Born approximation for inverse scattering with high contrast media

Jari Kaipio[1,2]

[1]Department of Mathematics
University of Auckland

[2]Department of Applied Physics
University of Eastern Finland

Joint work with
Timo Lahivaara, Teemu Luostari, Tomi Huttunen and Peter Monk

Overview

- The Born approximation in the context of inverse scattering is used to obtain a linear approximation to the forward map from index of refraction to the far-field pattern.
- It works well for weak scatterers but tends to fail for strong scatterers:

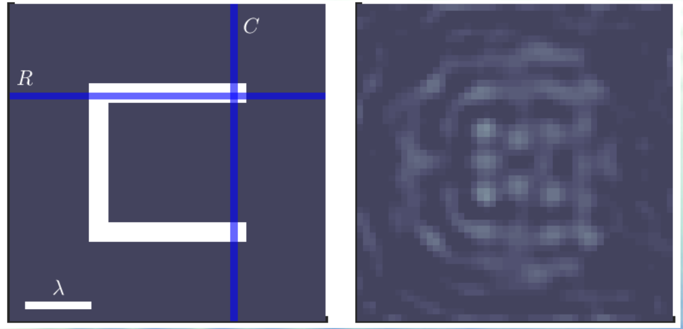

The forward transmission problem

- Let $u(x)$, $u^i(x,d)$ and $u^s(x)$ be the total wave, the incident wave and the scattered wave, respectively, and let $x, d \in \mathbb{R}^2$.
- The time-harmonic scattering problem can be written as

$$\begin{aligned} \Delta u + k^2 n(x) u &= 0 \text{ in } \mathbb{R}^2, \\ u &= u^i + u^s \text{ in } \mathbb{R}^2, \\ r^{1/2}\left(\frac{\partial u^s}{\partial r} - iku^s\right) &\to 0 \text{ as } r := |x| \to \infty. \end{aligned}$$

  where $n(x)$ is the (squared) refraction index, $k = \omega/c$ is the wave number, d is the direction of the incident wave, ω is the frequency and c is the wave speed where $n = 1$.
- Let the scatterer/contrast $m(x) = 1 - n(x)$ have a bounded support such that supp $m \subset D$.

The far-field pattern

- The total field u satisfies the Lippmann-Schwinger equation

$$u(x) = u^i(x) - k^2 \int_D \frac{i}{4} H_0^{(1)}(k|x-y|) m(y) u(y)\, dA(y), \qquad \forall x \in \mathbb{R}^2$$

where $H_0^{(1)}$ is the Hankel function of first kind and order zero.

- The far-field pattern can be computed as

$$u_\infty(\hat{x}, d) = -\frac{k^2 \exp(i\pi/4)}{\sqrt{8\pi k}} \int_D \exp(-ik\hat{x}\cdot y) m(y) u(y, d)\, dA(y)$$

where $\hat{x}$ is the direction of the out-going wave.

- The forward problem is to compute $u_\infty(\hat{x}, d)$ for a collection of $(d, \hat{x})$.

"Accurate" numerical approximation of the forward problem

- We can write

$$\Delta u^s + k^2 n(x) u^s = k^2(1 - n(x))u^i = k^2 m(x) u^i$$

  where the right hand "source" term has the same support as m.
- We use the finite element method with a perfectly matching layer with a computational domain $\bar{\Omega}$ that is $\sim$ twice the size of the scatterer.
- The far field pattern is then computed as above.
- Let us denote the respective model predictions as $\mathbb{R}^M \ni \bar{F}(\bar{m})$ where the barred entities refer to "accurate" representations.

The Born approximation

- The Born approximation in this context is simply the first term of the Neumann series

$$u(x) \approx u^i(x) - k^2 \int_D \frac{i}{4} H_0^{(1)}(k|x-y|) m(y) u^i(y)\, dA(y), \qquad \forall x \in \mathbb{R}^2$$

- This yields the approximate forward map

$$u_\infty(\hat{x}, d) \approx -\frac{k^2 \exp(i\pi/4)}{\sqrt{8\pi k}} \int_D \exp(ik(d-\hat{x})\cdot y) m(y)\, dA(y).$$

which (linear map) we denote by $F(m)$

- Note: the Neumann series converges usually only for very low contrasts $m \sim 1.07$.
- Weak scatterer: $(ka)^2 \|m\|_{L^\infty(D)}$ is "sufficiently small", where a is the size of the scatterer (D).

Bayesian models with auxiliary variables

- All information of the random variables χ is decoded in the joint density $\pi(\chi)$, in our case $\chi = (Y, m, e, \varepsilon)$ where
  - Measurements: Y
  - Primary unknowns: m
  - Measurement noise: e
  - Other secondary unknows: ε
- So, one is interested in the RV's m, the RV's Y have been measured and the RV's (e, ε) are uninteresting. Then the task is to model the conditional distribution

$$\pi(m \mid Y) \propto \int \int \pi(m, e, \varepsilon \mid Y) \, de \, d\varepsilon$$

which expresses the uncertainty of m given Y.

The measurement model

- We pose the additive measurement error model

$$\begin{aligned} Y &= \bar{F}(\bar{m}) + e \\ &= F(m) + e + \underbrace{\bar{F}(\bar{m}) - F(m)}_{\varepsilon(\bar{m})} \end{aligned}$$

  with $e \sim \mathcal{N}(e_*, \Gamma_e)$.
- One can marginalize over e to yield

$$\pi(Y \mid m) = \int \pi_e(Y - F(m) - \varepsilon)\pi_{\varepsilon|m}(\varepsilon \mid m)\, d\varepsilon$$

- At this stage, we approximate the joint density $\pi(\varepsilon, m)$ with a normal model

$$\pi(\varepsilon, m) \propto \exp\left\{-\frac{1}{2}\begin{pmatrix} \varepsilon - \varepsilon_* \\ m - m_* \end{pmatrix}^{\mathrm{T}} \begin{pmatrix} \Gamma_{\varepsilon\varepsilon} & \Gamma_{\varepsilon m} \\ \Gamma_{m\varepsilon} & \Gamma_{mm} \end{pmatrix}^{-1} \begin{pmatrix} \varepsilon - \varepsilon_* \\ m - m_* \end{pmatrix}\right\}$$

- For the approximate conditional density $\pi(\varepsilon \mid m) = \mathcal{N}(\varepsilon_{*|m}, \Gamma_{\varepsilon|m})$ we can then write

$$\begin{aligned} \varepsilon_{*|m} &= \varepsilon_* + \Gamma_{\varepsilon m}\Gamma_{mm}^{-1}(m - m_*), \\ \Gamma_{\varepsilon|m} &= \Gamma_{\varepsilon\varepsilon} - \Gamma_{\varepsilon m}\Gamma_{mm}^{-1}\Gamma_{m\varepsilon}. \end{aligned}$$

- Define the normal random variable ν so that $\nu \mid m = e + \varepsilon \mid m$ then

$$\nu \mid m \sim \mathcal{N}(\nu_{*|m}, \Gamma_{\nu|m})$$

where

$$\nu_{*|m} = e_* + \varepsilon_* + \Gamma_{\varepsilon m}\Gamma_{mm}^{-1}(m - m_*), \quad (1)$$

$$\Gamma_{\nu|m} = \Gamma_e + \Gamma_{\varepsilon\varepsilon} - \Gamma_{\varepsilon m}\Gamma_{mm}^{-1}\Gamma_{m\varepsilon} \quad (2)$$

- Which yields the approximate likelihood

$$\pi(Y \mid m) = \mathcal{N}(Y - F(m) - \nu_{*|m}, \Gamma_{\nu|m})$$

The posterior model

- For the inversion/inference, we adopt “a normal approximation for the prior model”

$$\pi(m) = \mathcal{N}(m_*, \Gamma_{mm}) \tag{3}$$

- The approximation for the posterior distribution can thus be written as

$$\pi(m \mid Y) \propto \pi(Y \mid m)\pi(m) \propto \exp\left(-\frac{1}{2}V(m \mid Y)\right)$$

where $V(m \mid Y)$ is the posterior potential that can be written in the form

$$V(m \mid Y) = \|L_{\nu|m}(Y - F(m) - \nu_{*|m})\|^2 + \|L_m(m - m_*)\|^2 \tag{4}$$

where $\Gamma_{\nu\,|\,m}^{-1} = L_{\nu|m}^{\mathrm{T}} L_{\nu|m}$ and $\Gamma_{mm}^{-1} = L_m^{\mathrm{T}} L_m$.

- We aim to compute the conditional mean (minimizer of the posterior potential) to obtain a "precomputed" estimator since we can write

$$\mathbb{E}(m \,|\, Y) = BY + c$$

- The posterior covariance can also be precomputed

$$\Gamma_{m|Y} = \left(\tilde{F}^{\mathrm{T}} \Gamma_{\nu|m}^{-1} \tilde{F} + \Gamma_{mm}^{-1} \right)^{-1} \tag{5}$$

  where $\tilde{F} = F + \Gamma_{\varepsilon m} \Gamma_{mm}^{-1}$. A further approximation, that is referred to as the enhanced error model, is obtained by setting $\Gamma_{\varepsilon m} = 0$.

- The overall approach is called the Bayesian approximation error approach.

The priors $\bar{\pi}(\bar{m})$, $\pi(m)$, $\bar{m}$ and m

- The actual prior $\bar{\pi}(\bar{m})$: 1-3 ellipses (inclusions) with random centers, orientations, eccentricities and contrasts $m \in (0, 1)$.
- Draw samples $\bar{m}^{(\ell)}$ from this prior model and compute projections $m^{(\ell)}$ and the approximation errors $\epsilon^{(\ell)}$
- Compute the joint second order statistics of (ε, m)
- The prior for the inversion/inference is $\pi(m) = \mathcal{N}(0, \Gamma_{mm})$ where Γ_{mm} is an isotropic homogeneous Ornstein-Uhlenback covariance with characteristic length λ (wavelength outside scatterer) and marginal variances $\mathrm{var}\,(m_k) = 0.4^2$ for all k.
- m is discretized in a 50×50 rectangular grid.

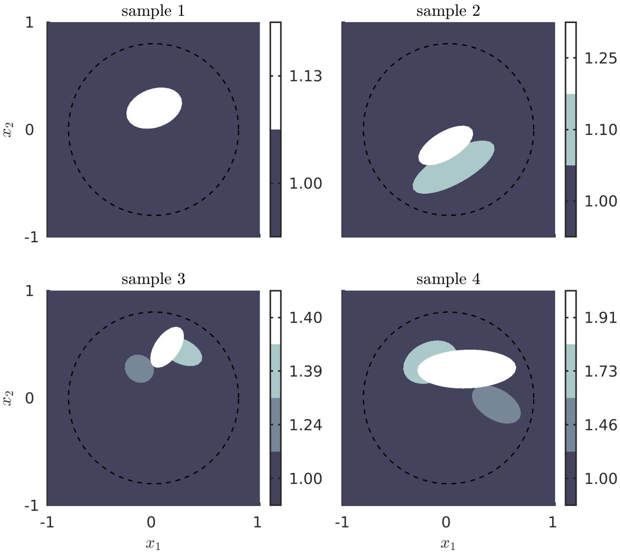
sample 1
sample 2
sample 3
sample 4
1
0
-1
x_2
1.13
1.00
1.25
1.10
1.00
1.40
1.39
1.24
1.00
1.91
1.73
1.46
1.00
-1
0
1
x_1

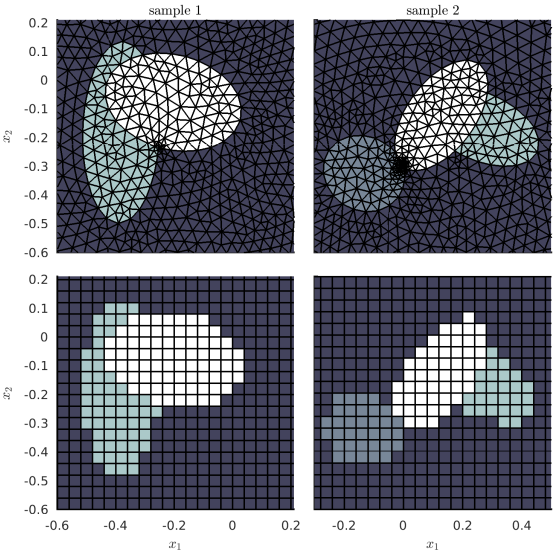
sample 1
sample 2
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
-0.5
-0.6
x_2
-0.6
-0.4
-0.2
0
0.2
0.4
x_1

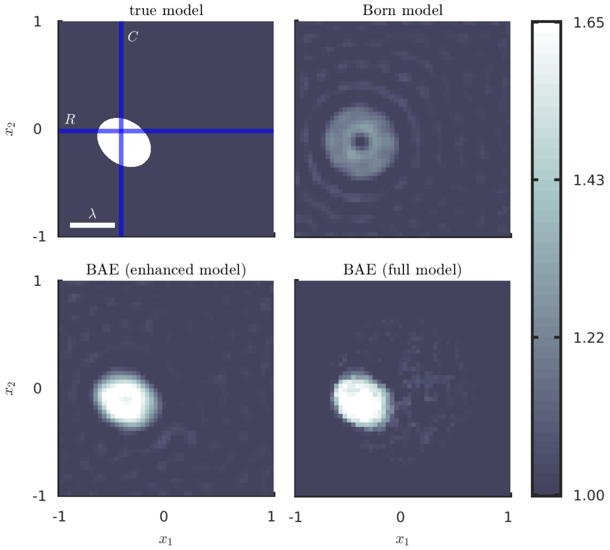
true model
Born model
BAE (enhanced model)
BAE (full model)
C
R
λ
1
0
-1
x_2
x_1
1.65
1.43
1.22
1.00

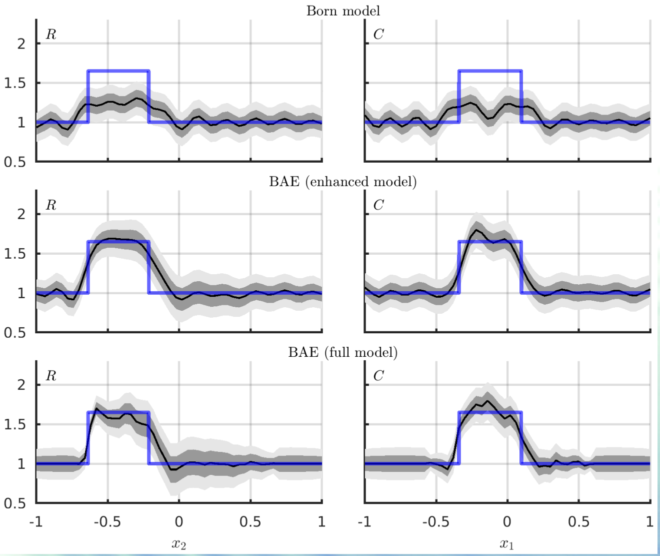
Born model
BAE (enhanced model)
BAE (full model)
R
C
2
1.5
1
0.5
-1
-0.5
0
0.5
1
x_2
x_1

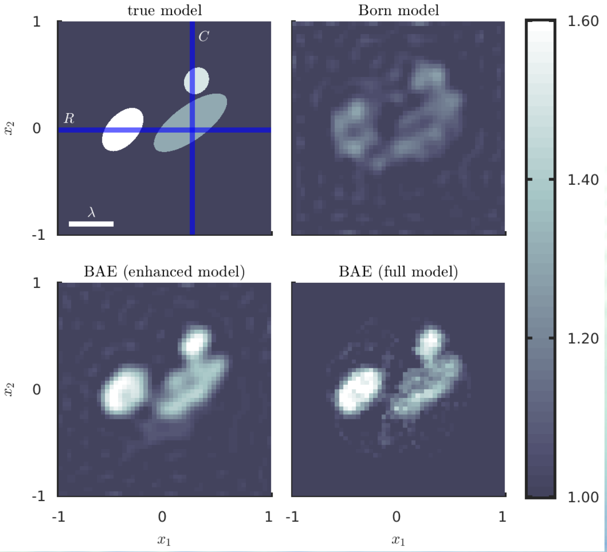
true model
Born model
BAE (enhanced model)
BAE (full model)
C
R
λ
1
0
-1
x_2
x_1
1.60
1.40
1.20
1.00

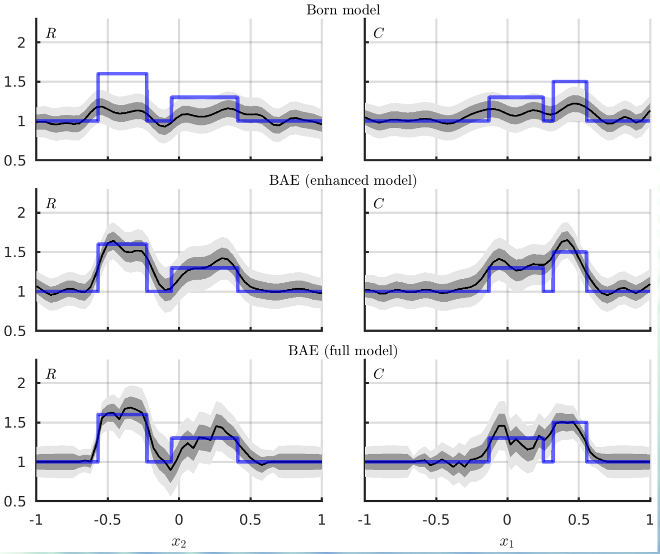
Born model
R
C
BAE (enhanced model)
R
C
BAE (full model)
R
C
2
1.5
1
0.5
-1
-0.5
0
0.5
1
x_2
x_1

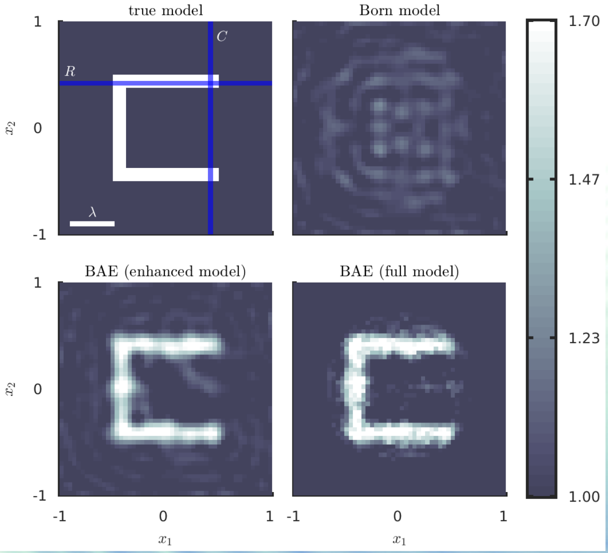
true model
Born model
BAE (enhanced model)
BAE (full model)
C
R
λ
1
0
-1
x_2
x_1
1.70
1.47
1.23
1.00

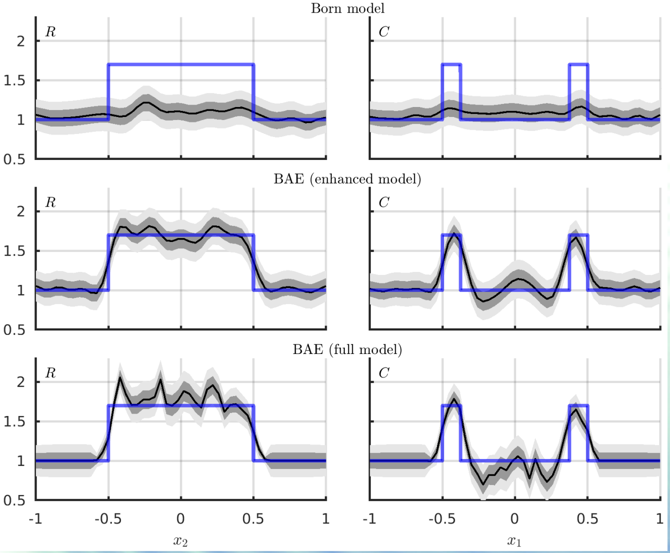
Born model
BAE (enhanced model)
BAE (full model)
R
C
2
1.5
1
0.5
-1
-0.5
0
0.5
1
x_2
x_1

Comments

- By carrying out simulations (model predictions) with an accurate and an approximate (Born) forward model, one can compute the approximate statistics of the related approximation/modelling errors
- The approach yields a computational scheme with essentially the same complexity as the (standard) Born approximation
- The simulations suggest that this approach can be feasible in the sense that the posterior error estimates are . . . feasible.