Scalable Bayesian seismic wavelet estimation

Guillermina Senn[1]

jointly with Matt Walker[2], Håkon Tjelmeland[1], and Andrew Holbrook[3]

[1]Norwegian University of Science and Technology (NTNU)
[2]bp · [3]University of California, Los Angeles

Lofotseminaret i anvendt geofysikk
August 21, 2025

Scalable Bayesian seismic wavelet estimation

Guillermina Senn[1]

jointly with Matt Walker[2], Håkon Tjelmeland[1], and Andrew Holbrook[3]

[1]Norwegian University of Science and Technology (NTNU)
[2]bp · [3]University of California, Los Angeles

Lofotseminaret i anvendt geofysikk
August 21, 2025

Summary

1. Data

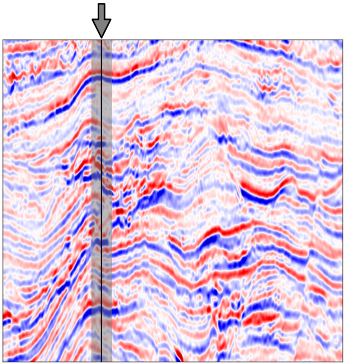

(a) Seismic data

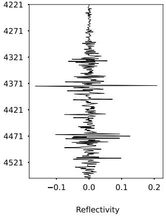

(b) Well-log

2. Model

1D convolutional model

$$\boldsymbol{d}_j = \boldsymbol{c}_j \star \boldsymbol{w} + \boldsymbol{e}_j \quad \text{for column } j.$$

+ Bayesian probabilistic model, on a cyclic lattice .

3. Scalable wavelet estimation

Sample posterior distribution

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \sigma_d^2 \mid \boldsymbol{d})$$

with Markov chain Monte Carlo (MCMC).

Summary

1. Data

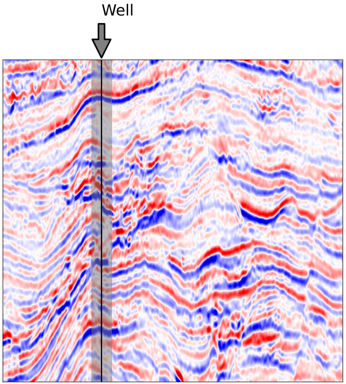

(a) Seismic data

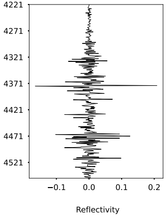

(b) Well-log

2. Model

1D convolutional model

$$\boldsymbol{d}_j = \boldsymbol{c}_j \star \boldsymbol{w} + \boldsymbol{e}_j \quad \text{for column } j.$$

+ Bayesian probabilistic model, on a cyclic lattice .

3. Scalable wavelet estimation

Sample posterior distribution

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \sigma_d^2 \mid \boldsymbol{d})$$

with Markov chain Monte Carlo (MCMC).

Summary

1. Data

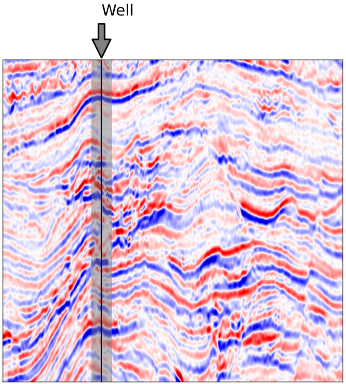

(a) Seismic data

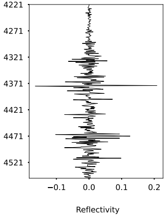

(b) Well-log

2. Model

1D convolutional model

$$\boldsymbol{d}_j = \boldsymbol{c}_j \star \boldsymbol{w} + \boldsymbol{e}_j \quad \text{for column } j.$$

+ Bayesian probabilistic model, on a cyclic lattice .

3. Scalable wavelet estimation

Sample posterior distribution

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \sigma_d^2 \mid \boldsymbol{d})$$

with Markov chain Monte Carlo (MCMC).

Summary

1. Data

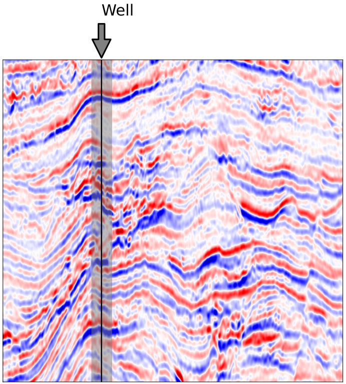

(a) Seismic data

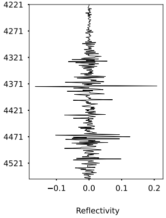

(b) Well-log

2. Model

1D convolutional model

$$\boldsymbol{d}_j = \boldsymbol{c}_j \star \boldsymbol{w} + \boldsymbol{e}_j \quad \text{for column } j.$$

+ Bayesian probabilistic model, on a cyclic lattice .

3. Scalable wavelet estimation

Sample posterior distribution

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \sigma_d^2 \mid \boldsymbol{d})$$

with Markov chain Monte Carlo (MCMC).

Summary

1. Data

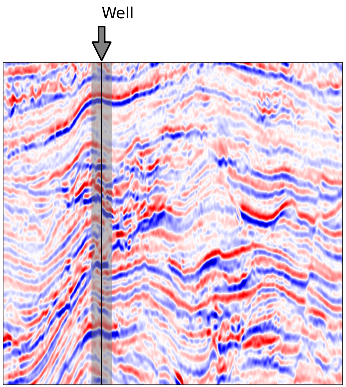

(a) Seismic data

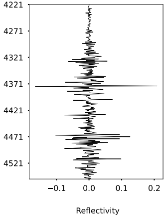

(b) Well-log

2. Model

1D convolutional model

$$\boldsymbol{d}_j = \boldsymbol{c}_j \star \boldsymbol{w} + \boldsymbol{e}_j \quad \text{for column } j.$$

+ Bayesian probabilistic model, on a cyclic lattice .

3. Scalable wavelet estimation

Sample posterior distribution

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \sigma_d^2 \mid \boldsymbol{d})$$

with Markov chain Monte Carlo (MCMC).

The 1D convolutional model

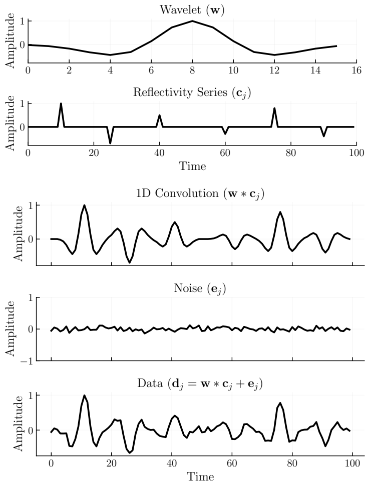

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \tag{1}$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

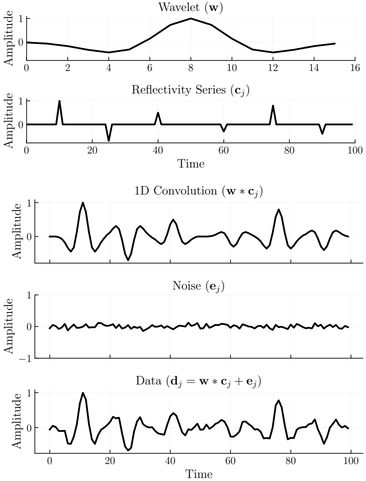

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \tag{1}$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

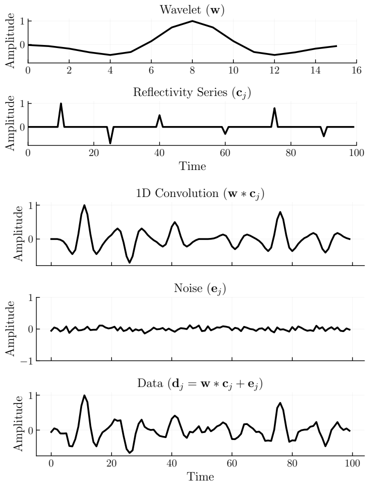

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \tag{1}$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

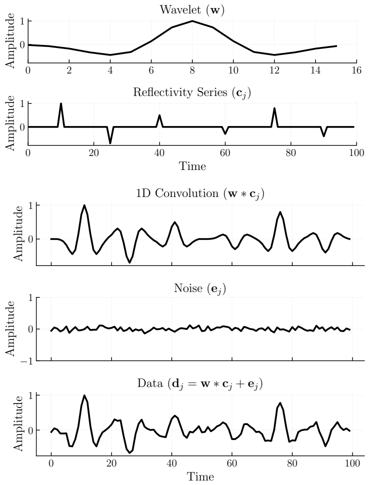

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \quad (1)$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

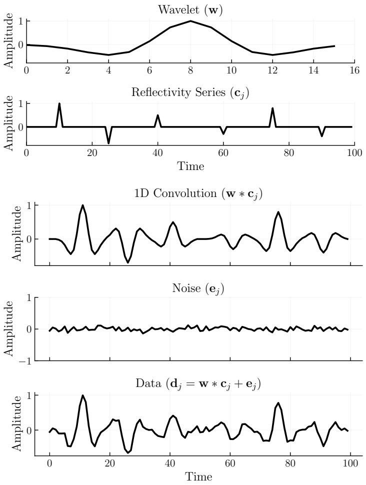

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \tag{1}$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

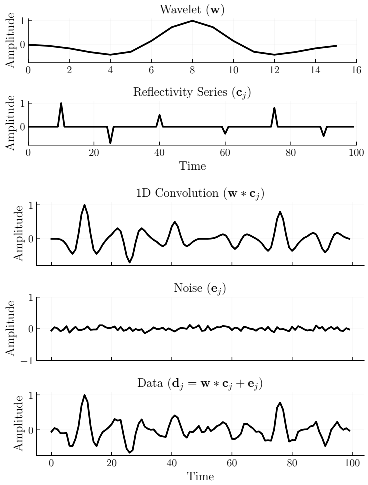

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \tag{1}$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

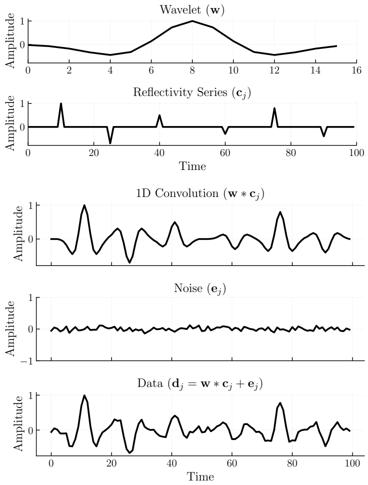

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \tag{1}$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The 1D convolutional model

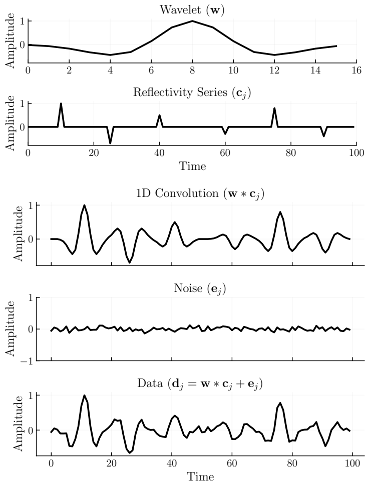

Single trace j:

- $\boldsymbol{d}_j, \boldsymbol{c}_j, \boldsymbol{e}_j \in \mathcal{R}^n$
- $\boldsymbol{w} = \{w_0, \cdots, w_{k-1}\} \in \mathcal{R}^k$
- 1D convolutional model is

$$\boldsymbol{d}_j = \boldsymbol{w} \star \boldsymbol{c}_j + \boldsymbol{e}_j \qquad (1)$$

Image with m traces:

- $\boldsymbol{d}, \boldsymbol{c}, \boldsymbol{e} \in \mathcal{R}^{nm}$.
- $\boldsymbol{W}$ is the wavelet convolutional matrix for the image.
- Convolutional model is

$$\boldsymbol{d} = \boldsymbol{W}\boldsymbol{c} + \boldsymbol{e}.$$

The Bayesian model

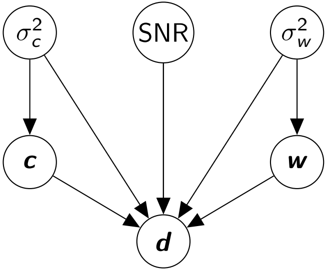

DAG of the model.

- Gaussian wavelet with inverse-gamma variance:

$$\boldsymbol{w} \sim N_k(\boldsymbol{\mu}_w, \sigma_w^2 \boldsymbol{R}_w), \quad \sigma_w^2 \sim IG(\alpha_w, \beta_w)$$

- Endpoints are constrained to zero.

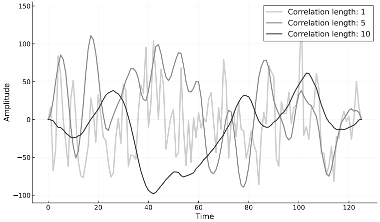

Samples from the wavelet prior.

The Bayesian model

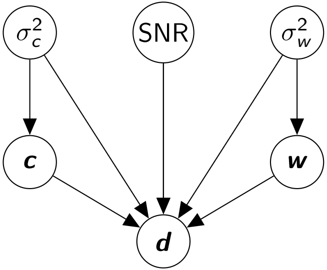

DAG of the model.

- Gaussian wavelet with inverse-gamma variance:

$$\boldsymbol{w} \sim N_k(\boldsymbol{\mu}_w, \sigma_w^2 \boldsymbol{R}_w), \quad \sigma_w^2 \sim IG(\alpha_w, \beta_w)$$

- Endpoints are constrained to zero.

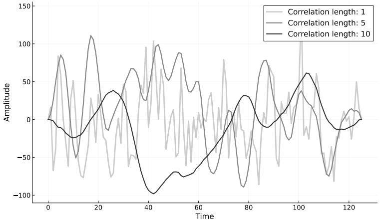

Samples from the wavelet prior.

The Bayesian model

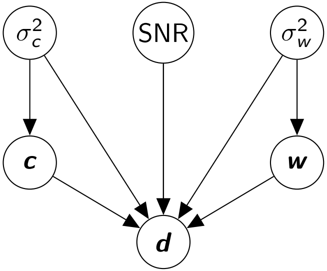

DAG of the model.

- Gaussian reflectivity with inverse-gamma variance:

$$\boldsymbol{c} \sim N_{nm}(\boldsymbol{\mu}_c, \sigma_c^2 \boldsymbol{R}_c), \quad \sigma_c^2 \sim IG(\alpha_c, \beta_c).$$

- Hard-constrained with well values.

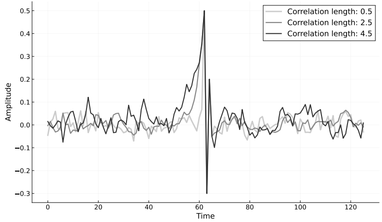

Samples from the reflectivity prior (one trace).

The Bayesian model

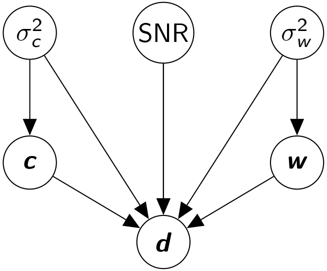

DAG of the model.

- Gaussian reflectivity with inverse-gamma variance:

$$\boldsymbol{c} \sim N_{nm}(\boldsymbol{\mu}_c, \sigma_c^2 \boldsymbol{R}_c), \quad \sigma_c^2 \sim IG(\alpha_c, \beta_c).$$

- Hard-constrained with well values.

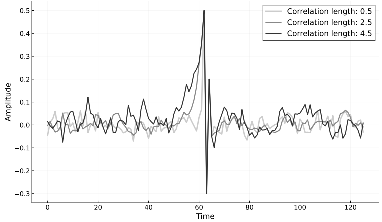

Samples from the reflectivity prior (one trace).

The Bayesian model

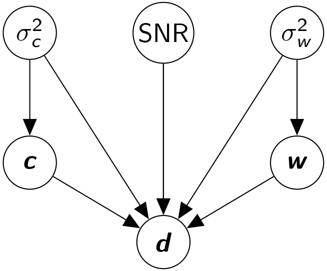

DAG of the model.

- Model the signal-to-noise ratio as stochastic

$$SNR = \frac{\mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c})}{\sigma_d^2} \iff \sigma_d^2 = \frac{\mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c})}{SNR}; \quad \mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c}) = \psi\, \sigma_c^2\, \sigma_w^2$$

- Gaussian observational noise with Gamma SNR:

$$\boldsymbol{d} | \boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \mathrm{SNR} \sim N_{nm}(\boldsymbol{W}\boldsymbol{c}, \sigma_d^2 \boldsymbol{R}_d), \quad \mathrm{SNR} \sim \mathrm{Gamma}(\alpha_{SNR}, \beta_{SNR})$$

The Bayesian model

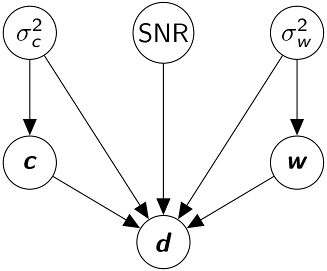

DAG of the model.

- Model the signal-to-noise ratio as stochastic

$$SNR = \frac{\mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c})}{\sigma_d^2} \iff \sigma_d^2 = \frac{\mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c})}{SNR}; \quad \mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c}) = \psi\, \sigma_c^2\, \sigma_w^2$$

- Gaussian observational noise with Gamma SNR:

$$\boldsymbol{d} | \boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \mathrm{SNR} \sim N_{nm}(\boldsymbol{W}\boldsymbol{c}, \sigma_d^2 \boldsymbol{R}_d), \quad \mathrm{SNR} \sim \mathrm{Gamma}(\alpha_{SNR}, \beta_{SNR})$$

The Bayesian model

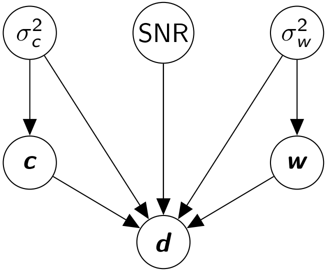

DAG of the model.

- Model the signal-to-noise ratio as stochastic

$$SNR = \frac{\mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c})}{\sigma_d^2} \iff \sigma_d^2 = \frac{\mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c})}{SNR}; \quad \mathrm{Var}(\boldsymbol{w} \star \boldsymbol{c}) = \psi\, \sigma_c^2\, \sigma_w^2$$

- Gaussian observational noise with Gamma SNR:

$$\boldsymbol{d}|\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \mathrm{SNR} \sim N_{nm}(\boldsymbol{Wc}, \sigma_d^2 \boldsymbol{R}_d), \quad \mathrm{SNR} \sim \mathrm{Gamma}(\alpha_{SNR}, \beta_{SNR})$$

Posterior distribution of the model

- Bayes' rule

$$p(\boldsymbol{\theta}|\boldsymbol{d}) = \frac{p(\boldsymbol{\theta}, \boldsymbol{d})}{p(\boldsymbol{d})} \propto p(\boldsymbol{d}, \boldsymbol{\theta})$$

gives

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta|\boldsymbol{d}) \propto p(\boldsymbol{d}|\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta)p(\boldsymbol{c}|\sigma_c^2)p(\boldsymbol{w}|\sigma_w^2)p(\sigma_c^2)p(\sigma_w^2)p(\zeta).$$

- The model is proposed on an extended cyclic lattice for scalability [1]
- We sample this intractable density with MCMC.

[1] Senn et al. (2025)

Posterior distribution of the model

- Bayes' rule

$$p(\boldsymbol{\theta}|\boldsymbol{d}) = \frac{p(\boldsymbol{\theta}, \boldsymbol{d})}{p(\boldsymbol{d})} \propto p(\boldsymbol{d}, \boldsymbol{\theta})$$

gives

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta|\boldsymbol{d}) \propto p(\boldsymbol{d}|\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta)p(\boldsymbol{c}|\sigma_c^2)p(\boldsymbol{w}|\sigma_w^2)p(\sigma_c^2)p(\sigma_w^2)p(\zeta).$$

- The model is proposed on an extended cyclic lattice for scalability [1]
- We sample this intractable density with MCMC.

[1] Senn et al. (2025)

Posterior distribution of the model

- Bayes' rule

$$p(\boldsymbol{\theta}|\boldsymbol{d}) = \frac{p(\boldsymbol{\theta}, \boldsymbol{d})}{p(\boldsymbol{d})} \propto p(\boldsymbol{d}, \boldsymbol{\theta})$$

gives

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta|\boldsymbol{d}) \propto p(\boldsymbol{d}|\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta)p(\boldsymbol{c}|\sigma_c^2)p(\boldsymbol{w}|\sigma_w^2)p(\sigma_c^2)p(\sigma_w^2)p(\zeta).$$

- The model is proposed on an extended cyclic lattice for scalability [1]
- We sample this intractable density with MCMC.

[1] Senn et al. (2025)

Posterior distribution of the model

- Bayes' rule

$$p(\boldsymbol{\theta}|\boldsymbol{d}) = \frac{p(\boldsymbol{\theta}, \boldsymbol{d})}{p(\boldsymbol{d})} \propto p(\boldsymbol{d}, \boldsymbol{\theta})$$

gives

$$p(\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta|\boldsymbol{d}) \propto p(\boldsymbol{d}|\boldsymbol{c}, \boldsymbol{w}, \sigma_c^2, \sigma_w^2, \zeta) p(\boldsymbol{c}|\sigma_c^2) p(\boldsymbol{w}|\sigma_w^2) p(\sigma_c^2) p(\sigma_w^2) p(\zeta).$$

- The model is proposed on an extended cyclic lattice for scalability [1]
- We sample this intractable density with MCMC.

[1] Senn et al. (2025)

MCMC on the cyclic lattice

Algorithm A general MCMC scheme

Initialize $\cdots$
for $t = 1, 2, \ldots, T$ **do**
  Update wavelet
  Update reflectivity
  Update reflectivity variance
  Update wavelet variance
  Update SNR

- A: **Update wavelet** from its full conditional: Gibbs sampler
  - Inefficient exploration due to high correlation in posterior samples.
- B: **Update wavelet** with Hamiltonian Monte Carlo.

MCMC on the cyclic lattice

Algorithm A general MCMC scheme

Initialize $\cdots$
for $t = 1, 2, \ldots, T$ **do**
  Update wavelet
  Update reflectivity
  Update reflectivity variance
  Update wavelet variance
  Update SNR

- A: **Update wavelet** from its full conditional: Gibbs sampler
  - Inefficient exploration due to high correlation in posterior samples.
- B: **Update wavelet** with Hamiltonian Monte Carlo.

MCMC on the cyclic lattice

Algorithm A general MCMC scheme

Initialize $\cdots$
for $t = 1, 2, \ldots, T$ **do**
  Update wavelet
  Update reflectivity
  Update reflectivity variance
  Update wavelet variance
  Update SNR

- A: **Update wavelet** from its full conditional: Gibbs sampler
  - Inefficient exploration due to high correlation in posterior samples.
- B: **Update wavelet** with Hamiltonian Monte Carlo.

MCMC on the cyclic lattice

Algorithm A general MCMC scheme

Initialize $\cdots$
for $t = 1, 2, \ldots, T$ **do**
  Update wavelet
  Update reflectivity
  Update reflectivity variance
  Update wavelet variance
  Update SNR

- A: **Update wavelet** from its full conditional: Gibbs sampler
  - Inefficient exploration due to high correlation in posterior samples.
- B: **Update wavelet** with Hamiltonian Monte Carlo.

Wavelet update with Hamiltonian Monte Carlo

- We want to avoid random walk behaviour in high dimensions.

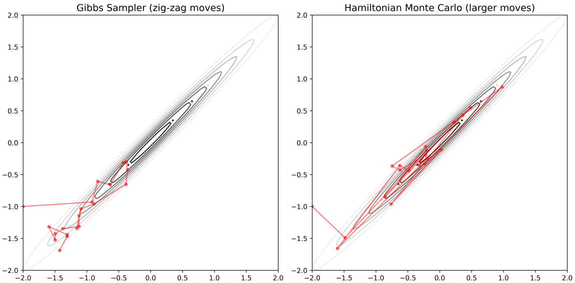

Posterior exploration of a correlated Gaussian.

- HMC explores regions with high mass (density x volume).
- HMC moves between points following a trajectory of length εL, defined by a Hamiltonian system.
- Must tune step size ε and number of steps L in the numerical integrator.

Wavelet update with Hamiltonian Monte Carlo

- We want to avoid random walk behaviour in high dimensions.

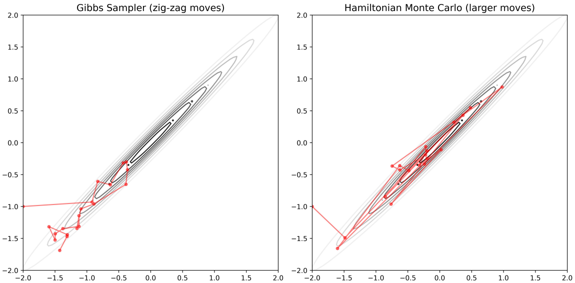

Posterior exploration of a correlated Gaussian.

- HMC explores regions with high mass (density x volume).
- HMC moves between points following a trajectory of length εL, defined by a Hamiltonian system.
- Must tune step size ε and number of steps L in the numerical integrator.

Wavelet update with Hamiltonian Monte Carlo

- We want to avoid random walk behaviour in high dimensions.

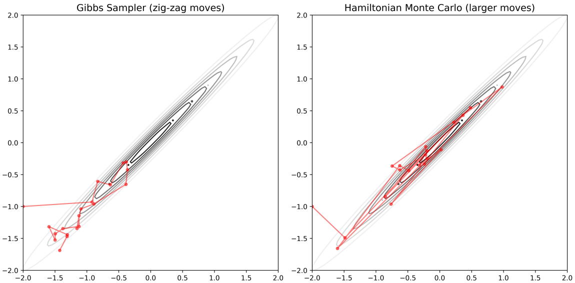

Posterior exploration of a correlated Gaussian.

- HMC explores regions with high mass (density x volume).
- HMC moves between points following a trajectory of length εL, defined by a Hamiltonian system.
- Must tune step size ε and number of steps L in the numerical integrator.

Wavelet update with Hamiltonian Monte Carlo

- We want to avoid random walk behaviour in high dimensions.

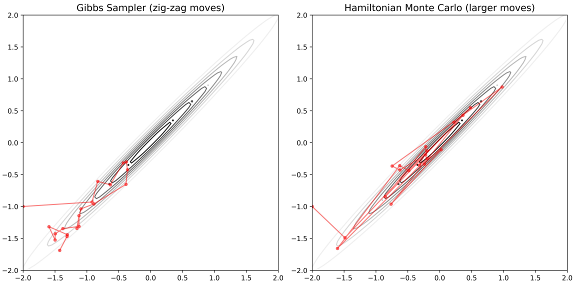

Posterior exploration of a correlated Gaussian.

- HMC explores regions with high mass (density x volume).
- HMC moves between points following a trajectory of length $\varepsilon\, L$, defined by a Hamiltonian system.
- Must tune step size ε and number of steps L in the numerical integrator.

Wavelet update with Hamiltonian Monte Carlo

- We want to avoid random walk behaviour in high dimensions.

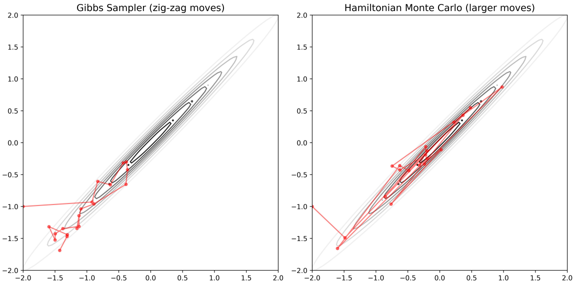

Posterior exploration of a correlated Gaussian.

- HMC explores regions with high mass (density x volume).
- HMC moves between points following a trajectory of length εL, defined by a Hamiltonian system.
- Must tune step size ε and number of steps L in the numerical integrator.

Scalability

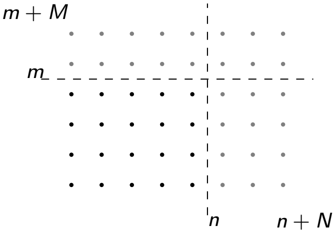

(a) The lattice is extended with M rows and N columns

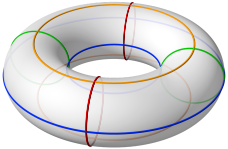

(b) The cyclic lattice.

- The covariance and convolutional matrices become circulant.
- We only need to store their first row.
- Circulant matrix algebra is done with the FFT on the first row.
- Fast evaluation of full conditional parameters in Gibbs and gradients in HMC.

Scalability

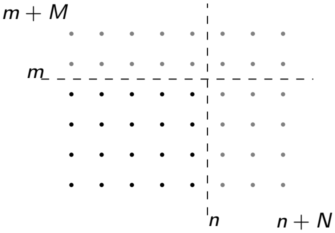

(a) The lattice is extended with M rows and N columns

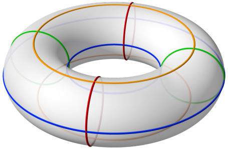

(b) The cyclic lattice.

- The covariance and convolutional matrices become circulant.
- We only need to store their first row.
- Circulant matrix algebra is done with the FFT on the first row.
- Fast evaluation of full conditional parameters in Gibbs and gradients in HMC.

Scalability

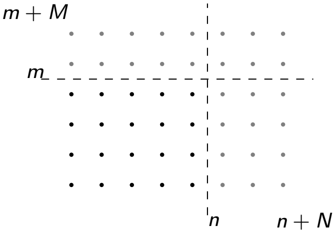

(a) The lattice is extended with M rows and N columns

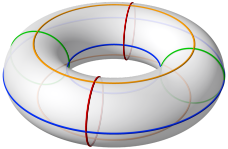

(b) The cyclic lattice.

- The covariance and convolutional matrices become circulant.
- We only need to store their first row.
- Circulant matrix algebra is done with the FFT on the first row.
- Fast evaluation of full conditional parameters in Gibbs and gradients in HMC.

Scalability

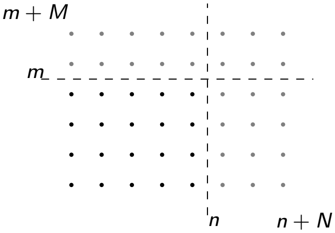

(a) The lattice is extended with M rows and N columns

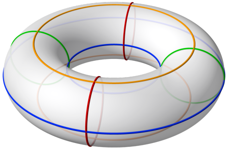

(b) The cyclic lattice.

- The covariance and convolutional matrices become circulant.
- We only need to store their first row.
- Circulant matrix algebra is done with the FFT on the first row.
- Fast evaluation of full conditional parameters in Gibbs and gradients in HMC.

Scalability

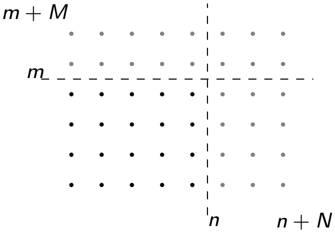

(a) The lattice is extended with M rows and N columns

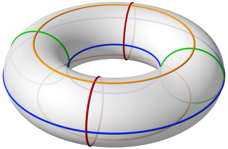

(b) The cyclic lattice.

- The covariance and convolutional matrices become circulant.
- We only need to store their first row.
- Circulant matrix algebra is done with the FFT on the first row.
- Fast evaluation of full conditional parameters in Gibbs and gradients in HMC.

Scalability

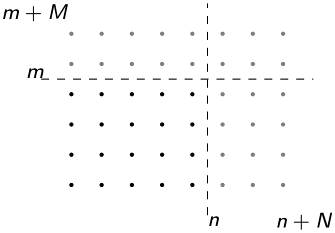

(a) The lattice is extended with M rows and N columns

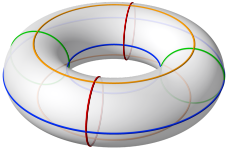

(b) The cyclic lattice.

- The covariance and convolutional matrices become circulant.
- We only need to store their first row.
- Circulant matrix algebra is done with the FFT on the first row.
- Fast evaluation of full conditional parameters in Gibbs and gradients in HMC.

Real application: Gas reservoir in offshore Egypt

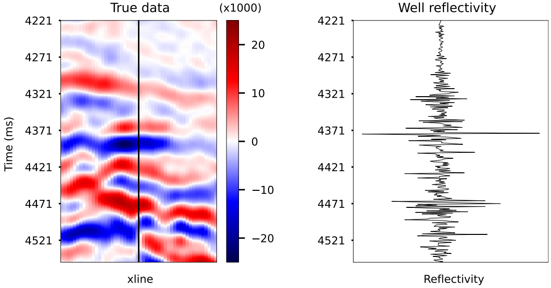

330x50ms AVO window and well reflectivity.

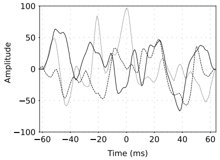

(a) Wavelet prior samples.

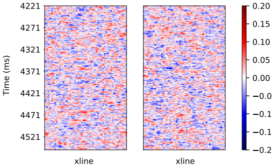

(b) Reflectivity prior samples.

Real application: Gas reservoir in offshore Egypt

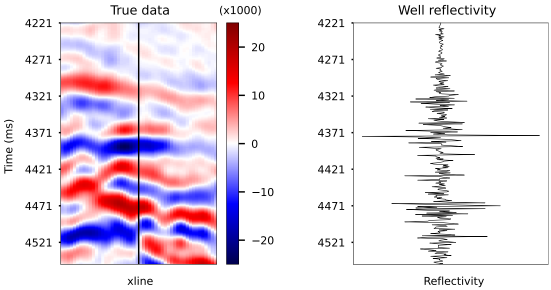

330x50ms AVO window and well reflectivity.

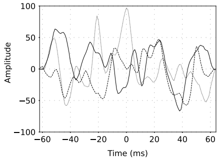

(a) Wavelet prior samples.

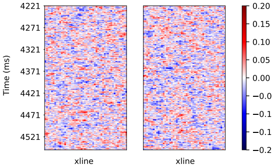

(b) Reflectivity prior samples.

Preliminary wavelet estimation results

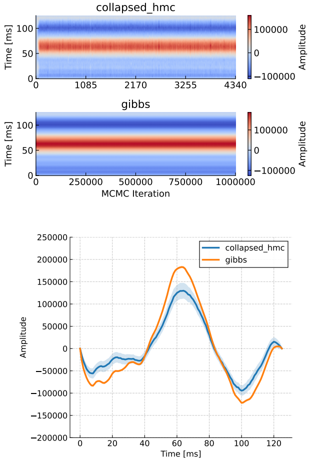

(a) Well trace.

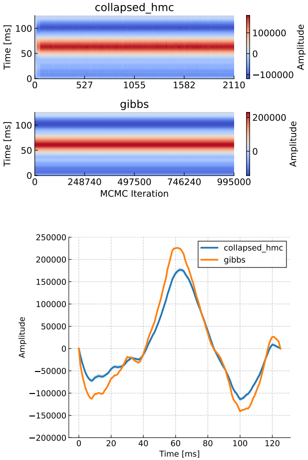

(a) 330x10ms.

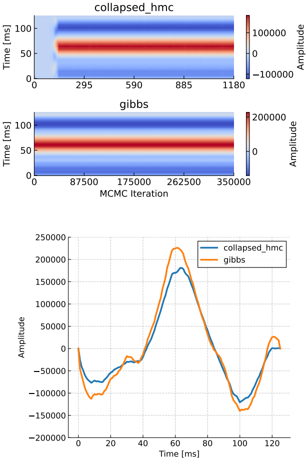

(a) 330x50ms.

Preliminary wavelet estimation results

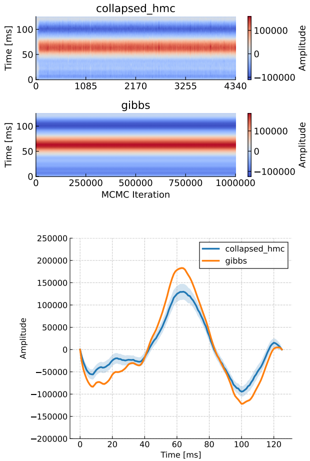

(a) Well trace.

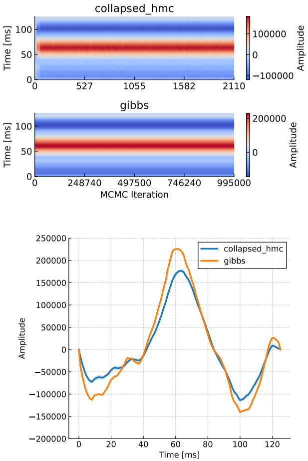

(a) 330x10ms.

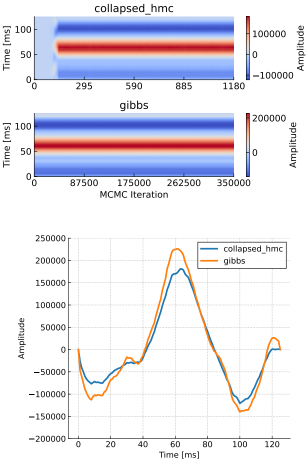

(a) 330x50ms.

Preliminary wavelet estimation results

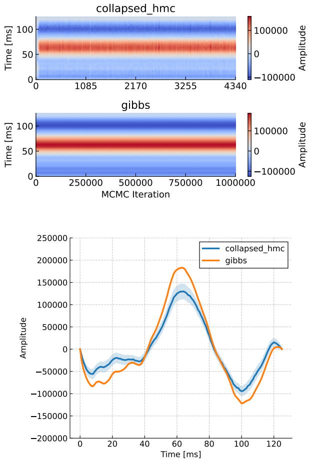

(a) Well trace.

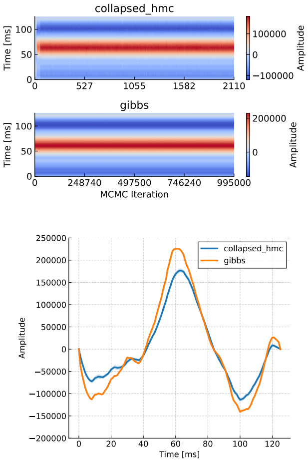

(a) 330x10ms.

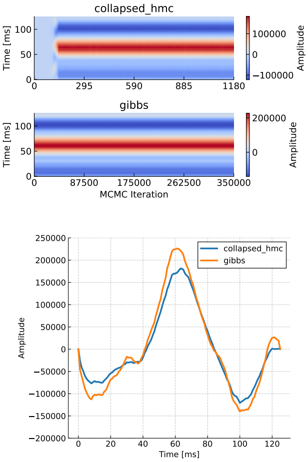

(a) 330x50ms.

Model fit

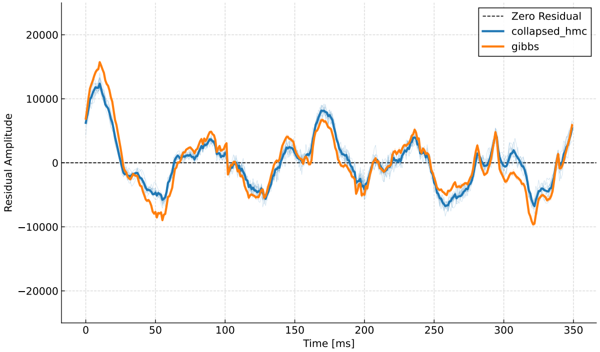

(a) Well trace. RMSE: 4340 (G) vs 3611 (HMC).

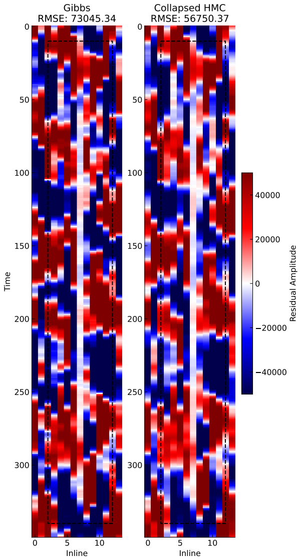

(a) 330x10ms. RMSE: 73000 (G) vs 57000 (HMC).

Conclusions

- Gibbs and HMC might be exploring different regions of the parameter space.
- Including data outside the well changes the wavelet estimation.

References

- Buland, A. and Omre, H. (2003). Bayesian wavelet estimation from seismic and well data. *Geophysics*, 68, 2000–2009.
- Buland, A. and Omre, H. (2003). Joint AVO inversion, wavelet estimation and noise-level estimation using a spatially coupled hierarchical Bayesian model. *Geophysical Prospecting*, 51(6).
- Senn, G., Walker, M. and Tjelmeland, H. (2025). Scalable Bayesian seismic wavelet estimation. *Geophysical Prospecting*, 73(5), 1635–1650.

Conclusions

- Gibbs and HMC might be exploring different regions of the parameter space.
- Including data outside the well changes the wavelet estimation.

References

- Buland, A. and Omre, H. (2003). Bayesian wavelet estimation from seismic and well data. *Geophysics*, 68, 2000–2009.
- Buland, A. and Omre, H. (2003). Joint AVO inversion, wavelet estimation and noise-level estimation using a spatially coupled hierarchical Bayesian model. *Geophysical Prospecting*, 51(6).
- Senn, G., Walker, M. and Tjelmeland, H. (2025). Scalable Bayesian seismic wavelet estimation. *Geophysical Prospecting*, 73(5), 1635–1650.

Conclusions

- Gibbs and HMC might be exploring different regions of the parameter space.
- Including data outside the well changes the wavelet estimation.

References

- Buland, A. and Omre, H. (2003). Bayesian wavelet estimation from seismic and well data. *Geophysics*, 68, 2000–2009.
- Buland, A. and Omre, H. (2003). Joint AVO inversion, wavelet estimation and noise-level estimation using a spatially coupled hierarchical Bayesian model. *Geophysical Prospecting*, 51(6).
- Senn, G., Walker, M. and Tjelmeland, H. (2025). Scalable Bayesian seismic wavelet estimation. *Geophysical Prospecting*, 73(5), 1635–1650.