A Probabilistic Framework for Three-Dimensional Equilibrium Reconstruction

Y. Zhang[1], H. Wu[1], S. Kwak[2], K. Rahbarnia[2], G. Fuchert[2], K. J. Brunner[2], E. V. Hausten[2], G. Verdoolaege[1] and the W7-X team

[1]*Department of Applied Physics, Ghent University, 9000 Ghent, Belgium*

[2]*Max Planck Institute for Plasma Physics, D-17491, Greifswald, Germany*

A probabilistic framework for three-dimensional ideal-MHD equilibrium reconstruction is presented. The equilibrium is parameterised by flux-space Fourier coefficients and flux functions; Gaussian process (GP) priors encode smoothness, axis regularity, and gauge constraints analytically. The posterior distribution over these latent variables, conditioned on magnetic and kinetic diagnostic data, couples real-space observations to flux-space quantities through coordinate-invariant likelihood functions. Force balance enters as a virtual-observation (soft constraint) term, interpolating between a hard-equilibrium and a purely data-driven regime. The framework naturally propagates measurement and model uncertainty throughout the inferred equilibrium.

Introduction

Magnetic equilibrium is the macroscopic state of a magnetically confined plasma, governed by ideal-MHD force balance $\mathbf{J}\times\mathbf{B}=\nabla p$ together with Maxwell's equations. In three-dimensional (3-D) configurations—stellarators, perturbed tokamaks, or devices with significant bootstrap-current coupling—the equilibrium is intrinsically non-axisymmetric and the standard Grad–Shafranov reduction does not apply. Deterministic solvers such as VMEC [1] or DESC [2] compute single equilibrium configurations, but provide no direct statistical description of uncertainty or identifiability when equilibria are inferred from experimental data.

We reformulate 3-D equilibrium reconstruction as a Bayesian inverse problem, replacing single deterministic solutions with a full posterior over geometry and plasma profiles, inferred from heterogeneous diagnostics with calibrated uncertainty.

Three-Dimensional Equilibrium Representation

Flux coordinates and geometry We adopt the flux-coordinate system (ρ,θ,ζ), where $\rho=\sqrt{\Phi/\Phi_{\mathrm{LCFS}}}\in[0,1]$ is the normalized toroidal-flux label (LCFS: last closed flux surface) and θ, ζ are 2π-periodic poloidal and toroidal angles. The label ρ is linear in the minor radius near the magnetic axis, ensuring that all equilibrium quantities admit analytic (integer-power) near-axis expansions.

The plasma geometry is represented by $\mathbf{X}(\rho,\theta,\zeta) = (R, \phi = \zeta, Z)$. Under stellarator symmetry, the geometry and field-line label are expanded in real Fourier series,

$$R = \sum_{m,n} R_{mn}(\rho)\cos(m\theta - n\zeta), \qquad Z = \sum_{m,n} Z_{mn}(\rho)\sin(m\theta - n\zeta), \tag{1}$$

and analogously for $\lambda(\rho,\theta,\zeta)$ (sine series). The radial coefficient functions $\{R_{mn}, Z_{mn}, \lambda_{mn}\}$ constitute the geometric degrees of freedom.

Magnetic field and current The equilibrium field takes the Clebsch form

$$\mathbf{B} = \nabla\zeta \times \nabla\Phi + \nabla\Psi(\rho) \times \nabla\theta^* = B^\theta\, \mathbf{e}_\theta + B^\zeta\, \mathbf{e}_\zeta, \tag{2}$$

where $2\pi\Phi(\rho)$ and $2\pi\Psi(\rho)$ are the toroidal and poloidal fluxes, $\theta^* = \theta + \lambda$ is the straight-field-line angle, and $\iota(\rho) = \Psi'/\Phi'$ ($' \equiv d/d\rho$) is the rotational transform. The current density is obtained via Ampère's law; covariant components are $j_i = \frac{1}{\mu_0\sqrt{g}}\,\varepsilon_{ijk}\,\frac{\partial B_k}{\partial u^j}$, and contravariant components follow by raising with the metric, $J^i = g^{ij} j_j$.

Boundary conditions and spectral gauge Axis regularity and LCFS boundary conditions are imposed analytically. It is worth noting that, besides the boundary conditions for R and Z ($R_{mn}(0) = Z_{mn}(0) = 0$, $m \neq 0$), a minor correction on λ is presented here: $\lim_{\rho\to 0}\lambda_\theta = \lim_{\rho\to 0}\left[(\sqrt{g}/g_{\zeta\zeta})/\langle\sqrt{g}/g_{\zeta\zeta}\rangle - 1\right] = 0$ ($m \neq 0$), where the poloidal average normalizing this limit corrects the original VMEC near-axis treatment of λ. A spectrum-condensation constraint [5] removes the poloidal-angle gauge freedom and selects a unique coefficient set.

Bayesian Formulation

Latent variables and prior The latent state collecting all equilibrium degrees of freedom is

$$\xi = \left\{R_{mn}(\rho), Z_{mn}(\rho), \lambda_{mn}(\rho), p(\rho), \iota(\rho), n_e(\rho), T_e(\rho)\right\}.$$

Each radial function is assigned a GP prior, $f(\rho) \sim \mathcal{GP}(m(\rho), k(\rho,\rho'))$, where $m(\rho)$ is the prior mean and $k(\rho, rho')$ the covariance kernel. Axis regularity and boundary constraints are incorporated directly into the prior through GP conditioning and constrained covariance constructions [3, 4].

Two-domain structure A central architectural feature is the separation between two spaces (Figure 1). The *flux-coordinate space* is the latent domain, where all priors are defined and no diagnostic data appear directly. The *real-space domain* is the observation domain, where diagnostics measure physical quantities at fixed (R, Z, ϕ) locations. The geometric embedding

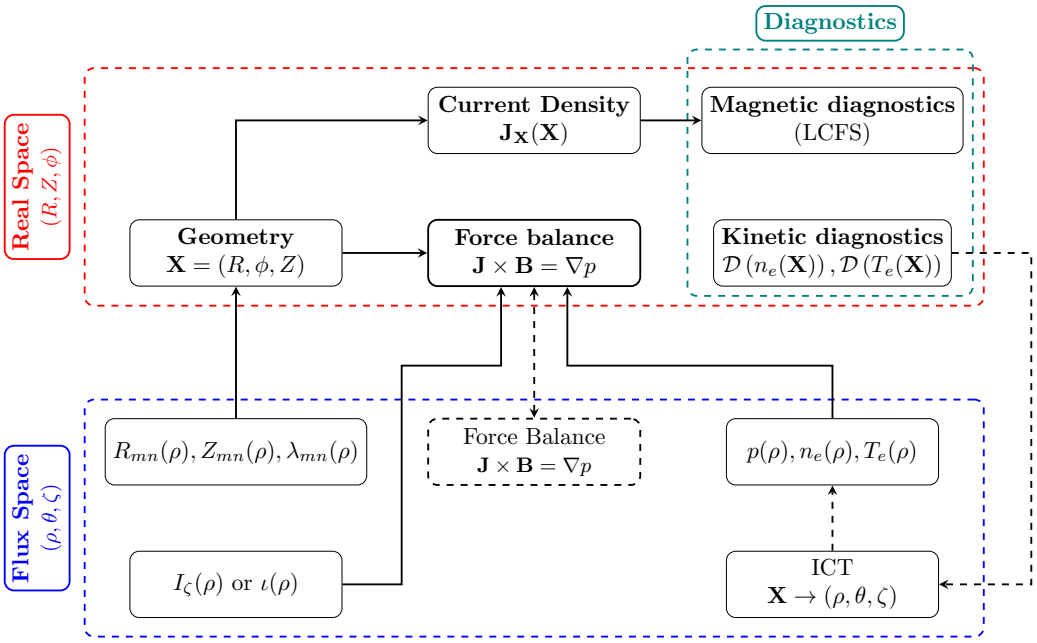

Figure 1: Overview of the Bayesian 3-D equilibrium reconstruction framework: latent quantities are defined in flux coordinates, while diagnostics act in real space.

$\mathbf{X}(\rho, \theta, \zeta; \xi)$ is the deterministic forward map connecting them; its functional inverse, the *inverse coordinate transform* (ICT), maps real-space diagnostic locations back to flux coordinates for given ξ (detailed below for kinetic diagnostics).

Likelihood Functions

Magnetic diagnostics For pick-up coils and flux loops, the predicted signal $\mu_i^{\text{mag}}(\xi)$ is a deterministic functional of the magnetic field implied by ξ. In a free-boundary setting this requires solving the vacuum Laplace problem via Green's third identity. A Gaussian noise model yields

$$p(\mathscr{D}_{\text{mag}}|\xi) = \mathscr{N}\left(\mathscr{D}_{\text{mag}} \mid \boldsymbol{\mu}_{\text{mag}}(\xi), \Sigma_{\text{mag}}\right).$$

Kinetic diagnostics and the ICT For kinetic diagnostics (interferometry, Thomson scattering, ECE) the measurement location $\mathbf{X}_i$ maps to a flux coordinate (ρ_i, θ_i) through the (deterministic) ICT. Moreover, if $\mathbf{X}_i$ itself carries positioning uncertainty, it is represented directly in real space, e.g. $R \sim \mathscr{N}(R_i, \sigma_R^2)$, $Z \sim \mathscr{N}(Z_i, \sigma_Z^2)$, and $\phi \sim \mathscr{N}(\phi_i, \sigma_\phi^2)$. These uncertainties are marginalized over in flux coordinates with the correct Jacobian to ensure coordinate invariance:

$$p(d_i|\xi) = \int \mathscr{N}(d_i|q(\rho), \sigma_i^2)\, \mathscr{N}(\delta\mathbf{r}|\mathbf{0}, \Sigma_{X_i}) \times \left|\det \mathscr{J}(\rho, \theta; \xi)\right| d\rho\, d\theta\, d\zeta, \tag{3}$$

where $q(\rho)$ is the profile quantity predicted from ξ, $\delta\mathbf{r} = \mathbf{X}(\rho, \theta, \zeta; \xi) - \mathbf{X}_i$, and $\Sigma_{X_i} = \text{diag}(\sigma_R^2, \sigma_Z^2, \sigma_\phi^2)$. Because $\phi \equiv \zeta$, the 3×3 Jacobian reduces to the in-plane block $\mathscr{J} = \partial(R, Z)/\partial(\rho, \theta)$, so

the 2-D form stays exact. In the limit of small Σ_{X_i} the Jacobian factors cancel and $p(d_i|\xi) \approx \mathcal{N}(d_i|q(\rho_i(\xi)), \sigma_i^2)$.

Physics virtual observations Force-balance constraints enter as virtual observations. Collecting the MHD residuals at quadrature nodes into $\mathbf{r}_{\text{eq}}(\xi)$ (VMEC-style (F_R, F_Z, F_λ) or DESC-style (F_ρ, F_β); both project the same physics into different residual spaces), the soft constraint is

$$p(\mathscr{D}_{\text{eq}}|\xi) = \mathcal{N}\big(\mathbf{0} \mid \mathbf{r}_{\text{eq}}(\xi), \Sigma_{\text{eq}}\big),$$

where $\Sigma_{\text{eq}} \to \mathbf{0}$ recovers a hard equilibrium constraint and finite Σ_{eq} allows controlled deviation representing model mismatch.

Posterior and Outlook

Full posterior structure Combining prior and likelihood factors, the posterior is

$$p(\xi|\mathscr{D}) \propto p_{\text{mag}}(\mathscr{D}_{\text{mag}}|\xi)\, p_{\text{kin}}(\mathscr{D}_{\text{kin}}|\xi)\, p_{\text{phys}}(\mathscr{D}_{\text{eq}}|\xi) \times p_{\mathscr{G}\mathscr{P}}(\xi)\, p_{\text{axis}}(\xi)\, p_{\text{SC}}(\xi). \tag{4}$$

Here $\mathscr{D}_{\text{kin}} = \{d_i\}$, $\mathscr{D}_{\text{eq}}$ are the kinetic and virtual-observation data (Sec. 4), and $p_{\text{axis}}, p_{\text{SC}}$ the axis-regularity and spectrum-condensation priors (Sec. 2). Diagnostic data constrain latent flux-space variables, while physical and gauge constraints are encoded in the prior. The VMEC nested-surface model is adopted as the equilibrium representation.

GP kernel hyperparameters $\boldsymbol{\vartheta}$ (length scales, amplitudes) are either fixed from physics knowledge, estimated by maximising the marginal likelihood, or marginalised jointly with ξ.

Outlook This formulation closes the gap between deterministic 3-D equilibrium solvers and statistically rigorous inference: every latent quantity carries a calibrated uncertainty, and force balance is enforced exactly in the limit $\Sigma_{\text{eq}} \to 0$. Validation proceeds in three stages — synthetic 2-D equilibria with known ground truth, synthetic 3-D equilibria, and finally W7-X experimental data.

References

[1] S. P. Hirshman and J. C. Whitson, Phys. Fluids **26**, 3553 (1983).

[2] D. W. Dudt and E. Kolemen, Phys. Plasmas **27**, 102513 (2020).

[3] R. Fischer *et al.*, Fusion Sci. Technol. **76**, 879–893 (2020).

[4] M. A. Chilenski *et al.*, Nucl. Fusion **55**, 023012 (2015).

[5] S. P. Hirshman, W. I. van Rij and P. Merkel, Comput. Phys. Commun. **43**, 143–155 (1986).