A COURSE FOR PHD STUDENTS

Numerical Methods for Inverse Problems

Carlos Castro
Universidad Politécnica de Madrid

October 14, 2026

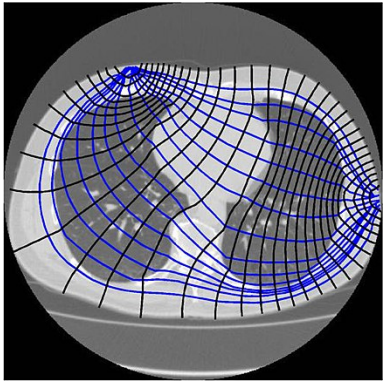
Inverse problems for PDE models

Problem	Model	Data	Unknown
Calderón (EIT)	$\nabla \cdot (\gamma \nabla u) = 0$ in Ω	DtN map on $\partial\Omega$	conductivity γ
Scattering (Helmholtz)	$\Delta u + k^2 u = q\,u$ in $\mathbb{R}^3$	far-field pattern	potential q
Scattering (elasticity)	$\mu\Delta u + (\lambda+\mu)\nabla(\nabla\cdot u) + \omega^2\rho\, u = 0$	far field of P, S waves	density, Lamé

The data depend nonlinearly on the unknown, and the inverse map is discontinuous: stability is typically only logarithmic.

Motivation

Electrical Impedance Tomography (wikipedia)



Stream lines (blue) and equipotential lines (black)

Inject currents at the boundary, measure voltages, image what is inside.

The Calderón problem

Consider the following system in the unit ball

$$(P1)\begin{cases} -\operatorname{div}\big(\gamma(x)\nabla u(x)\big) = 0, & x \in \Omega, \\ u = f, & x \in \partial\Omega. \end{cases}$$

Given $\gamma \in L^\infty(\Omega)$ we define the Dirichlet to Neumann map (DtN) as

$$\begin{aligned} \Lambda_\gamma : H^{1/2}(\partial\Omega) &\to H^{-1/2}(\partial\Omega) \\ f &\mapsto \gamma \tfrac{\partial u}{\partial n}\big|_{\partial\Omega} \end{aligned}$$

The reconstruction problem: can we invert the map

$$\gamma \longmapsto \Lambda_\gamma \, ?$$

Equivalent potential problem

$$(P2)\begin{cases} -\Delta u(x) + q(x)\, u(x) = 0, & x \in \Omega, \\ u = f, & x \in \partial\Omega. \end{cases}$$

Given $q \in L^\infty(\Omega)$ we define the Dirichlet to Neumann map (DtN) as

$$\begin{array}{rccc} \Lambda_q : H^{1/2}(\partial\Omega) & \to & H^{-1/2}(\partial\Omega) \\ f & \mapsto & \frac{\partial u}{\partial n} \end{array}$$

The reconstruction problem: can we find q from Λ_q?

Remark. For smooth γ, (P1) is equivalent to (P2) with

$$q = \frac{\Delta \gamma^{1/2}}{\gamma^{1/2}}$$

Is the problem well-posed?

Uniqueness: does Λ_γ determine γ?

- J. Sylvester and G. Uhlmann (1987). Dimension $d \geq 3$ and smooth γ.
- A. Nachman (1996). Dimension $d = 2$ and $\gamma \in C^2$.
- M. Lassas, L. Päivärinta and G. Uhlmann (2003). Dimension $d \geq 3$ and γ Lipschitz.
- K. Astala and L. Päivärinta (2006). Dimension $d = 2$ and $\gamma \in L^\infty$.

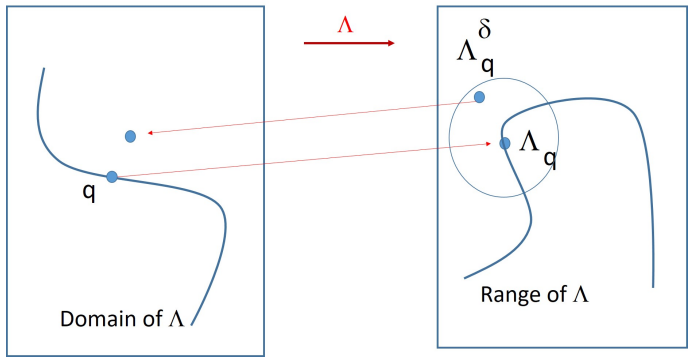
Stability:

- G. Alessandrini (1988) proved logarithmic stability for $d \geq 3$.
- N. Mandache (2001) proved the optimality of this log-stability result for $d \geq 2$.

Yes, the data determine the potential. But only logarithmically stably.

Ill-posed problem

Close DtN maps may be associated with very different q; also, a perturbed DtN map may be outside the range of Λ [Siltanen and Mueller 2012].



The Calderón problem in Schrödinger form

Let $\Omega \subset \mathbb{R}^n$ be a bounded domain and $q \in L^\infty(\Omega)$ (with 0 not a Dirichlet eigenvalue):

$$\begin{cases} -\Delta u + q\,u = 0 & \text{in } \Omega \\ u = f & \text{on } \partial\Omega \end{cases} \qquad \Lambda_q f = \partial_\nu u\,\big|_{\partial\Omega}$$

Inverse problem: recover q from the DtN map Λ_q

Choose an orthonormal basis $\varphi_1, \varphi_2, \dots$ of $L^2(\partial\Omega)$ and measure

$$M_{ij}[q] = \big\langle (\Lambda_q - \Lambda_0)\varphi_i, \varphi_j \big\rangle_{\partial\Omega}, \qquad i, j \le N.$$

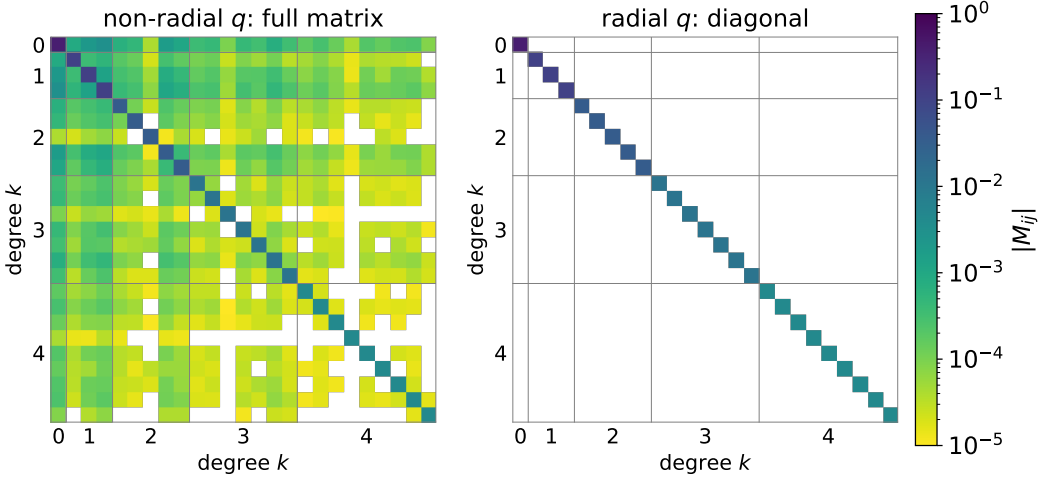
Green's formula (Alessandrini's identity) turns each entry into an integral of q:

$$M_{ij}[q] = \int_\Omega q\; u_i\, w_j\; dx, \qquad \begin{array}{l} u_i\colon\ (-\Delta+q)u_i = 0,\ u_i = \varphi_i \\ w_j\colon\ \Delta w_j = 0,\ w_j = \varphi_j \end{array}$$

- Forward map $F : q \mapsto M[q] \in \mathbb{R}^{N\times N}$, symmetric, nonlinear ($u_i$ depends on q).

The radial case in the ball: the matrix becomes a sequence

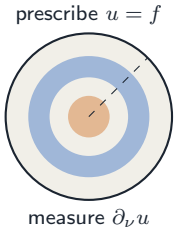
$\Omega = B \subset \mathbb{R}^3$, basis $\varphi = Y_k^m$ (spherical harmonics), harmonic extension $w = r^k Y_k^m$.



Simplified model: a radial potential in the 3D ball

$\Omega = B \subset \mathbb{R}^3$, $q = q(|x|)$. For $\varphi = Y_k^m$ the solution separates, $u = r^k\, V_k(r)\, Y_k^m$ with $V_k(1) = 1$:

$$\big(r^{2k+2}\, V_k'\big)' = r^{2k+2}\, q\, V_k, \qquad V_k \text{ regular at } 0.$$



So $\partial_\nu u = \big(k + V_k'(1)\big) Y_k^m$: spherical harmonics are eigenfunctions,

$$\Lambda_q Y_k^m = \lambda_k[q]\, Y_k^m, \qquad \lambda_k[q] = k + V_k'(1).$$

$$\mu_k = \lambda_k[q] - k = V_k'(1)$$

Forward map $F : q \mapsto \mu = (\mu_k)_k$; noisy data $\mu^\delta = F(q) + e$.

The data are moments of q

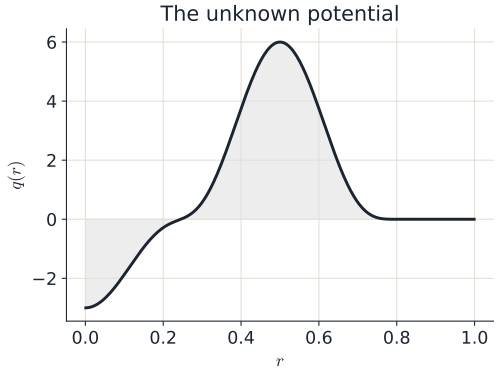
General identity (Alessandrini): $M_{ij}[q] = \displaystyle\int_\Omega q\, u_i\, w_j\, dx.$

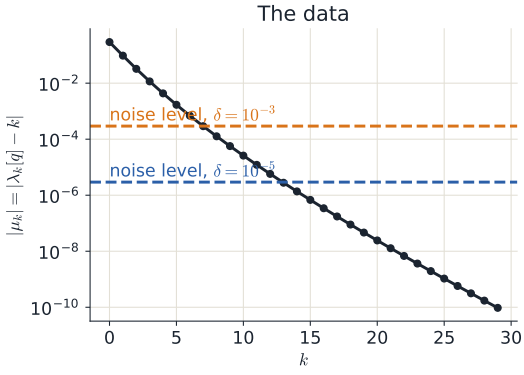
Take $\varphi_i = \varphi_j = Y_k^m$: $u_i = r^k V_k\, Y_k^m$ (previous slide) and the harmonic one is $w_j = r^k Y_k^m$. In polar coordinates $dx = r^2\, dr\, d\sigma$, with $\int_{S^2} (Y_k^m)^2\, d\sigma = 1$:

$$\mu_k = M_{ii} = \int_0^1 \int_{S^2} q(r)\, \big(r^k V_k\, Y_k^m\big)\big(r^k Y_k^m\big)\, r^2\, d\sigma\, dr$$

$$\mu_k = \int_0^1 q(r)\, r^{2k+2}\, V_k(r)\, dr$$

The running example



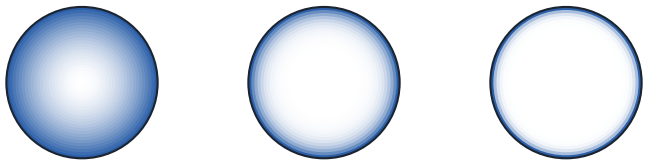


The data decay at least like a^{2k}: with noise 10^{-3} only 7 of the 30 eigenvalues are above the noise level, with 10^{-5} only 13.

Why so little information? High harmonics stay near the boundary

$$\mu_k = \int_0^1 q(r)\, r^{2k+2}\, V_k(r)\, dr$$

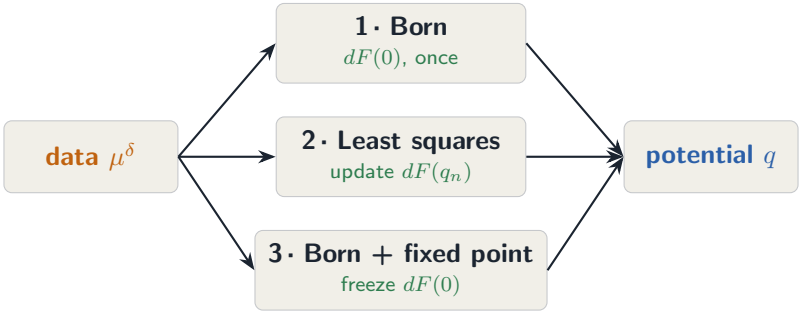
sensitivity r^{2k+2} of the datum μ_k to q



$k = 0$ $k = 3$ $k = 10$

Each datum sees q only near the boundary, and the higher k, the thinner the layer. The centre of the ball is almost invisible: this is where the core of the example lives.

Three different strategies



Colour code: potential data noise operators

FIRST IDEA

The Born approximation

Linearize and invert

Born approximation

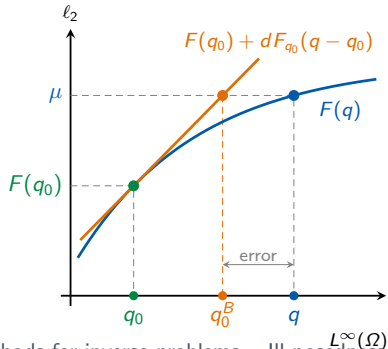
$$\begin{aligned} F : L^\infty(\Omega) &\to \ell_2 \\ q &\mapsto \mu = (\lambda_k[q] - k)_{k\geq 0} \quad (\lambda_k \text{ eigenvalues of } \Lambda_q) \end{aligned}$$

$$\mu_k = \int_0^1 q(r)\, r^{2k+2}\, V_k(r)\, dr$$

Take a reference background potential q_0 and look for q_0^B s.t.

$$dF_{q_0}\left(q_0^B - q_0\right) = \mu - \mu_0$$

Main issue: Does q_0^B exists?



The Born approximation for a radial q

$$\mu_k = \int_0^1 q(r)\, r^{2k+2}\, V_k(r)\, dr$$

$$\left(r^{2k+2}\, V_k'\right)' = r^{2k+2}\, q\, V_k, \qquad V_k \text{ regular at } 0, \qquad V_k(1) = 1$$

When $q = 0$ it is solved by $V_k \equiv 1$; for small q, $V_k = 1 + O(\|q\|)$. Therefore:

$$\mu_k = \int_0^1 q\, r^{2k+2}\, V_k\, dr = \underbrace{\int_0^1 q\, r^{2k+2}\, dr}_{(Aq)_k = (dF_0\, q)_k} + O(\|q\|^2).$$

$$\mu \approx A\, q, \qquad \mu = A\, q^B$$

The linear map $A = dF_0$ is explicit: a moment problem. Recovering q from its moments is a well-known ill-posed problem!

Computing it exactly

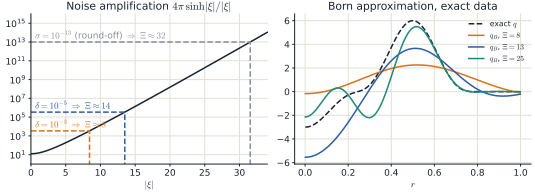
The moment problem $A\,q_B = \mu$ has an explicit solution in Fourier (Barceló, C., Macià, Meroño, 2024):

$$\widehat{q_B}(\xi) = \sum_k c_k(\xi)\,\mu_k, \qquad c_k(\xi) = 2\pi^{3/2}\frac{(-1)^k}{k!\,\Gamma(k+\frac{3}{2})}\Big(\frac{|\xi|}{2}\Big)^{2k}$$

Noise ε in each μ_k reaches frequency ξ amplified by $|c_k(\xi)|$.

Keep the frequencies where this stays below the signal:

δ	10^{-3}	10^{-5}	round-off
ξ	8	13	32

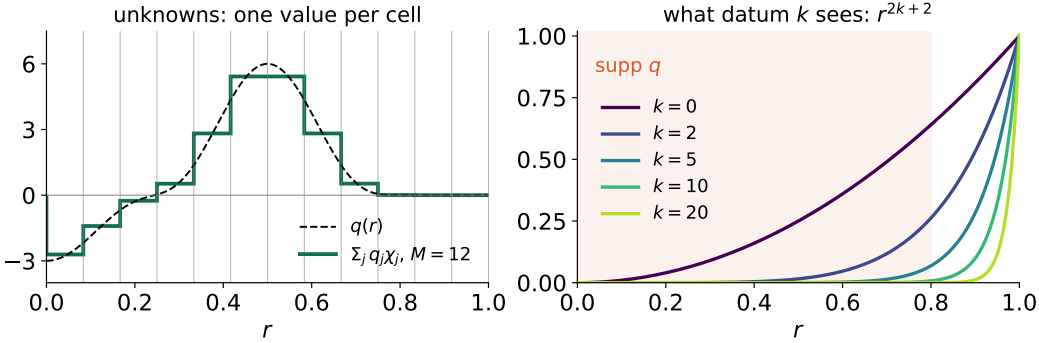


The cut-off is the regularisation, fixed by the noise level alone.

Discretisation: K data, M unknowns

Piecewise constant potential on M cells, $q = \sum_{j=1}^{M} q_j \chi_j$. The Born map becomes a matrix:

$$\mu = A\, q^B, \qquad A_{kj} = \int_{\text{cell } j} r^{2k+2}\, dr, \qquad A \in \mathbb{R}^{K \times M}.$$



Inverting A through its SVD

Reduced SVD ($K \geq M$): $A = U\Sigma V^\top = \sum_{j=1}^{M} \sigma_j \, u_j v_j^\top$, $\quad \sigma_1 \geq \cdots \geq \sigma_M > 0$, $A\, v_j = \sigma_j \, u_j$.

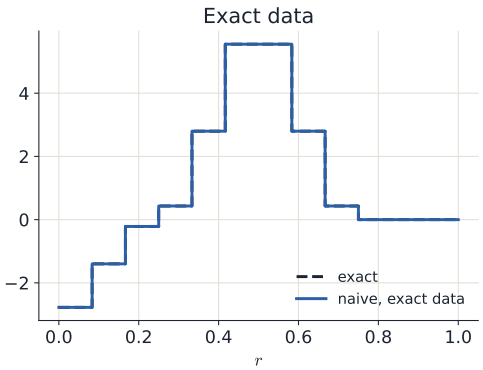
- $V \in \mathbb{R}^{M \times M}$: columns v_j are orthonormal potential modes.
- $U \in \mathbb{R}^{K \times M}$: columns u_j are orthonormal data modes.
- $\Sigma = \mathrm{diag}(\sigma_1, \ldots, \sigma_M)$: square and invertible.

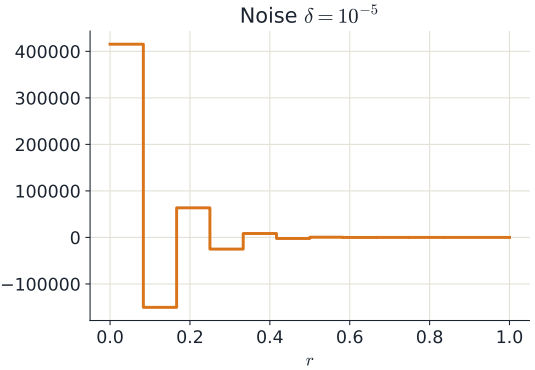
$$q^B_{\text{naive}} = V\,\Sigma^{-1}\,U^\top \mu^\delta = \sum_{j=1}^{M} \frac{u_j{}^\top \mu^\delta}{\sigma_j}\, v_j$$

Read from right to left: project the data on u_j, divide by σ_j, rebuild with v_j.

A first attempt: compute q_{naive}

Running example, $K = 30$ data, $M = 12$ cells, with data generated by A itself, so there is no model error:





Why it fails: dividing by σ_j

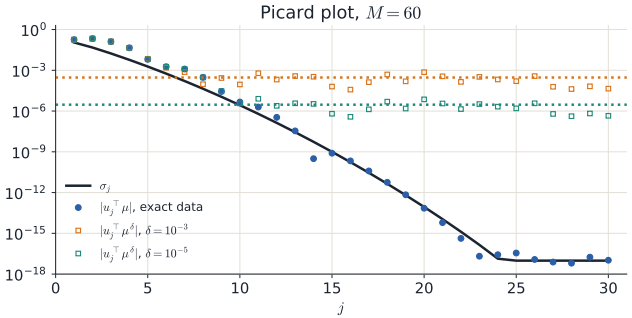
Back to the formula: split the data into signal and noise, $\mu^\delta = \mu + e$:

$$q^B_{\text{naive}} = \sum_j \Big(\underbrace{\frac{u_j^\top \mu_j}{\sigma_j}}_{\text{signal}} + \underbrace{\frac{u_j^\top e}{\sigma_j}}_{\text{amplified noise}} \Big) v_j$$

- Signal: the coefficients of the true q, of size $O(1)$.
- Noise: for example $\delta \sim 10^{-5}$.
- σ_j decays exponentially: $\sigma_{12} \approx 10^{-10}$.

Amplified noise: $\sim 10^5$ for $j = 12$.

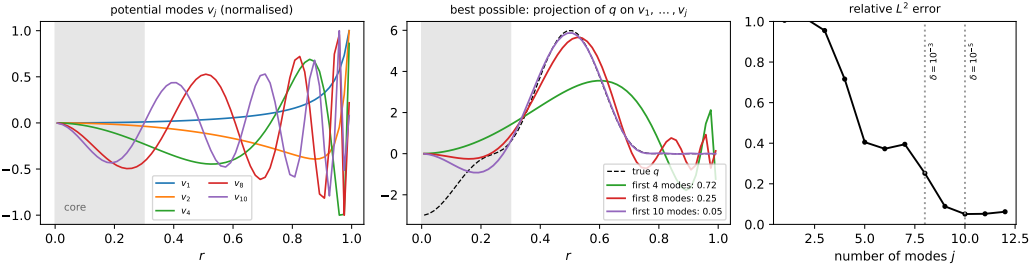
How many modes can we recover? The Picard plot



Discrete Picard condition: data coefficients must decay faster than σ_j. Clean data do (circles); noisy data stop at the noise floor (squares). Beyond the crossing, mode j is pure noise: about 8 modes at $\delta = 10^{-3}$, 10 at 10^{-5}.

What can 8 or 10 modes represent?

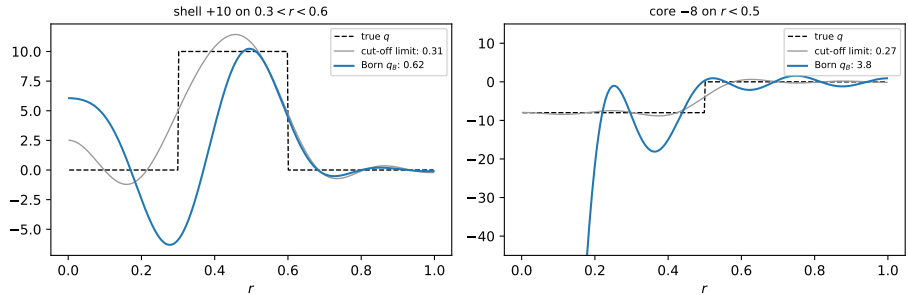
The data determine only the first modes $v_1, \dots, v_j$. The best any reconstruction can do from them alone is the projection of the true potential, $P_j q = \sum_{i \le j} (v_i^\top q)\, v_i$.



Left: the modes live near the boundary and vanish in the interior. Centre: 10 modes ($\delta = 10^{-5}$). Right: two modes more make the difference.

The Born approximation is not always close to q

Two discontinuous potentials, exact data (60 eigenvalues, cut-off 25):



Discontinuities are not the problem; size and sign are. The gap $q - q_B$ is the nonlinear part, and it must be recovered by a second step.

STRATEGY 2

Least squares

Fit the data, penalise roughness, iterate.

The idea: fit the data instead of inverting the map

Forward map and data:

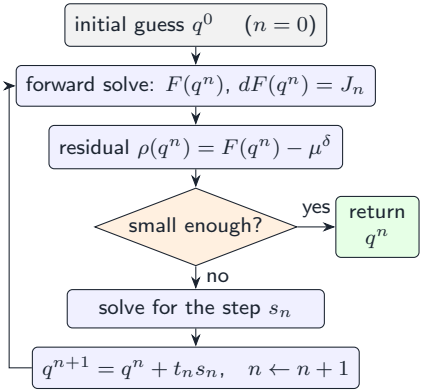
$$F:\ q \longmapsto \big(\mu_0(q), \dots, \mu_{K-1}(q)\big), \qquad \mu^\delta = F(q^\dagger) + e, \quad e_k \sim \mathcal{N}(0, \varepsilon^2)$$

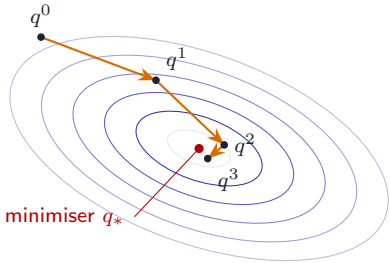
($q^\dagger$: the true potential). We look for the potential whose *predicted* data are closest to the *measured* data:

$$\min_q \quad \Phi(q) \;=\; \frac{1}{2}\left\|F(q) - \mu^\delta\right\|^2$$

- Only needs a **forward solver** (and ideally its derivative).
- Works with any amount of data, noisy or redundant.

Iterative methods: improve a guess, step by step





level sets of Φ: each step goes downhill

Three ingredients: a **starting guess** ($q^0 = 0$), an **descent direction** (s_n), a **stopping test**.

The Jacobian: differentiating the ODE

The forward problem is

$$\mu_k = \int_0^1 q(r)\, r^{2k+2}\, V_k(r)\, dr$$

$$(r^{2k+2}V_k')' = r^{2k+2}\, q\, V_k, \qquad V_k(1) = 1, \qquad \mu_k = V_k'(1).$$

1. **Perturb** $q \to q + t\,h$, differentiate in t: $\delta\mu_k = \delta V_k'(1)$, where
$$(r^{2k+2}\,\delta V_k')' = r^{2k+2}\big(q\,\delta V_k + h\,V_k\big), \qquad \delta V_k(1) = 0.$$
2. **Multiply by V_k, integrate by parts twice:** boundary terms vanish and the q-terms cancel.
3. **Result:**

$$\delta\mu_k[h] = \int_0^1 h\, r^{2k+2}\, V_k^2\, dr \quad \Longrightarrow \quad (J_n)_{kj} = \int_{\text{cell } j} r^{2k+2}\, V_k^2\, dr$$

Minimising Φ: Newton method

A minimiser of $\Phi(q) = \frac{1}{2}\|F(q) - \mu^\delta\|^2$ solves

$$G(q) \;=\; \nabla\Phi(q) \;=\; J(q)^\top \rho(q) = 0, \qquad \rho(q) = F(q) - \mu^\delta.$$

Newton's method for $G = 0$ needs the Hessian

$$\nabla^2\Phi(q) \;=\; \underbrace{J(q)^\top J(q)}_{\text{derivative of } \rho} \;+\; \underbrace{\textstyle\sum_k \rho_k(q)\, \nabla^2\mu_k(q)}_{\text{derivative of } J}.$$

Gauss–Newton: drop the second term, i.e. $\nabla^2\Phi \approx J^\top J$,

$$J_n{}^\top J_n\, s_n = -J_n{}^\top \rho(q^n), \qquad q^{n+1} = q^n + t_n s_n$$

Cheap. Only first derivatives: J_n comes with the forward solve.

Safe. $J_n^\top J_n$ is positive semi-definite: s_n is a descent direction.

The same step, read geometrically

Linearise F, not Φ, at q^n:

$$F(q^n + s) \approx F(q^n) + J_n s$$

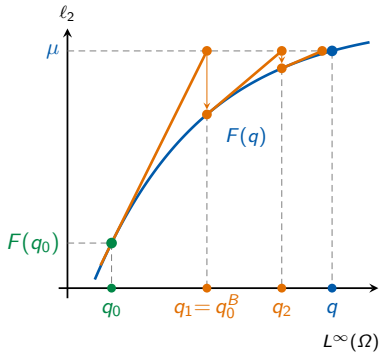
and search for $s = q^{n+1} - q^n$ such that

$$\mu^\delta = F(q^{n+1}) \sim F(q^n) + J_n s$$

which is equivalent to fit the *linear* model:

$$\min_s \ \tfrac{1}{2}\|J_n s + \rho(q^n)\|^2$$

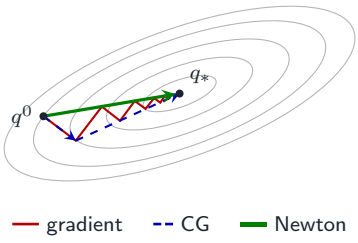
$$J_n{}^\top J_n \, s_n = -J_n{}^\top \rho(q^n)$$



follow the tangent until it hits μ^δ.

Why Gauss–Newton?

If F were linear, $F(q) = Aq$: then $\Phi(q) = \frac{1}{2}\|Aq - \mu^\delta\|^2$ is *exactly* quadratic, with Hessian $A^\top A$.



Error reduction per step on this quadratic, $\kappa = \mathrm{cond}(A^\top A)$:

- gradient: $(\kappa - 1)/(\kappa + 1)$
- CG: $(\sqrt{\kappa} - 1)/(\sqrt{\kappa} + 1)$
- Newton: exact in one step

F **nonlinear:** replace A by the linearisation J_n at q_n and take that Newton step:

$$s_n = -(J_n{}^\top J_n)^{-1} J_n{}^\top \rho_n.$$

Gauss–Newton redraws this picture at every iterate.

Regularise: the Tikhonov functional

The first Gauss–Newton step from $q^0 = 0$ gives q^B, but fails the same way as SVD. Add a size penalty, natural for $q \in L^\infty \subset L^2$:

$$\Phi_\alpha(q) = \frac{1}{2}\left\|F(q) - \mu^\delta\right\|^2 + \frac{\alpha}{2}\int_0^1 |q|^2\, dr$$

with $\int_0^1 |q|^2\, dr \approx q^\top L q$, $L = h\, I$. Same Gauss–Newton, one extra term:

$$\left(J_n{}^\top J_n + \alpha L\right) s_n = -J_n{}^\top \rho(q^n) - \alpha L\, q^n, \qquad q^{n+1} = q^n + t_n s_n$$

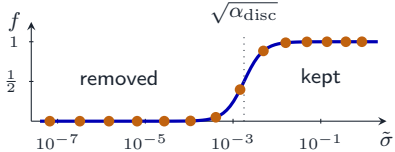
One parameter, α, balances fitting the data against keeping q small.

What α does: a filter on the modes

Born model, $L = I$: $(A^\top A + \alpha I)\, q = A^\top \mu^\delta$, i.e. mode by mode $(\sigma_j^2 + \alpha)\, c_j = \sigma_j\, u_j^\top \mu^\delta$:

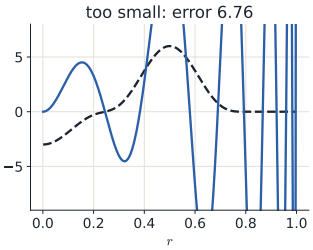
$$q_\alpha = \sum_j \underbrace{\frac{\sigma_j^2}{\sigma_j^2 + \alpha}}_{f_j(\alpha)} \underbrace{\frac{u_j^\top \mu^\delta}{\sigma_j}\, v_j}_{\text{term } j \text{ of } q_{\text{naive}}}$$

- $\sigma_j \gg \sqrt{\alpha}$: $f_j \approx 1$, mode **kept**.
- $\sigma_j \ll \sqrt{\alpha}$: $f_j \approx 0$, mode **removed**: no noise amplified, but that part of $q^\dagger$ is lost.
- α too small: noisy; too large: over-smoothed.

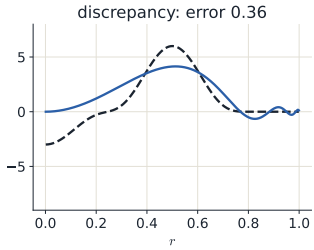


dots: $\tilde\sigma_j$ of A; 6 kept at $\delta = 10^{-3}$

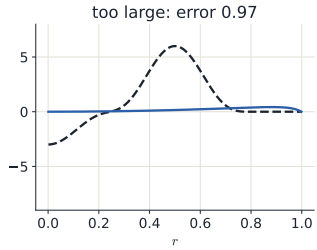
Under-, well- and over-regularised



$\alpha = 10^{-4}\,\alpha_{\text{disc}}$
oscillations, noise fitted



$\alpha = \alpha_{\text{disc}}$
shell found, core lost



$\alpha = 10^{5}\,\alpha_{\text{disc}}$
shrunk towards 0

Running example, $\delta = 10^{-3}$, dashed: true q.

Too little regularisation fits the noise; too much fits nothing. The discrepancy principle picks the balance, but even then the core is out of reach.

Two penalties: size or slope

$$\Phi_\alpha(q) = \frac{1}{2}\left\|F(q) - \mu^\delta\right\|^2 + \frac{\alpha}{2}\,\mathcal{R}(q)$$

$\mathcal{R}(q) = \int_0^1 |q|^2\,dr$ **($L = h\,I$)**

Penalises the **size** of q.

Large α: $q \to 0$, the background.

Filter diagonal in the SVD of A.

Assumes only $q \in L^2$: jumps allowed.

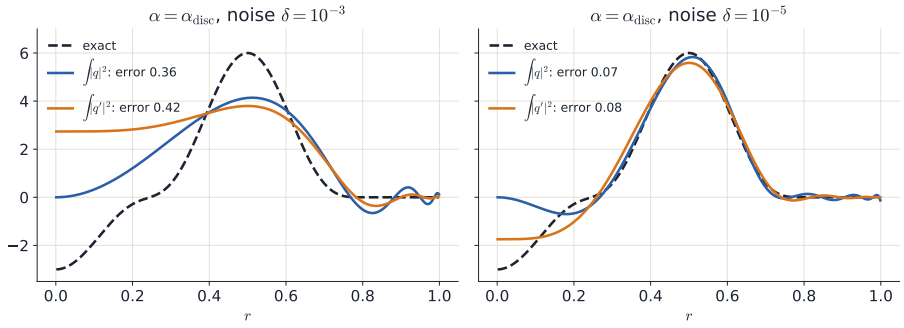
$\mathcal{R}(q) = \int_0^1 |q'|^2\,dr$ **($L = \frac{1}{h}D^\top D$)**

Penalises the **oscillation** of q; constants are free.

Large α: $q \to$ the best constant fit.

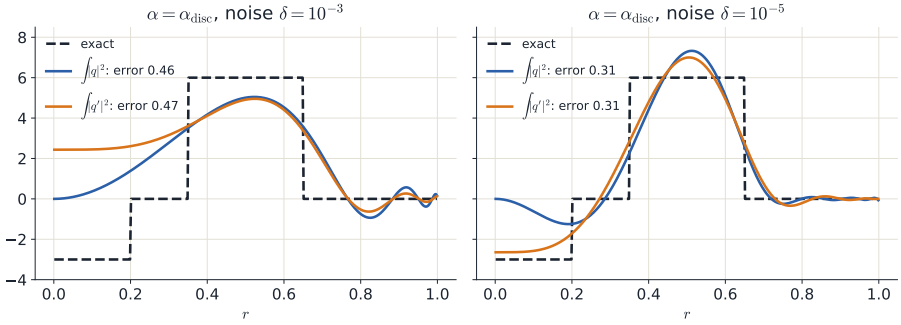
Assumes $q \in H^1$: no jumps.

Same data, two penalties



Near $r = 0$ the L^2 penalty returns 0 and the H^1 one a flat value

A discontinuous potential



What least squares costs

- Every iteration linearises again: a new Jacobian and a new linear solve.
- Outside the radial case J_n is not free: each product $J_n v$ or $J_n^\top w$ costs one PDE solve (linearised or adjoint), and J_n is too large to form.

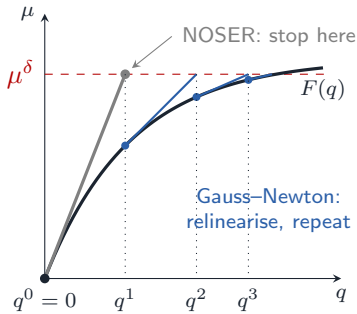
NOSER is one Gauss–Newton step

NOSER, Newton's One-Step Error Reconstructor (Cheney, Isaacson, Newell, Simske, Goble, 1990): in EIT, one Gauss–Newton step from the background.

From $q^0 = 0$: $J_0 = A$, $F(0) = 0$, so

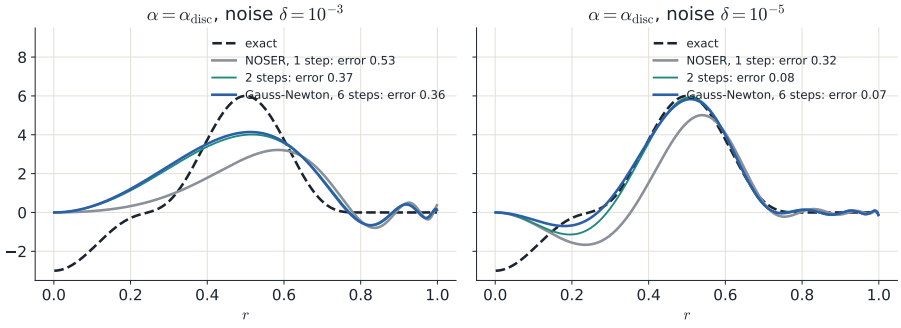
$$\left(A^\top A + \alpha L\right) q^1 = A^\top \mu^\delta : \quad q^1 = R_\alpha \mu^\delta = q_B.$$

- The regularised Born approximation.
- Linear in the data: one precomputed matrix.
- Keeps the linearisation error.



each step follows the tangent at the current point.

One step or several



Most of the gain is in the second step, and it is larger at low noise: there the linearisation error, not the noise, is what limits NOSER.

Another way to regularise: mollifiers

Tikhonov lets α decide how much is lost. The mollifier method (Louis and Maass, 1990) changes the question: reconstruct a *local average* of q, chosen in advance,

$$q_\gamma(r_0) = \int_0^1 e_\gamma(r_0, r)\, q(r)\, dr, \qquad e_\gamma(r_0, \cdot) \text{ a bump of width } \gamma \text{ at } r_0, \qquad \int e_\gamma = 1.$$

Each datum is already an average of q: $(Aq)_k = \int_0^1 q\, r^{2k+2}\, dr$. So if

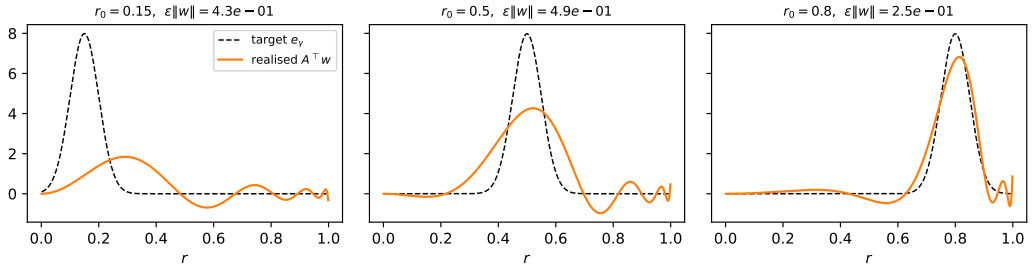
$$\sum_{k<K} w_k(r_0)\, r^{2k+2} \;\approx\; e_\gamma(r_0, r) \qquad (\text{i.e. } A^\top w \approx e_\gamma)$$

then, by linearity,

$$\boxed{q_\gamma(r_0) \;\approx\; \sum_k w_k(r_0)\, \mu_k^\delta \;=\; w(r_0)^\top \mu^\delta}$$

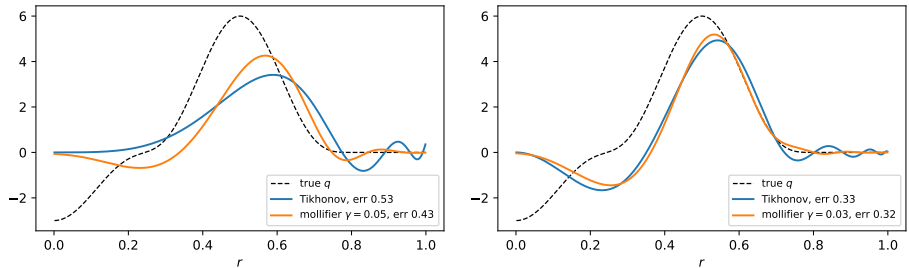
We never invert A: we only write the bump as a combination of the measurement kernels. The weights do not depend on the data.

What can be seen at each depth



Same target width $\gamma = 0.05$, $\delta = 10^{-3}$. Near the boundary the target is reached. The interior is out of reach.

Mollifier against Tikhonov on the same data



Reconstructions from the same nonlinear data. Left: $\delta = 10^{-3}$. Right: 10^{-5}: similar quality.

STRATEGY 3

Linearise, then correct

The Born approximation and a fixed-point iteration.

The Born approximation is a Hausdorff moment problem

Recall: $F(q) = Aq + O(\|q\|^2)$, $(Aq)_k = \int_0^1 q\, r^{2k+2}\, dr$. The Born approximation q^B solves $Aq^B = \mu$: its moments are the data. With $s = r^2$ and $g(s) = \frac{1}{2}\sqrt{s}\, q^B(\sqrt{s})$,

$$\int_0^1 s^k\, g(s)\, ds = \mu_k, \qquad k = 0, 1, 2, \ldots$$

- **Uniqueness:** the powers s^k are complete (Weierstrass), so μ determines q^B.
- **Existence:** not every sequence is a moment sequence, and $\mu = F(q) \neq Aq$.
- **Instability:** $(AA^\top)_{kl} = \frac{1}{2k+2l+5}$, a Hilbert-type matrix: condition number $\approx 10^{14}$ already for $K = 10$.

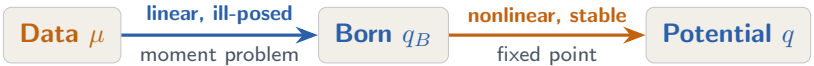
Born turns the nonlinear problem into a classical one: unique, but severely ill-posed.

Factorise the inverse problem

Define the Born approximation q_B as the function whose moments are the data:

$$\int_0^1 q_B(r)\, r^{2k+2}\, dr \;=\; \lambda_k[q] - k \quad \text{for all } k \geq 0, \qquad \text{i.e.} \quad A\, q_B = F(q)$$

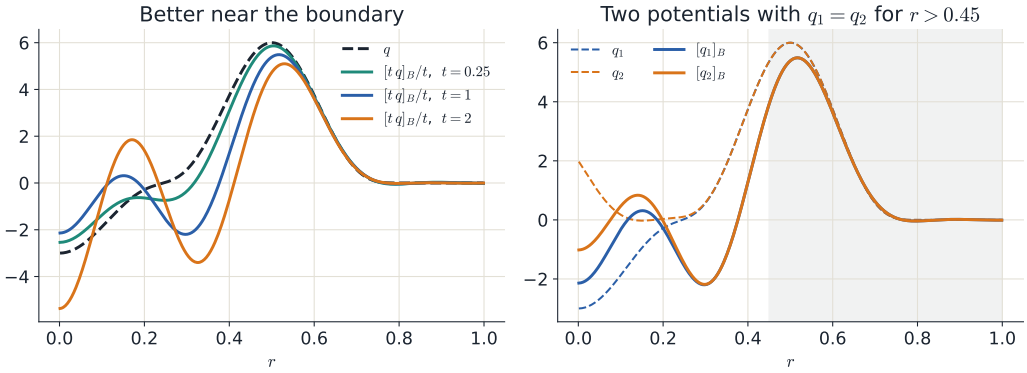
The Born approximation splits the inverse problem in two steps:



- **First arrow:** solve $Aq_B = \mu^\delta$. The operator A is explicit, but inverting it needs regularisation, exactly as in Part 1: keep the modes above the noise.
- **Second arrow:** recover q from q_B. The map $q_B \mapsto q$ is stable: small errors in q_B give small errors in q.

All the instability sits in the first arrow, and that arrow is linear: the regularisation of Part 1 applies as it is.

What the Born approximation sees



Near the boundary it is accurate even for fairly large potentials, and it depends only on the potential between that point and the boundary. It also keeps the support and the jumps of q (BCMM, 2024).

From the Born approximation to the potential

Fixed point (Barceló, Castro, Macià, Meroño, 2024):

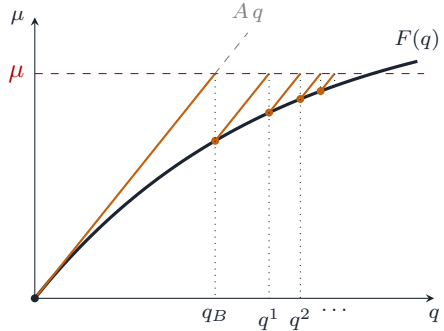
$$q^0 = q_B, \qquad q^{n+1} \;=\; q_B + q^n - [q^n]_B$$

With $[q]_B = R\,F(q)$, $R \approx A^{-1}$:

$$q^{n+1} \;=\; q^n + R\,\big(\mu - F(q^n)\big)$$

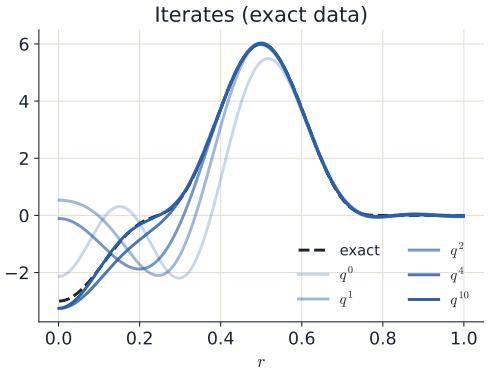
Newton-type, with the slope **frozen** at $q = 0$: no Jacobians, no new linear solves, one forward solve per step.

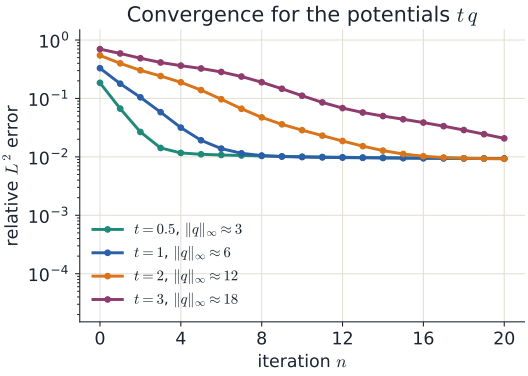
Converges when q is not too large, so that $F'(q)$ stays close to A.



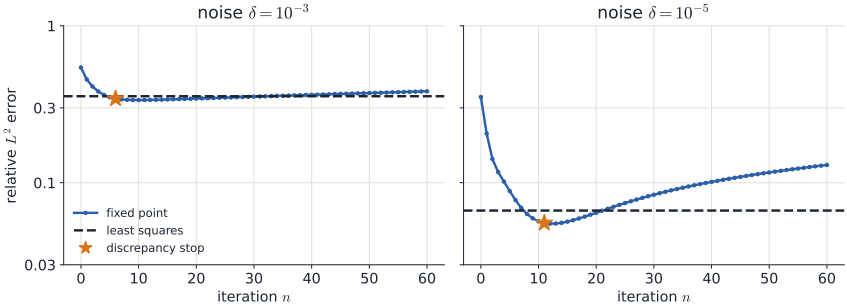
every step uses the same slope $A = F'(0)$,
not the tangent at q^n (compare Gauss–Newton).

Fixed point with exact data



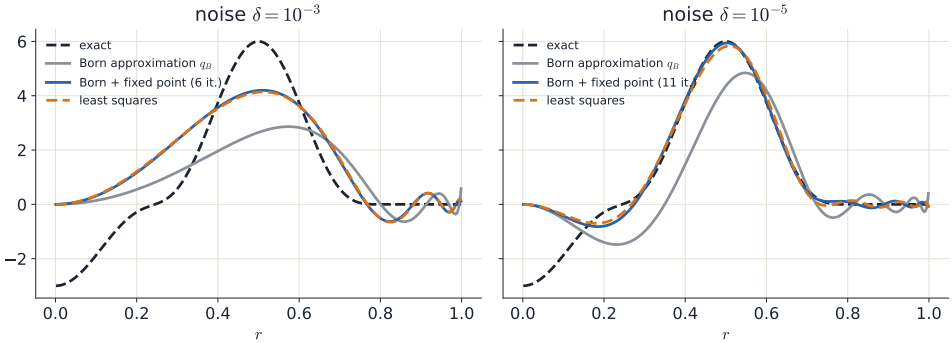


Fixed point with noisy data: when to stop



The error falls and then grows: the late steps fit the noise (semi-convergence). Stop when the misfit reaches the noise.

Comparison on the same data



Born q_B: only $A = F'(0)$. Fixed point: $F(q^n)$ at each step, derivative only at $q = 0$. Least squares: $F(q^n)$ *and* $J_n = F'(q^n)$ at each step.

Accuracy and cost

	NOSER = Born q_B	Born + fixed point	Gauss–Newton
Error, $\delta = 10^{-3}$	0.53	**0.34**	0.36
Error, $\delta = 10^{-5}$	0.32	**0.055**	0.066
Error, $\|q\|_\infty = 3\,/\,12\,/\,24$	0.19 / 0.52 / 0.78	0.06 / 0.13 / 0.15	**0.06 / 0.06 / 0.09**
Forward solves	0	7	19
Jacobians (30×400)	none $(A = F'(0))$	none	6

COMPARISON AND LEARNING

Same data, two algorithms

Then the Born approximation as the input of a neural network.

The Born approximation as network input

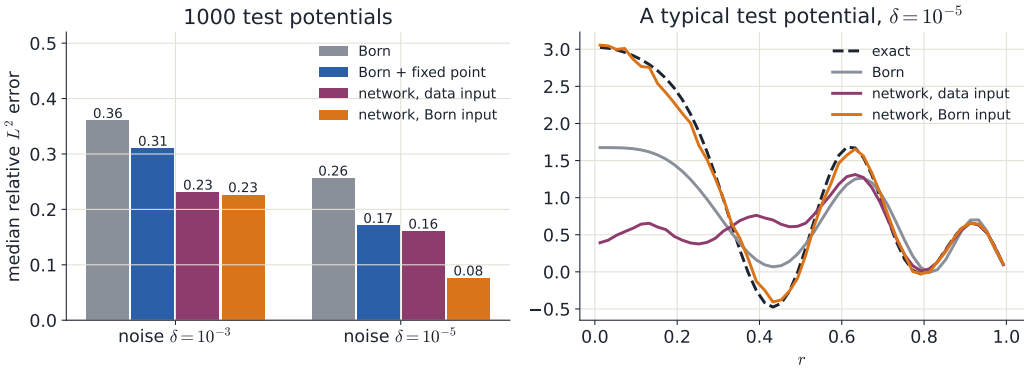
The linear step has already unpacked the data; the network only has to learn the nonlinear, stable part.

(a) data μ → network → potential q

(b) data μ → linear map R → Born approximation q_B → network → correction $q - q_B$

- Same network in both: two hidden layers of 128 units.
- Same training set: 5000 random potentials (cosine series, size up to about 7) with noisy data.
- Test on 1000 new potentials; classical Born and Born + fixed point as baselines.

Learning: what the Born input buys



At $\delta = 10^{-5}$ the Born input halves the error (0.08 against 0.16), and more so with only 500 training samples (0.10 against 0.27). At $\delta = 10^{-3}$ both networks reach 0.23: the noise is the limit, not the method.

A RELATED EXAMPLE

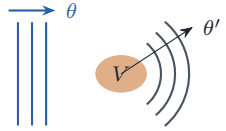
Inverse scattering

Born approximation.

Inverse scattering: the same structure in $\mathbb{R}^2$

A plane wave hits a potential V with compact support:

$$-\Delta u + Vu = k^2 u, \qquad u = e^{ik\theta\cdot x} + u_s,$$



u_s outgoing. Far away, $u_s \approx \frac{e^{ikr}}{\sqrt{r}}\, u_\infty(\theta', \theta, k)$: the data.

Direct problem: the Lippmann–Schwinger equation, with $\Phi_k = \frac{i}{4} H_0^{(1)}(k|\cdot|)$,

$$u(x,\theta,k) = e^{ik\theta\cdot x} + \int \Phi_k(x-y)\, V(y)\, u(y,\theta,k)\, dy$$

Exact identity (the far-field limit of the same equation):

$$u_\infty(\theta', \theta, k) = \int e^{-ik\theta'\cdot x}\, V(x)\, u(x,\theta,k)\, dx$$

J. A. Barceló, C. Castro, J. M. Reyes, Inverse Problems 32 (2016) 015006 .

Born: the data are the Fourier transform of V

Split $u = e^{ik\theta\cdot x} + u_s$ in the identity, and write $\xi = k(\theta' - \theta)/2\pi$:

$$u_\infty(\theta', \theta, k) = \underbrace{\widehat{V}(\xi)}_{\text{linear}} + \underbrace{\int e^{-ik\theta'\cdot x}\, V\, u_s\, dx}_{\text{nonlinear}}.$$

Born approximation: drop the nonlinear term,

$$\widehat{V_B}(\xi) = u_\infty(\theta', \theta, k) \qquad \Longrightarrow \qquad V_B = \text{one inverse FFT}$$

- Each data point gives one Fourier coefficient: the map $(\theta', \theta, k) \mapsto \xi$ fixes which one.
- Fixed energy: $|\theta' - \theta| \leq 2$, so only $|\xi| \leq k/\pi$ is measured: a low-pass filter.

Here A is the Fourier transform: no small singular values, but a band limit. Larger k, finer details.

Factorization in practice: freeze the Fourier transform, iterate

Keep the nonlinear term, evaluated at the previous iterate: $V^0 = 0$ and

$$\widehat{V^{n+1}}(\xi) = u_\infty(\theta', \theta, k) - \int e^{-ik\theta'\cdot y}\, V^n(y)\, u_s^n(y, \theta, k)\, dy$$

where u_s^n solves Lippmann–Schwinger with the previous potential:

$$u_s^n = \int \Phi_k(|\cdot - y|)\, V^n(y)\, \left(e^{ik\theta\cdot y} + u_s^n(y)\right) dy.$$

- $V^1 = V_B$; the linear operator never changes (freeze A): one inverse FFT per step.
- The nonlinearity enters only through u_s^n, like $V_k[q^n]$ in the ball.

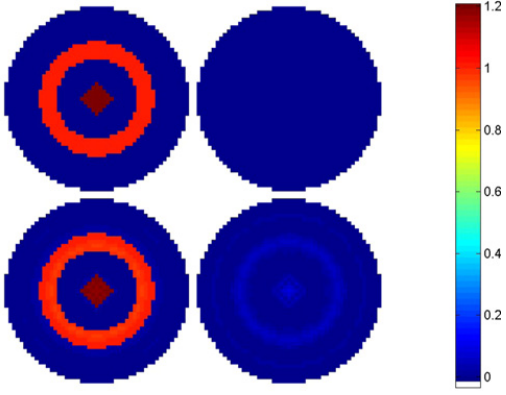
Same fixed point as route 3: invert the linear part once and for all, put the nonlinearity in the data.

The two problems side by side

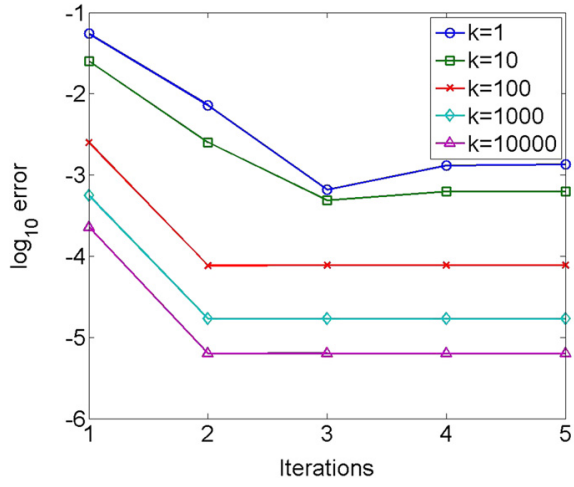
	Ball, radial q	**Scattering in $\mathbb{R}^2$**
Direct problem	ODE for V_k	Lippmann–Schwinger for u
Zero potential	$V_k \equiv 1$	plane wave $e^{ik\theta\cdot x}$
Exact identity	$\mu_k = \int_0^1 q\, r^{2k+2}\, V_k\, dr$	$u_\infty = \int e^{-ik\theta'\cdot x}\, V\, u\, dx$
Born	$V_k \to 1$	$u \to e^{ik\theta\cdot x}$
Operator A	moments: σ_j decay fast	Fourier: band $\lvert\xi\rvert \le k/\pi$
Fixed point	update $V_k[q^n]$	update $u_s^n[V^n]$

Same three ingredients: a direct solver, an exact identity, and a linear operator that is frozen.

Results: fixed energy

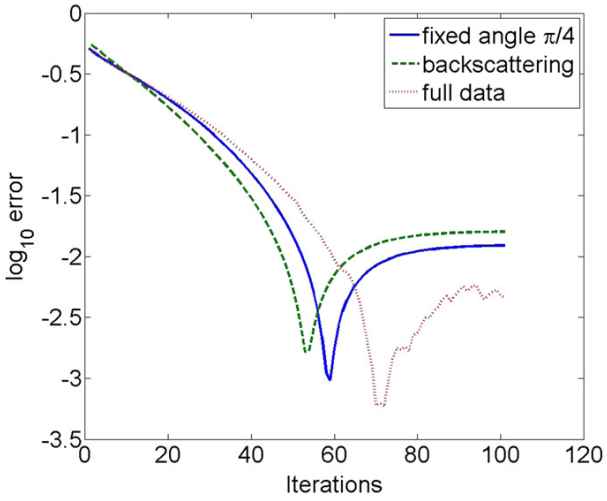


True V (top) and Born V_B (bottom), $k = 10$



$\log_{10}$ error against iterations, $k = 1, \ldots, 10^4$

Results: stopping is regularisation



Fixed angle, backscattering, full data; $N^2 = 2^{12}$

- The error decreases for 50–70 iterations,
- then grows again: the iteration starts fitting the noise.
- Semi-convergence, as for gradient and CG: the number of iterations is the regularisation parameter.

J. A. Barceló, C. Castro, J. M. Reyes, Inverse Problems 32 (2016) 015006 .

References: the Born approximation and its iteration

- J. A. Barceló, C. Castro, F. Macià, C. J. Meroño. The Born approximation in the three-dimensional Calderón problem. *J. Funct. Anal.* 283 (2022).
- J. A. Barceló, C. Castro, F. Macià, C. J. Meroño. The Born approximation in the three-dimensional Calderón problem II: numerical reconstruction in the radial case. *Inverse Probl. Imaging* 18 (2024).
- C. Castro, F. Macià, C. Meroño, D. Sánchez-Mendoza. Characterization of Dirichlet-to-Neumann maps via the Born approximation. arXiv:2601.11975.
- T. Daudé, F. Macià, C. Meroño, F. Nicoleau. Stable factorization of the Calderón problem via the Born approximation. arXiv:2402.06321.
- J. A. Barceló, C. Castro, J. M. Reyes. Numerical approximation of the potential in the two-dimensional inverse scattering problem. *Inverse Problems* 32 (2016) 015006.
- R. G. Novikov. An iterative approach to non-overdetermined inverse scattering at fixed energy. *Sb. Math.* 206 (2015) 120–134.
- A. P. Calderón. On an inverse boundary value problem. *Seminar on Numerical Analysis and its Applications to Continuum Physics*, Rio de Janeiro (1980).