Seismic uncertainty quantification beyond the normal – machine-learning-based variational inference with Moreau-enveloped priors

Gabrio Rizzuti, Ivan Vasconcelos, Lorenzo Casasanta, Shearwater GeoServices*

ABSTRACT

In many applications in seismic inversion, the setup of the inverse problem typically relies on the minimization of an objective function, which is the linear combination of a differentiable term and a non-smooth regularization term. Hinging on *proximal solvers*, a whole stew of optimization strategies has been devised to overcome the mathematical challenges of non-smooth regularization (Parikh, 2014). Here, we look at the non-trivial question of uncertainty quantification (UQ) for such problems - generalizing UQ beyond the use of normally-distributed (Gaussian) priors. While convex optimization is widely applied to deterministic inverse problems, there are no established strategies for estimating the associated posterior distribution, especially when considering *variational inference* (VI, Blei, 2017). To address this need for VI-based UQ in general convex optimization problems we aim at a notable solution based on the smoothing of such non-smooth objectives via infimal convolution, also known in the literature as the *Moreau envelope*.

We start by considering the log-posterior distribution associated with a generic inverse problem:

$$-\log p(\boldsymbol{x}|\boldsymbol{y}) = f(\boldsymbol{x}) + \lambda g(\boldsymbol{x}) + \text{const},$$

where $\boldsymbol{x}$ is a variable of interest (e.g. acoustic impedance) and $\boldsymbol{y}$ are measured data (e.g. a seismic stack). The functional f is convex and differentiable, while g is a convex—though generally not differentiable—function. The parameter λ weights the relative strength of the regularization. We approximate the target posterior with a candidate distribution implicitly represented by a transport map T. Samples from the candidate distribution are thus obtained by evaluating T with random inputs. Our choice of T belongs to the class of invertible networks known as *normalizing flows* (NF, Kobyzev, 2021). Conventional gradient-based optimization is, however, not directly applicable to this problem since g is not differentiable.

Our main goal is to replace g with a differentiable perturbation thereof. The main tool will be the proximal operator of a convex function, defined as:

$$\text{prox}_{\lambda,g}(\boldsymbol{x}) = \arg\min_{\boldsymbol{x}'} \|\boldsymbol{x}' - \boldsymbol{x}\|^2/2 + \lambda g(\boldsymbol{x}').$$

The Moreau envelope of g with "smoothing" level ρ is:

$$g_\rho(\boldsymbol{x}) = \min_{\boldsymbol{x}'} \|\boldsymbol{x}' - \boldsymbol{x}\|^2/(2\rho) + g(\boldsymbol{x}').$$

Formally, we have

$$\lim_{\rho\to 0^+} g_\rho(\boldsymbol{x}) = g(\boldsymbol{x}),$$

which ensures that, for small ρ, the Moreau envelope is close to the original functional. Furthermore, g_ρ is differentiable and:

$$\nabla g_\rho(\boldsymbol{x}) = (\boldsymbol{x} - \overline{\boldsymbol{x}})/\rho, \qquad \text{s.t. } \overline{\boldsymbol{x}} = \text{prox}_{\rho,g}(\boldsymbol{x}).$$

In conclusion, we can swap the non-smooth term g with its Moreau envelope, while still employing gradient-based optimization in the variational inference framework. Similar techniques can be applied to hard-constraint regularization as well (Peters, 2022).

By relying on our approach to envelope total-variation constraints in acoustic impedance UQ (Figure 1), we show the capability of inferring posterior information from non-Gaussian priors – as evidence by the displayed asymmetry between point-wise percentiles.

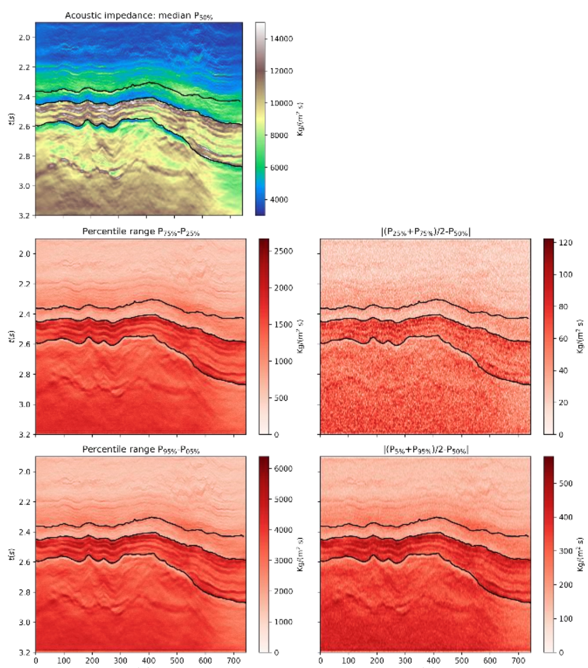

Figure 1: Post-stack acoustic inversion results with Moreau-enveloped total variation regularization ($\rho = 0.1$), obtained via NF-based VI. We show the point-wise median and point-wise percentile to highlight the asymmetry of the posterior distribution (cf. 25% vs 75%, and 5% vs 95% percentiles).

Acknowledgements

We thank Prof. Tristan van Leeuwen for many insightful discussions on convex optimization.