Statistical Estimation of the Parameters of a PDE

Colin Fox, Geoff Nicholls
(University of Auckland)

- Nomenclature for image recovery
- Statistical model for inverse problems
- Traditional approaches – deconvolution example
- Recovering electrical conductivity via inference

Image Recovery nomenclature for Inverse Problems

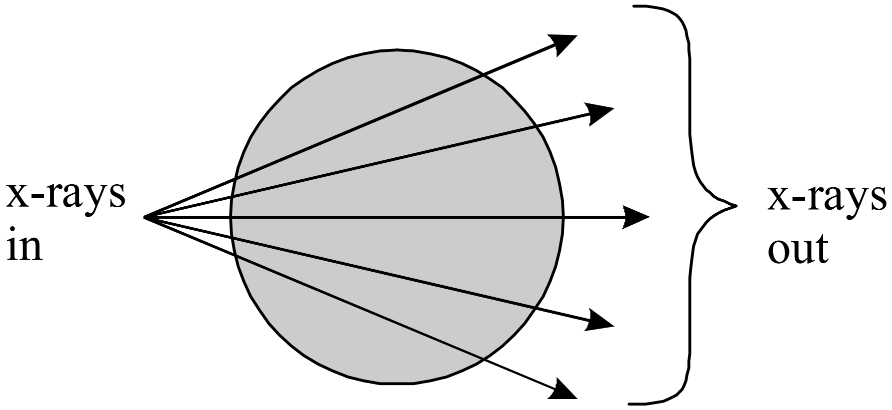

X-ray tomography

Image: spatially varying quantity of interest

- optical reflectance of a scene
- optical or radio brightness of sky
- sound speed in tissue / ocean / earth
- electrical conductivity of tissue / mud

Recovery: estimate image from indirect data

Forward Problem	**Inverse Problem**
image ⟶ data	data ⟶ image
physical model (PDE)	implicit
direct computation	indirect
well posed	ill posed
unique	never unique

Conductivity Imaging Measurement Set

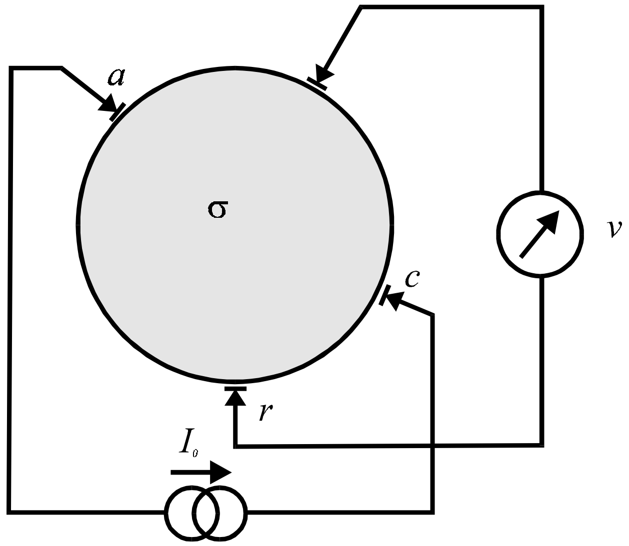

- Electrodes at $x_1, x_2, \cdots, x_E$
- Assert currents at electrodes $j = (j(x_1), j(x_2), \cdots, j(x_E))^T$
- Measure voltages $v = (\phi(x_1), \phi(x_2), \cdots, \phi(x_E))^T$.

Unknown $\sigma(x)$ related to measurements via Neumann BVP

$$\begin{aligned} \nabla \cdot \sigma(x) \nabla \phi(x) &= 0 \qquad && x \in \Omega \\ \sigma(x) \frac{\partial \phi(x)}{\partial n(x)} &= j(x) \qquad && x \in \partial\Omega \end{aligned}$$

- Set of measurements is current-voltage pairs

$$\{j^n, v^n\}_{n=1}^N$$

Inverse problem is to find σ from these measurements (non linear)

Statistical model of Inverse Problem

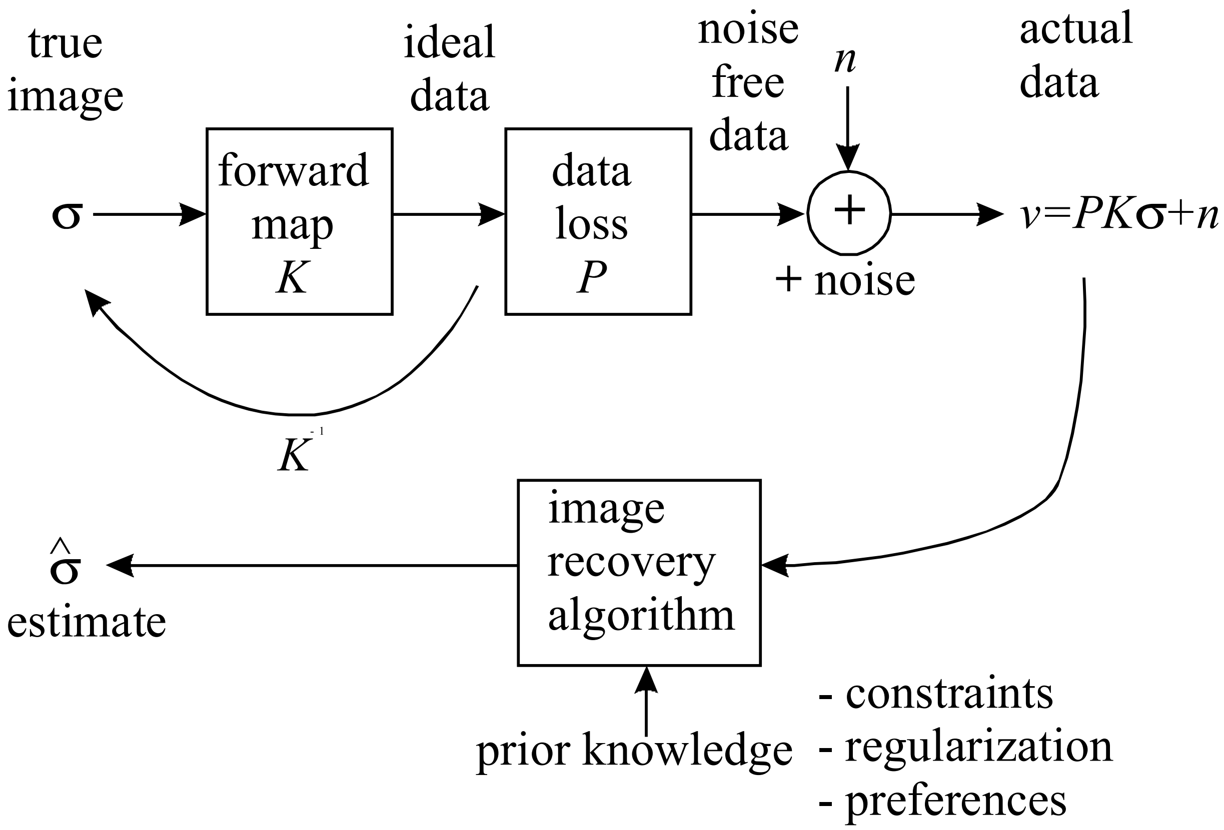

If $n \sim N(0, s^2)$, $v \sim N(PK\sigma, s^2)$

Given measurements v, the likelihood for σ is

$$L_v(\sigma) \equiv \Pr(v|\sigma) \propto \exp(|v - \phi(\sigma)|^2/2s^2)$$

Posterior distribution for σ conditional on v

$$\Pr(\sigma|v) = \frac{\Pr(v|\sigma)\Pr(\sigma)}{\Pr(v)} \qquad \text{(Bayes rule)}$$

$\Pr(\sigma)$ is the prior distribution

Solutions = Summary Statistics

All information contained in posterior distribution $\Pr(\sigma|v)$

Traditional Solutions - modes

$$\begin{aligned}\hat{\sigma}_{\text{MLE}} &= \arg\max L_v(\sigma) \equiv \arg\max \Pr(v|\sigma)\\ \hat{\sigma}_{\text{MAP}} &= \arg\max \Pr(\sigma|v) \equiv \arg\max \Pr(v|\sigma)\Pr(\sigma)\end{aligned}$$

e.g. simple Gaussian prior: $\Pr(\sigma) \propto \exp\left(-|\sigma|^2/2\lambda^2\right)$

$$\hat{\sigma}_{\text{MAP}} = \arg\min |v - \phi(\sigma)|^2 + \alpha|\sigma|^2 \qquad \alpha = s^2/\lambda^2$$

- Tikhonov regularization, Kalman filtering, Backus-Gilbert
- $\alpha \to 0$ Moore-Penrose inverse, $\alpha = 0$ least-squares

Inferential Solutions

"Answers" are expectations over the posterior

$$\mathrm{E}[f(\sigma)] = \int \Pr(\sigma|v)\, f(\sigma)\, d\sigma$$

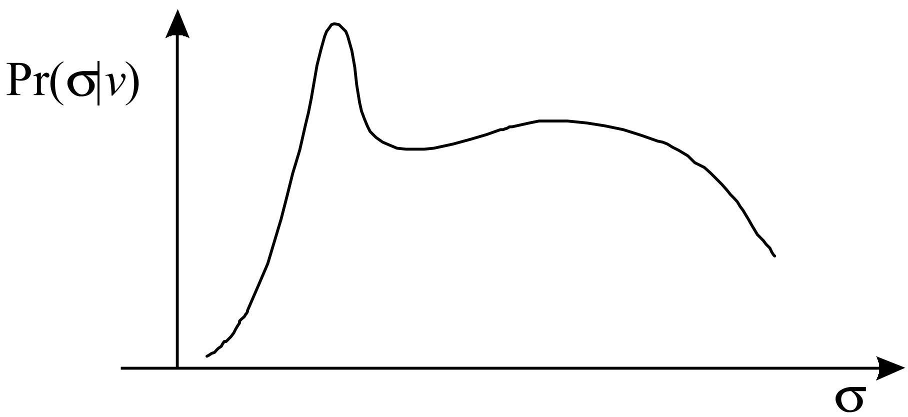

Traditional Solutions – Fourier Deconvolution

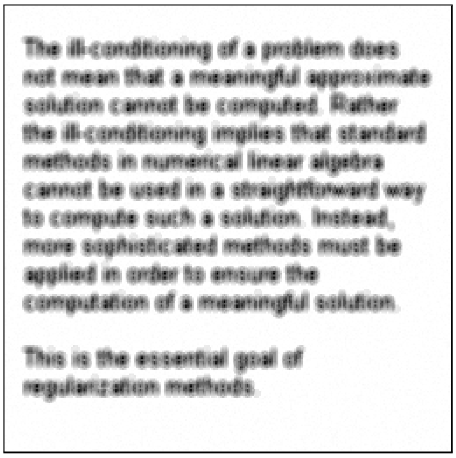

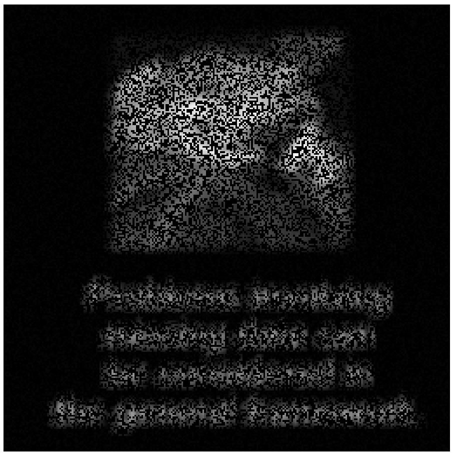

Noisey blurred image

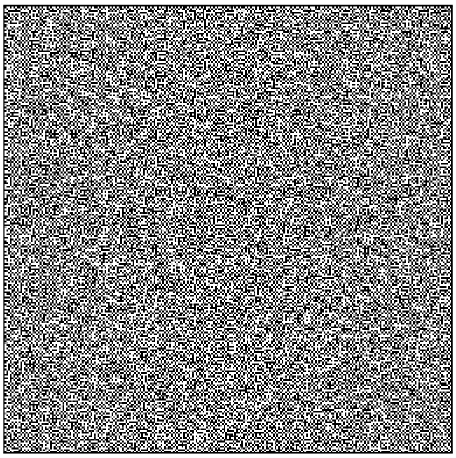

Exact inverse

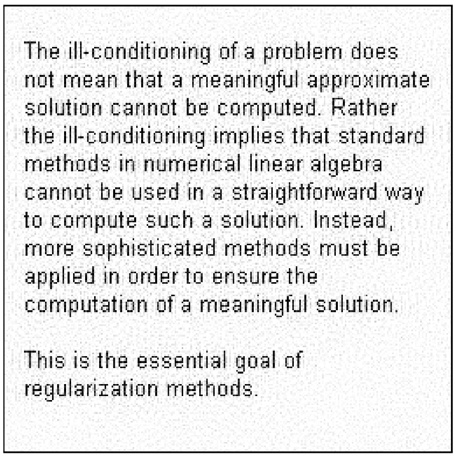

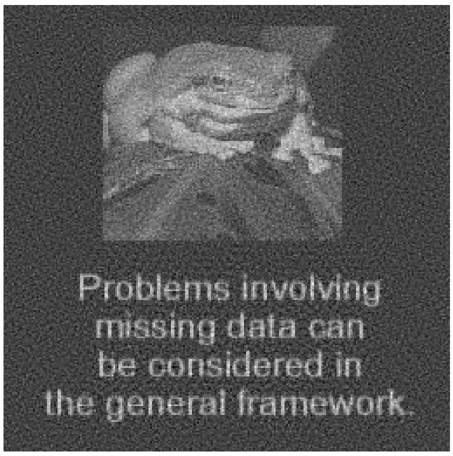

MAP solution

Bayesian Formulation for Conductivity Imaging

	current in Ω	potential in Ω	voltage electrode	current electrode	conductivity
r.v.	R	Φ	V	J	Σ
value	ρ	ϕ	v	j	σ

Posterior

$$\begin{aligned} &\Pr\left\{\Sigma=\sigma, \Phi^n=\phi^n, R^n=\rho^n \middle| \left\{J^n, V^n\right\}=\left\{j^n, v^n\right\}\right\} \\ &\quad = \Pr\left\{\left\{J^n, V^n\right\}=\left\{j^n, v^n\right\} \middle| \Sigma=\sigma, \Phi^n=\phi^n, R^n=\rho^n\right\} \\ &\qquad \times \Pr\left\{\Sigma=\sigma, \Phi^n=\phi^n, R^n=\rho^n\right\} \end{aligned}$$

$R=-\Sigma\nabla\Phi$, $\phi=\Gamma_\sigma\left(\rho|_{\partial\Omega}\right)$ and $\rho=-\sigma\nabla\phi$

$$\Pr\left\{\Sigma=\sigma, \Phi^n=\phi^n, R^n=\rho^n\right\}=\Pr\left\{\Sigma=\sigma, R^n=\rho^n\right\}$$

Stipulate $\Pr\left\{\Sigma=\sigma\right\}$ only – usually a MRF

$$\begin{aligned} &L\left(\sigma, \phi^n, \rho^n\right) \\ &\quad = \Pr\left\{\left\{J^n, V^n\right\}=\left\{j^n, v^n\right\} \middle| \Sigma=\sigma, \Phi^n=\phi^n, R^n=\rho^n\right\} \\ &\quad = \Pr\left\{\left\{V^n\right\}=\left\{v^n\right\} \middle| \Phi^n=\phi^n\left(\sigma, \rho^n\right)\right\} \\ &\qquad \times \Pr\left\{\left\{J^n\right\}=\left\{j^n\right\} \middle| R^n=\rho^n\right\} \end{aligned}$$

Errors i.i.d.

$$\begin{aligned} L\left(\sigma, \phi^n, \rho^n\right) &= \Pi_{n=1}^N \Pr\left\{V^n=v^n \middle| \Phi^n=\Gamma_\sigma\left(\rho^n|_{\partial\Omega}\right)\right\} \\ &\quad \times \Pr\left\{J^n=j^n \middle| R^n=\rho^n\right\}. \end{aligned}$$

Noise is normal (say)

$$\Pr\left\{J=j|R=\rho\right\} \sim \mathrm{N}\left(\left(\rho\left(\mathrm{x}_1\right), \rho\left(\mathrm{x}_2\right), \cdots, \rho\left(\mathrm{x}_\mathrm{k}\right)\right)^\mathrm{T}, \mathrm{s}_\rho^2\right)$$

Samples from the Prior

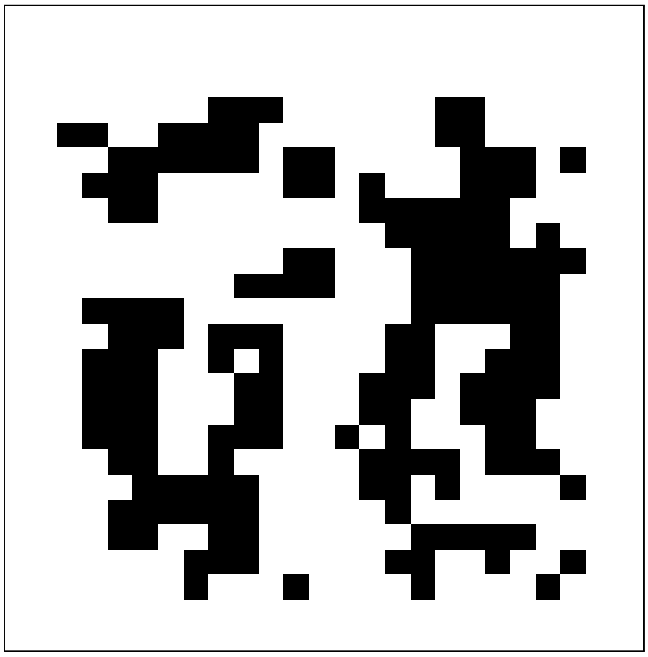

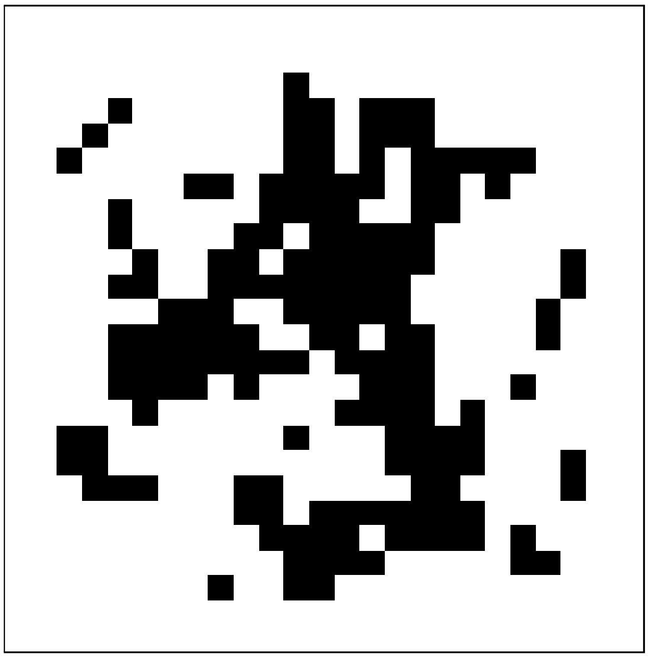

Markov chain Monte Carlo

- Monte Carlo integration
  If $\{X_t, t = 1, 2, \ldots, n\}$ are sampled from $\Pr(\sigma|v)$
  $$\mathrm{E}[f(\sigma)] \approx \frac{1}{n}\sum_{t=1}^{n} f(X_t)$$
- Markov chain
  Generate $\{X_t\}_{t=0}^{\infty}$ as a Markov chain of random variables $X_t \in \Sigma_\Omega$, with a t-step distribution $\Pr(X_t = \sigma|X_0 = \sigma^{(0)})$ that tends to $\Pr(\sigma|v)$, as $t \to \infty$.

Metopolis-Hastings algorithm

(1) given state $\sigma_t$at time t generate candidate state σ' from a proposal distribution $q(.|\sigma_t)$

(2) Accept candidate with probability
$$\alpha(X|Y) = \min\left(1, \frac{\Pr(Y|v)q(X|Y)}{\Pr(X|v)q(Y|X)}\right)$$

(3) If accepted, $X_{t+1} = \sigma'$ otherwise $X_{t+1} = \sigma_t$

(4) Repeat

$q(.|\sigma_t)$ can be any distribution that ensures the chain is:

– irreducible

– aperiodic

Three-Move Metropolis Hastings

Choose one of 3 moves with probability ζ_p, $p = 1, 2, 3$

Transition probabilities $\{\Pr^{(p)}(X_{t+1} = \sigma_{t+1}|X_t = \sigma_t)\}_{p=1}^{3}$ (reversible w.r.t. $\Pr(\sigma|v)$).

Overall transition probability is

$$\begin{aligned}\Pr(X_{t+1} &= \sigma_{t+1}|X_t = \sigma_t) \\ &= \sum_{p=1}^{3} \zeta_p \Pr{}^{(p)}(X_{t+1} = \sigma_{t+1}|X_t = \sigma_t).\end{aligned}$$

If at least one of the moves is irreducible on Σ_Ω, then the equilibrium distribution is $\Pr(\sigma|v)$.

A pixel n is a *near-neighbour* of pixel m if their lattice distance is less than $\sqrt{8}$.

An *update-edge* is a pair of near-neighbouring pixels of unequal conductivity. ($\mathcal{N}^*(\sigma)$, $\mathcal{N}_m^*(\sigma)$)

Move 1*Flip a pixel*. Select a pixel m at random and assign σ_m a new conductivity σ'_m chosen uniformly at random from the other $\mathcal{C} - 1$ conductivity values.

Move 2*Flip a pixel near a conductivity boundary*. Pick an update-edge at random from $\mathcal{N}^*(\sigma)$. Pick one of the two pixels in that edge at random, pixel m say. Proceed as in Move 1.

M**ove 3***Swap conductivities at a pair of pixels*. Pick an update-edge at random from $\mathcal{N}^*(\sigma)$. Set $\sigma'_m = \sigma_n$ and $\sigma'_n = \sigma_m$.

Experiment 1

(discrete variables – three conductivity levels)

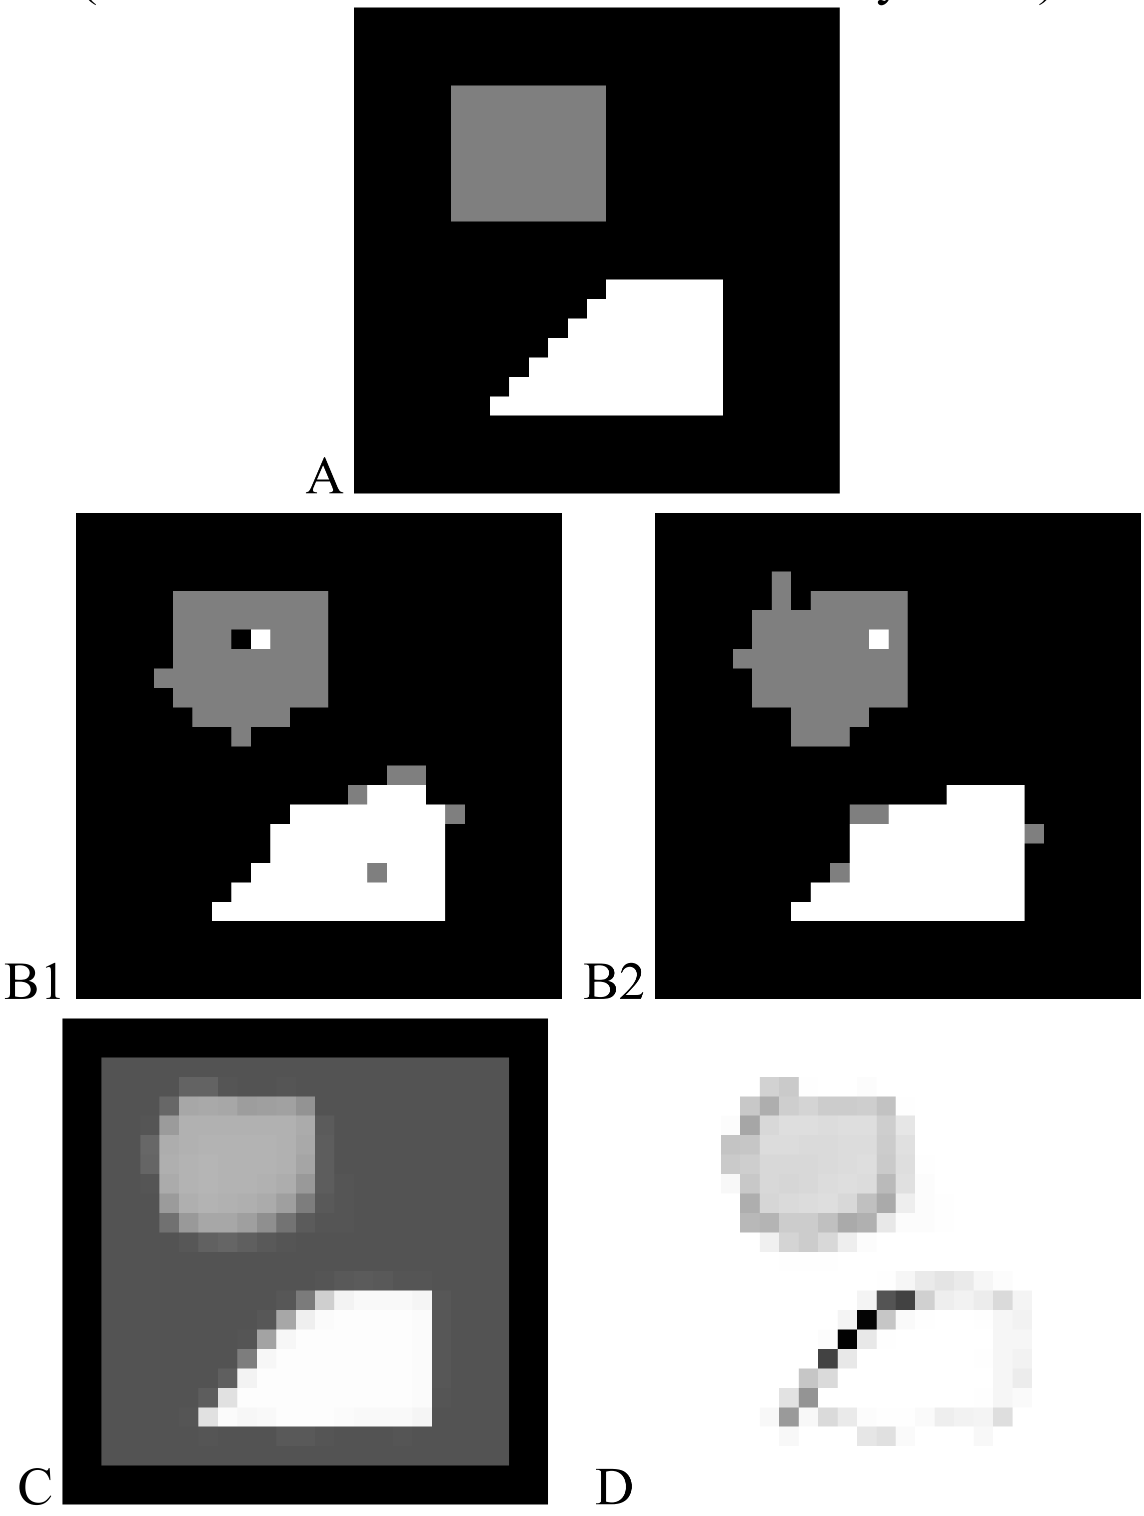

Experiment 2
(continuous variables – three conductivity types)

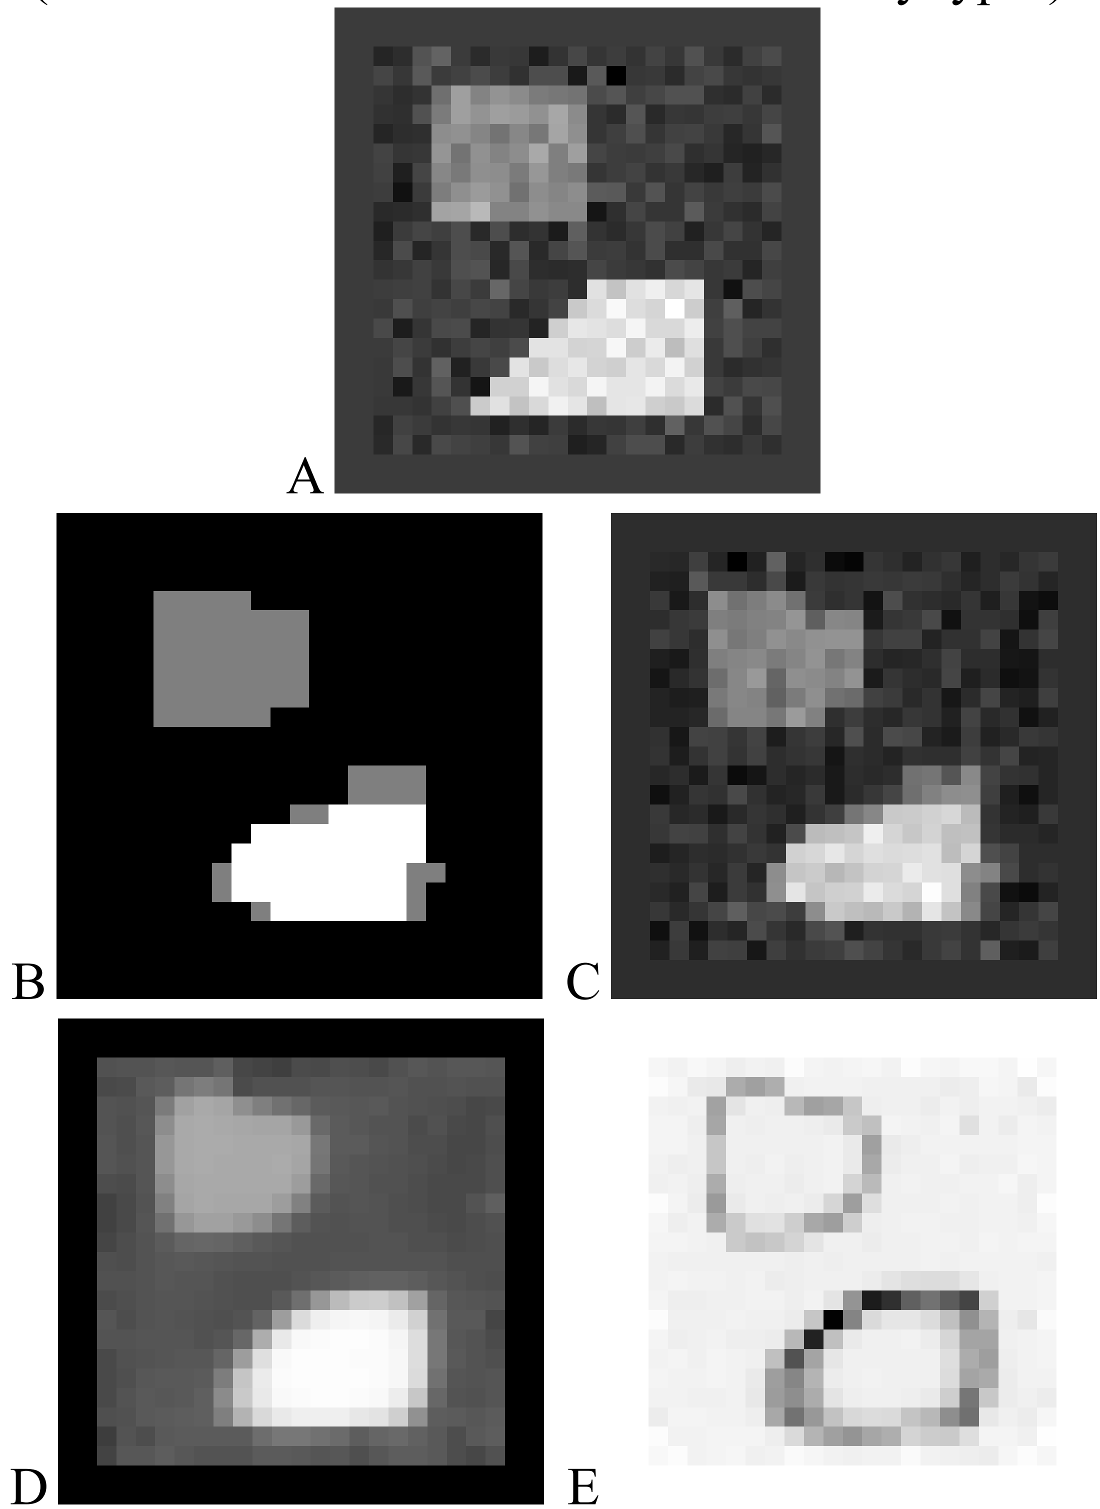

Experiment 3
(shielding)

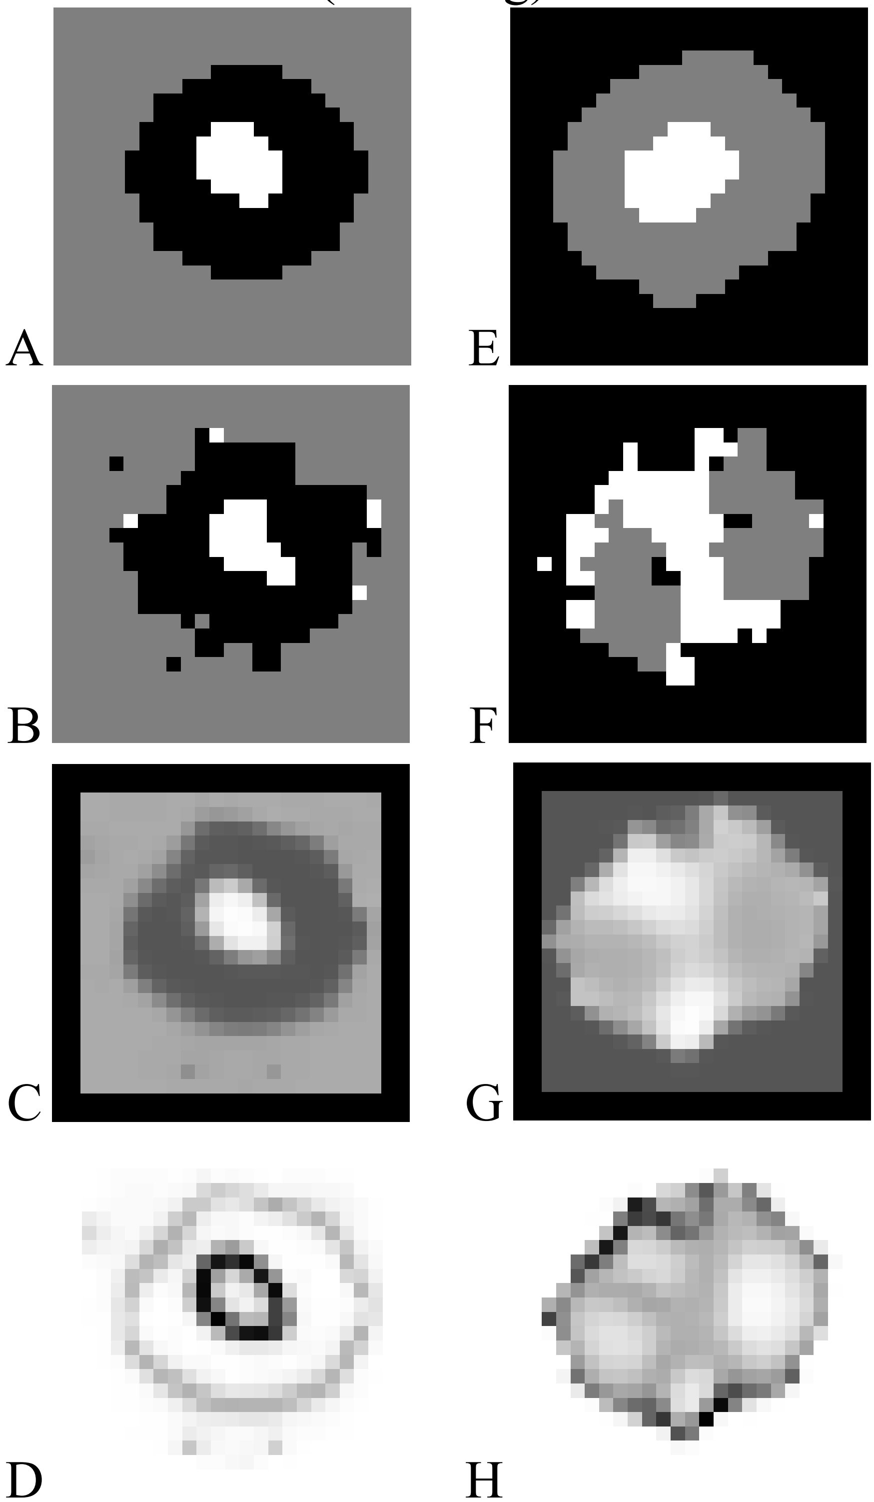

Summary

- If you can simulate the forward map then you can sample and calculate expectations over the posterior, i.e., ‘solve’ the inverse problem
- Statistical inference provides a unifying framework for inverse problems