Constraining gradient-based inversion with a variational autoencoder to reproduce geological patterns

Jorge Lopez-Alvis[1,3], Eric Laloy[2], Thomas Hermans[3], and Frédéric Nguyen[1]

[1]Applied Geophysics, Urban and Environmental Engineering, University of Liege, Liege, Belgium
[2]Engineered and Geosystems Analysis, Institute for Environment, Health and Safety, Belgian Nuclear Research Centre, Mol, Belgium
[3]Department of Geology, Ghent University, Ghent, Belgium

Given the sparsity of geophysical data it is useful to rely on prior information on the expected geological patterns to constrain the inverse problem and obtain a realistic image of the subsurface. By using several examples of such patterns (e.g. those obtained from a training image), deep generative models learn a low-dimensional latent space that can be seen as a reparameterization of the original high-dimensional parameters and then inversion can be done in this latent space. Examples of such generative models are the variational autoencoder (VAE) and the generative adversarial network (GAN). Both usually include deep neural networks within their architecture and have shown good performance in reproducing high-dimensional structured subsurface models. However, they both use a highly nonlinear function to map from latent space to the original high-dimensional parameter space which hinders the optimization of the objective function during inversion. Particularly, such nonlinearity may give rise to local minima where gradient-based inversion gets trapped and therefore fails to reach the global minimum. GAN has been previously used with gradient-based inversion in a linear traveltime tomography synthetic test where it was shown to often fail in reaching a consistent RMSE (compared to the added noise) because optimization converges to local minima. On the other hand, inversion with MCMC and GAN was shown to reach acceptable RMSE values. When applicable, however, a gradient-based inversion is preferred because of its lower computational demand. We propose using VAE together with gradient-based inversion and show that optimization reaches lower RMSE values on average compared to GAN in a linear traveltime tomography synthetic case. We also compare the subsurface models that are generated during the iterations of the optimization to explore the effect of the different latent spaces used by GAN and VAE. We identify a trade-off between a strict following of the patterns and getting trapped in local minima during optimization, i.e. VAE seems to be able to break some continuous channels in order to not get trapped in local minima whereas GAN does not break channels. Finally, we perform some synthetic tests with nonlinear traveltime tomography and show that gradient-based inversion with VAE is able to recover a similar global structure to the true model but its final RMSE values are still far from the added noise level.