Fifth Technical Meeting on Divertor Concepts
28 Oct 2025

Bayesian Experimental Design for Divertor Heat Loads

Michael Battye[a], Cyd Cowley[b], Dan Greenhouse[b], Ethan Hargrove[b]

[a] University of York
[b] digiLab

Motivation

We don't yet have a single, systematic method to design/position divertor diagnostics

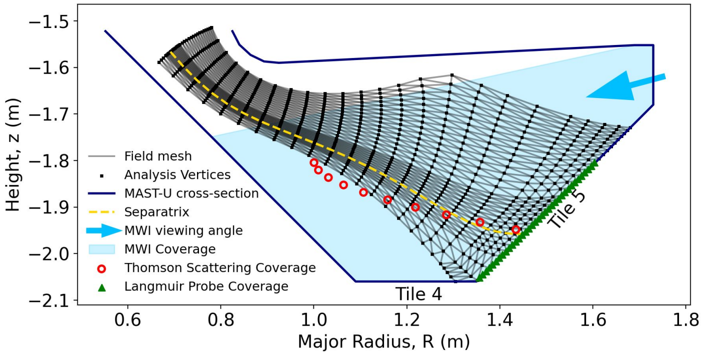

D Greenhouse et al 2025 Plasma Phys. Control. Fusion 67 035006

What is BED?

Bayesian Experimental Design (BED): choose experiments/sensors that *maximise expected information gain* about specific quantities of interest.

Design Goal

Given candidate sensor locations and expected plasma states, choose the fewest, most informative sensors for our QoIs (e.g., peak q, heat-flux width, tile hotspots).

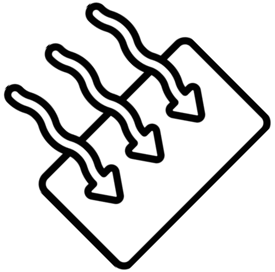

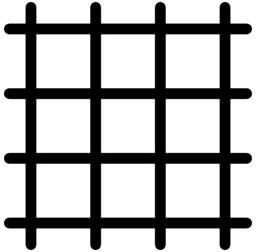

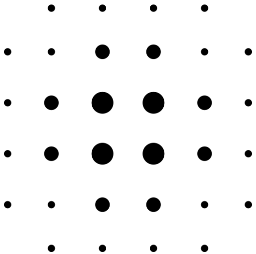

2D Toy Workflow

Data: **25 SOLPS heat-flux profiles → GP fit → 100+ sampled profiles** for **2D MAST-U slice**.

Thermal model: **Laplace** ($\nabla^2 T = 0$) solved by **Gauss–Seidel**; Neumann BC from $q = -k\partial T/\partial n$ at the PF surface and Dirichlet BC of T_{coolant}.

Purpose: *validate BED plumbing & scoring at low cost.*

What is a GP?

Gaussian Process (GP): a non-parametric regression giving mean + uncertainty over functions.

MAST-U State 1

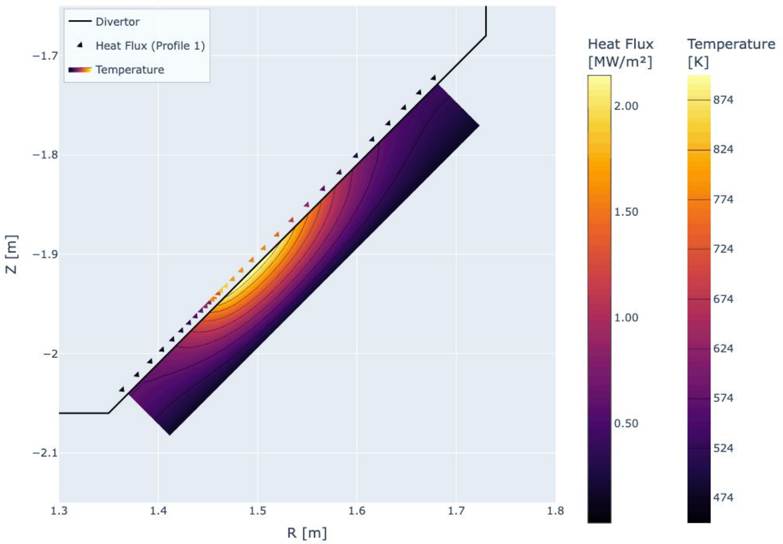

Results: 2D prototype

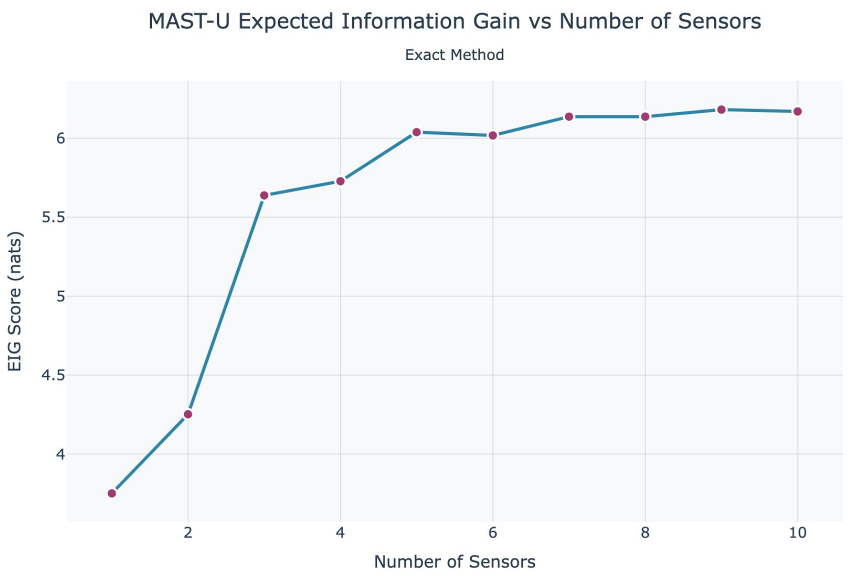

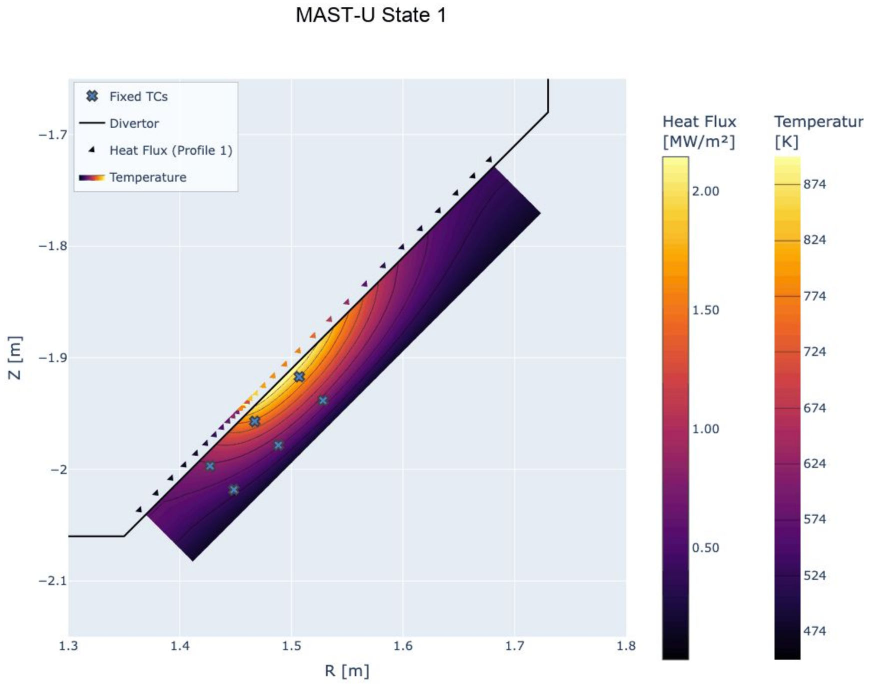

BED Workflow

Priors → Forward Model → Likelihood → Evidence → Posterior → $\sigma_g^{(\mathrm{post})}$

θ: upstream states → q profiles

$y(\theta, d)$: sensor temps

$p(y|\theta, d)$

$p(y|d)$

$p(\theta|y, d)$

probabilistic measure of misfit

distribution of upstream states

what we want after measuring

$$p(\theta \mid y, d) = \frac{p(y \mid \theta, d)\, p(\theta)}{p(y \mid d)}$$

normaliser

BED Workflow

Priors → Forward Model → Likelihood → Evidence → Posterior → $\sigma_g^{(\mathrm{post})}$

θ: upstream states → q profiles | $y(\theta, d)$: sensor temps | $p(y|\theta, d)$ | $p(y|d)$ | $p(\theta|y, d)$

In **BED**, we need to understand how the difference between the *evidence* and *likelihood* behaves across designs.
The **Expected Information Gain (EIG)** uses the *log of this evidence* to score designs:

$$\mathrm{EIG}(d) = \mathbb{E}_{y,\theta}[\log p(y \mid \theta, d) - \log p(y \mid d)]$$

expected information gain across all outcomes

how well the model predicts the data

how likely is the data overall

Uncertainty Engine: Sensor Designer

Inputs:

- **Sensor dataset** y: temps at candidate points
- **QoI dataset** $g(\theta)$: λ_q, S, peak q, tile hotspot temp, etc.
- **Uncertainties**: model/measurement noise, geometric tolerances
- **Scoring**: **EIG** via sensors (exact) and QoI-utility via **GP surrogates** (fast).

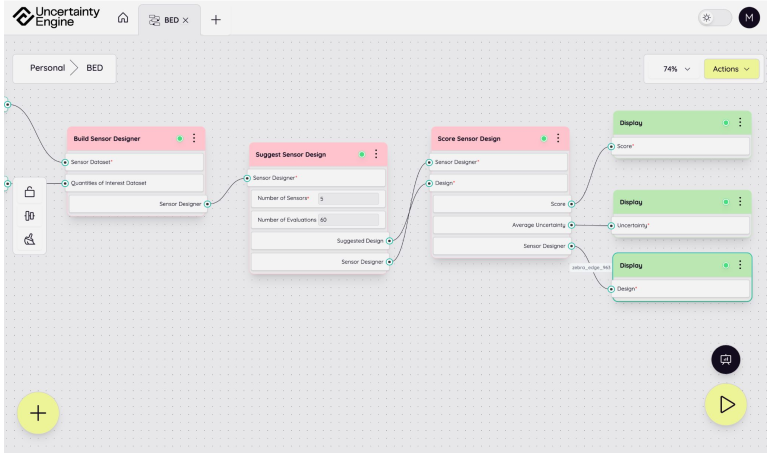

Optimisation: Genetic algorithm over sensor subsets.

3D Physics Chain

1. **HEAT** for optical approximation of divertor **heat-flux deposition** on detailed 3-D tiles.
2. **OpenFOAM** for **temperature field** under those loads.
3. **Surface grid "snap"** → candidate **thermocouple points** (temps become y).
4. **Bayesian experimental design** → suggested sensor sets
5. Select **best sensors** → convert to useful formats & visualise

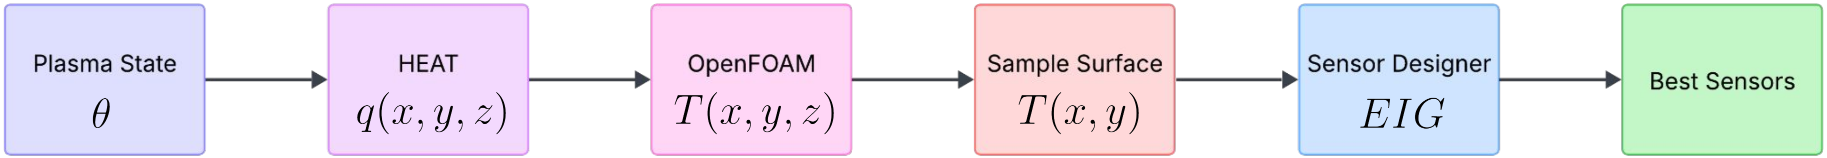

Looby T et al. 2022 *Fusion Science & Technology* 78 10-27

Results: 3D Pilot

1 2 3 4 5 6 7 8 9

Inputs:

NSTX-U Divertor Geometry [IBDH region]

$$P_{\mathrm{SOL}} = 30MW$$

Power radiated in SOL = 30%

$$\lambda_{q,\mathrm{CN}}\,[\mathrm{mm}] \in [2.0,\ 20.0]_{\Delta 0.5}$$

$$S\,[\mathrm{mm}] \in [0.0,\ 5.0]_{\Delta 0.2}$$

No. of Evals = 333

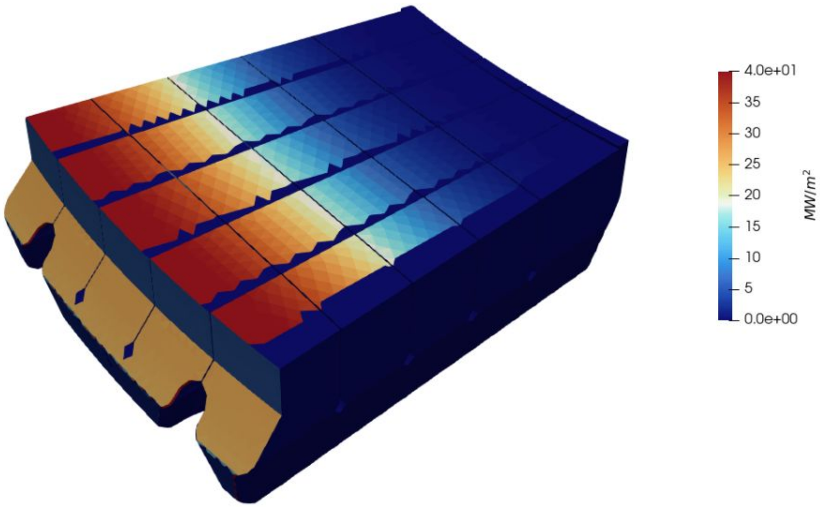

Goal: To mimic future power plant conditions to optimise surface thermocouple placement.

Results: 3D Pilot

1 2 3 4 5 6 7 8 9

$$\langle EIG \rangle = 12.7 \pm 1.3 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 12.9\% \qquad \frac{\sigma_S}{\langle S \rangle} = 56.5\%$$

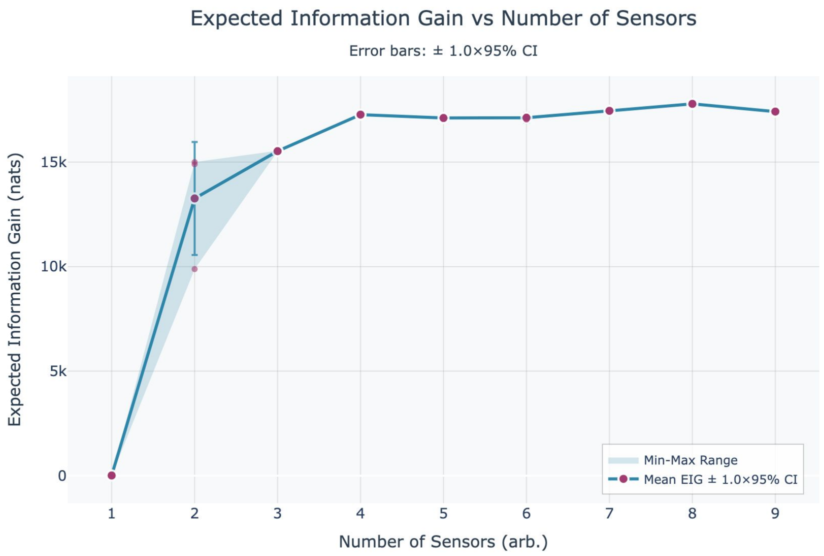

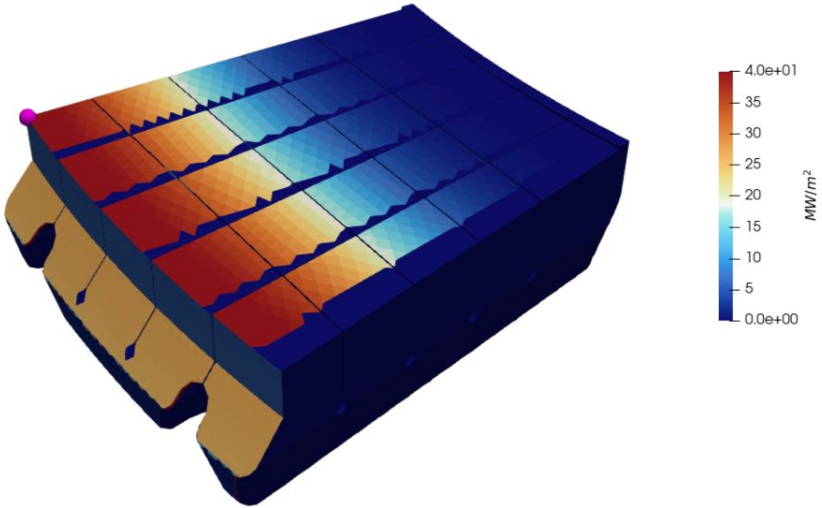

Results: 3D Pilot

1 2 3 4 5 6 7 8 9

$$\langle EIG \rangle = 13\,000 \pm 2\,000 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 1.5\% \qquad \frac{\sigma_S}{\langle S \rangle} = 5.6\%$$

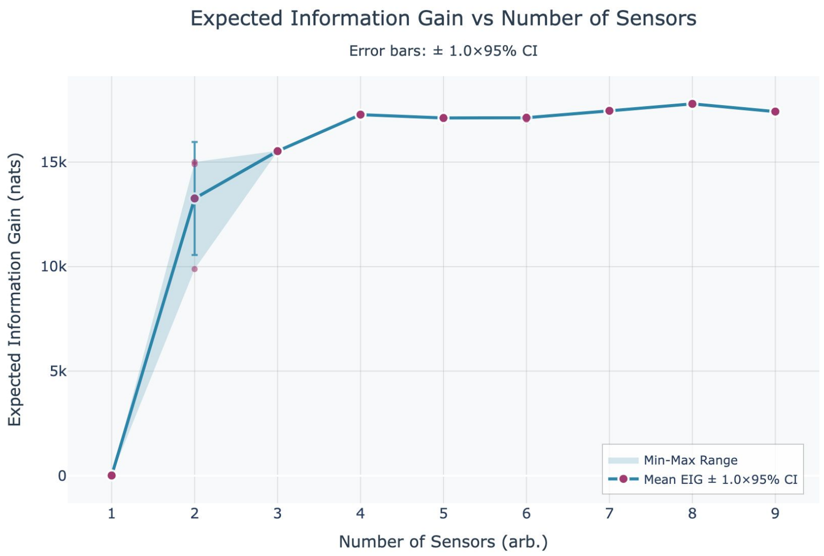

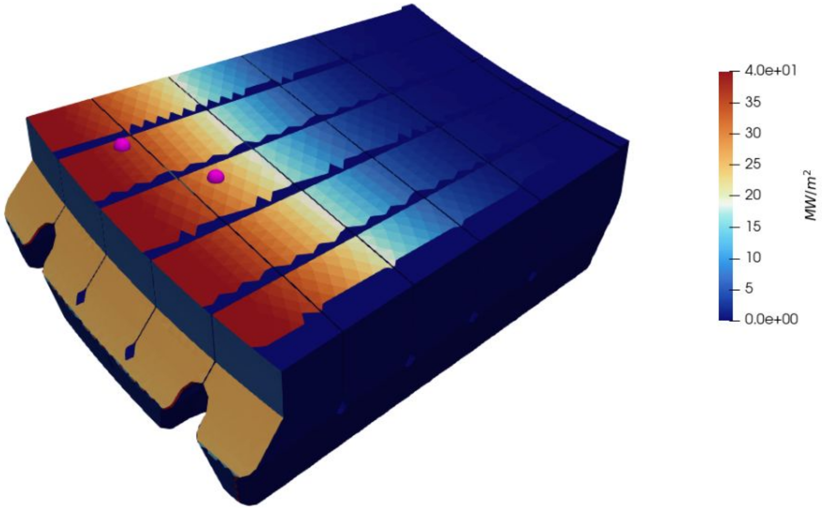

Results: 3D Pilot

1 2 **3** 4 5 6 7 8 9

$\langle EIG \rangle = 15\,520 \pm 10$ nats $\quad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 1.0\% \quad \frac{\sigma_S}{\langle S \rangle} = 0.9\%$

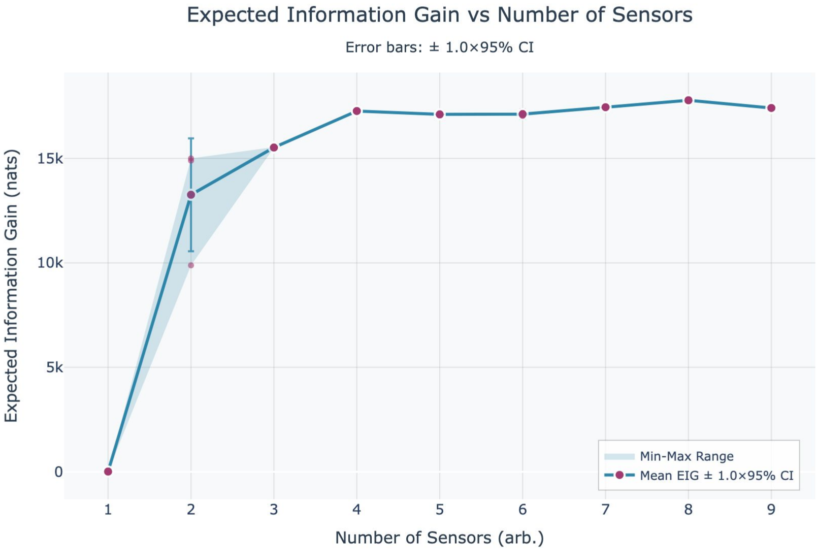

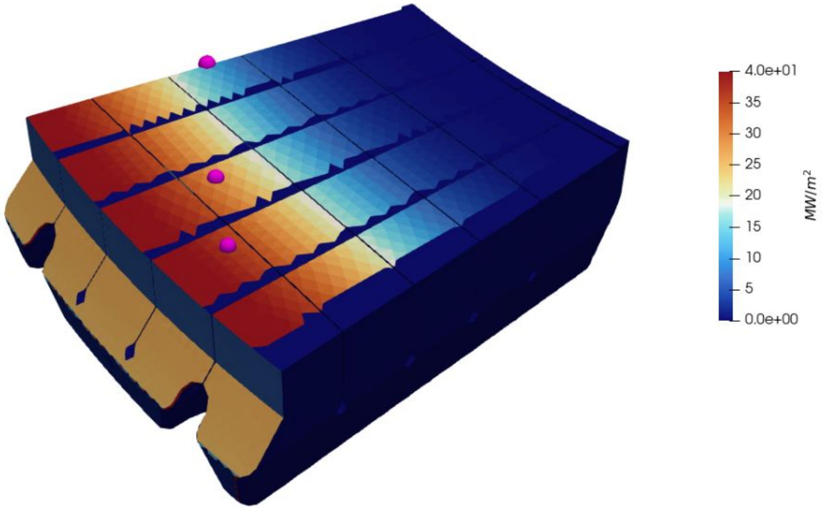

Results: 3D Pilot

1 2 3 **4** 5 6 7 8 9

$$\langle EIG \rangle = 17\,267 \pm 2 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 0.5\% \qquad \frac{\sigma_S}{\langle S \rangle} = 0.8\%$$

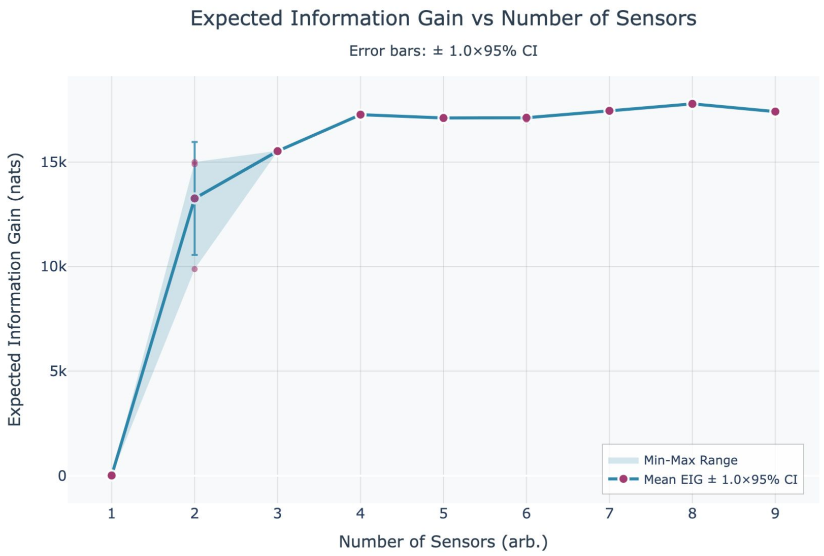

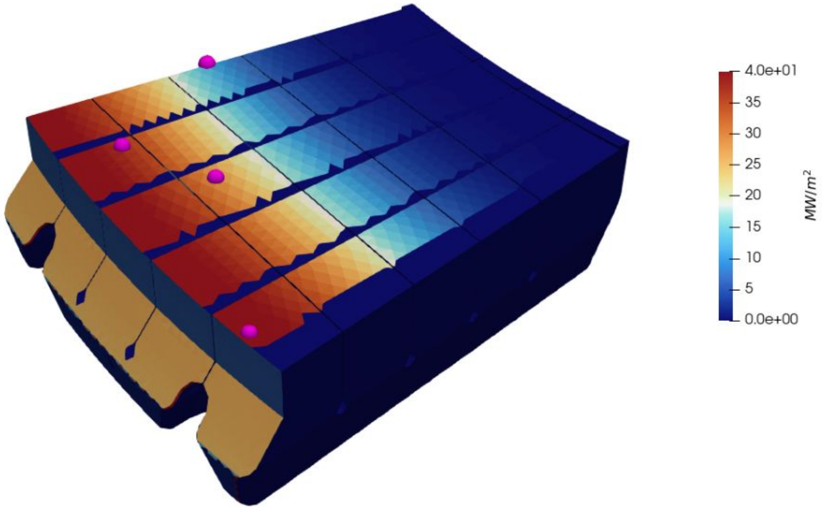

Results: 3D Pilot

1 2 3 4 **5** 6 7 8 9

$\langle EIG \rangle = 17\,110 \pm 20$ nats $\quad \dfrac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 0.5\%$ $\quad \dfrac{\sigma_S}{\langle S \rangle} = 0.8\%$

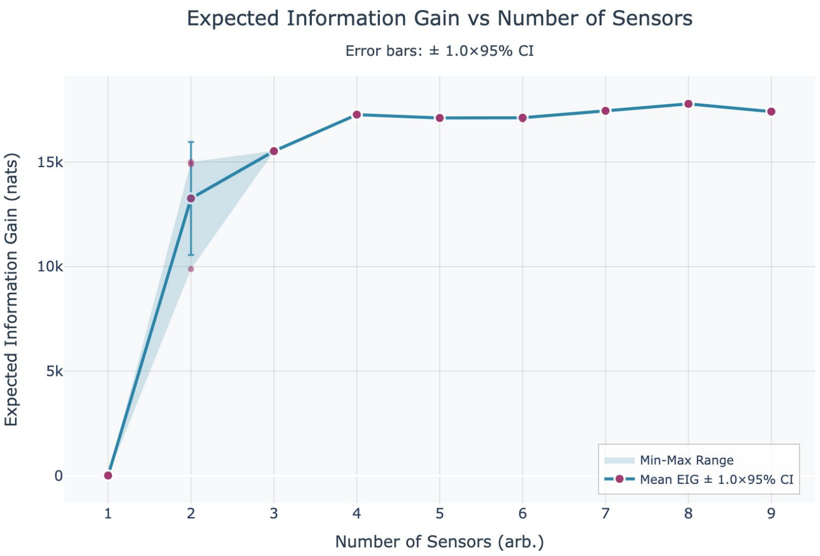

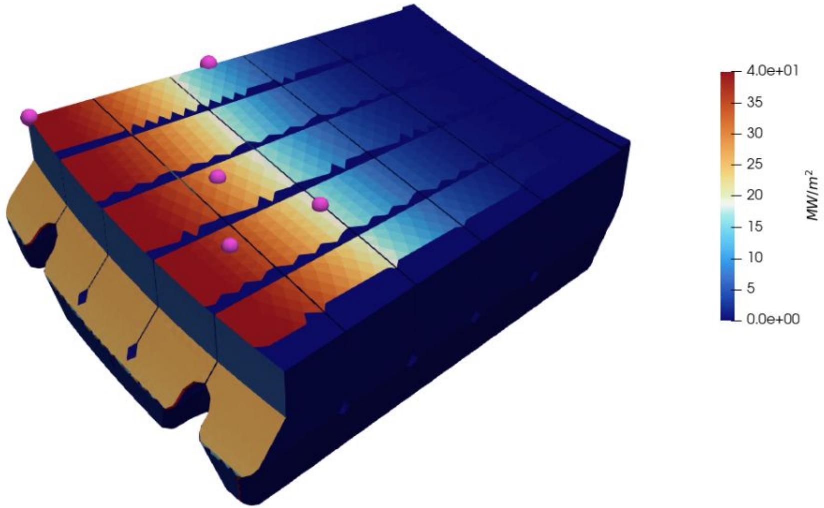

Results: 3D Pilot

1 2 3 4 5 **6** 7 8 9

$$\langle EIG \rangle = 17\,120 \pm 70 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 0.5\% \qquad \frac{\sigma_S}{\langle S \rangle} = 0.8\%$$

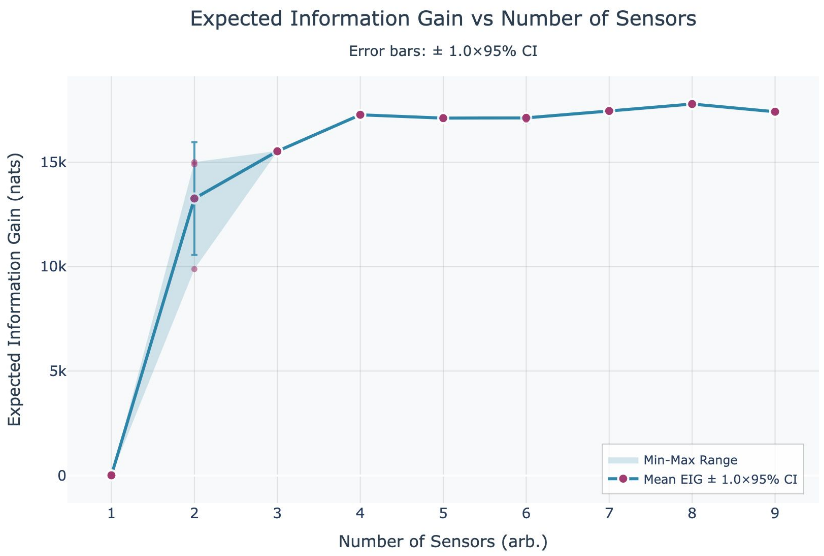

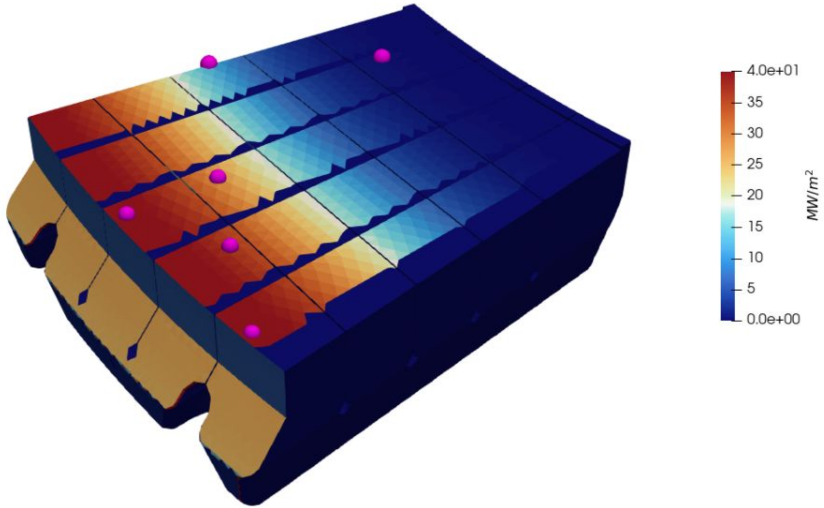

Results: 3D Pilot

1 2 3 4 5 6 7 8 9

$$\langle EIG \rangle = 17\,450 \pm 20 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 0.5\% \qquad \frac{\sigma_S}{\langle S \rangle} = 0.8\%$$

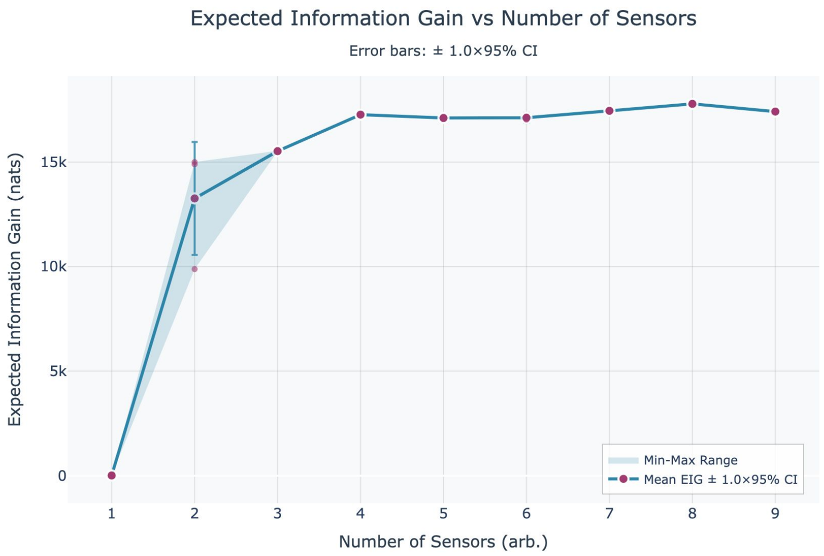

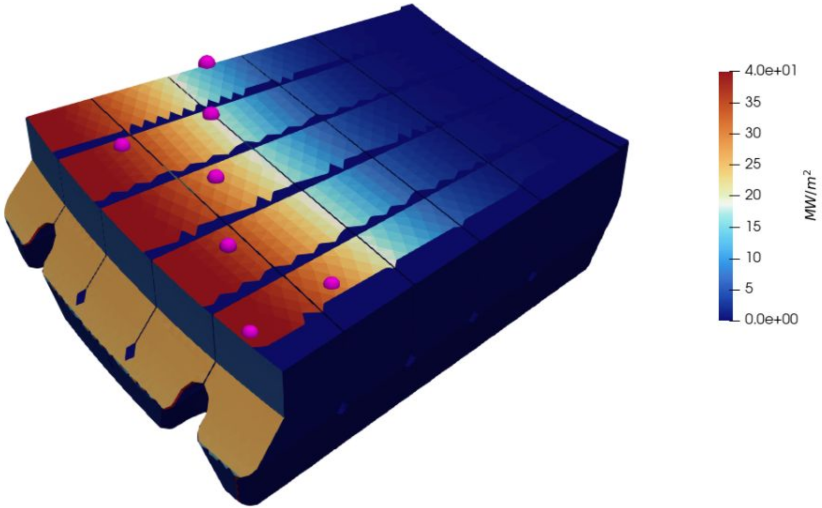

Results: 3D Pilot

1 2 3 4 5 6 7 **8** 9

$$\langle EIG \rangle = 17\,780 \pm 20 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 0.5\% \qquad \frac{\sigma_S}{\langle S \rangle} = 0.7\%$$

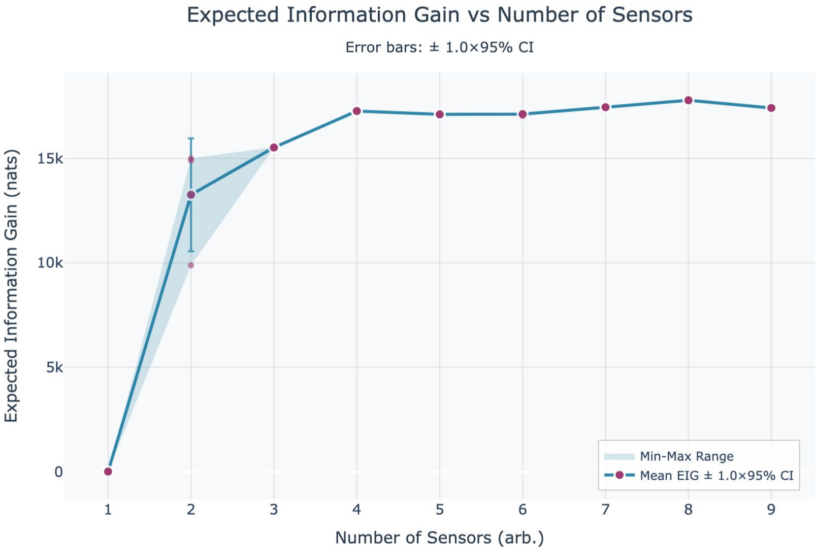

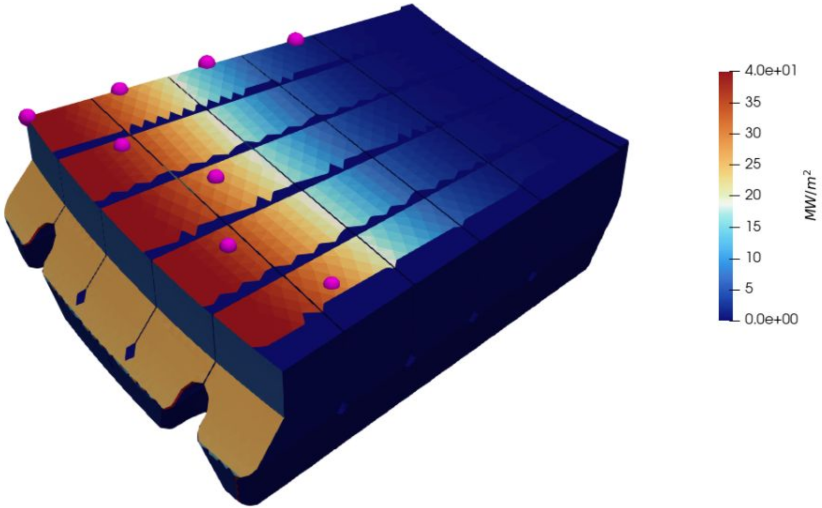

Results: 3D Pilot

1 2 3 4 5 6 7 8 9

$$\langle EIG \rangle = 17\,450 \pm 20 \text{ nats} \qquad \frac{\sigma_{\lambda_{q,CN}}}{\langle \lambda_{q,CN} \rangle} = 0.5\% \qquad \frac{\sigma_S}{\langle S \rangle} = 0.7\%$$

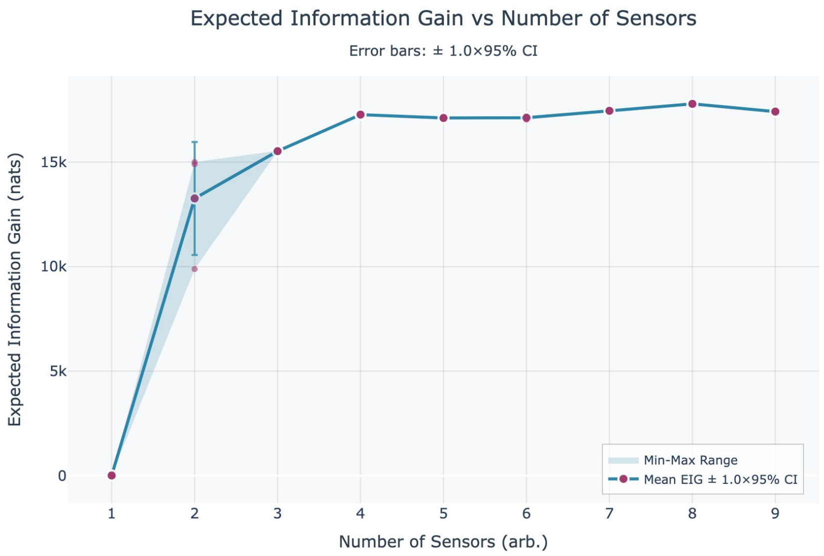

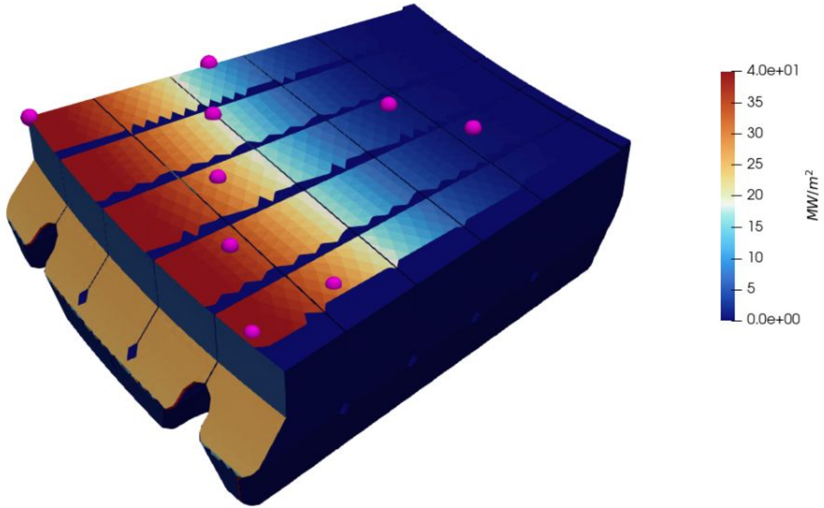

What next?

Multi-diagnostic fusion

(TC+IR+bolometry)

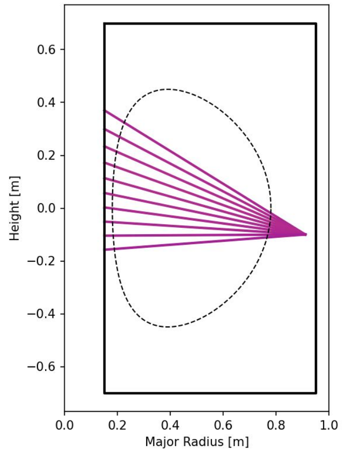

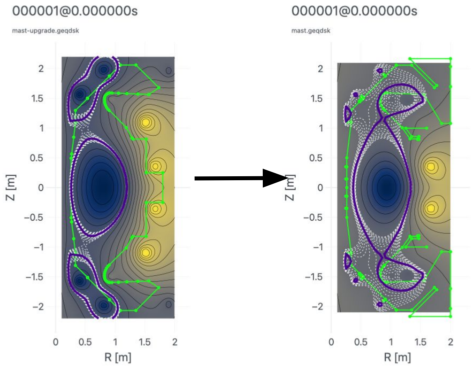

IR camera placement/FOV optimisation

Extensive MAST-U geometry

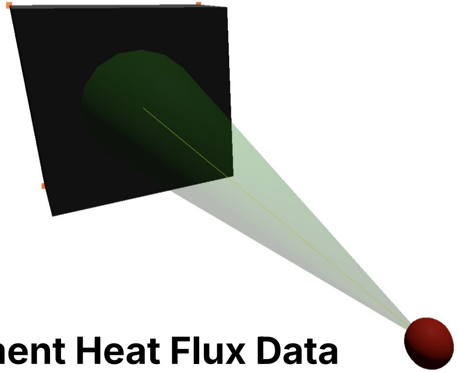

Add Toroidal Ripple

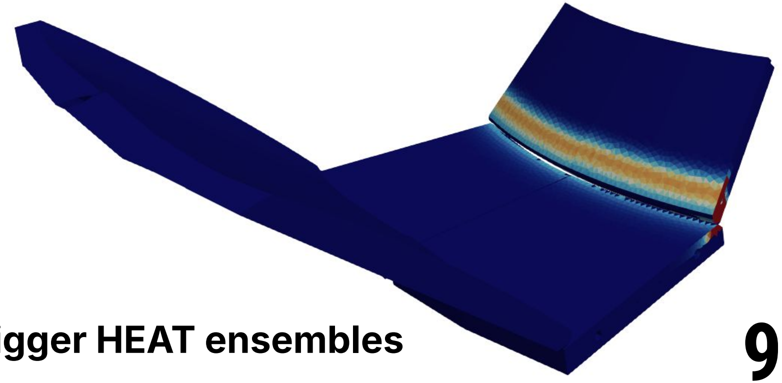

Detachment Heat Flux Data

Bigger HEAT ensembles

Thank you for your time.

Michael Battye
mike.battye@york.ac.uk

What is BED?

BED turns sensor placement into an optimisation problem: maximise EIG so the posterior over heat-flux parameters is as tight and decision-useful as possible.

1. State your goal (QoI)
   What do you ultimately care about? (e.g., heat flux width, Gaussian spreading, peak heat flux, or full heat flux profile.
2. Choose a design space
   Ξ**:** admissible **sensor layouts** ξ (locations, depth, count, wiring limits).
3. Specify a prior $p(\theta)$
   Physics/simulation-informed uncertainty over heat-flux parameters
4. Forward model $F(\theta, \xi)$
   Maps parameters → temperatures at sensor sites (FEM/FOAM or a GP surrogate).
5. Noise model
   $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$ (homoscedastic) or $\mathcal{N}(0, \Sigma_\epsilon)$ (correlated).
6. Likelihood $$p(y \mid \theta, \xi) = \mathcal{N}(F(\theta, \xi), \Sigma_\epsilon)$$
7. Predictive (marginal)
   $p(y|\xi)$: integrate out θ. If using a GP + PCA, this is approximately Gaussian with known μ,Σ.
8. $$\mathrm{EIG}(\xi) = \underbrace{\mathrm{E}_{p(\theta)\,p(y|\theta,\xi)}\big[\log p(y \mid \theta, \xi)\big]}_{\text{joint}} - \underbrace{\mathrm{E}_{p(y|\xi)}\big[\log p(y \mid \xi)\big]}_{\text{marginal}}$$
9. Estimator choice
   - Small n/simple F: nested MC
   - GP surrogate: compute entropies via **log-det** and **low-rank updates**; use small joint batches if needed.
10. Optimise ξ
    - **Discrete subsets**: GA/beam search with caching; proxy "exact" objective for large sets.
    - **Continuous placement**: gradient-based (SGD/Adam), possibly with variational bounds.
11. Pick a design & validate
    Simulate posterior (or run an end-to-end synthetic "reconstruction") to verify uncertainty reduction in QoIs.
12. Adaptive design

Equations

Forward Model

$$[y = F(\theta,\xi) + \varepsilon,\ \varepsilon \sim \mathcal{N}(0,\Sigma_\varepsilon)]$$

$$\underbrace{\theta \sim p(\theta)}_{\text{prior}} \Longrightarrow \underbrace{y|\theta,\xi \sim p(y|\theta,\xi)}_{\text{likelihood}}$$

$$\underbrace{p(y|\xi)}_{\text{marginal likelihood / evidence}} = \int p(y|\theta,\xi)\, p(\theta)\, d\theta$$

$$\underbrace{p(\theta|y,\xi)}_{\text{posterior}} = \frac{\underbrace{p(y|\theta,\xi)}_{\text{likelihood}}\ \underbrace{p(\theta)}_{\text{prior}}}{\underbrace{p(y|\xi)}_{\text{evidence}}}$$

Bayes' Rule

Gaussian entropy

$$H(\mathcal{N}(\mu,\Sigma)) = \tfrac{1}{2}\log\big((2\pi e)^n \det\Sigma\big)$$

EIG as mutual information (entropy drop)

$$[\mathrm{EIG}(\xi) = H(\theta) - \mathrm{E}_{\,y\sim p(y|\xi)}[\,H(\theta|y,\xi)\,] = H(y|\xi) - \mathrm{E}_{\,\theta\sim p(\theta)}[\,H(y|\theta,\xi)\,]]$$

EIG decomposition

$$\mathrm{EIG}(\xi) = \underbrace{\mathrm{E}_{p(\theta)p(y|\theta,\xi)}[\log p(y \mid \theta,\xi)]}_{\text{joint}} - \underbrace{\mathrm{E}_{p(y|\xi)}[\log p(y \mid \xi)]}_{\text{marginal}}$$

Monte Carlo estimator

$$\widehat{\mathrm{EIG}} = \frac{1}{N}\sum_{i=1}^{N}\left[\log p\big(y^{(i)} \mid \theta^{(i)},\xi\big) - \log\left(\frac{1}{M}\sum_{j=1}^{M} p\big(y^{(i)} \mid \tilde{\theta}^{(j)},\xi\big)\right)\right]$$

Dataset mean q_avg: 11.683
Dataset mean q_p95: 59.627

Fractional/relative mean avg_uncertainty_q_avg per number of sensors (relative to dataset mean q_avg value):
1 sensor(s): 0.163
2 sensor(s): 0.015
3 sensor(s): 0.006
4 sensor(s): 0.005
5 sensor(s): 0.005
6 sensor(s): 0.005
7 sensor(s): 0.005
8 sensor(s): 0.005
9 sensor(s): 0.005

Fractional/relative mean avg_uncertainty_q_p95 per number of sensors (relative to dataset mean q_p95 value):
1 sensor(s): 0.060
2 sensor(s): 0.016
3 sensor(s): 0.007
4 sensor(s): 0.007
5 sensor(s): 0.006
6 sensor(s): 0.006
7 sensor(s): 0.006
8 sensor(s): 0.007
9 sensor(s): 0.006