density function [1]. Central component of all Bayesian methods is Bayes' law, which is given by

$$\pi(x|\tilde{d}) = \frac{\pi(\tilde{d}|x)\pi(x)}{\pi(\tilde{d})} \propto \pi(\tilde{d}|x)\pi(x). \tag{1}$$

Hereby $\pi(\tilde{d}|x)$ is referred to as the likelihood function and $\pi(x)$ is referred to as the prior. The term $\pi(\tilde{d})$ is termed as evidence. The evidence has the role of a normalization constant and can be skipped for practical reasons leading to the right hand side term in equation (1). The likelihood function $\pi(\tilde{d}|x)$ incorporates knowledge about the measurement noise v. The prior $\pi(x)$ rates the probability about x being a solution. Thus, Bayes' law provides a probability measure for x being the solution for the problem given the data $\tilde{d}$ and the prior.

A class of algorithms out of the pool of the Bayesian framework are sequential Bayesian estimators with its most prominent member being the Kalman filter (KF) [2]. The KF is an iterative filter algorithm suitable for state estimation in dynamical systems. It considers all variables to be normal distributed random variables, i.e. x is completely determined by a mean μ_x and a covariance matrix Σ_x. Subsequently, the KF is less adequate for problems where non-Gaussian distributed random variables appear. An alternative filter able to handle arbitrary distributions is the Particle filter (PF) [3]. The PF utilizes a sampling procedure making it computationally expensive as an evaluation of the forward map F is required for every sample and in every iteration step. Hereby, $F : x \mapsto y$ presents a (computer) model of the measurement process P. This circumstance offers essential limitations for the real time performance of the PF on

systems like ECT.

Recently approximation or surrogate techniques have been proposed to speed up the computations in ECT [4]. These techniques replace the computationally expensive forward map F by a cheap approximation F^*. Subsequently an approximation error $E = F - F^*$ occurs. However, the Bayesian framework allows an incorporation of this type of error as well. Based on this fact we take advantage of the so called enhanced error model [1] to apply a PF approach capable for real time UQ in ECT.

This paper is structured as follows. In section 2 we briefly describe the inverse problem of ECT and the used framework. Section 3 gives details about the PF and how this filter can be used for UQ. In section 4 we describe the design of an affine surrogate model and how to combine knowledge about the approximation error into the inference process. Section 5 presents exemplary results and an analysis of the approach including simulated and measured data.

2. ELECTRICAL CAPACITANCE TOMOGRAPHY AND RBF SHAPE MODEL

Figure 1(a) schematically depicts an ECT sensor [5]. A number of N_{elec} electrodes are mounted on the exterior of the process pipe. The boundaries of the electrodes are referred to as Γ_i, $i = 1 \dots N_{\text{elec}}$. The problem domain Ω is enclosed by a screen which surface is denoted by $\partial\Omega$. The interior of the pipe is referred to as region of interest Ω_{ROI}. Electric fields in Ω are governed by the potential equation $\nabla \cdot (\varepsilon_0 \varepsilon_r \nabla V) = 0$, where V is the electric scalar potential. $\varepsilon_0 = 8.854 \times 10^{-12}\,\text{AsV}^{-1}\text{m}^{-1}$ and the dimensionless ε_r are the absolute and the relative permittivity, respectively. The boundary conditions on all electrodes and on the screen are of Dirichlet type and given by $V_{\partial\Omega} = 0$ on the screen, $V_{\Gamma_j} = V_0$ on the transmitter electrode and $V_{\Gamma_i} = 0$, $i \neq j$, on the remaining $(N_{\text{elec}} - 1)$ receiver electrodes. Given the N_{elec} solutions V_j, the $\frac{N_{\text{elec}}(N_{\text{elec}}-1)}{2}$ inter electrode capacitances are computed by Gauss's law

$$c_{i,j} = \frac{1}{V_0} \int_{\Gamma_i} \vec{n} \cdot \varepsilon_0 \varepsilon_r \nabla V_j \, ds, \tag{2}$$

to assemble the matrix $C = [c_{i,j}]$. $\vec{n}$ denotes the perpendicular vector on Γ_i. Each column of the symmetric capacitance matrix C corresponds to one transmitter electrode and each line to one receiver electrode. The computation steps follow the measurement process for low-Z receiver circuitry. In this scheme the displacement current at the receiver electrode to ground is measured, while the transmitter electrode is connected to an AC voltage. For the computation of the capacitance matrix C we take use of the finite element method (FEM). During the evaluation of the forward map $F : x \mapsto C$, N_{elec} solutions of the governing PDE are required. With the numerical scheme of the FEM this requires the solution of an equation system of form $Kv = r$ for N_{elec} times, making the evaluation of F computationally expensive and time consuming.

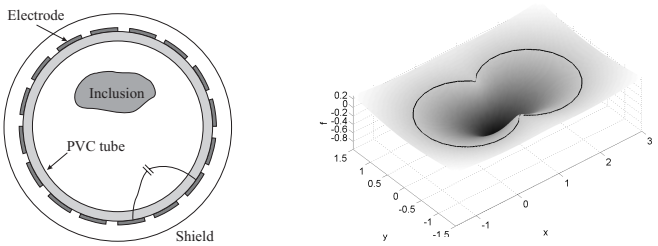

(a) ECT Scheme.

(b) Exemplary RBF for a shape (bold line).

Fig. 1. Scheme of an ECT sensor and example for the function Φ of the RBF shape description (bold line).

In this work we aim on the estimation of the shape $\partial\Omega_{\text{i}}$ of a homogeneous inclusion within a homogeneous background material (i.e. an oil bubble in air) with known permittivity values.

For the shape description a shape model based on radial basis functions (RBFs) [6] is used. Hereby the boundary $\partial\Omega_{\text{i}}$ is given by the set $\partial\Omega_{\text{i}} = \{z | \Phi(z) = 0\}$, where z presents the cartesian coordinates. The design of the function Φ is presented in [6]. Figure 1(b) depicts an example shape and the corresponding surface of Φ. Hereby, Φ or $\partial\Omega_{\text{ROI}}$ are parameterized by the vector x. As can be seen, the function Φ is designed to have a negative sign for points inside $\partial\Omega_{\text{i}}$. The sign of Φ is used to identify the finite elements in Ω_{ROI} which present the inclusion described by $\partial\Omega_{\text{i}}$. This step is referred to as mapping and is computationally efficient due to the sign property of Φ.

Summarizing the shape description in combination with the finite element approach to compute C we refer to this computation as forward map $F : \partial\Omega_{\text{i}} \mapsto C$. The inverse problem now consists in the estimation of $\partial\Omega_{\text{i}}$ from measured data $\tilde{d}$ and UQ about $\partial\Omega_{\text{i}}$.

3. THE PARTICLE FILTER

To apply the PF to the given problem, the data representation x and the model (forward map) have to be recast into the state space form [2]

$$x_k = f_k(x_{k-1}, w_{k-1}), \tag{3}$$
$$y_k = h_k(x_k, v_k), \tag{4}$$

where f_k presents the transition of the state over time, and equation (4) is referred to as measurement equation. v_k

and w_k are the measurement noise and the process noise, respectively. For the presented shape determination task the model output y_k corresponds to the measurements $\tilde{d}_k$ and x_k corresponds to the shape $\partial\Omega_{\mathrm{i}}$. The PF belongs to the class of Bayesian recursive filters. Hence, it estimates the current state x_k out of temporal measurement information $\tilde{d}_l, l = 0,1,\ldots,k$. For this, the PF uses a set of $m = 1,\ldots,M$ particles $x^{(m)}$ with their corresponding weights $w^{(m)}$.

In the first computation step of a PF iteration, equal weights $w^{(m)}$ are proposed to all particles $x_{k-1}^{(m)}$, before they propagate through the system (3). This prediction process of $x_k^{(m)}$ is determined by a deterministic drift through the dynamics of f_k and a stochastic diffusion based on the process noise w. Using the transition probability $\pi(x_k|x_{k-1})$, the probability $\pi(x_k|\tilde{D}_{k-1})$ of the predicted state follows the Chapman-Kolmogorov equation

$$\pi(x_k|\tilde{D}_{k-1}) = \int_{\mathbb{R}^N} \pi(x_k|x_{k-1})\pi(x_{k-1}|\tilde{D}_{k-1})dx_{k-1}, \quad (5)$$

where $\tilde{D}_{k-1} = \{\tilde{d}_1,\ldots,\tilde{d}_{k-1}\}$ collects the temporarily history of measurements. The second step uses the new measurement information $\tilde{d}_k$ to evaluate the posterior density

$$\pi(x_k|D_k) \propto \pi(d_k|x_k)\pi(x_k|D_{k-1}), \quad (6)$$

using Bayes' law on the predicted particles. This corresponds to the weights $w_k^{(m)}$. The term $\pi(\tilde{d}_k|x_k)$ is the likelihood. For additive Gaussian measurement noise it is given

$$\pi(\tilde{d}_k|x_k) \propto \exp\left(-\frac{1}{2}(y_k - \tilde{d}_k)^T\Sigma^{-1}(y_k - \tilde{d}_k)\right), \quad (7)$$

where y_k contains the simulated elements corresponding to the data $\tilde{d}$ and $\Sigma = \Sigma_v$ is the covariance matrix of the measurement noise.

To ensure the properties of a probability density function for the weights $w_k^{(m)}$, they subsequently have to be normalized by their sum. Further, a resampling scheme has to be applied to avoid the degeneration of the algorithm. In this work residual resampling was used [2]. The expectation of x_k can then be estimated by

$$\hat{x}_k = E(x_k|D_k) \approx \sum_{m=1}^{M} w_k^{(m)} x_k^{(m)}. \quad (8)$$

By similar schemes of equation (8) arbitrary statistics about x can be computed for UQ.

3.1. Particle Filter for ECT

To apply the particle filter for shape estimation and UQ in ECT we combine the RBF shape model with the finite element model in order to compute y (corresponds to the elements of C) to form the measurement equation h_k. For the state evolution f_k of the particles in the ECT example, we used the four moves translation, rotation, scale and corner move [7]. The effect on $\partial\Omega_{\mathrm{ROI}}$ of the four moves is depicted in figure 2. The selection of a specific move for a particle is done randomly by means of a uniform distribution. The manipulation of the particles by the selected move is done randomly and the deviation of the moves follows a normal distribution. Thus, we only consider a stochastic diffusion for the state evaluation and no deterministic propagation.

The corner move depicted in figure 2(d) is necessary to manipulate the shape of the inclusion. Without the corner move the PF is not able to reconstruct inclusions which differ from the initial shape by corners. I.e. for circular inclusions used in the initialization, the span of reconstructible inclusions only includes other circles and ellipsoidal objects. However, the setup of the corner move is critical and it was reported, that the statistical efficiency of the corner move is low [7]. Alternative moves like a beveling move for the shape might be considered, but we decided to use the presented moves.

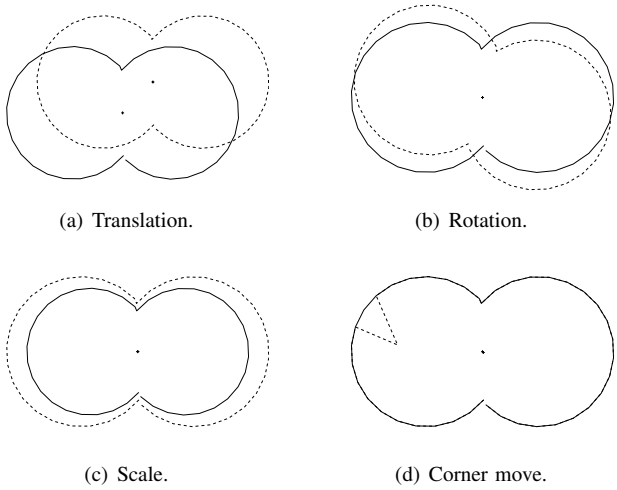

(a) Translation. (b) Rotation. (c) Scale. (d) Corner move.

Fig. 2. The four basic moves from the shape $\partial\Omega_{\mathrm{ROI,k}}$ (state x_{k-1}, bold line) to the state $\partial\Omega_{\mathrm{ROI,k+1}}$ (x_k dashed line).

4. AN AFFINE SURROGATE FOR ECT

The computational bottleneck for the application of the PF are the M evaluations of the forward map F for the determination of the likelihood of the particles in every iteration. As a certain number of particles is required to achieve a certain quality of the output, M has a lower limit. The runtime of F can be decreased by efficient computation techniques, but remains the limiting part. Recently the idea of applying fast surrogate techniques to replace F by an approximation F^* has been proposed for ECT [4]. A

surrogate model F^* is an approximation of F which ideally has a considerable lower runtime than F but it introduces a deterministic model error $e = F - F^*$.

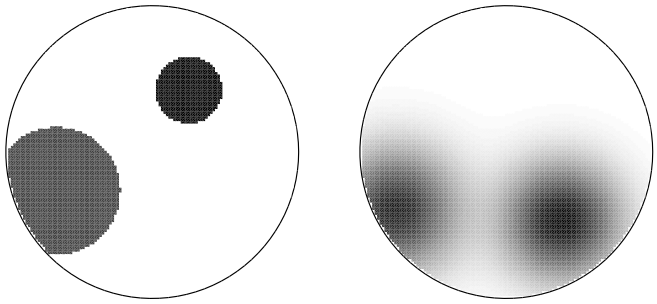

(a) Rod like distribution.

(b) Gaussian bump like distribution.

Fig. 3. Examples for random patterns to generate the approximation F^*.

In this work we choose the approach to replace the computation of $c_{i,j}$ by $c^*_{i,j} = p^T_{i,j}\hat{\varepsilon}$, where $\hat{\varepsilon}$ holds the permittivities of the finite elements in Ω_{ROI}. For the computation of the individual vectors $p_{i,j}$ random patterns about the expected inclusion scenario are generated [**?**]. Figure 3 depicts two exemplary patterns for expected material distributions. Figure3(a) depicts a scenario where rod like inclusions are places inside Ω_{ROI}. Figure 3(b) depicts a scenario where Gaussian bumps are assumed. Thus, a smoothness assumption about $\pi(x)$ can be applied. For each of the patterns the corresponding charges are computed [**?**]. Then $p_{i,j}$ can be found by solving the equation system $\mathcal{E}p_{i,j} = c_{i,j}$ using the pseudo inverse. The rows of the matrix $\mathcal{E}$ contain the individual vectors ε^T of the patterns. In order to achieve an affine approximation a leading 1 is added to ε^T.

Figure 4 depicts a validation of the approximation error by means of a Gaussian distribution. This means that we computed the mean and the standard deviation of the approximation error on a second set of sampled patterns.

The approximation was generated for inclusions with a permittivity of $\varepsilon_r = 3.5$. As can be seen, the mean μ_e stays almost zero. The variance σ_e increases for measurements where the receiver electrode is close to the transmitter electrode. In terms of the signal to noise ratio (SNR), the approximation error e leads to a lower SNR value of about 35 dB. Without the Bayesian approach one would have to accept this additional error. However, the Bayesian framework allows to incorporate knowledge about e by means of its probability density function (pdf).

Figure 5 depicts the pdf of one individual error $e_{i,j}$. As can be seen the distribution is nonsymmetric and has a tail. In order to keep $\pi(e)$ a simple function we applied the idea of the enhanced error model [1] and modeled e by means of a Gaussian distribution $e \propto \mathcal{N}(\mu_e, \Sigma_e)$. By this, the covariance matrix Σ becomes $\Sigma = \Sigma_v + \Sigma_e$. Thus, the likelihood (7) considers the approximation error in form of an increased measurement noise.

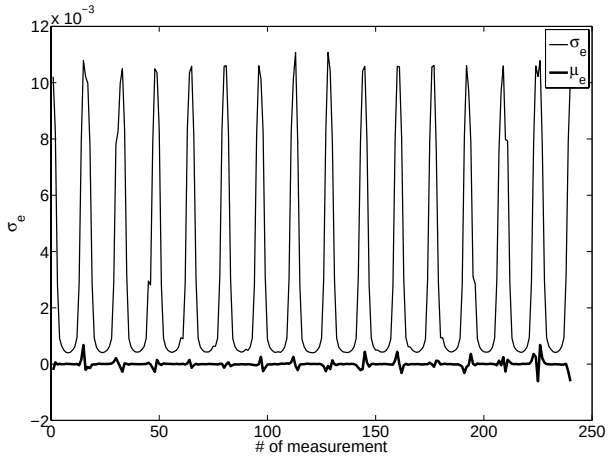

Fig. 4. Standard deviation and mean of the approximation error for the 240 measurements.

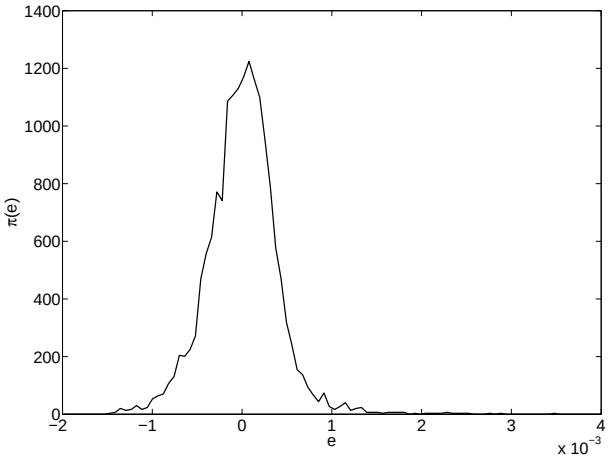

Fig. 5. Probability density function of the approximation error $e_{i,j}$.

5. RESULTS

In this section we will provide simulation and measurement results about the applicability of the presented approach. A simulation study is presented in subsection 5.1. Results using measured data are presented in subsection 5.2. For both types of experiments we considered the presence of one inclusion, although we generated the approximation also for scenarios of several inclusions as depicted in figure 3. For the prior $\pi(x)$ we just used a simple $1/0$ prior, meaning that $\pi(x)$ was set to zero for particles outside Ω_{ROI}.

5.1. Simulation Results

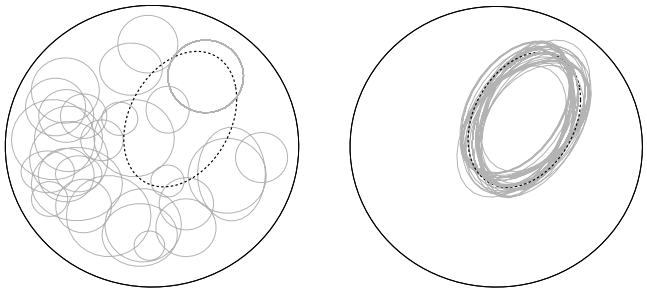

(a) Initialization. (b) Different particles.

Fig. 6. Typical images during the PF operation.

For the experiments we computed the capacitances for an ellipsoidal inclusion ($\varepsilon_r = 3.5$) on an $N_{\text{elec}} = 16$ electrode ECT system. To obtain $\tilde{d}$ the computed data was corrupted with normal distributed noise with a standard deviation of $\sigma = 1 \times 10^{-3}$. The number of particles was set to $M = 20$. Figure 6 depicts some scenarios during the PF operation. Figure 6(a) depicts the initial particles $x_0^{(m)}$. The bold ellipsoidal contour presents the true material distribution. Figure 6(b) depicts the particles after some iterations. As can be seen, the bulk of the particles cover the true inclusion. By the use of the surrogate model F^* a frame rate of about 15 frames/s could be achieved on a standard PC. For comparison, the use of F results in a frame rate of less than 4 frames/s. Figure 7 depicts the trend of the estimated center coordinates and the area of the inclusion using the conditional mean over the samples for some iterations. Figure 8 depicts the distribution of relative error for one component of the center coordinates. Due to the availability of several particles and their weight UQ becomes possible.

5.2. Measurement Results

Figure 10 and **??** depict two reconstruction results using measured data. The two test scenarios considered a large inclusion placed close to the boundary of the pipe and a small inclusion placed right to the center. The inclusions were PVC rods with a relative permittivity of $\varepsilon_r = 3.5$. Figure 10 and figure 11 depict the reconstruction results. For the left hand side figures the accurate forward map F was used. The right hand side pictures depict the results where the fast surrogate approximation F^* has been used. The dashed shapes in figure 10 and 11 depict the position of the rod as depicted in the corresponding photographs. As can be seen, both approaches provide a reasonable estimate of the position and the shape of the contour - yet the version of the PF using F^* provides real time capability.

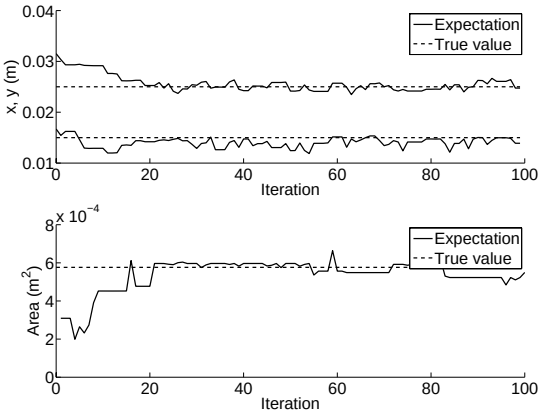

Fig. 7. Trend of the expectation of the center position and the area of the inclusion. The dashed lines mark the true values.

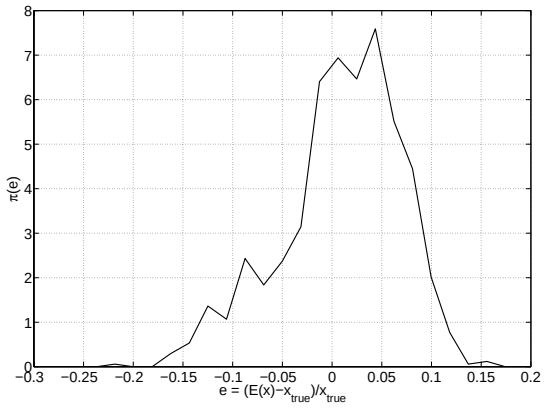

Fig. 8. Distribution of the relative estimation error or the x-component of the center coordinate.

The approximation works for both, the large and the small inclusion, well. Due to the availability of several particles and their weight $w_i^{(m)}$ UQ becomes possible. The fact that the approximation was generated for a prior where several inclusions are considered, does not influence the results. The result of the small reconstruction using the accurate model F (figure 11(a)) appears to be biased. We addressed this bias to be causes by a problem in the calibration data of the sensor. Beside this the result is acceptable.

6. CONCLUSION

In this work we presented an accelerated particle filtering for ECT using an affine surrogate approximation of the forward map. Within the Bayesian framework the approximation error $e = F - F^*$ can be incorporated into the estimation process. The results demonstrate the real time applicability of sampling methods for Bayesian inference in inverse problems for accelerated estimation and UQ. The applicability of

(a) Photography of the large inclusion.

(b) Photography of the small inclusion.

Fig. 9. Photographes of the two inclusions used for the experiments with measured data.

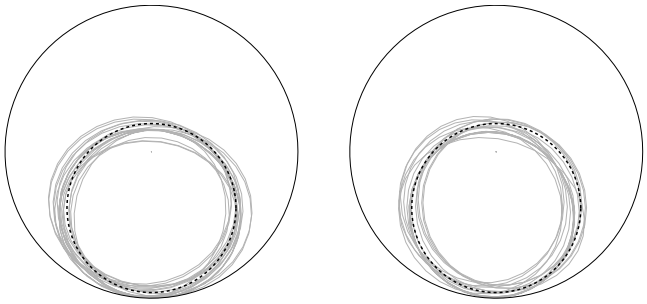

(a) Result using F.

(b) Result using F^*.

Fig. 10. Reconstruction results with measured data for the large inclusion (see photo 9(a)).

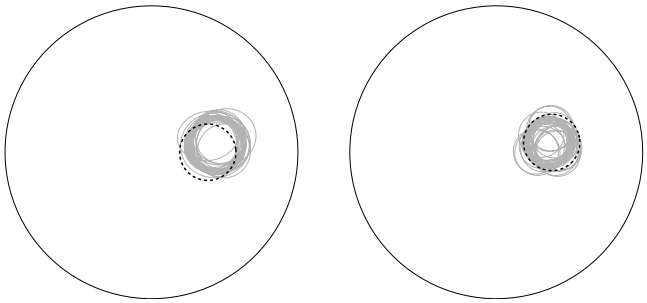

(a) Result using F.

(b) Result using F^*.

Fig. 11. Reconstruction results with measured data for the small inclusion (see photo 9(b)).

the approach was demonstrated for simulated and measured data from an ECT sensor. It was demonstrated that a different prior for the setup of the surrogate approximation is not critical with respect to the scenario of the reconstruction. The approach can be applied to different inverse problems like electrical impedance tomography (EIT) or similar other problems.

7. REFERENCES

[1] J. P. Kaipio and E. Somersalo, "Statistical and computational inverse problems", New York: Applied Mathematical Sciences, Springer, 2004.

[2] A. Doucet, N. de Freitas, and N. J. Gordon, "Sequential Monte Carlo methods in practice", New York: Springer, 2001.

[3] M. S. Arulampalam, S. Maskell, N. Gordon, T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking", IEEE Transactions on Signal Processing, vol.50, no.2, pp.174-188, Feb 2002.

[4] D. Watzenig, M. Neumayer, and C. Fox, "Accelerated Markov chain Monte Carlo sampling in Electrical Capacitance tomography", in the int. Jnl. of Comp. and Math. in elect. and electronic eng. (Compel), vol. 30, no. 6, 2011

[5] M. Neumayer, H. Zangl, D. Watzenig, and A. Fuchs, " Current reconstruction algorithms in Electrical Capacitance Tomography", in New Developments and Applications in Sensing Technology, Springer, 2011.

[6] K. Uhliir, J. Patera, and V. Skala, "Radial basis function method for iso-line extraction", in Elect. Comp. and Informatics, 2004, pp. 439-444.

[7] D. Watzenig and C. Fox, "A review of statistical modelling and inference for electrical capacitance tomography", in Jnl. of Meas. Sci. and Techn., vol. 20, no. 5, 2009, 052002.