

Contents

| | | |
|-------|-------------------------------------------------------------------------------|---|
| 0.1 | Interferometry Results - Density Profile and Current Inversions | 2 |
| 0.1.1 | Fixed equilibrium ψ_N with simple prior | 2 |
| 0.1.2 | Smoothing priors | 3 |
| 0.1.3 | Interferometry and current tomography - Free ψ_N non-linear inversion. . | 5 |

0.1 Interferometry Results - Density Profile and Current Inversions

As discussed in section ??, the interferometry likelihood $P(\underline{D}_i | n_e(R, Z))$ is sufficiently modelled as a fixed width Gaussian distribution, centred on the line integrated electron density along each line of sight. If the density at any spatial point $n_e(R, Z)$, is a linear function of the parameters used, the determination of those parameters from the (fringe-corrected) interferometry data \underline{D}_i satisfies the conditions of the truncated linear Gaussian inversion (TLGI - see section ??).

The 8 lines of sight are available on JET are clearly insufficient to obtain any certainty on a complete 3D or even 2D (e.g. assuming axisymmetry) parametrisation. Since, for low-flow equilibrium electron density should be approximately constant on each closed flux surface (see section ??), the density can be parametrised as a 1D function of normalised poloidal flux $n_e(\psi_N)$. For the SOL surfaces, the assumption may be slightly less valid, but since all but one of the lines of sight are very insensitive to the SOL density, the uncertainty will always be far larger than any systematic error due to this assumption.

0.1.1 Fixed equilibrium ψ_N with simple prior

To maintain linearity, the single fixed flux map $\psi_N(R, Z)$ taken from the standard equilibrium analysis (EFIT) is used and to ensure $n_e(\psi_N)$ is linear with respect to its parameters, a simple linear interpolation of a series of knots at fixed ψ_N is used. The values of n_e at the knots serve as the parameters. To allow the use of the TLGI, the prior must also be expressed as a Gaussian. The simplest form to use is an independent normal distribution on each mode, centred at zero with $+1\sigma$ slightly higher than JETs maximum operating regime and truncated so that $P(n_e < 0) = 0$ since densities cannot be negative.

Figure 1a shows the described inversion to $P(n_e | \underline{D}_i, \psi_N)$ for a typical JET O-mode plasma, given the fixed flux surfaces of the standard magnetics-only equilibrium and the simple Gaussian prior.

These profiles look immediately 'incorrect' to any even mildly experienced Tokamak physicist, since it is usual to expect fairly smooth profiles throughout the plasma core, where there is no




Figure 1: a) Profiles at maximum (thick) and at 20 samples (thin) from posterior $P(n_e(\psi_N) | D_i, \psi_N^{EFIT}(R, Z))$ using a simple uncorrelated Gaussian prior $P(n_e) = \mathcal{G}(n_e; 0, 10^{21})$ for $n_e > 0$. b) Values of two neighbouring knots for 10000 samples showing degeneracy due to integrated measurements.

reason for transport or particle source rates to change significantly over short length scales. The posterior here does not in-fact exclude such profiles and shows that the short length scale information cannot be determined by the line integral measurements. Figure 1b illustrates this, showing that this degeneracy is correctly described by the posterior. The average n_e in the region of the two neighbouring nodes shown is fairly well determined, but the difference not. The posterior is correctly showing that with no other information, such large amplitude oscillations with short-wavelength could exist in the plasma.

0.1.2 Smoothing priors

Clearly, this claimed knowledge that transport and source rates do not change over short length scales should be included in the prior. To maintain applicability of the TLGI, this belief must be represented as a multivariate Gaussian over the parameters. A common approach is to include a Gaussian prior probability on the difference between each neighbouring pair of nodes. Each Gaussian is usually centred on 0, with a fixed width across the whole profile, or across separated regions such as the plasma core and edge. The effect is to smooth the profiles by constraining the first differential and hence it is known throughout this work as the *first differential smoothing*

prior. The general expression is:

$$1^{st} \text{ diff: } \log P(y(x)) = -\frac{1}{2\sigma_{\frac{dy}{dx}}^2} \sum_{i=1}^{N-1} \left[\frac{(y_{i+1} - y_i)}{(x_{i+1} - x_i)} - 0 \right]^2 \quad (1)$$

An alternative is to apply a Gaussian to the change in this difference between each pair of nodes and the next, giving the *second differential smoothing prior*:

$$2^{nd} \text{ diff: } \log P(y(x)) = -\frac{1}{2\sigma_{\frac{d^2y}{dx^2}}^2} \sum_{i=2}^{N-1} \left[\frac{(y_{i+1} - y_i)}{(x_{i+1} - x_i)} - \frac{(y_i - y_{i-1})}{(x_i - x_{i-1})} - 0 \right]^2 \quad (2)$$

In this case, these can be applied either to the spatial derivatives or to the derivative with respect to the normalised flux.

The electron density gradient in steady-state is determined by the local electron particle transport and the local electron source rate and since there is no prior information about the magnitude of these, the first differential constraint, which prefers small gradients is hard to justify (other than by pessimism about controlled fusion). However, the second differential constraint fits the belief that these gradients should not change rapidly over short scales. Since the parametrisation in this case is over normalised flux, it is most practical to constrain $d^2n_e/d\psi_N^2$. (The constraints are sufficiently weak that applying it to the true spatial differentials d^2n_e/dr^2 instead, has little effect on the results).

Figure 2 shows the posterior obtained with the second differential smoothing prior with $\sigma_{d^2n_e/d\psi_N^2} = 1.0 \times 10^{21}$.

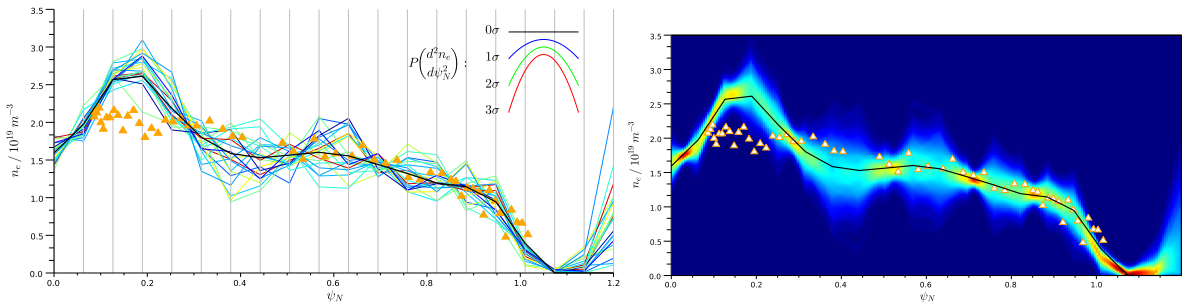


Figure 2: a) Profiles at maximum (thick) and at 20 samples (thin) from posterior $P(n_e(\psi_N) | D_i, \psi_N^{EFIT}(R, Z))$ using second differential smoothing prior. Also shown are curves at 1, 2 and 3σ of the prior used. b) Profile of marginal PDFs from the same posterior.

Considering that only 8 data values relating to n_e have been used in the inversion and that the smoothing prior is relatively weak, the results of the inversion are surprisingly close to those of the HRTS diagnostic shown for comparison. For several values of ψ_N , the certainty is reasonably high, shown by the 'hot spots' in figure 2b. These correspond to the values of ψ_N for which one of the lines of sight is tangential to the surface, making that data value very sensitive to the density on that surface.

The remaining discrepancy is only in the very core of the plasma ($\psi_N < 0.4$) where the posterior samples all show structure that is not seen in the HRTS profile, and which is difficult to see any cause of in this stable and quiescent O-mode plasma. Again, the posterior is in fact correct. It does show all reasonably probable n_e profiles, given ψ_N . If the profiles are considered improbable, the implication is that the fixed ψ_N is improbable (or possibly one of the other assumptions).

0.1.3 Interferometry and current tomography - Free ψ_N non-linear inversion.

If the electron density is the object of interest, then the posterior really desired is $P(n_e | D_i)$ which is equal to that found in the last section $P(n_e | \underline{D}_i, \psi_N)$, integrated over all possible flux surface arrangements. This would be as uninformative as assuming nothing of the 2D geometry and parameterising $n_e(R, Z)$ directly and so the PDF would give almost any possible electron density. However, some information about the flux surface geometry can be gained from the magnetic diagnostics by including their data D_m to give the posterior $P(n_e | D_i, D_m)$. Practically, this is a combination of the current tomography (CT) from section ?? and the density inversion from the last section with ψ_N calculated from the toroidal current beams \underline{J} . The toroidal currents \underline{J} and electron density knots \underline{n}_e form the complete set of parameters.

This unfortunately introduces the non-linear relationship between the interferometry data D_i and the currents J , since a change to J changes both the normalisation and geometry of ψ_N . This non-linearity turns out to be weak enough that the LGI procedure can still be performed and while the resulting Gaussian is not a good enough approximation to the real PDF to draw samples directly, it does give the approximate global shape, especially for the strongest correlations.

After appropriate rescaling, it performs excellently as a proposal distribution for the Metropolis Hastings MCMC (see section ??), something that is otherwise extremely hard to obtain in the high dimensionality necessary to correctly describe the toroidal currents J . Practically, the use of the approximate LGI reduces the execution time of the MCMC by over an order of magnitude.

Figure 3 shows the density profiles from the full posterior. The inversion is to a parametrisation of 50 knots of $n_e(\psi_N)$, 217 plasma current beams spread inside the firstwall and 48 Iron-core current parameters. The priors used were the second differential prior on n_e with $\sigma_{d^2 n_e / d\psi_N^2} = 1.0 \times 10^{21}$ and a very weak Conditional-Auto-Regressive (CAR) prior over the plasma currents.




Figure 3: a/c) n_e Profiles for 20 samples and b/d) profile marginals from posterior $P(n_e(\psi_N), J | D_i, D_m)$ using second differential smoothing prior and free plasma currents, versus a/b) ψ_N and b/d) R_{mag} .

As well as the density profile, each sample of $P(n_e, J | D_i, D_m)$ contains a complete description of the toroidal current. Comparing this to the posterior from current tomography alone $P(J | D_m)$, indicates the information that the interferometry data (and the n_e prior) provides about the currents and hence the magnetic geometry. Figure 4 shows the separatrix and magnetic axis position from samples of the two posteriors to demonstrate the principal. In this case,

the apparent accuracy might be easily obtained simply by increasing the prior assumptions (e.g. using a strong CAR prior, or assuming equilibrium) but the improvement here has been derived directly from the data of a diagnostic not normally considered capable of measuring anything about the plasma current.

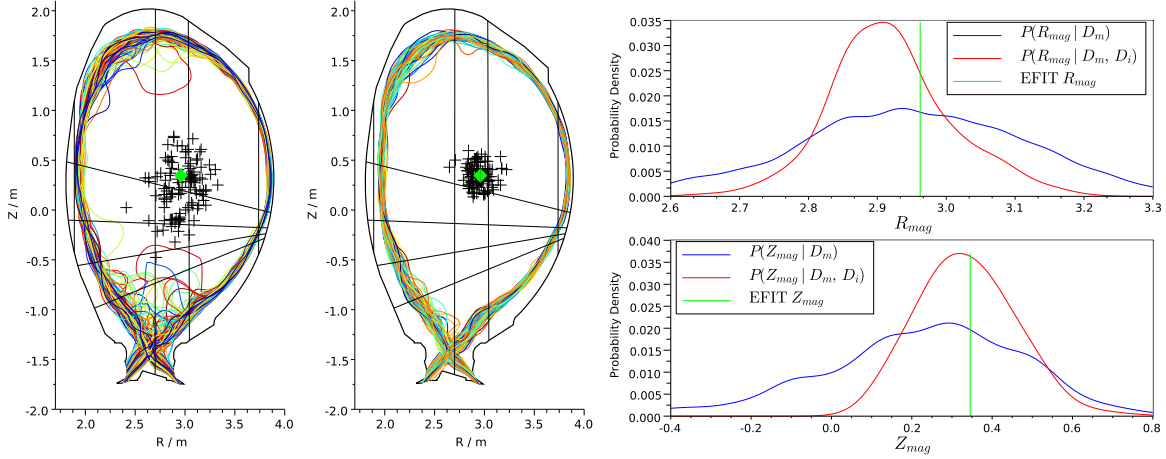


Figure 4: Comparison of magnetic configurations of the posterior from Current Tomography (CT) alone and including interferometry data and electron density profile prior, with the same parameterisation and prior for toroidal currents. a) Separatrix and magnetic axis position for 100 samples from posterior for left) CT only and right) CT+interferometry. Interferometry lines of sight and first wall (black) and standard magnetics EFIT axis position (green diamond) are also shown. b) PDFs for top) R_{mag} and bottom) Z_{mag} each showing CT-only (blue), CT+interferometry (red) and EFIT (green).

At first glance the results are concerning, since it is hard to believe that 8 numbers can provide information about the plasma currents at the same time as electron density profiles that appear a similar resolution as the HRTS system. It should be emphasised though, that the electron density priors (the smoothing and assumption of constancy on flux surfaces) are providing much of the information and the results are accurate because the priors are good priors for O-mode and L-mode plasma.

For H-Mode plasmas, the transport rates change rapidly across the plasma edge and hence the smoothing prior must be relaxed, even if only in this region. Unfortunately, JET H-mode shots are usually positioned in such a way that the interferometry line of sight which measures only SOL density (channel 1) is further from the plasma edge, as seen in figure 5b. The pedestal information is effectively provided by this channel and the difference between the two core channels

(channels 2 and 3) which couples the density near the X-point to the plasma core. To illustrate this, figure 5a shows the integration weights and density profile (at the posterior maxima) for the O-mode plasma of figure 3 and for a typical H-mode plasma. It is clear that the pedestal region of the H-mode pulse is illuminated by even less information than in the O-mode case. The posterior profiles become dominated by the smoothing priors which favour less steep profiles and this has a knock-on effect on the core, due to the differences between channels 2 and 3, which can be seen by comparing the posterior profile marginals to the HRTS standard analysis in figure 6.

Clearly, for a useful posterior on H-mode pulses, more information is required and to avoid imposing stronger priors, more diagnostic data needs to be added.

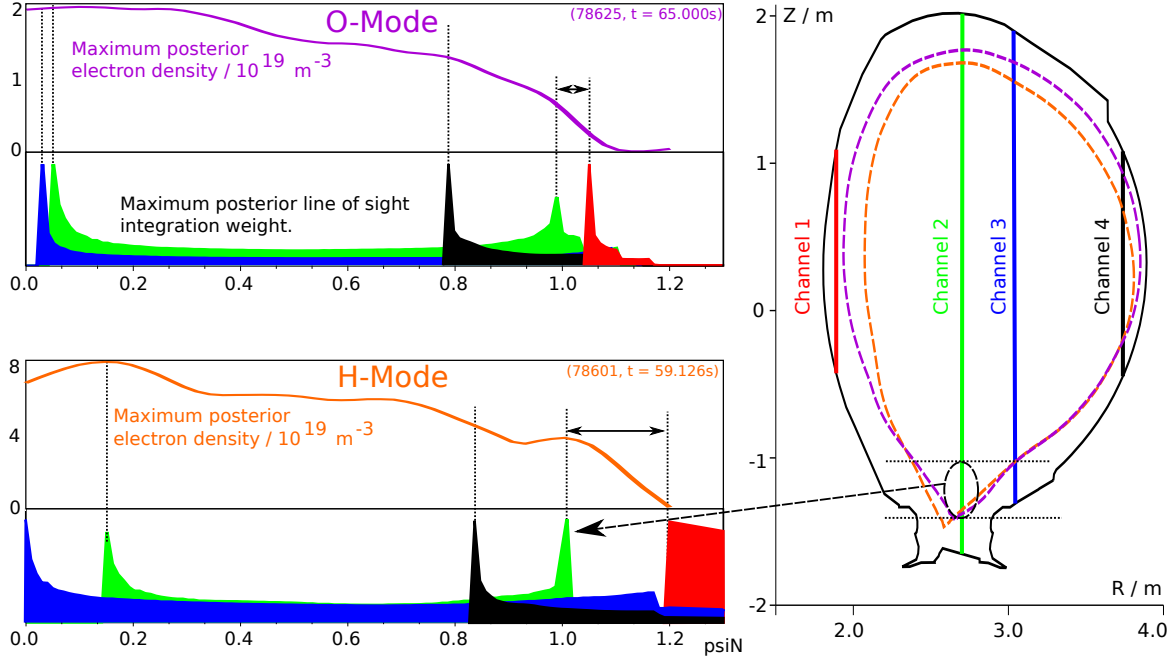


Figure 5: a) Integration weights over ψ_N and $n_e(\psi_N)$ at posterior maximum for typical O-mode (top) and H-Mode (bottom) plasmas. b) Separatrix ($\psi_N = 1$) for both plasmas relative to lines of sight.




Figure 6: Profile marginals from posterior $P(n_e(\psi_N), J | D_i, D_m)$ for H-Mode plasma using second differential constraint 'smoothing' prior and free plasma currents.

Bibliography