Bayesian Inference with Embedded Model Error in Radiographic Material Reconstruction

Khachik Sargsyan, Joy Mueller, Logan Somoza (SNL-CA)
Peter Lalor, Curtis Michels, Ben S. McDonald, Andrew J. Gilbert (PNNL)
Chengkun Huang (LANL)







SIAM Conference on Imaging Science (IS26)
Salt Lake City, UT
November 17, 2026

Acknowledgements

This work is an interlab collaboration supported by the
Laboratory Directed Research and Development (LDRD) programs of
Sandia National Laboratories,
Pacific Northwest National Laboratory,
Los Alamos National Laboratory.



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Material Identification from Radiographs

- A radiograph records the areal density integrated along each ray
- Shape is easy, *material* is hard: attenuation curves look alike across materials
  - a second beam helps: two X-ray spectra, or X-ray and neutron
- High-energy (MeV) material reconstruction is typically ill-posed
- Estimates rarely come with uncertainties, at best Cramér–Rao bounds

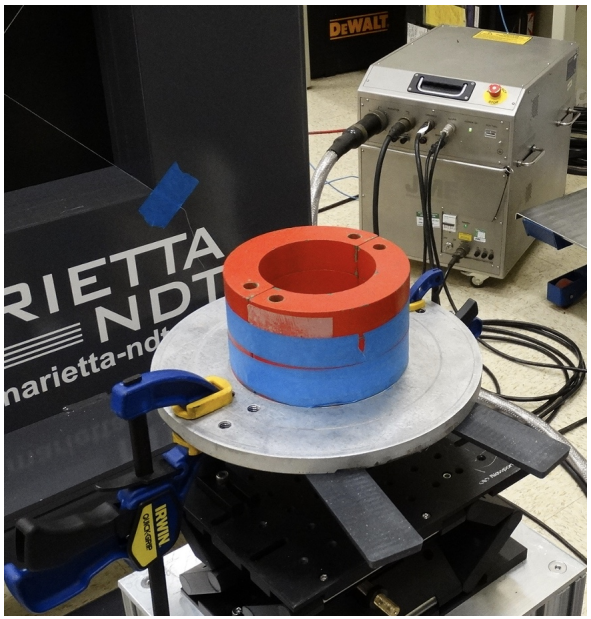
Builds on [Cunningham et al., 1995; Gilbert et al., 2025]

This talk: Bayesian inference over a *geometric* description of the object

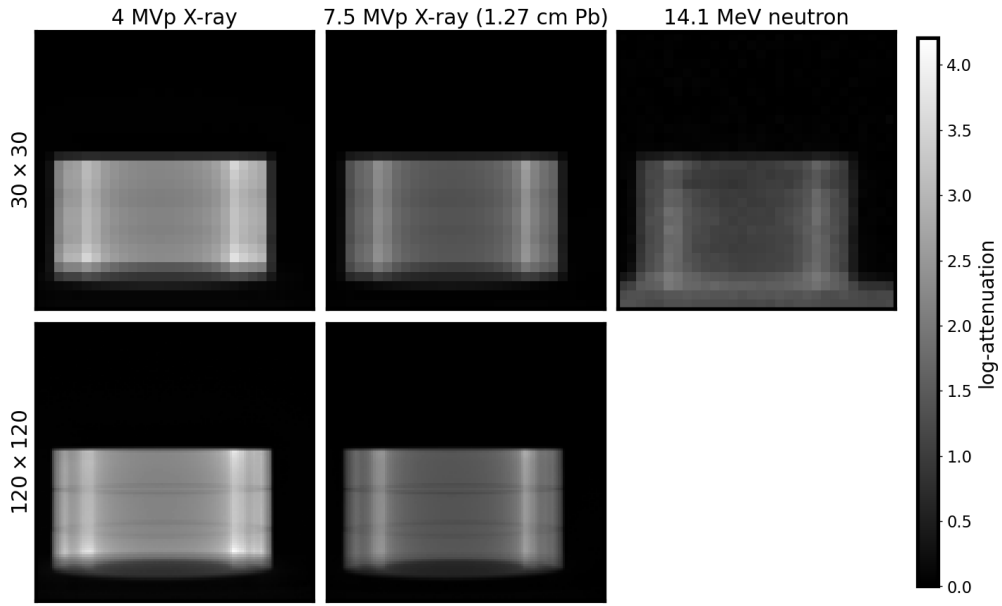
What the Bayesian view buys

- Identifiability analysis
- Correlations, not only error bars
- Explicit data noise and weighting
- Model error, embedded in the physics
- Uncertainty in areal density
- What more pixels can and cannot buy

Data: Two X-Ray and One Neutron Beam



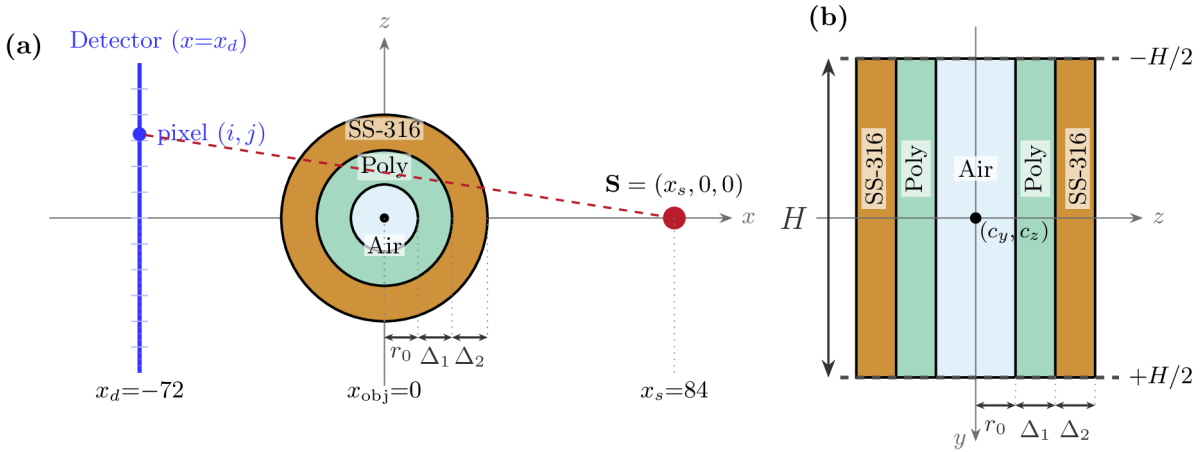
steel annulus on its mounting plate



- X-ray: betatron at 4 MVp (unfiltered) and 7.5 MVp (1.27 cm Pb), up to 120×120
- Neutron: DT generator, 14.1 MeV, associated-particle imaging; 30×30 PVT array
- Targets: stainless-steel and polyethylene annuli, radii 4.2 and 6.6 cm
- Inference on log-attenuation data $\mathbf{d}$

[Gilbert et al., 2025]

Geometric Model $\mathcal{G}$: 28,800 $\Rightarrow$ 6



		range (cm)
r_0	core radius	$[0, 10]$
Δ_1	poly shell	$[0, 5]$
Δ_2	SS-316 shell	$[0, 5]$
H	height	$[0, 20]$
c_y	offset in y	$[-5, 5]$
c_z	offset in z	$[-5, 5]$

$$\boldsymbol{\theta} = [r_0, \Delta_1, \Delta_2, H, c_y, c_z], \qquad r_1 = r_0 + \Delta_1, \quad r_2 = r_1 + \Delta_2$$

- Pixel field: $120 \times 120 \times 2 = 28{,}800$ densities, potentially TV-regularized
- Geometry: six interpretable geometric parameters, ordering $r_0 \leq r_1 \leq r_2$ by construction
- Effectively regularizes the problem, and makes sampling feasible: adaptive MCMC learns a $d \times d$ proposal covariance
  - 21 entries at $d = 6$ vs. $> 4 \times 10^8$ for the field

Radiographic Model $\mathcal{M}$: Areal Density $\Rightarrow$ Radiograph

- Polychromatic Beer–Lambert transmission, inspection k, pixel (i, j):

$$T_k(i,j) = \frac{\sum_E \phi_{0,k}(E)\, \exp\!\left[-\sum_m \mu_m(E)\, \rho_m(i,j)\right]}{\sum_E \phi_{0,k}(E)}$$

  - $\phi_{0,k}(E)$: measured source spectra; $\mu_m(E)$: mass attenuation coefficients

- Scatter removed from the data in preprocessing
- Detector blur, a calibrated mixture of Gaussians, acting on the transmission deficit:

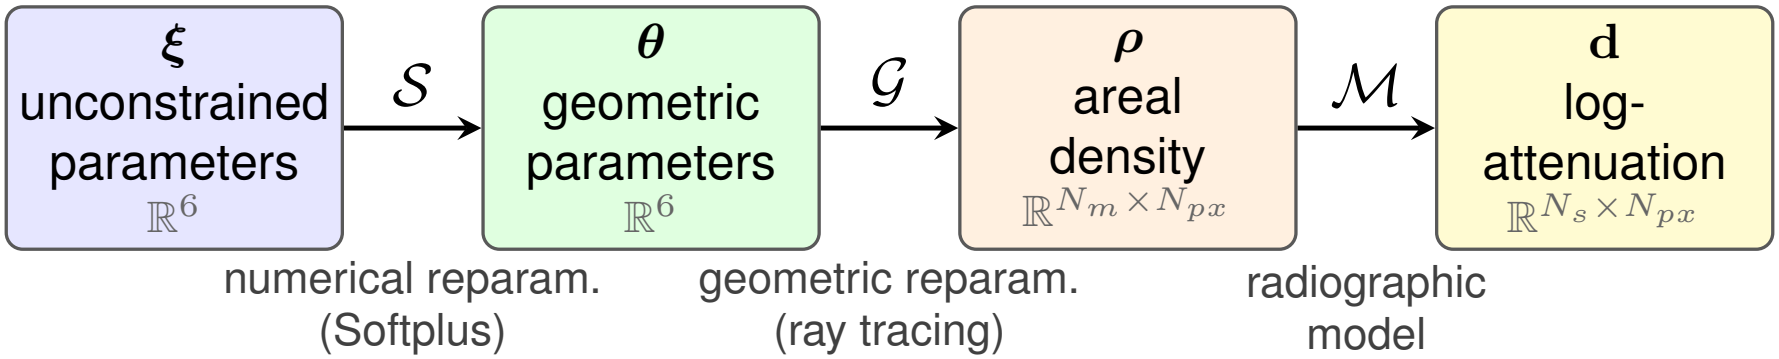
$$T_k^{\mathrm{blr}} = C_k\,(T_k - 1)\,C_k^{\top} + 1$$

- Log-attenuation, matching the flat-field-corrected measurement:

$$\left[\mathcal{M}(\boldsymbol{\rho})\right]_k(i,j) = -\log T_k^{\mathrm{blr}}(i,j)$$

- Evaluated through precomputed look-up tables in $\boldsymbol{\rho}$ for fast repeated calls

Full Forward Model $\mathcal{F} = \mathcal{M} \circ \mathcal{G} \circ \mathcal{S}$



- Numerical reparameterization $\mathcal{S}$: unconstrained $\boldsymbol{\xi} \in \mathbb{R}^6$, positivity by construction using Softplus function

$$\boldsymbol{\theta} = \mathcal{S}(\boldsymbol{\xi}) = \left[S_+(\xi_1), \ldots, S_+(\xi_4),\, \xi_5,\, \xi_6\right], \qquad S_+(u) = \ln(1 + e^u)$$

- Geometric reparameterization $\mathcal{G}$: analytic ray tracing through the concentric shells gives the areal densities $\rho_m(i, j)$
- Radiographic model $\mathcal{M}$: previous slide
- Analytic vector–Jacobian products $J^\top \mathbf{u}$, $J = \partial \mathcal{F} / \partial \boldsymbol{\xi}$, by the chain rule through $\mathcal{M}$, $\mathcal{G}$ and $\mathcal{S}$

Deterministic Formulation

- Point estimate by weighted least squares over the unconstrained parameters:

$$\hat{\boldsymbol{\xi}} = \arg\min_{\boldsymbol{\xi}\in\mathbb{R}^6} \ell(\boldsymbol{\xi}), \qquad \ell(\boldsymbol{\xi}) = \tfrac{1}{2}\left\|\Sigma^{-1/2}\big(\mathbf{d} - \mathcal{F}(\boldsymbol{\xi})\big)\right\|^2, \qquad \hat{\boldsymbol{\theta}} = \mathcal{S}(\hat{\boldsymbol{\xi}})$$

  - $\Sigma = \mathrm{diag}(\boldsymbol{\sigma}^2)$ weights the residuals; its overall scale does not move the optimum

- L-BFGS-B [Byrd et al., 1995] with the analytic gradient

$$\nabla_{\boldsymbol{\xi}}\ell = J^\top \Sigma^{-1}\big(\mathcal{F}(\boldsymbol{\xi}) - \mathbf{d}\big),$$

- Six parameters instead of a $28{,}800$-unknown field: well-posed without explicit penalty regularization

But a single point carries no confidence assessment. Additionally, different geometries can produce indistinguishable projections
$\Rightarrow$ Bayesian inference

Bayesian Inference

$$\underbrace{\pi(\boldsymbol{\xi} \mid \mathbf{d})}_{\text{posterior}} = \frac{\overbrace{\pi(\mathbf{d} \mid \boldsymbol{\xi})}^{\text{likelihood}} \overbrace{\pi(\boldsymbol{\xi})}^{\text{prior}}}{\underbrace{\pi(\mathbf{d})}_{\text{evidence}}}$$

- Likelihood: independent Gaussian residuals, $\pi(\mathbf{d} \mid \boldsymbol{\xi}) \propto \exp\big(-\ell(\boldsymbol{\xi})\big)$ with the deterministic ℓ
- Prior: uniform on a physical box $\mathcal{B}$ for $\boldsymbol{\theta}$, transported to $\boldsymbol{\xi}$ via the softplus Jacobian (happens to be the sigmoid function)
- Posterior:

$$\pi(\boldsymbol{\xi} \mid \mathbf{d}) \;\propto\; \exp\big(-\ell(\boldsymbol{\xi})\big) \prod_{k=1}^{4} \sigma(\xi_k), \qquad \mathcal{S}(\boldsymbol{\xi}) \in \mathcal{B}$$

  Maximum a Posteriori (MAP) $=$ deterministic optimum

- Adaptive MCMC [Haario et al., 2001] in PyTUQ, 10^5 samples, then pushed forward to credible bands for $\boldsymbol{\theta}$, $\boldsymbol{\rho} = \mathcal{G}(\boldsymbol{\theta})$ and $\mathbf{d} = \mathcal{F}(\boldsymbol{\xi})$

PyTUQ: https://github.com/sandialabs/pytuq

Likelihood Function and Residual Weighting

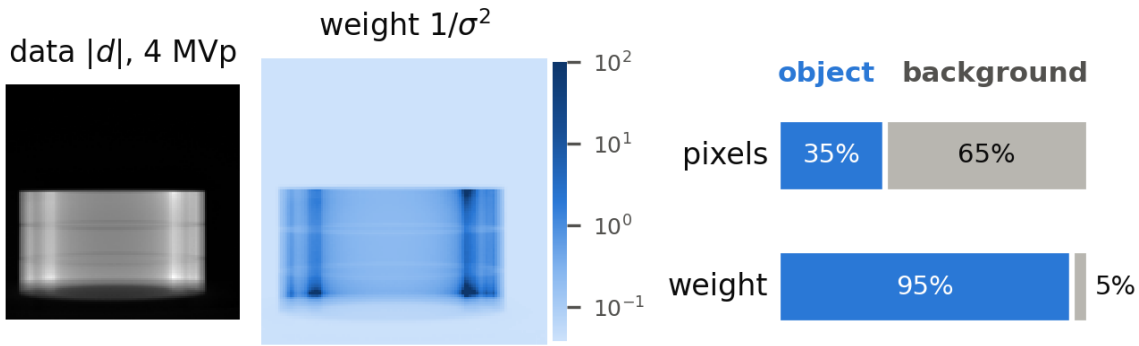
- Background has no geometric information, yet most of the pixels
- Gaussian likelihood over inspections k and pixels (i, j):

$$\pi(\mathbf{d} \,|\, \boldsymbol{\xi}) \propto e^{-\ell(\boldsymbol{\xi})}, \qquad \ell(\boldsymbol{\xi}) = \frac{1}{2} \sum_k \sum_{i,j} \left(\frac{d_k(i,j) - [\mathcal{F}(\boldsymbol{\xi})]_k(i,j)}{\sigma_k(i,j)} \right)^2$$

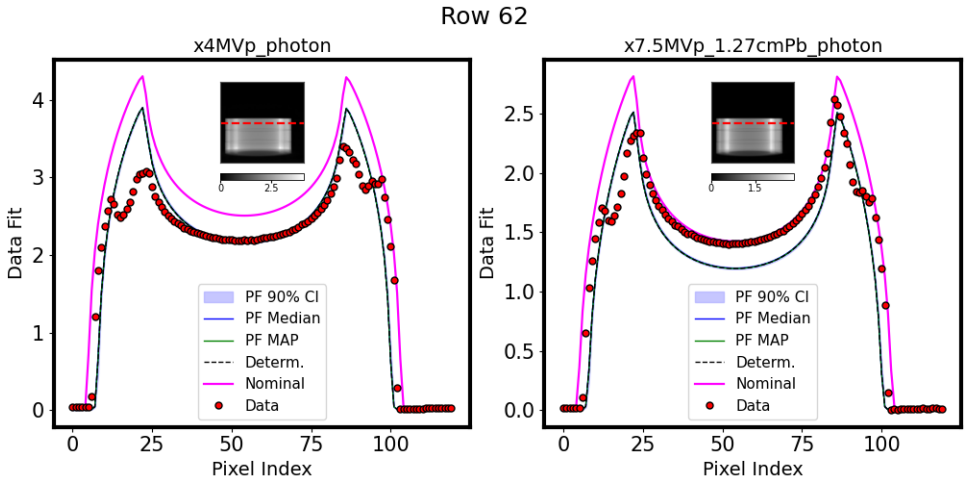
with a data-dependent scale that down-weights the background:

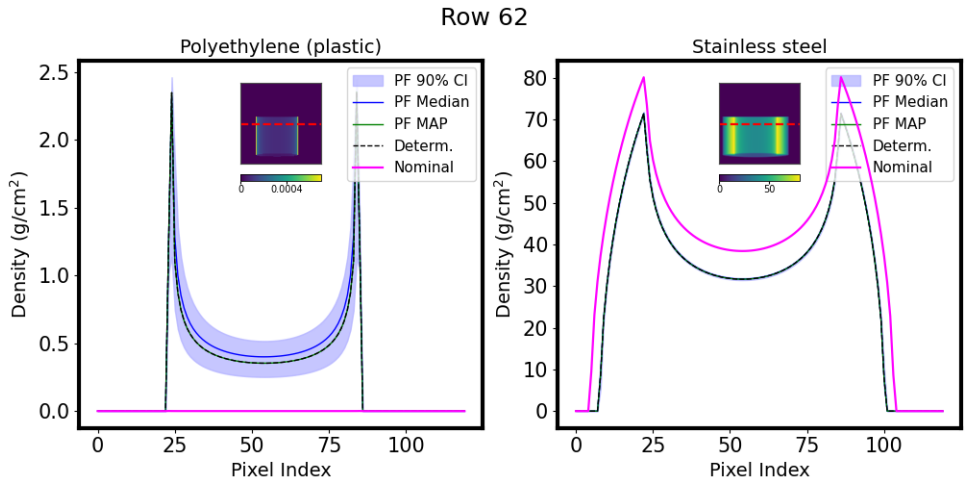
$$\sigma_k(i,j) = a + b\Big(1 - \min\Big(\frac{|d_k(i,j)|}{d_{\text{ref}}}, 1\Big)\Big), \qquad a = 0.1,\ b = 5$$

- d_{ref}: 0.999 quantile of $|\mathbf{d}| \Rightarrow \sigma = a$ at peak, $a + b$ in background
- weight ratio $\big((a+b)/a\big)^2 = 2601$ between peak and background



Demo: Two X-Ray Beams at 120×120





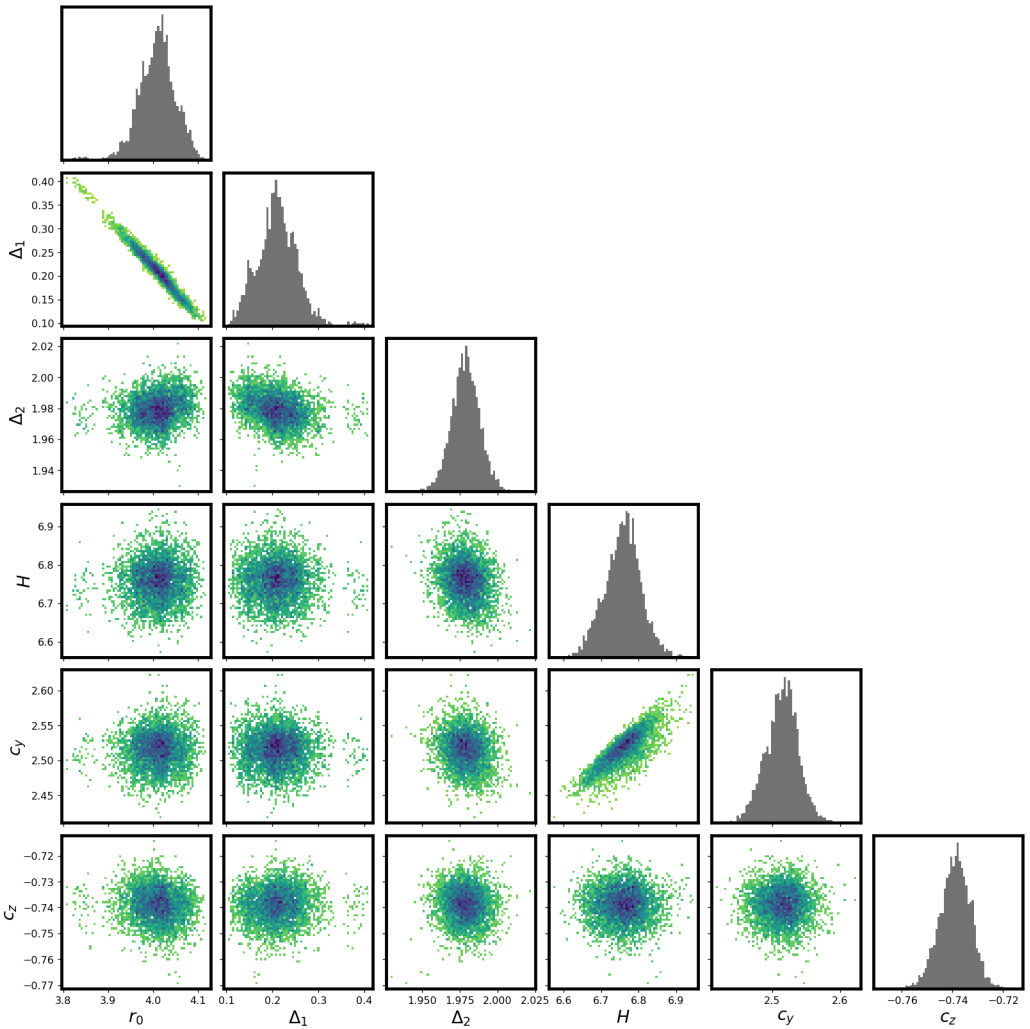
Log-attenuation
(push forward of $\mathcal{M} \times \mathcal{G} \times \mathcal{S}$)

- The prediction tracks the data; the credible band is too narrow to see
- The visible gap is between prediction and *nominal* object; mounting holes drive part of it (model error not accounted for)

Areal density
(push forward of $\mathcal{G} \times \mathcal{S}$)

- Steel band: about 1% of the median at the peak
- Polyethylene (nominally absent): about 50%

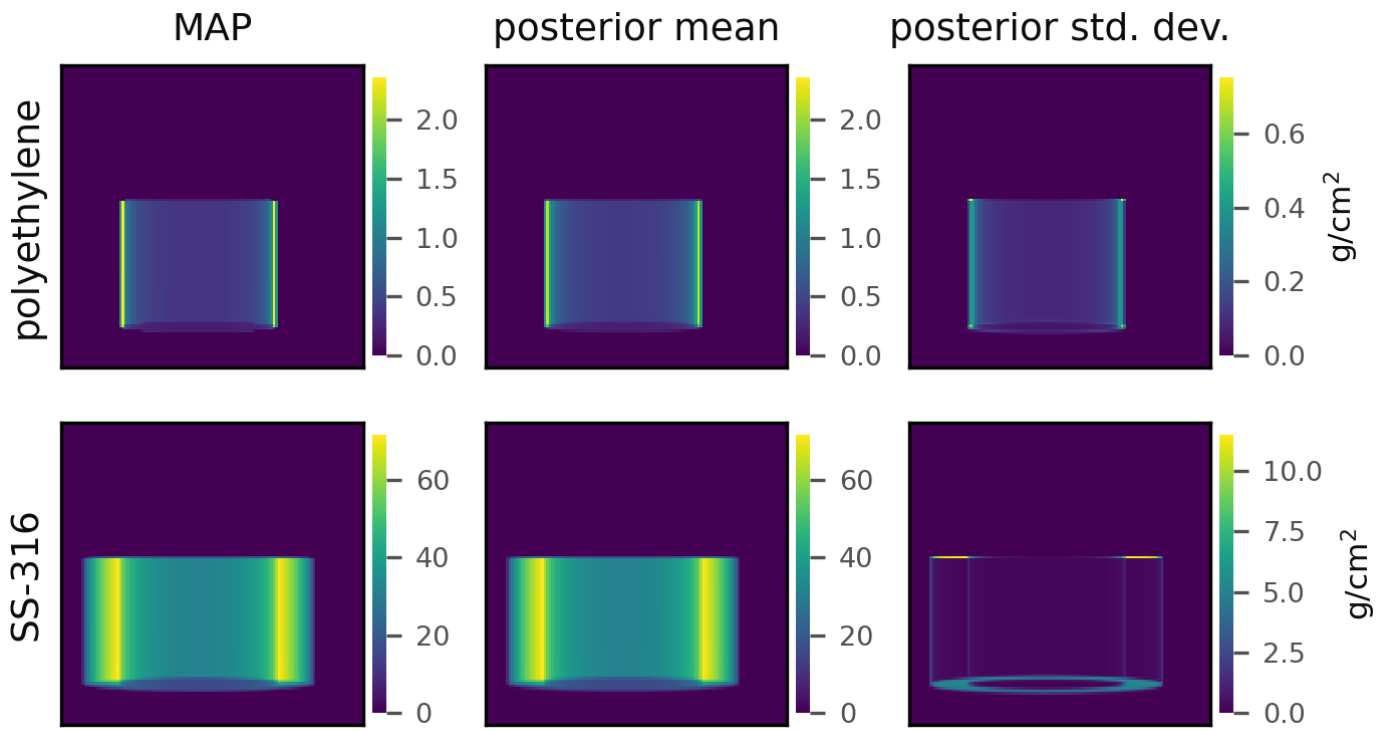
Uncertainty in Input Space: Joint Posterior



- (r_0, Δ_1): corr $= -0.98$
  - only the sum r_1 is determined
- (H, c_y): corr $= +0.81$
  - two near-collinear X-ray views cannot tell a taller cylinder from a shifted one

The Bayesian posterior exposes couplings set by the geometry; a deterministic inversion returns one point and hides joint impact

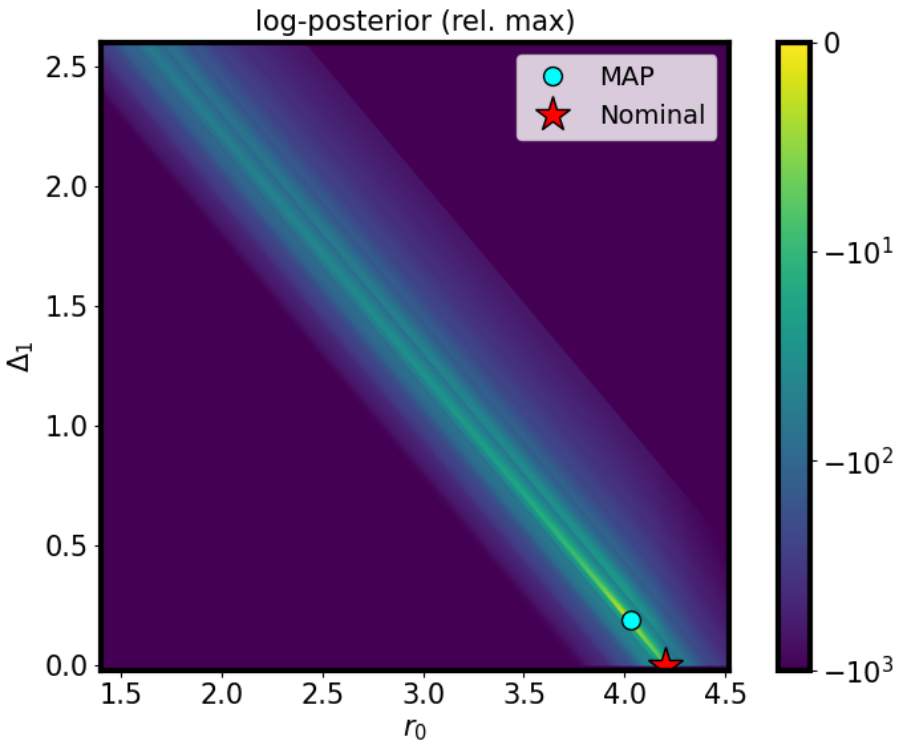
Uncertainty in Physical Space: Outline vs. Footprint



- Steel (SS-316) blocks most of the beam: the *amount* of steel is well known; only *where* it ends is uncertain, so the uncertainty outlines the object
- Plastic (polyethylene) hides behind the steel: its *amount* is uncertain, so the uncertainty fills the whole region it occupies

The Bayesian posterior delineates both where a material ends and how much of it there is; a deterministic result shows neither

Non-identifiability of the Air/Polyethylene Split

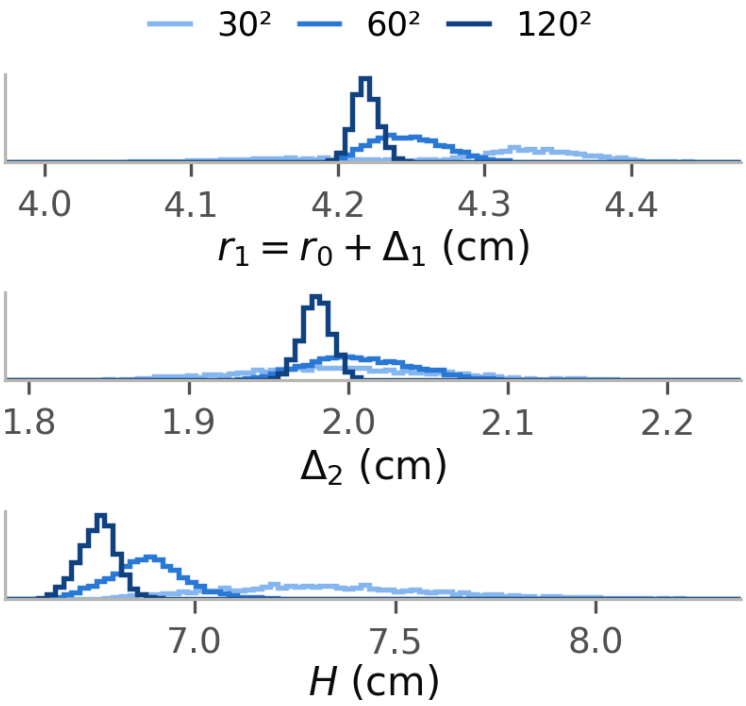


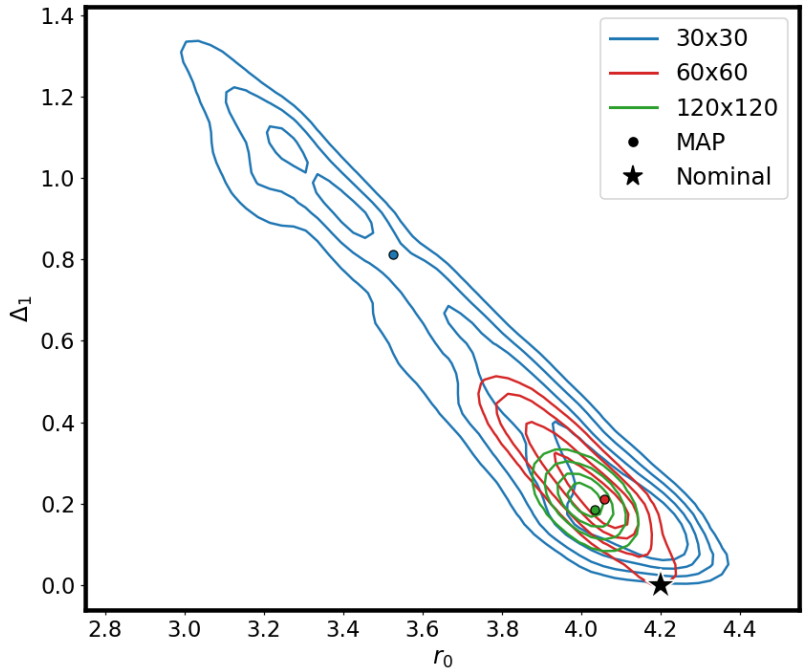
exact log-posterior on an (r_0, Δ_1) grid, other parameters at the MAP

- The data points to where the steel begins, $r_1 = r_0 + \Delta_1$, but barely constrains how the space inside splits between air and polyethylene: a long, narrow ridge along $r_0 + \Delta_1 = \text{const}$
- Why: the steel blocks most of the beam, and the air/polyethylene interface sits in its shadow
- Practical non-identifiability [Raue et al., 2009]: report r_1, not the split (r_0, Δ_1).

The Bayesian posterior highlights what data can and can not determine;
a deterministic solve does not reveal the ridge

Resolution Dependence





joint (r_0, Δ_1)

- More pixels shrink the posterior for all parameters, and *along* the ridge for (r_0, Δ_1)
- The accuracy does not improve: the gap to the nominal geometry does not shrink

Posterior narrows with every added pixel, the model–data gap does not: model error sets the floor $\Rightarrow$ it has to enter the inference

Why Model Error, and Why Inside the Model

- The forward model is an idealization: spectrum, scatter and blur; uniform shells; holes, blocks and pedestal not modeled
- Ignoring it leads to overconfident and biased predictions
  - more pixels, narrower uncertainty, same misfit: 90% band covers $40\% \to 20\% \to 7\%$ with resolution increase
  - increasingly sure about the wrong values of the geometric parameters

How to handle model error

- Outside [Kennedy and O'Hagan, 2001]: $\mathbf{d} = \mathcal{F}(\boldsymbol{\theta}) + \boldsymbol{\delta} + e$
  - can violate physics; confounded with noise and $\boldsymbol{\theta}$
  - lives on the measured data space: no impact on areal density
- Inside, embedded [Sargsyan et al., 2019]: $\mathbf{d} = \mathcal{F}(\boldsymbol{\theta} + \boldsymbol{\delta}) + e$
  - keeps physics and positivity; error sits at the assumptions
  - propagates to parameters, areal densities, any model outputs

Model Error, Embedded in the Geometry

- Embedded model error [Sargsyan et al., 2015, 2019]: a random correction *inside* the model, $\boldsymbol{\theta} \to \boldsymbol{\theta} + \boldsymbol{\delta}$, with

$$\boldsymbol{\delta} = \big(a_1 z_0,\ a_2 z_0,\ s_{\Delta_2} z_1,\ s_H z_2,\ s_{c_y} z_3,\ s_{c_z} z_4\big), \qquad \boldsymbol{z} \sim \mathcal{N}(0, I)$$

$$\mathbf{d} = \mathcal{F}(\hat{\boldsymbol{\theta}} + \boldsymbol{\delta}) + b_k + e$$

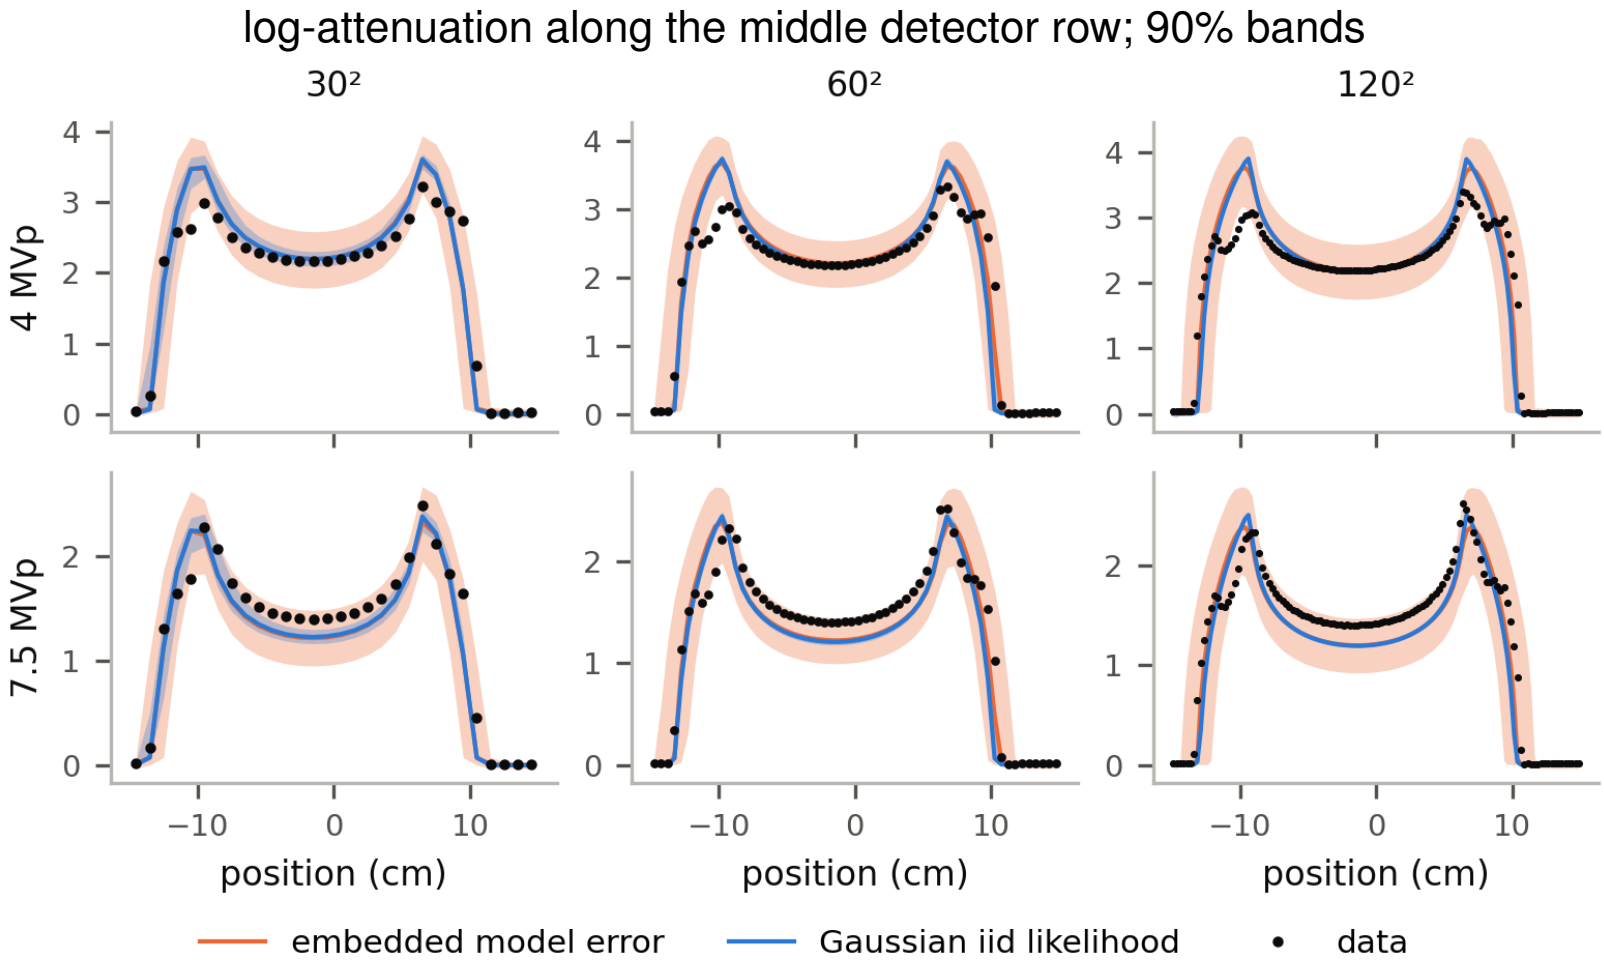
  - one shared factor z_0 for (r_0, Δ_1): perfectly correlated along the previously discovered posterior ridge
  - b_k: background offset of beam k; e: pixel noise, plus a small floor away from the object
- Two stages: $\boldsymbol{\theta}$ fixed at its MAP $\hat{\boldsymbol{\theta}}$, then the scales (a, s) by MAP only
  - no MCMC: the corrections carry the predictive uncertainty
- Marginal Gaussian: pixels independent given their embedded mean and variance

$$\pi(\mathbf{d} \,|\, a, s) = \prod_p \mathcal{N}\big(d_p;\ \mu_p + b_k,\ \sigma_p^2 + n_k^2 + f_p^2\big)$$

$$\mu_p = \mathbb{E}_{\boldsymbol{z}}\big[\mathcal{F}_p(\hat{\boldsymbol{\theta}} + \boldsymbol{\delta})\big], \qquad \sigma_p^2 = \mathbb{V}_{\boldsymbol{z}}\big[\mathcal{F}_p(\hat{\boldsymbol{\theta}} + \boldsymbol{\delta})\big]$$

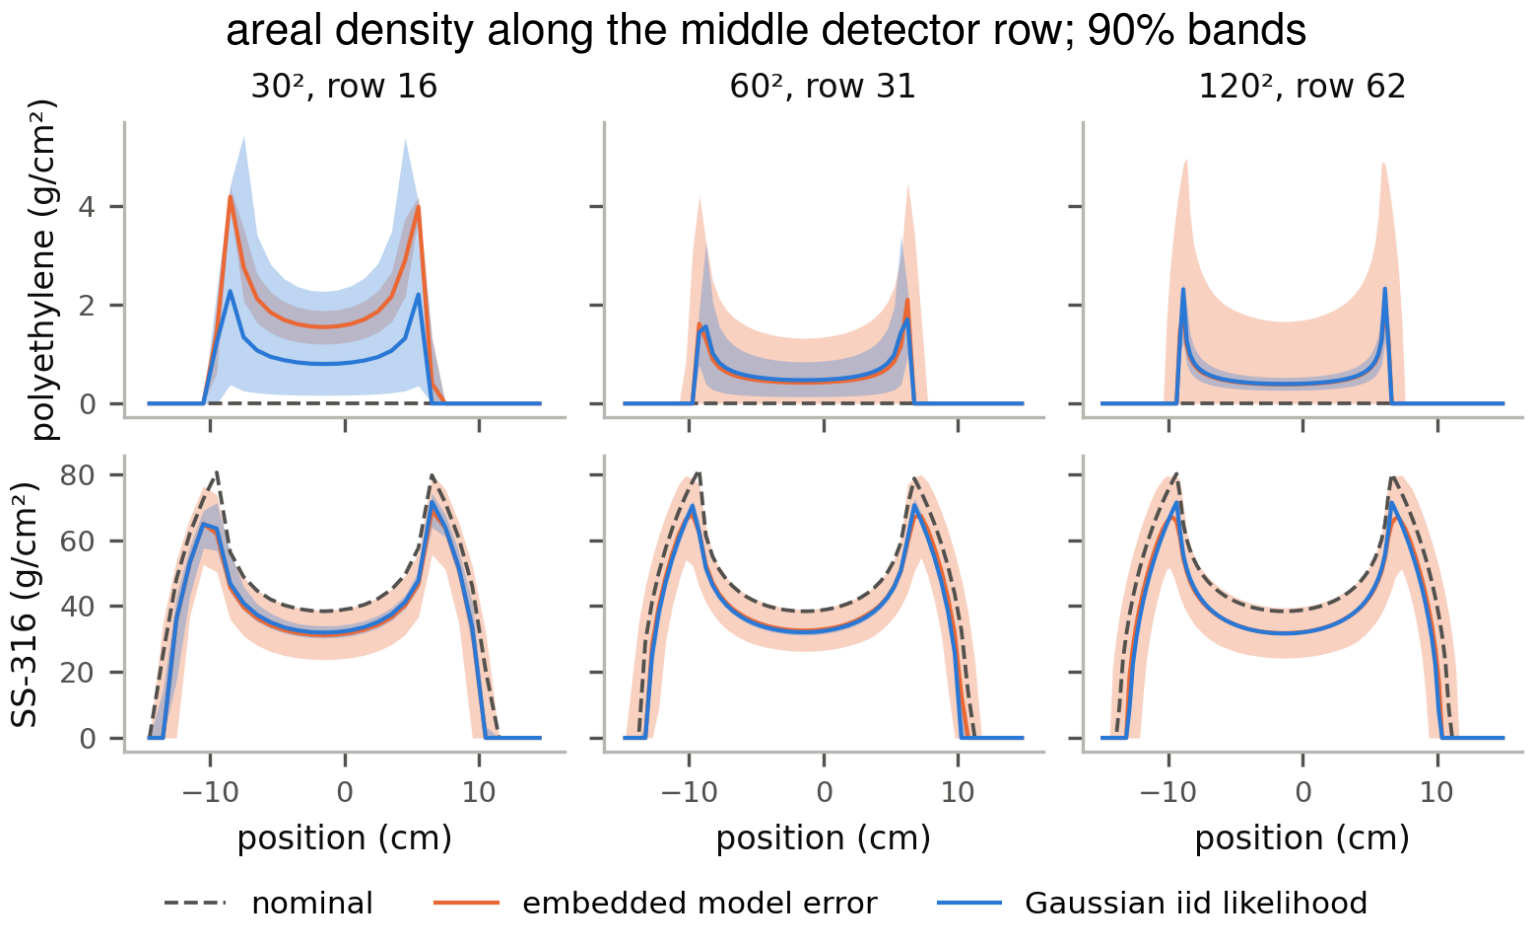
  - moments computed by sampling the exact model
  - n_k: pixel noise; f_p: floor away from the object (zero within 1.5 cm)

Data Fit Across Resolutions



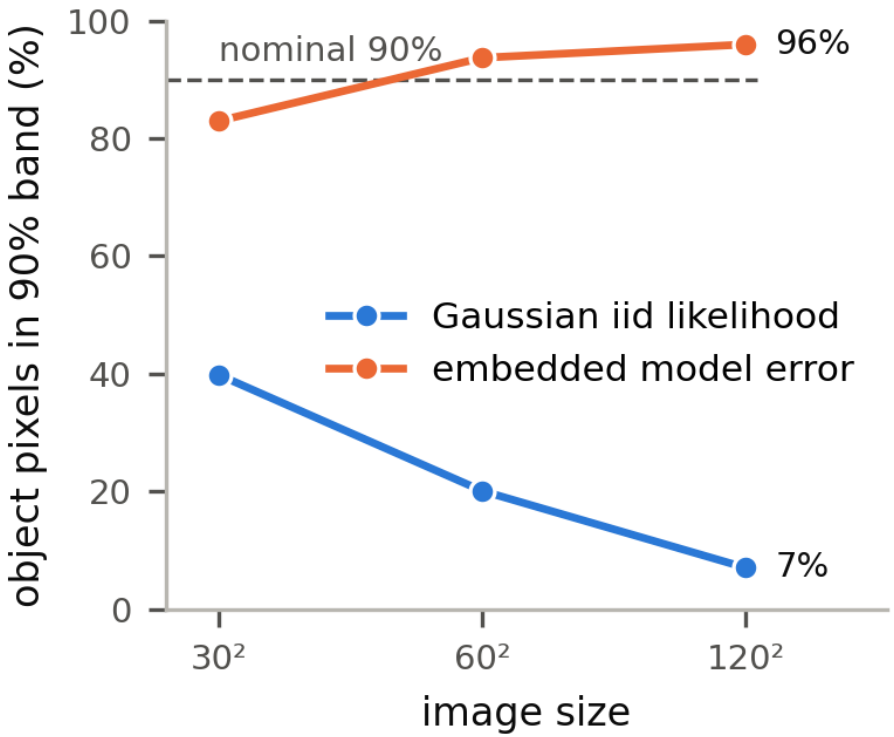
- *Gaussian iid likelihood:* the band collapses onto its median as pixels are added
- *Embedded model error likelihood:* about the same width at every resolution, covering the dips at holes and walls

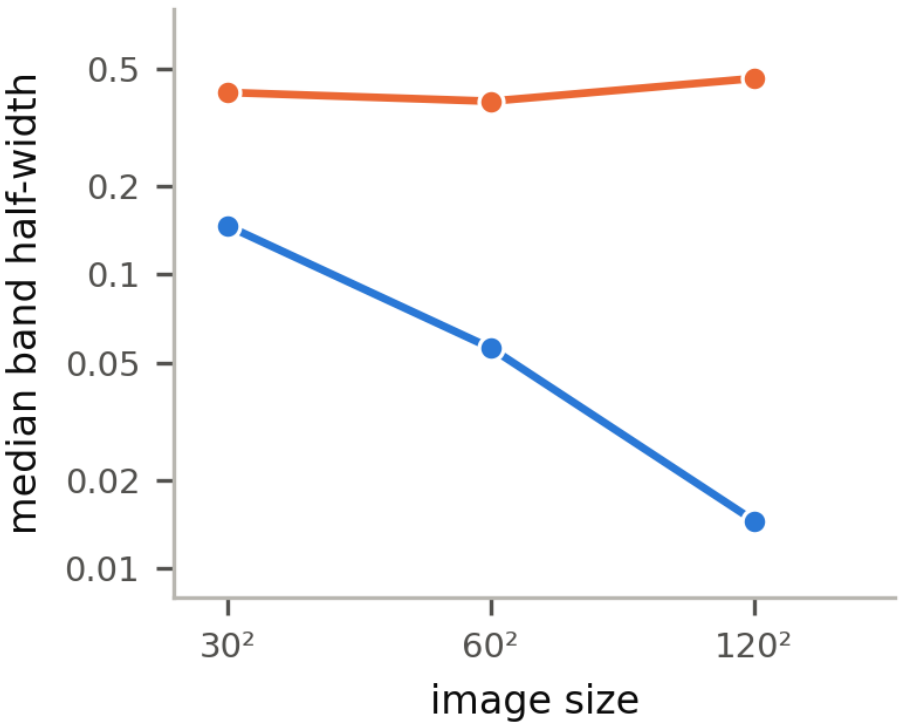
Areal Density Across Resolutions



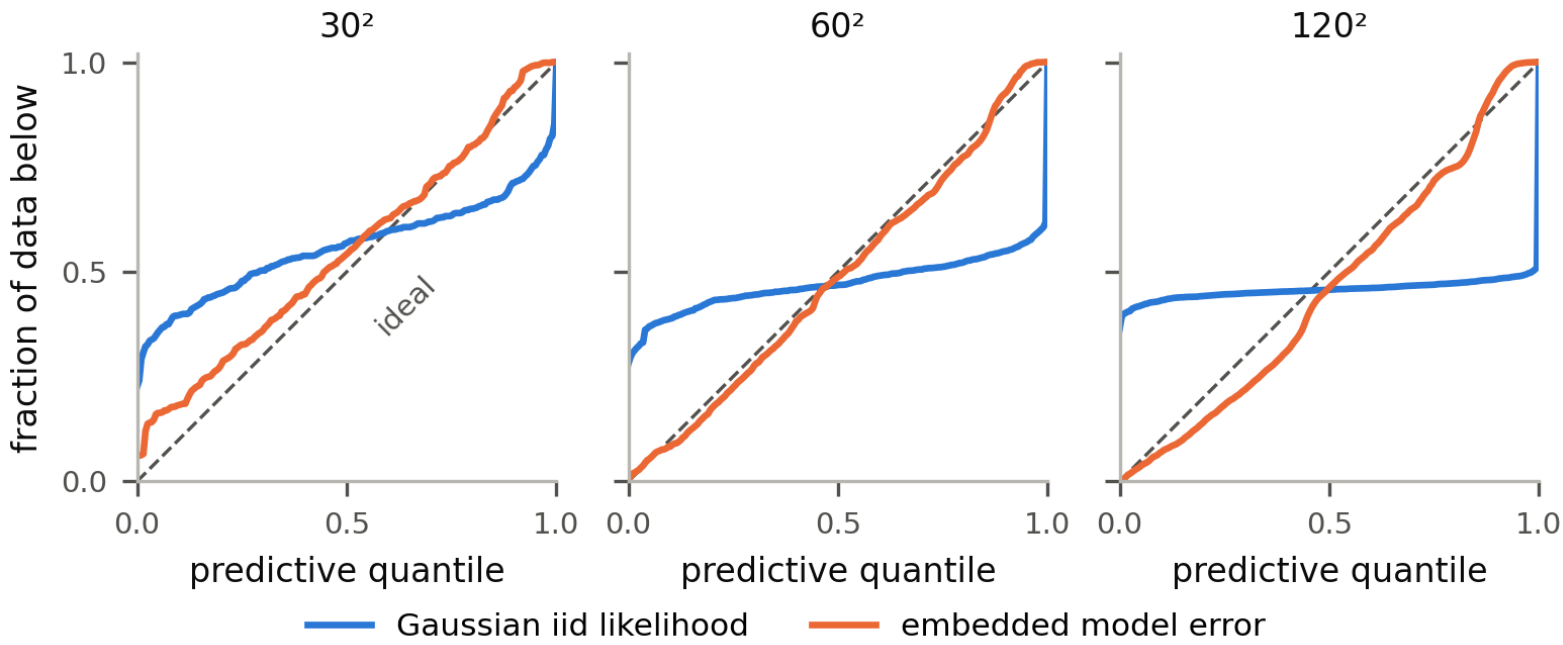
- SS-316: the embedded band keeps its width and reaches the nominal profile in most places; the current one shrinks onto its median
- Polyethylene: the embedding allows a phantom shell of zero to a few g/cm^2

Uncertainty Band and Data Coverage Remain Stable with More Pixels





Uncertainty Is Better Calibrated for All Quantiles



- Fraction of object data below each predictive quantile; ideal: the diagonal
- Gaussian iid likelihood: a step, nearly every datum outside its spread
- Embedded model error: near the diagonal at every resolution

Preliminary, but the uncertainty bands stay honest as the data grows

Summary

- Geometric reparameterization makes Bayesian inference tractable: tens of thousands of unknowns $\Rightarrow$ six physical parameters

- The Bayesian posterior shows what a deterministic point estimate does not
  - which combinations the data is informative on
  - how the parameters trade off, as set by the viewing geometry
  - where a material ends vs. how much of it there is
  - how modeling choices, e.g. the weighting, move the answer
  - explicit treatment of model error

J. Mueller, P. Lalor, L. Somoza, C. Michels, B. S. McDonald, C. Huang, A. J. Gilbert, K. Sargsyan, "Radiographic Material Reconstruction via Geometric Inference and Bayesian Uncertainty Quantification," to be submitted to *SIAM Journal on Imaging Sciences*, 2026.

- Model error: intervals narrow with more data, the misfit does not
  - embedded in the model, it keeps the physics/geometry, reaches every quantity of interest, and keeps the bands calibrated at every resolution

Trunk

Literature

Radiography and material reconstruction

Roessl, E., Herrmann, C., "Cramér–Rao lower bound of basis image noise in multiple-energy x-ray imaging", *Phys. Med. Biol.*, 54(5), 2009.

Gilbert, A. J., *et al.*, "X-Ray and Neutron Radiography for Quantitative Material Reconstructions", *IEEE Trans. Nucl. Sci.*, 72(7), 2025.

Campbell, L. W., Gilbert, A. J., Wittman, R. S., "Scatter and Blur Corrections for High-Energy X-Ray Radiography", *IEEE Trans. Nucl. Sci.*, 2025.

Nguyen-Fotiadis, N. T. T., *et al.*, "Physics-aware Discrete Reparameterization with Symmetry-aware Bayesian Fusion for Multimodal Radiography", *NeurIPS ML4PS Workshop*, 2025.

Geometric parameterizations

Aghasi, A., Kilmer, M., Miller, E. L., "Parametric Level Set Methods for Inverse Problems", *SIAM J. Imaging Sci.*, 4(2), 2011.

Iglesias, M. A., Lu, Y., Stuart, A. M., "A Bayesian Level Set Method for Geometric Inverse Problems", arXiv:1504.00313, 2015.

Meles, G. A., Linde, N., Marelli, S., "Bayesian tomography with prior-knowledge-based parametrization and surrogate modelling", *Geophys. J. Int.*, 231(1), 2022.

Software

Baillo, E., Sargsyan, K., Debusschere, B., "PyTUQ: Python Toolkit for Uncertainty Quantification (v1.0.0z)", Zenodo, 2025.

Bayesian inference, sampling, optimization

Cunningham, G. S., Hanson, K. M., Jennings, G. R., Wolf, D. R., "An interactive tool for Bayesian inference", *Review of Progress in Quantitative Nondestructive Evaluation*, 1995.

Kaipio, J., Somersalo, E., "Statistical and Computational Inverse Problems", Springer, 2005.

Stuart, A. M., "Inverse problems: A Bayesian perspective", *Acta Numerica*, 19, 2010.

Haario, H., Saksman, E., Tamminen, J., "An adaptive Metropolis algorithm", *Bernoulli*, 7(2), 2001.

Byrd, R. H., Lu, P., Nocedal, J., Zhu, C., "A limited memory algorithm for bound constrained optimization", *SIAM J. Sci. Comput.*, 16(5), 1995.

Raue, A., Kreutz, C., Maiwald, T., *et al.*, "Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood", *Bioinformatics*, 25(15), 2009.

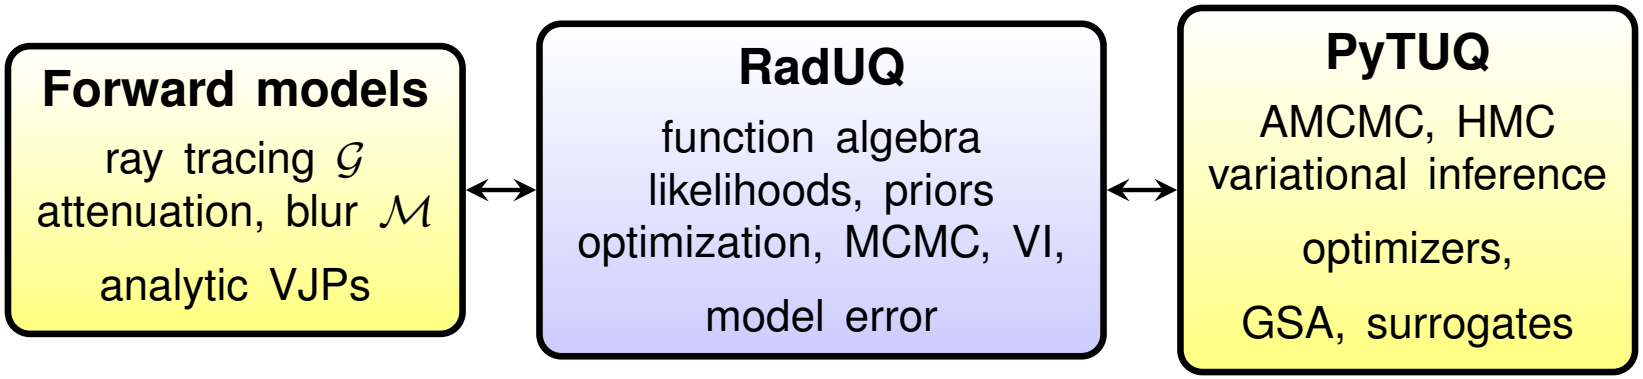
Model error

Kennedy, M. C., O'Hagan, A., "Bayesian calibration of computer models", *J. R. Stat. Soc. B*, 63(3), 2001.

Sargsyan, K., Najm, H. N., Ghanem, R., "On the statistical calibration of physical models", *Int. J. Chem. Kinet.*, 47(4), 2015.

Sargsyan, K., Huan, X., Najm, H. N., "Embedded model error representation for Bayesian model calibration", *Int. J. Uncertain. Quantif.*, 9(4), 2019.

Software: RadUQ Bridges Forward Models and UQ



- RadUQ: radiography UQ library on top of PyTUQ
  - forward model as a composable function $\mathcal{F} = \mathcal{M} \circ \mathcal{G} \circ \mathcal{S}$, gradients through the chain
  - weighted Gaussian and Poisson likelihoods; priors with the softplus Jacobian
  - one interface, many solvers: multistart optimization, adaptive MCMC, VI, randomized maximum likelihood, embedded model error
- PyTUQ, the Python offspring of UQTk
  `https://github.com/sandialabs/pytuq`
- A new forward model or a new UQ method plugs in without touching the other side

Geometric Map $\mathcal{G}$: Ray Tracing Through the Shells

- Ray to pixel (i, j) in the cylinder frame (axis $\hat{y}$): $\mathbf{P}(t) = \mathbf{S} + t\,\mathbf{v},\ t \in [0, 1]$, $\mathbf{S} = (x_s, -c_y, -c_z),\ \mathbf{v} = (x_d - x_s,\, y_{ij},\, z_{ij})$
- Wall of radius R: roots $t^{\pm}_{\text{cyl}}$ of $At^2 + Bt + C = 0$ ($\perp$: normal to the axis)

$$A = \|\mathbf{v}_\perp\|^2, \quad B = 2\,\mathbf{S}_\perp \cdot \mathbf{v}_\perp, \quad C = \|\mathbf{S}_\perp\|^2 - R^2, \quad D = B^2 - 4AC$$

- Caps at $y = \pm H/2$: $t^{\pm}_{\text{cap}} = (\pm H/2 + c_y)/y_{ij}$; clamp to the segment,

$$t_{\text{in}} = \max\big(t^-_{\text{cyl}},\, \min t_{\text{cap}},\, 0\big), \qquad t_{\text{out}} = \min\big(t^+_{\text{cyl}},\, \max t_{\text{cap}},\, 1\big)$$

- Chord, then onion peeling over the shells (d_m: mass densities)

$$L(R) = \|\mathbf{v}\|\,\big(t_{\text{out}} - t_{\text{in}}\big)_+ \ \text{ if } D > 0, \qquad \rho_m = d_m\big[L(r_{m+1}) - L(r_m)\big]$$

- Unclipped, $L = \|\mathbf{v}\|\sqrt{D}/A$: cusps at tangency ($D = 0$) $\Rightarrow$ local minima

Sampling Makes the Point Estimate Trustworthy

- Blind L-BFGS-B multistart (100–400 restarts) misses the best optimum in 95% of 80 runs
  - median shortfall 1.5 log-posterior units, worst 16.8
  - largest at the highest information: 10.2 at 120×120 vs. 2.0 at 30×30
- Not a smooth narrow valley, but genuine local minima. Chord through a cylinder:

$$L(R) \;\propto\; \sqrt{(R^2 - b^2)_+}\,, \qquad \frac{\partial L}{\partial R} \to \infty \;\text{ as } R \downarrow b$$

  - a cusp wherever a boundary crosses a ray's tangent point
  - curvature undefined $\Rightarrow$ (quasi-)Newton does not help; gradient-free search occasionally escapes
- Restarting from chain samples: optimum $=$ MAP in every production run

Sampling is usually presented as error bars added *after* the optimum. Here it is what makes the optimum trustworthy in the first place.

A Flat Prior Is Not Uninformative

- Expert knowledge lives in $\boldsymbol{\theta}$ (cm), not in $\boldsymbol{\xi}$: uniform $\pi_{\boldsymbol{\theta}}$ on a physical box $\mathcal{B}$
- Transported to $\boldsymbol{\xi}$, the softplus Jacobian appears:

$$\pi_{\boldsymbol{\xi}}(\boldsymbol{\xi}) \propto \prod_{k=1}^{4} \sigma(\xi_k), \qquad \sigma(u) = \frac{1}{1+e^{-u}}$$

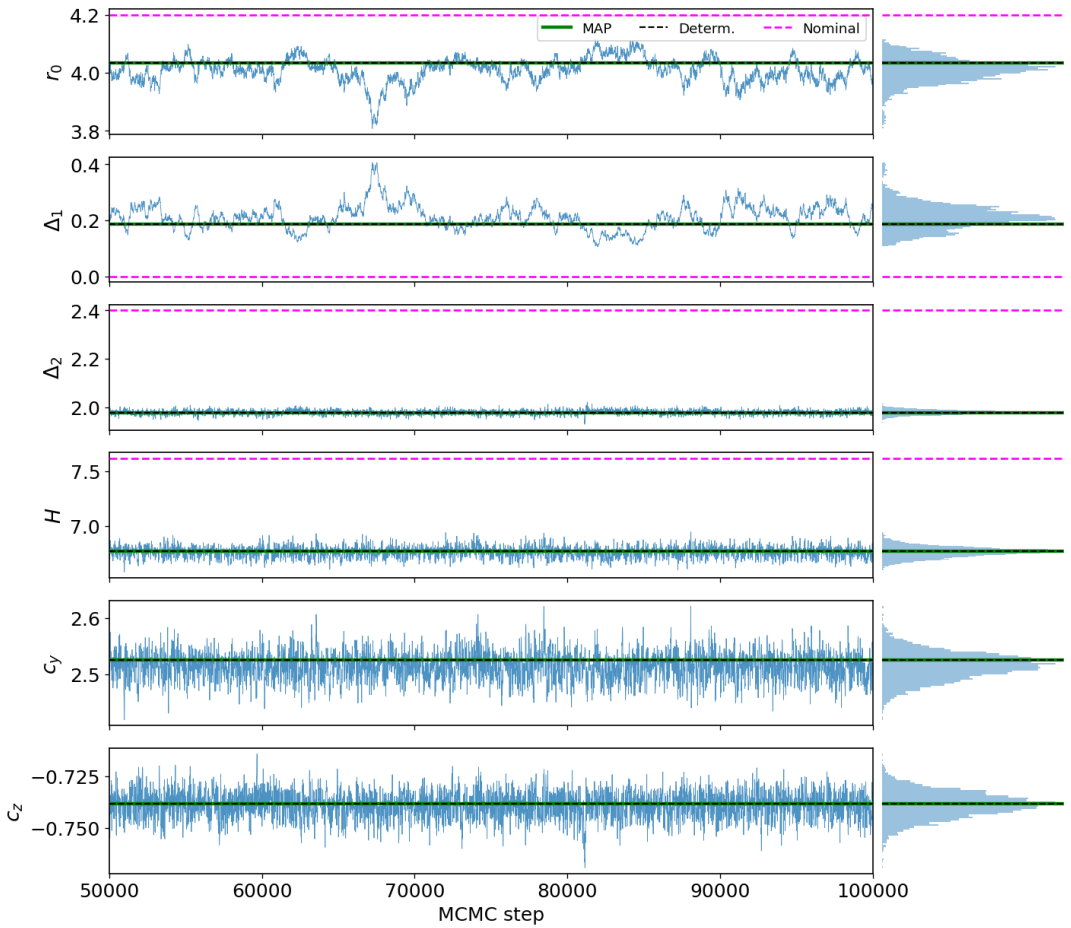
- A flat prior in $\boldsymbol{\xi}$ instead implies

$$\pi_{\boldsymbol{\theta}}(\theta_k) \;\propto\; \frac{1}{1-e^{-\theta_k}} \;\sim\; \frac{1}{\theta_k} \quad \text{as } \theta_k \to 0$$

  - piles prior mass at zero thickness
- It matters exactly where the data are weak: along the unidentified (r_0, Δ_1) ridge

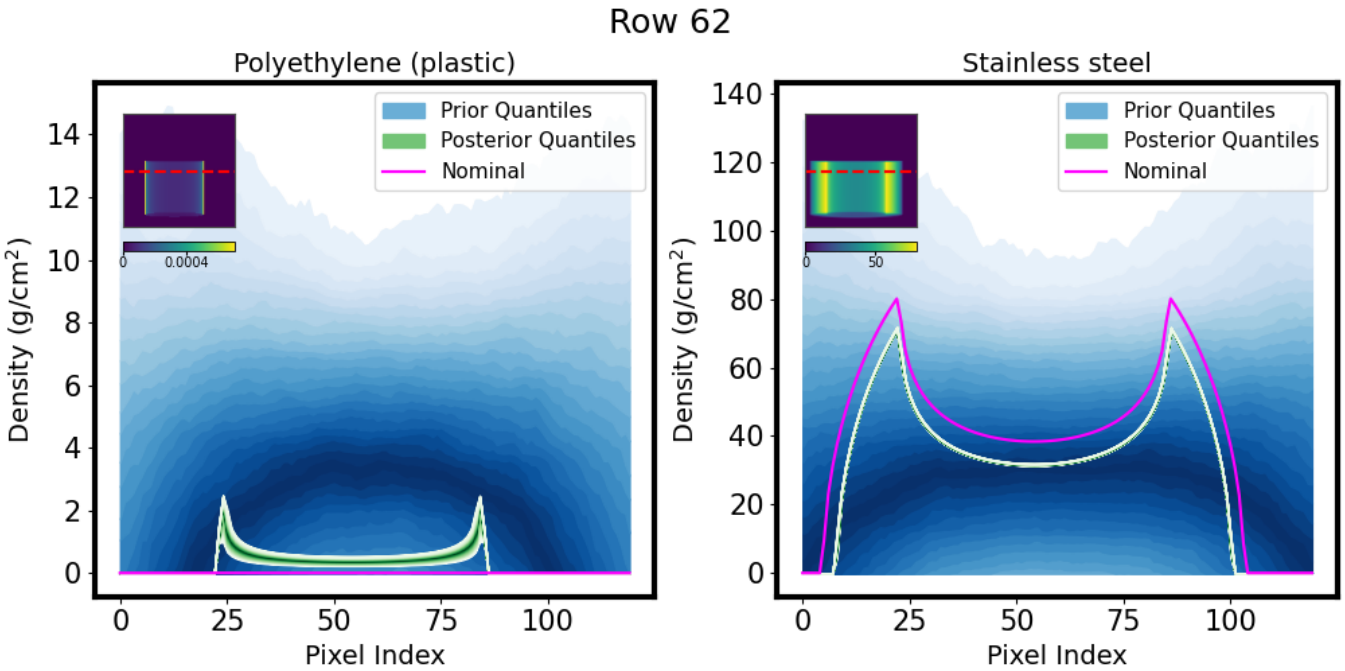
	flat in $\boldsymbol{\xi}$	uniform in $\boldsymbol{\theta}$	shift in r_1
two X-ray, 120×120	$\Delta_1 = 0.19$	$\Delta_1 = 0.20$	< 0.02
all three, 30×30	$\Delta_1 = 0.10$	$\Delta_1 = 0.62$	< 0.02

MCMC Traces, Two X-Ray Beams at 120×120



- Retained samples: second half of the chain, every tenth step
- Stationary over the retained samples; slower, larger excursions in r_0 and Δ_1 along the ridge

Prior to Posterior

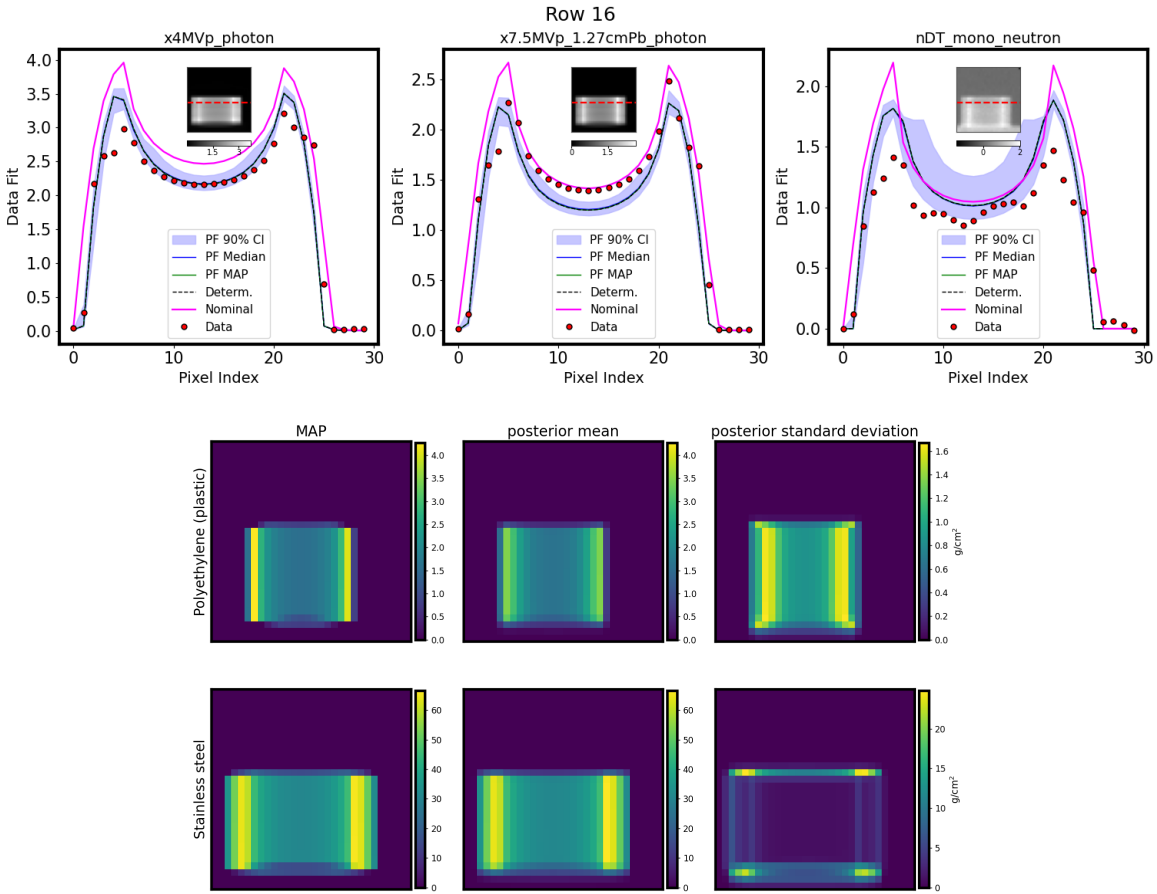


- Prior: up to 15 g/cm^2 polyethylene, 140 g/cm^2 steel along this row
- Contraction of the posterior standard deviation relative to the prior: a measure of 'learnability'

parameter	contraction
$r_1 = r_0 + \Delta_1$	$384\times$
Δ_2	$154\times$
H	$118\times$
r_0	$69\times$
Δ_1	$33\times$

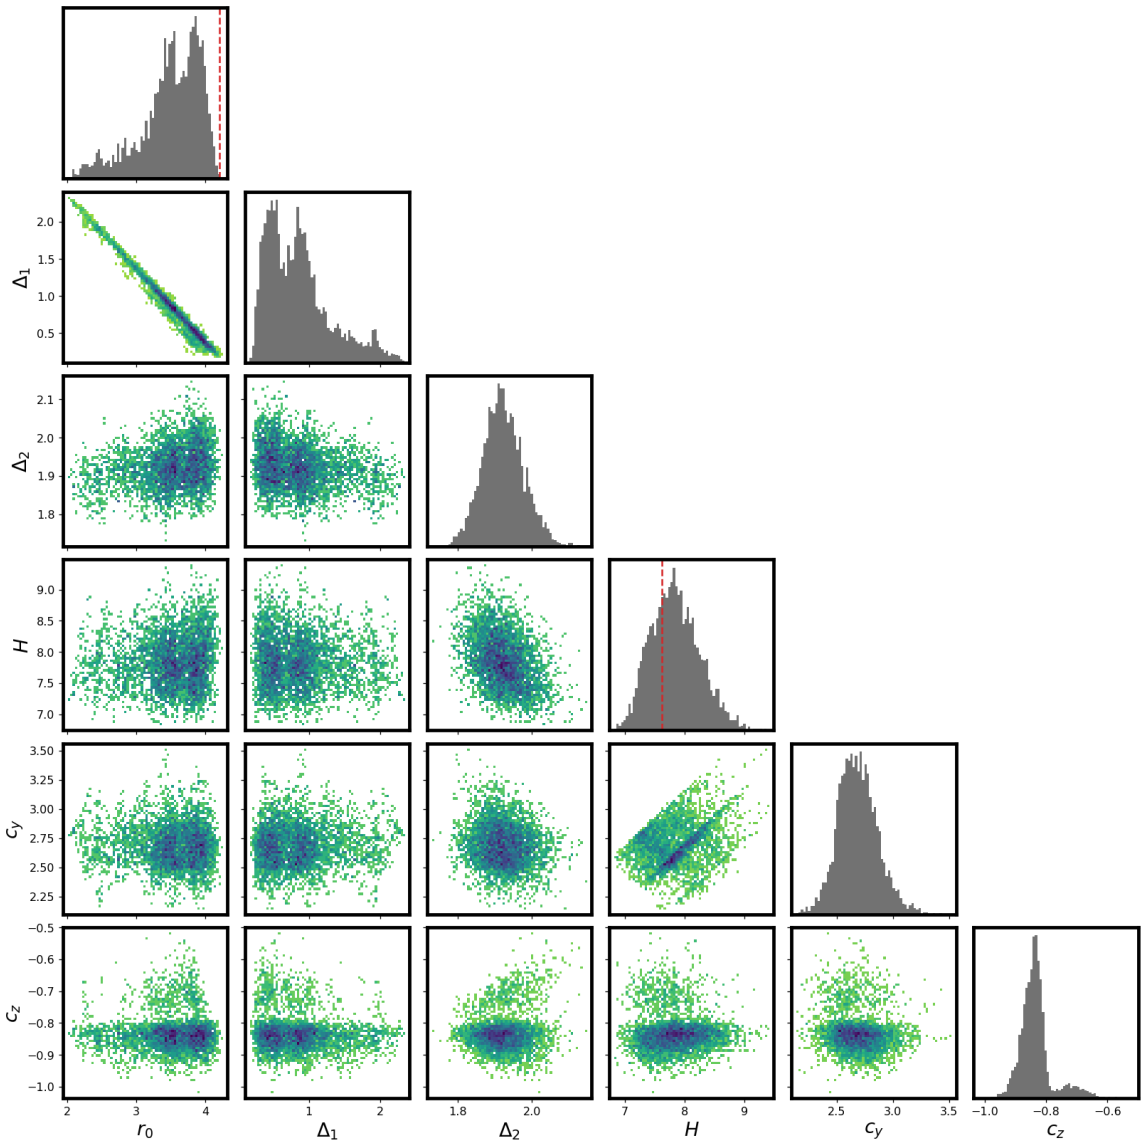
The data locate the inner radius of the steel, not its split into air and polyethylene

All Three Beams at 30×30



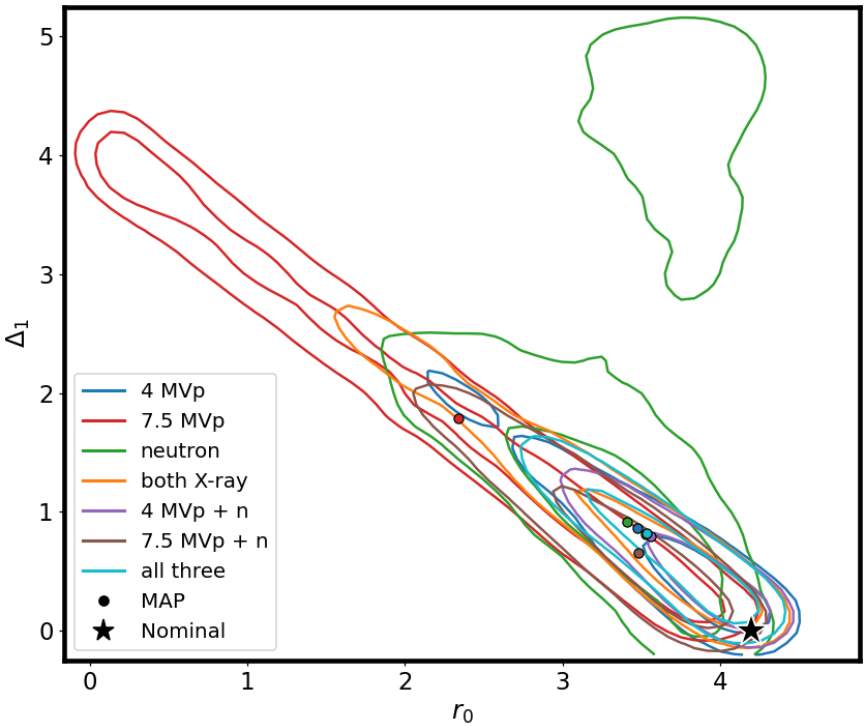
- Bands widen with $11\times$ fewer data: SS-316 $1\% \rightarrow 18\%$ of median, polyethylene $50\% \rightarrow 80\%$
- The neutron constrains the shell interior least: RMSE 0.44 vs. 0.41 (4 MVp) and 0.23 (7.5 MVp)

Joint Posterior, All Three Beams at 30×30



- (r_0, Δ_1) ridge persists at -0.99: not a resolution artifact
- (H, c_y) coupling drops $+0.81 \to +0.28$
- (Δ_2, H) emerges at -0.38
- Multimodal Δ_1 and c_z marginals

Neutron Data (Steel)



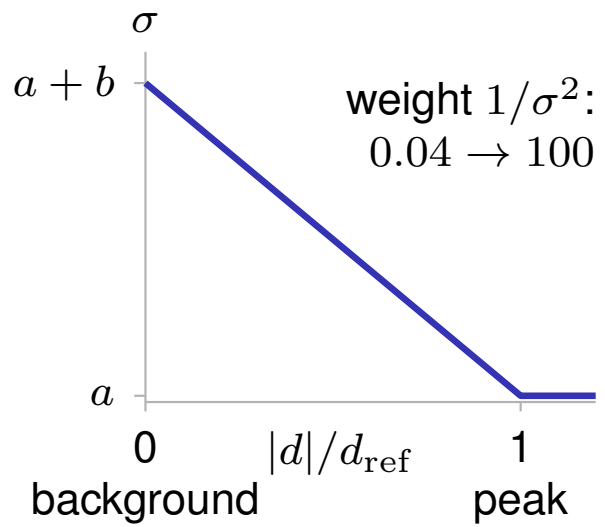
seven inspection subsets at 30×30; 50 and 90% credible levels

- All seven subsets lie on the same (r_0, Δ_1) ridge: more beams shorten it, but do not rotate it
- 7.5 MVp alone stretches to the all-poly end, $r_0 \to 0$
- Neutron alone is the broadest, with a detached island: $r_1 \in [3.99,\ 8.41]$ cm
- X-rays carry the radii: the neutron narrows r_1 by only about a fifth (for polyethylene it halves it)
- Height (not shown) rises with the neutron, which also sees the pedestal

Weighting Matters, per Beam

$$\sigma_k(i,j) = a + b\Big(1 - \min\Big(\frac{|d_k(i,j)|}{d_{\mathrm{ref}}}, 1\Big)\Big), \qquad a = 0.1,\ b = 5$$

		r_1 (90% CI), nominal 4.200	
beam	res.	unweighted	weighted
4 MVp	30^2	$4.788\ [4.557, 5.012]$	$4.320\ [4.124, 4.415]$
	120^2	$4.941\ [4.901, 4.960]$	$4.210\ [4.194, 4.272]$
7.5 MVp	30^2	$4.249\ [4.044, 4.484]$	$4.123\ [3.855, 4.384]$
	120^2	$4.240\ [4.188, 4.297]$	$4.101\ [4.062, 4.166]$
both	30^2	$4.554\ [4.434, 4.707]$	$4.312\ [4.110, 4.394]$
	120^2	$4.668\ [4.615, 4.686]$	$4.219\ [4.205, 4.234]$

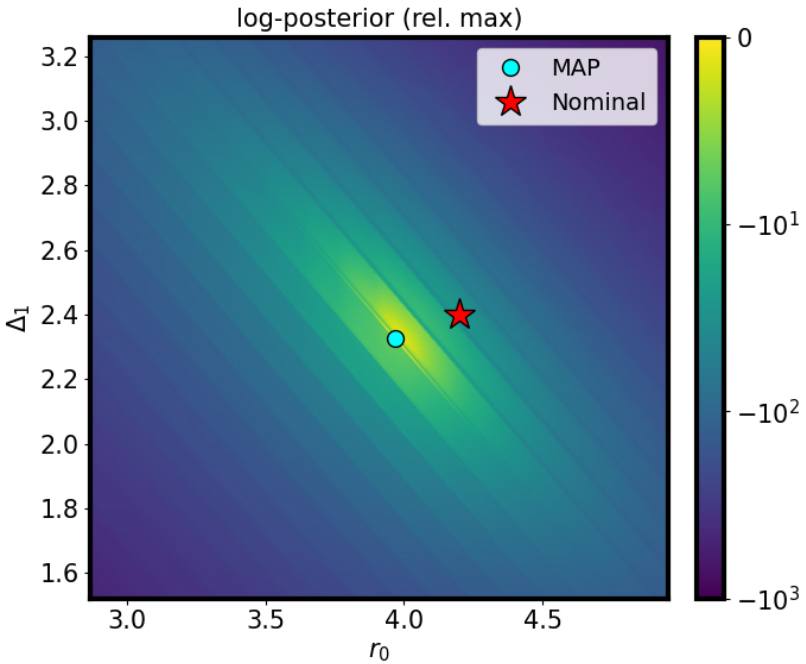


- 4 MVp, unweighted (constant σ): more pixels give a *more confident wrong answer*; weighted, it approaches nominal
- 7.5 MVp, less attenuated by the steel: both stay within 0.13 cm of nominal
- The pair inherits the 4 MVp failure $\Rightarrow$ choose the weighting per target, even per beam

Choosing the Weighting (a, b)

- Chosen empirically on the two X-ray beams at 120×120, among weight ratios from 9 to $10{,}201$: $a = 0.1$, $b = 5$ (ratio 2601)
- Stronger weighting, $a = 0.05$ and 0.02 (ratios up to $63{,}001$), over all configurations
  - no consistent improvement
  - a harder deterministic problem
- a and b encode an expected *contrast*, not a noise level: $\sigma = a$ at the brightest pixels, so what matters is a relative to the peak signal
- Future: hierarchical treatment of (a, b), with σ tied to $\mathcal{F}(\boldsymbol{\xi})$ so the likelihood stays a proper density

Second Target: The Polyethylene Annulus

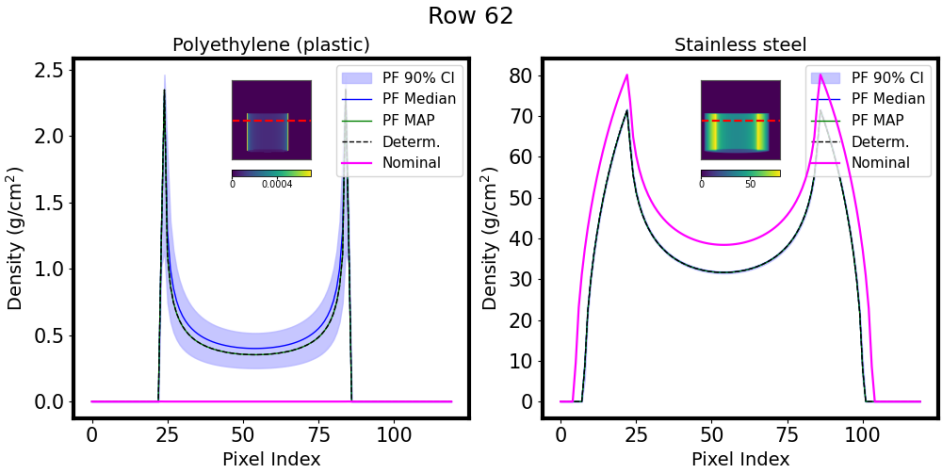


exact log-posterior, a single maximum

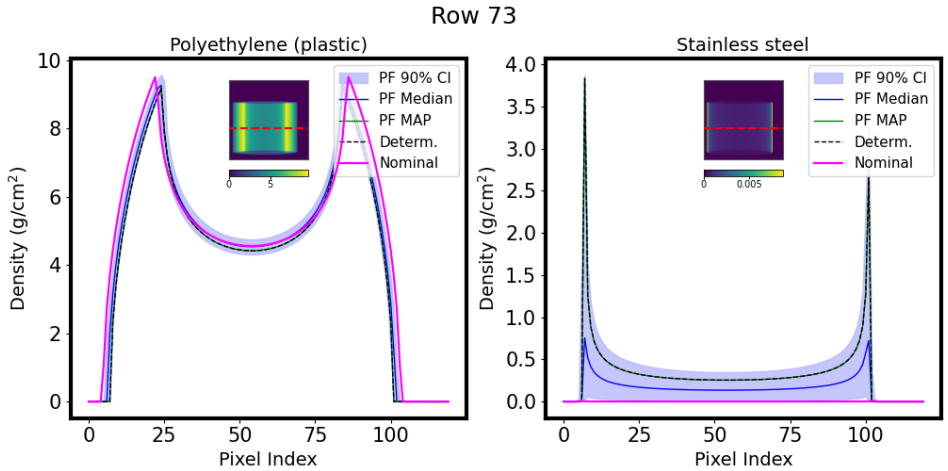
- Far more transparent: 4 MVp peak 0.82 vs. 3.53 for steel; the neutron image is now the brightest
- Degeneracy gone: κ median $11.4 \to 2.2$, range 1.8–2.8
- It appears when a weakly attenuating interface is masked by a strongly attenuating one
- Weighting does not carry over: $a = 0.1$ is 3% of the steel peak but 12% of the poly peak $\Rightarrow$ unweighted likelihood is better here

Roles Reverse; Confusion Runs One Way

steel target: polyethylene (left), SS-316 (right)



poly target: polyethylene (left), SS-316 (right)



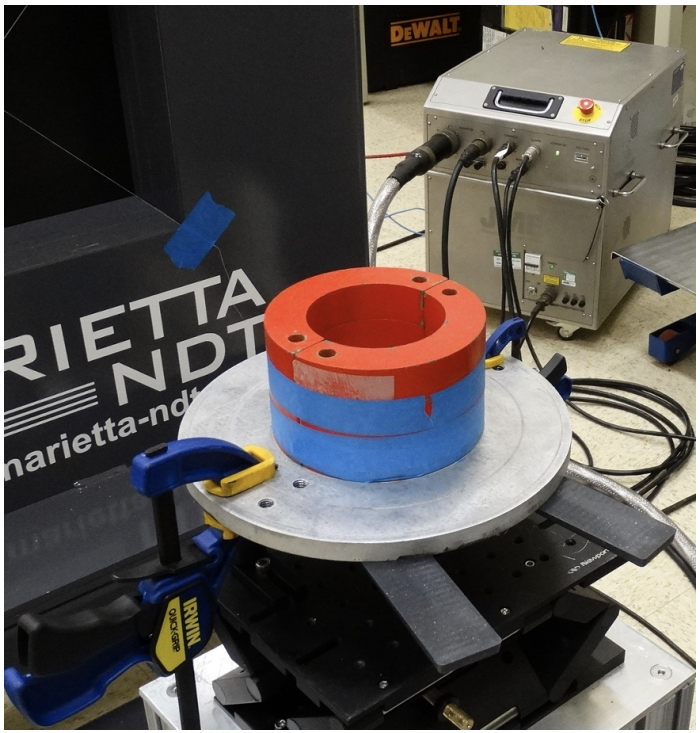
- The dominant attenuator is determined to about 1% (steel) or a few percent (poly)
- Steel target: phantom polyethylene, $\Delta_1 = 0.06$–1.8 cm where there is none
- Poly target: no phantom steel, $\Delta_2 = 0.007$–0.15 cm against zero
- A thin dense layer can imitate a thick light one; the converse needs more light material than the outer radius allows

Ruling out a dense material is better supported by these data than ruling out a light one

Poly, Weighted Likelihood Hits the Box

- With neutron data the weighted likelihood returns $\Delta_1 \approx 4.7$ cm against a nominal 2.4
- Posterior mass within 2% of a bound on the four shape parameters
  - unweighted: 0.0% in all seven configurations
  - weighted: above 5% in six of seven, up to 29% (all three) and 32% (neutron alone)
- Values set by the prior box, not by the data $\Rightarrow$ not quoted as results
- Two further schemes, mildly favoring the signal $(36\times)$ or the background, are not usable either

What the Posterior Reveals About the Model

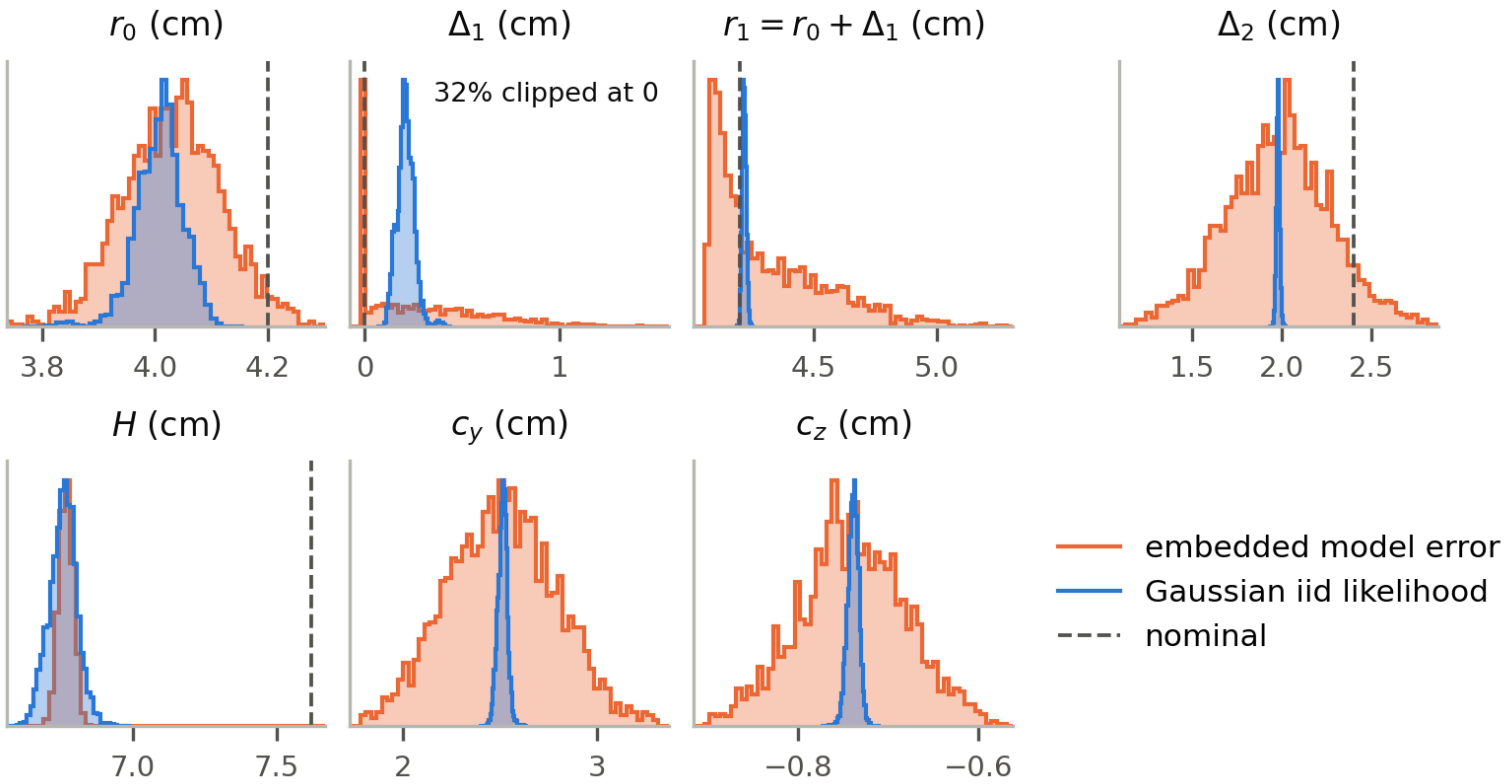


mounting holes, block-wise construction, plate

- Nominal dimensions are a *reference*, not a grade
- The object as measured is not the object as modeled
  - pedestal seen by the neutron $\Rightarrow$ taller cylinder
  - empty mounting holes, block-wise assembly
  - simplified spectrum, scatter and blur models
- H and Δ_2 miss nominal by far more than the credible intervals, and the intervals narrow faster than the discrepancy
- A point estimate cannot raise this question; a posterior can

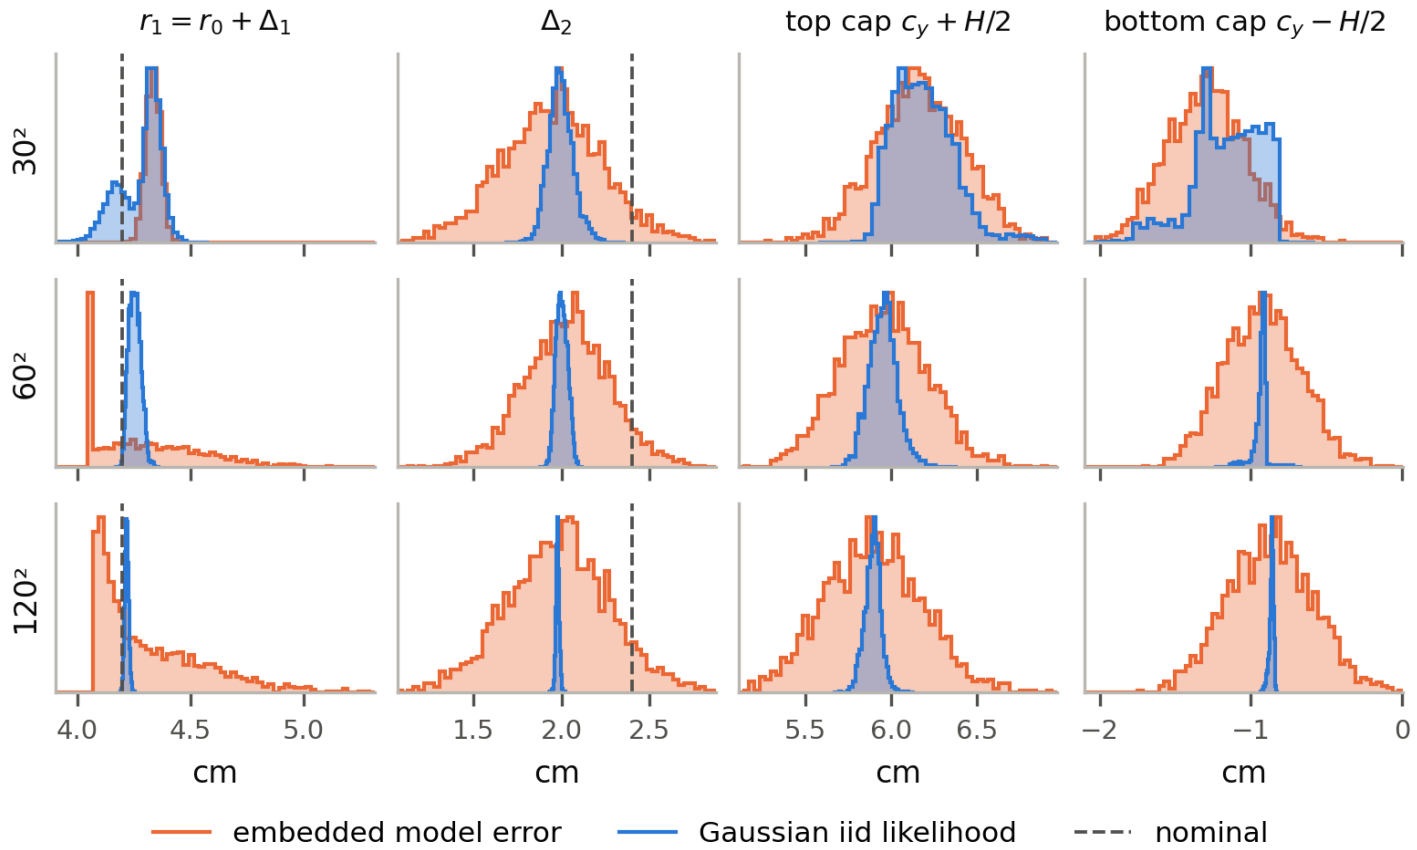
Next: model-error likelihoods that keep the intervals from shrinking below the model-error floor

Parameter PDFs with Model Error, 120^2



- Embedded: $\hat{\boldsymbol{\theta}} + \boldsymbol{\delta}$, whose spread is model error (curves scaled to unit peak)
- Δ_2 widens from 0.009 to 0.29 cm and covers the nominal 2.4; r_0, c_y, c_z widen too
- H stays narrow (the cap correction went into c_y); Δ_1 is clipped at zero, which skews r_1

Parameter PDFs Across Resolutions



- The iid posterior narrows with resolution; the embedded spread stays at 0.24–0.29 cm for Δ_2 and both caps
- Caps replace H and c_y, which trade the cap correction (H at 30^2, c_y above)
- At 30^2 the shared factor mostly moves the split (r_0, Δ_1), so r_1 stays narrow