

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Using optimal transport to mitigate cycle-skipping in ultrasound computed tomography

Christian Boehm, Lion Krischer, Ines Ulrich, Patrick Marty, Michael Afanasiev, et al.

Christian Boehm, Lion Krischer, Ines Ulrich, Patrick Marty, Michael Afanasiev, Andreas Fichtner, "Using optimal transport to mitigate cycle-skipping in ultrasound computed tomography," Proc. SPIE 12038, Medical Imaging 2022: Ultrasonic Imaging and Tomography, 1203809 (4 April 2022); doi: 10.1117/12.2605894

SPIE.

Event: SPIE Medical Imaging, 2022, San Diego, California, United States

Using optimal transport to mitigate cycle-skipping in ultrasound computed tomography

Christian Boehm^a, Lion Krischer^b, Ines Ulrich^a, Patrick Marty^a, Michael Afanasiev^b, and Andreas Fichtner^a

^aETH Zurich, Switzerland

^bMondaic Ltd., Zurich, Switzerland

ABSTRACT

Full-waveform inversion (FWI) for ultrasound computed tomography is an advanced method to provide quantitative and high-resolution images of tissue properties. Two main reasons hindering the widespread adoption of FWI in clinical practice are (1) its high computational cost and (2) the requirement of a good initial model to mitigate the non-convexity of the inverse problem. The latter is commonly referred to as “cycle-skipping”, which occurs for phase differences between synthetic and observed signals and usually traps the inversion in a local minimum. Source-encoding strategies, which simultaneously activate several emitters and have been proposed to reduce the simulation cost, further contribute to this issue due to the multiple arrivals of the wavefronts.

We present a time-domain acoustic full-waveform inversion strategy utilizing a recently proposed misfit functional based on optimal transport. Using a graph-space formulation, the discrepancy between simulated and observed signals can be computed efficiently by solving an auxiliary linear program. This approach alleviates the common need for either a good initial model and / or low-frequency data. Furthermore, combining this misfit functional with random source-encoding and a stochastic trust-region method significantly reduces the computational cost per FWI iteration.

In-silico examples using a numerical phantom for breast screening ultrasound tomography demonstrate the ability of the proposed inversion strategy to converge to the ground truth even when starting from a weak prior and cycle-skipped data.

Keywords: Optimal transport, ultrasound computed tomography, full-waveform inversion, cycle-skipping

1. INTRODUCTION

Full-waveform inversion (FWI) is a powerful imaging technique providing quantitative models of tissue properties in 2D slices or 3D volumes of the human body. Modeling the entire space-time evolution of the ultrasonic wavefield enables FWI to infer high-resolution images from both transmission and reflection data. Previous works on several applications including ultrasound computed tomography (USCT) for soft tissue,¹⁻³ transcranial ultrasound^{4,5} and photoacoustic imaging⁶ have demonstrated the ability of FWI to precisely reconstruct properties such as sound speed, density, or attenuation. However, challenges related to the computational cost and the non-convexity of FWI need to be overcome in order to pave the way toward widespread adoption of FWI in clinical systems.

Compared to geophysical data, ultrasound transducers typically operate within a narrow bandwidth, and striking a balance between lower and higher frequencies is challenging. On the one hand, high frequencies are desirable to improve the expected spatial resolution of the reconstruction, but the computational cost of the simulations increases substantially. On the other hand, low frequency contents in the waveforms are crucial to extend the basin of attraction around the initial model and avoid getting trapped in a local minimum.⁷

It is well-known that the conventionally used least-squares misfit between observed and simulated waveforms has only a very narrow range in which it is locally convex with respect to phase shifts of the signals. This

Further author information: (Send correspondence to C.B.)

C.B.: E-mail: christian.boehm@erdw.ethz.ch, Telephone: +41 44 633 3332

issue – commonly referred to as “cycle-skipping” – necessitates an accurate prior model or the inversion will fail. Other misfits such as cross-correlations⁸ are robust with respect to time shifts in the signal, but require identifying distinct phases, which would rely on careful picking of arrival times from the A-scans. Time- and frequency-dependent phase misfits⁹ are routinely used in continental-scale seismology, but limited to small phase shifts in pre-selected time windows. In addition to adaptive full-waveform inversion^{10,11} using the Wasserstein metric as a misfit functional has recently gained substantial attention in seismic imaging, because of its ability to mitigate cycle-skipping.^{12,13}

In this contribution, we adapt the graph-spaced approach proposed by Métivier et al.¹³ to ultrasound computed tomography and demonstrate its ability to recover models close to the ground truth when starting from a poor prior model or in the absence of low-frequency data for in-silico examples using numerical phantoms. This is particularly interesting when combined with source-encoding strategies^{14,15} to simultaneously simulate the superposition of the pressure field for multiple emitters.

2. METHODS

2.1 Full-waveform inversion with encoded sources

We consider time-domain (visco-)acoustic wave propagation, where the tissue properties $\mathbf{m} = (c, \rho, Q)$ are parameterized in terms of space-dependent sound speed c , density ρ and quality factor Q . The wave equation can then compactly be written as

$$L(\mathbf{u}, \mathbf{m}) = \mathbf{f}, \quad (1)$$

with initial conditions

$$\mathbf{u}_{t=0} = 0, \quad \partial_t \mathbf{u}_{t=0} = 0, \quad (2)$$

and where L denotes the wave operator, \mathbf{u} is a scalar displacement potential and \mathbf{f} is the energy injected by an external source. For fixed \mathbf{m} and \mathbf{f} , the solution $\mathbf{u} = \mathbf{u}(\mathbf{m}; \mathbf{f})$ to the wave equation is unique.¹⁶ While the parameter-to-solution map $\mathbf{m} \mapsto \mathbf{u}(\mathbf{m})$ is non-linear, the dependence on the right-hand-side is linear, i.e., for fixed \mathbf{m} the solution \mathbf{u} depends linearly on \mathbf{f} . A typical USCT dataset consists of a collection of A-scans recorded by a multitude of transducers distributed across the sidewalls of a scanning device. Emitting transducers are activated in sequence producing scans for each transducer location. FWI can be formulated as a non-linear minimization problem governed by a vector of wave equations with right-hand-sides \mathbf{f}_i for each emitter,

$$\min_{\mathbf{m}} \sum_{i=1}^{n_s} \chi(\mathbf{u}_i(\mathbf{m}); \mathbf{u}_i^{\text{obs}}) \quad \text{subject to} \quad L(\mathbf{u}_i, \mathbf{m}) = \mathbf{f}_i, \quad i = 1 \dots, n_s, \quad (3)$$

where $\mathbf{u}_i^{\text{obs}}$ denotes the observed data for emitter i , and χ is a misfit functional to compare synthetic and observed ultrasound scans. When using checkpointing to reduce the memory burden of time-domain simulations of the wave equation, each iteration of FWI requires at least three simulations per emitter to evaluate both the misfit functional χ , and its gradient with respect to changes in the tissue material. The latter is the outcome of a so-called adjoint simulation in conjunction with a check-pointed forward simulation. The setup in eq. (3) using a wave equation with a different source term for each emitter thus results in $3n_s$ simulations per iteration for n_s independent sources. Although the simulations for different emitters are embarrassingly parallel, reducing the computational cost is a key target to make FWI more widespread affordable.

Source-encoding^{14,15,17} exploits the linearity of the wave equation with respect to the source term, which enables us to simulate the superposition of multiple emitters and comparing those signals to the stacked recordings. Using weights $\beta = (\beta_i)_i$ for each emitter, we obtain

$$\mathbf{u}_\beta = \sum_{i=1}^{n_s} \beta_i \mathbf{u}_i \quad \text{for} \quad L(\mathbf{u}_\beta, \mathbf{m}) = \sum_{i=1}^{n_s} \beta_i \mathbf{f}_i \quad \text{and} \quad L(\mathbf{u}_i, \mathbf{m}) = \mathbf{f}_i. \quad (4)$$

A common strategy is to (1) only activate a subset of emitters in each encoded sources, i.e., $\beta_i = 0$ for inactive emitters, (2) consider a small number of encoded sources in the FWI objective, eq. (3), using different weights, and (3) choose the non-zero weights from $\{-1, 1\}$. Additionally, the onset time of the source injection can vary,

assuming that either data recordings are available for a longer time interval or padding with zeros is acceptable because of attenuation.

In the FWI algorithm, we combine source-encoding with a stochastic trust-region method¹⁸ that extracts curvature information from L-BFGS with a diffusion-based Hessian initialization. Here, the initial guess of the Hessian operator is implicitly obtained by solving diffusion equations for each parameter field in \mathbf{m} as initial values

$$\partial_t q(\mathbf{x}, t) - \nabla \cdot (\mathcal{D}(\mathbf{x}) \nabla q(\mathbf{x}, t)) = 0, \quad (5)$$

with a symmetric second-order diffusion tensor \mathcal{D} . This can be related to a Wiener process using

$$\mathcal{D}(\mathbf{x}) = \frac{1}{2} \sigma(\mathbf{x}) \sigma(\mathbf{x})^T, \quad (6)$$

with direction- and space-dependent smoothing lengths σ_i . Incorporating the smoothing operator into the L-BFGS approximation is an effective alternative (or addition) to explicit regularization terms in the objective. The smoothing lengths can be chosen proportional to the wavelength to mitigate small-scale artifacts in the gradients which are frequently obtained in FWI.

Each batch consists of multiple randomly encoded sources, using random weights for the source amplitude and onset time. The initial batch selection and time shifts of the encoded sources are determined by an auxiliary optimization problem using approximate time-of-flights of the prior to maximize coverage of the encoded wavefield. This is particularly important for ring-shaped apertures (cf. section 3) to avoid focusing of the encoded wavefield at the center of the ring.

The main challenges of source-encoding are twofold. First, interference in the stacked wavefields leads to a loss of information. This can be mitigated by modifying the encoding weights in every iteration. Second, the stacked traces contain multiple primary and secondary arrivals from all emitting locations. This contributes to the issue of cycle-skipping and increases the risk of converging toward a local minimum, which is why a careful choice of the misfit functional is crucial.

2.2 Graph-space optimal transport misfit

To address the issue of cycle-skipping, we consider a misfit functional based on optimal transport, which is closely related to the Wasserstein distance.^{12,13} It has been observed that this metric is convex with respect to time-shifted patterns in the signal, and optimal transport has been applied successfully in many image processing applications. Because A-scans have positive and negative amplitudes and the total energy contained in synthetic and observed traces can differ, some adjustments are required to apply optimal transport as a misfit functional in FWI. Following the work of Métivier et al.,¹³ we consider an approach that is based on matching point clouds originating from the discretized signals. To this end we consider two time series $v, w : [0, T] \rightarrow \mathbb{R}$ defined on the same time interval $[0, T]$, as well as their graphs $G(v), G(w)$, where the graph G of a time series s is defined as:

$$G(s) := \{(t, s(t)), t \in [0, T]\}. \quad (7)$$

After discretizing v and w on the same computational grid using n_t equidistantly-spaced time samples, we obtain two point clouds $\mathbf{v}, \mathbf{w} \in \mathbb{R}^{n_t \times 2}$ approximating the graphs, where the i -th row is given by $\mathbf{v}_i = (t_i, v(t_i))$ and $\mathbf{w}_i = (t_i, w(t_i))$, respectively. The Wasserstein metric applied to \mathbf{v} and \mathbf{w} will then compute the space-time distance between both point clouds. This graph-space approach results in an auxiliary optimization problem to compute the misfit, which can be reformulated as a Linear Sum Assignment Problem (LSAP).¹³ Here, each

tuple $(\mathbf{v}_i, \mathbf{w}_j)$ is assigned a weight d_{ij} which encodes the mapping of both points in the graph-space.

$$\begin{aligned} \min_{d_{ij}} \quad & \sum_{i,j=1}^{n_t} d_{ij} \|\mathbf{v}_i - \mathbf{w}_j\|_2^2 \\ \text{subject to} \quad & d_{ij} \geq 0, \quad \forall i, j = 1, \dots, n_t, \\ & \sum_{i=1}^{n_t} d_{ij} = 1, \quad \forall j = 1, \dots, n_t, \\ & \sum_{j=1}^{n_t} d_{ij} = 1, \quad \forall i = 1, \dots, n_t. \end{aligned} \quad (8)$$

Problems of this structure are well studied, and it is known that the continuous relaxation of the LSAP has an integer solution, which means that efficient solvers for linear programs can be applied to compute the misfit. In particular, the solution of eq. (8) can be represented as a permutation σ^* of the index set $\{1, \dots, n_t\}$, which encodes the matching between the two point clouds. Hence, the misfit value χ is obtained from

$$\chi(\mathbf{v}) = \sum_{i=1}^{n_t} \|\mathbf{v}_i - \mathbf{w}_{\sigma^*(i)}\|_2^2. \quad (9)$$

Furthermore, Theorem 1 in Métivier et al.¹³ shows that the corresponding adjoint source is given by

$$\nabla \chi(\mathbf{v}) = 2(\mathbf{v} - \mathbf{w}_{\sigma^*}), \quad \text{where } \mathbf{w}_{\sigma^*} = (\mathbf{w}_{\sigma^*(i)})_{i=1, \dots, n_t}. \quad (10)$$

Evaluating the misfit and computing the adjoint source for optimal transport is significantly more expensive than their counterparts in conventional misfit functionals such as the L^2 -norm. Note that the distance matrix in eq. (8) is dense, although the maximum expected time shift can be used to prune portions of the graph. We solve the LSAP independently for each trace using `scipy`'s¹⁹ implementation of the Hungarian algorithm.²⁰ For computational efficiency, the signals are downsampled before extracting the discrete point clouds. Afterwards, the adjoint sources are upsampled again to the original time step using Lanczos interpolation. The error in the resulting gradients is negligible and does not slow down convergence.




Figure 1. Example of graph-space optimal transport for two time series. Synthetic and observed data are significantly different and subject to cycle-skipping. The connecting lines show the permutation of the observed signal that is identified as the optimal transport for matching the point cloud of the synthetic signal. The color indicates the associated cost of the assignment in the time-amplitude domain.

Fig. 1 illustrates the point cloud matching for two synthetic time series. The weighting of amplitude difference vs. distance in time gives an important tuning parameter of the graph-space approach. Emphasizing amplitude mismatches, i.e., assigning higher costs in the distance matrix to differences in time, yields an optimal permutation that mimics the L^2 -norm misfit. Vice versa, assigning higher costs to amplitude differences yields an optimal

permutation that moves energy along the time axis, and therefore shows similar behavior as a misfit based on cross-correlations of the signals. We can select this tuning parameter based on the maximum expected time-of-flight difference between synthetic and observed signals to ensure a sufficient basin of attraction even for poor starting models in the inversion. This effect of convexifying the misfit functional with respect to time shifts is shown in Fig. 2. Here, we consider the source wavelet of the dataset provided by Li et al.²¹ and compute the misfit value as a function of time shift in the wavelet. Different choices of the tuning parameter show the transition of the misfit shape between L^2 -norm and cross-correlation time shifts. Note that a misfit based on cross-correlation time-shifts only appears convex in this idealized case of two identical, but time-shifted signals. It would not work for more complex signals including multiple arrivals, whereas optimal transport is able to simultaneously match both direct arrivals as well as reflection and refraction data.

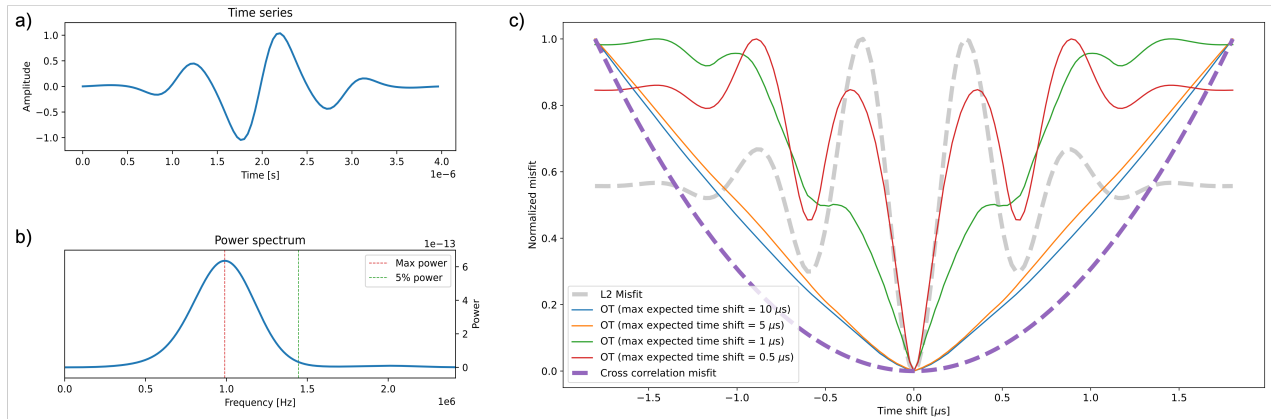


Figure 2. Optimal transport misfit for two time shifted signals. (a) Wavelet considered in this test. (b) Power spectrum of the wavelet. (c) Comparison of the optimal transport misfit for time shifted versions of the wavelet shown in (a) and different choices of the tuning parameter. Optimal transport widens the locally convex basin of attraction to mitigate cycle-skipping, while being able to maintain a steep slope for signals with small phase shifts.

3. RESULTS

We present an in-silico example for ultrasound computed tomography based on a numerical breast phantom.^{21,22} The example considers a coronal slice of a numerical breast phantom provided by Li et al.²¹ with reference *536035891B*. The aperture consists of a ring with 1024 equi-angularly spaced transducers that all act in emitting and receiving mode. Simulated data are generated using a Ricker wavelet with a dominant frequency of 1 MHz as source time function and without attenuation. Each encoded source activates 32 transducers with random spatial weights of ± 1 . The individual emitters fire at different starting times with a maximum time shift of 30 μ s. The simulation interval for each encoded source is 200 μ s.

The initial model is homogeneous with a sound speed of 1500 m/s, and density equal to 1000 kg/m³. The misfit only considers recordings starting at 100 μ s, thus partially excluding arrivals of the encoded sources that have not interacted with the phantom. The maximum expected time shift to tune the optimal transport distance is initially set to 5 μ s, and all traces are downsampled to 14 MHz before solving the LSAP. Fig. 3 shows an example of source-encoded signals for the true phantom and the homogeneous initial model. While the direct arrivals from near emitters match well, all secondary arrivals resulting from reflected or refracted waves are missing in the synthetic trace of the initial model, and we observe severe cycle-skipping toward the end of the observation interval.

A comparison of the ground truth and different inversion strategies are shown in Fig. 4. The sound speed map obtained from optimal transport includes major features and is in good agreement with the true model. However, small-scale details are missing and the outlines of the tissue are blurred. The reconstruction can be further improved by adaptively reducing the tuning parameter and smoothing lengths during the iterations to account for the improved waveform fit and smaller time shifts between the measurements and the improved synthetics.




Figure 3. Example of graph-space point cloud matching for an encoded source recording. Missing arrivals in the synthetic data and cycle-skipping of later arrivals increase the risk of converging to the wrong model.

Additionally, this strategy accelerates convergence, as the optimal transport misfit functional approaches the L^2 -norm. The density model (not shown here) exhibits sharper contours but underestimates the magnitude of the perturbations. For comparison, Fig. 4(f) depicts the inversion outcome using the conventional L^2 -norm misfit on the same data selection. Because of the poor initial model, the iterations quickly stall in a local minimum.

Fig. 5 shows the data fit for the initial model and the reconstruction using an encoded source and the signals at four different receiver locations. Even though the initial model is not capturing many arrivals and some parts of the signals are subject to heavy cycle-skipping, the recovered model produces traces that match remarkably well. Fig. 6 compares signal gathers for a single emitter generated with the ground truth, the initial model and the reconstructed model shown in Fig. 4(d), respectively. Note that those data have not been used for the inversion, which relied exclusively on randomly encoded sources. It thus serves as a validation check of the inversion. The initial model exhibits cycle-skipping in the first arrivals and misses all reflections and secondary arrivals. Although the fine-scale structure of the ground truth contains a variety of scattered and refracted waves, the main features of the recorded data are captured by the reconstructed model.

4. DISCUSSION

Using optimal transport as misfit functional significantly improves the local convexity of the inverse problem and therefore avoids cycle-skipping. The numerical results demonstrate convergence toward the ground truth even from a homogeneous background, very weak prior knowledge, and significantly different waveforms.

The price for the gain in robustness is twofold. First, computing the misfit value requires the solution of the auxiliary linear sum assignment problem, which is significantly more expensive compared to conventional misfit measures such as the L^2 -norm. While this overhead is not necessarily relevant as the computational cost are typically dominated by forward and adjoint simulations, it might be an issue when considering large point clouds; either from long time series or when considering several receivers simultaneously. The algorithm proposed by Guthe & Thuerck²³ is an interesting alternative for future research. The second cost increase is due to the observation that optimal transport requires more iterations than the L^2 misfit to converge in our numerical tests. Fig. 2 demonstrates that the maximum expected time shift between synthetic and observed waveforms can be used to balance amplitude and time shifts. This parameter can be adjusted adaptively during the inversion to speed up convergence. In particular, as time-of-flight differences between synthetic and observed data decrease during the inversion, the parameter can be decreased to allow for a smooth transition into the L^2 -norm misfit.

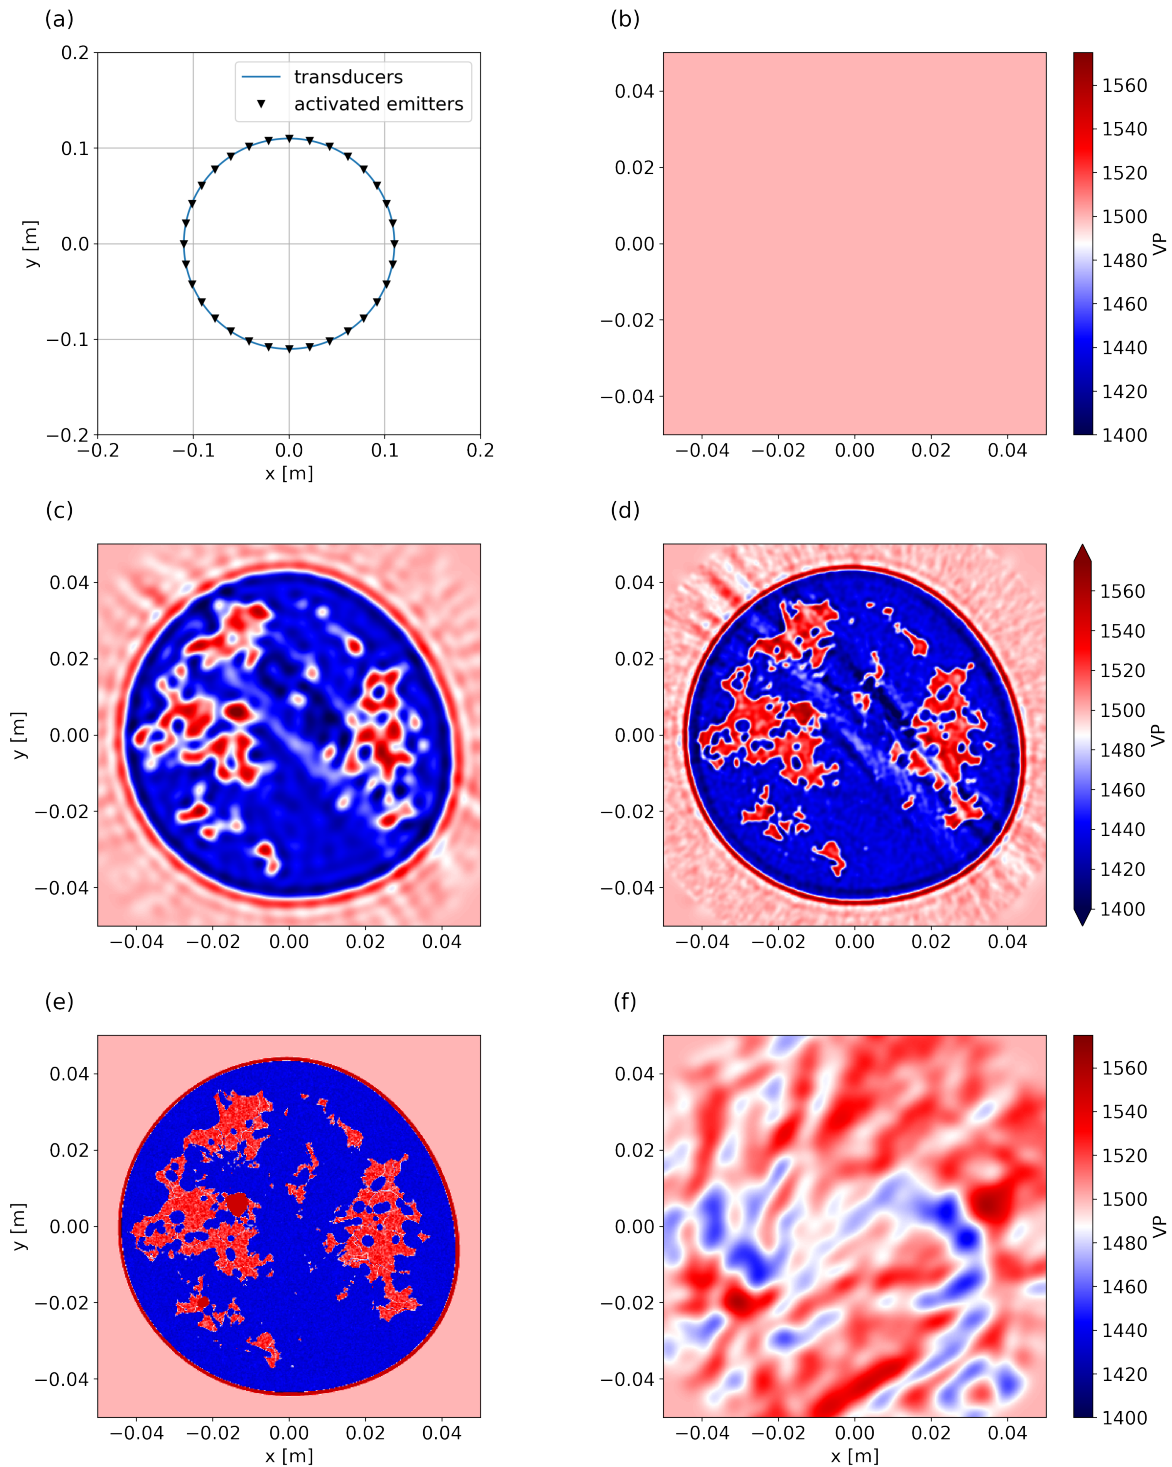


Figure 4. 2D FWI of a breast phantom using optimal transport. All plots show the speed of sound. (a) Synthetic aperture used for the inversion. As an example, the black triangles mark the emitters activated for a single encoded source. (b) Homogeneous initial model. (c) Reconstruction using optimal transport and a fixed tuning parameter of $5\ \mu\text{s}$ and smoothing lengths of 2 mm. (d) Reconstruction using optimal transport and an adaptive tuning parameter which is reduced to $0.5\ \mu\text{s}$ and smoothing lengths of 0.5 mm. (e) True phantom used for simulated measurements. (f) Reconstruction using the L^2 -norm. The iterations got trapped in a local minimum because of cycle-skipping.




Figure 5. Waveform fit for an encoded source with 32 emitting locations using either the initial model (left) or the final model of the inversion (right). The four transducer shown here are 90 degrees apart from each other.

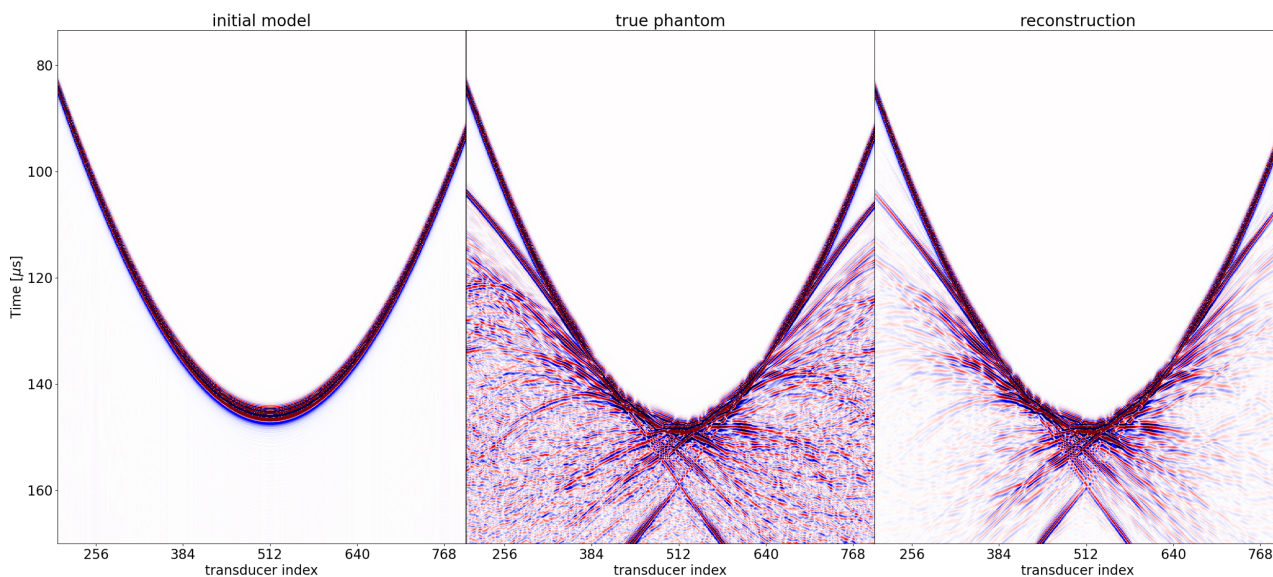


Figure 6. A comparison between the shot gathers for a single emitter and a ring of transducers in a plane with the emitter. The plots show A-scan data simulated for the initial model (left), the ground truth (middle), as well as the final model (right). The homogeneous initial model leads to cycle-skipped primary waves and is missing all secondary arrivals. Synthetic data retrieved from the FWI reconstruction with optimal transport are free of cycle skips and contain the later arrivals. The colorbar has been clipped to 10% of the maximum amplitude.

The initial models used in the numerical example could easily be improved at low computational cost prior to FWI, for instance, by applying a ray-based time-of-flight inversion or by obtaining the shape of the phantom from reflection imaging. It is important to note, however, that the graph-spaced optimal transport misfit used here is able to converge to the ground truth even from a very poor initial guess. This gain in robustness and reduced risk of converging to a local minimum is the primary purpose of the proposed inversion strategy.

The reconstruction quality would potentially improve by including an explicit regularization term in the objective. To this end, an adaptive tuning strategy using total variation regularization is currently under investigation.

5. CONCLUSIONS

We present an inversion method for acoustic full-waveform inversion using a graph-space formulation of optimal transport to quantify waveform mismatches. The numerical results indicate that this misfit functional mitigates cycle-skipping introduced by the superposition of several wavefields when using encoded sources, or in the absence of low frequencies in the measured ultrasound data. This helps to make FWI significantly more robust with respect to the choice of the initial model and alleviates the need for a good prior. The proposed strategy seems very promising and could expand the applicability of FWI to challenging datasets in USCT and photoacoustic imaging.

ACKNOWLEDGMENTS

This work was supported by the Swiss National Supercomputing Centre (CSCS) under project ID sm59.

REFERENCES

- [1] Pratt, R. G., Huang, L., Duric, N., and Littrup, P., “Sound-speed and attenuation imaging of breast tissue using waveform tomography of transmission ultrasound data,” in [*Medical Imaging 2007: Physics of Medical Imaging*], **6510**, 65104S, International Society for Optics and Photonics (Mar. 2007).
- [2] Pérez-Liva, M., Herraiz, J. L., Udías, J. M., Miller, E., Cox, B. T., and Treeby, B. E., “Time domain reconstruction of sound speed and attenuation in ultrasound computed tomography using full wave inversion,” *The Journal of the Acoustical Society of America* **141**(3), 1595–1604 (2017).
- [3] Wiskin, J., Malik, B., Pirshafiey, N., and Klock, J., “Limited view reconstructions with transmission ultrasound tomography: clinical implications and quantitative accuracy,” in [*Medical Imaging 2020: Ultrasonic Imaging and Tomography*], **11319**, 113190V, International Society for Optics and Photonics (Mar. 2020).
- [4] Guasch, L., Calderón Agudo, O., Tang, M.-X., Nachev, P., and Warner, M., “Full-waveform inversion imaging of the human brain,” *npj Digital Medicine* **3**, 28 (Feb. 2020).
- [5] Marty, P., Boehm, C., and Fichtner, A., “Acoustoelastic full-waveform inversion for transcranial ultrasound computed tomography,” in [*Medical Imaging 2021: Ultrasonic Imaging and Tomography*], Byram, B. C. and Ruiter, N. V., eds., **11602**, 210 – 229, International Society for Optics and Photonics, SPIE (2021).
- [6] Mitsuhashi, K., Poudel, J., Matthews, T. P., Garcia-Urbe, A., Wang, L. V., and Anastasio, M. A., “A Forward-Adjoint Operator Pair Based on the Elastic Wave Equation for Use in Transcranial Photoacoustic Computed Tomography,” *SIAM Journal on Imaging Sciences* **10**, 2022–2048 (Jan. 2017).
- [7] Symes, W. W., Chen, H., and E.Minkoff, S., [*Full-waveform inversion by source extension: Why it works*], 765–769, SEG (2020).
- [8] Korta Martiartu, N., Boehm, C., and Fichtner, A., “3-D Wave-Equation-Based Finite-Frequency Tomography for Ultrasound Computed Tomography,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **67**(7), 1332–1343 (2020).
- [9] Kristeková, M., Kristek, J., and Moczo, P., “Time-frequency misfit and goodness-of-fit criteria for quantitative comparison of time signals,” *Geophysical Journal International* **178**, 813–825 (08 2009).
- [10] Warner, M. and Guasch, L., “Adaptive waveform inversion: Theory,” *GEOPHYSICS* **81**(6), R429–R445 (2016).

- [11] Calderon Agudo, O., Stronge, G., Robins, T., and Guasch, L., “Overcoming cycle-skipping in full-waveform inversion of ultrasound data,” in [*Proceedings of the International Workshop on Medical Ultrasound Tomography: 14.-15. Oct. 2019, Wayne State University, Detroit, Michigan, USA*], Boehm, C., Hopp, T., Ruiter, N., and Duric, N., eds., 137–149, KIT Scientific Publishing, Karlsruhe (2021).
- [12] Engquist, B. and Froese, B. D., “Application of the wasserstein metric to seismic signals,” *Communications in Mathematical Sciences* **12**(5), 979–988 (2014).
- [13] Métivier, L., Brossier, R., Méridot, Q., and Oudet, E., “A graph space optimal transport distance as a generalization of l^p distances: application to a seismic imaging inverse problem,” *Inverse Problems* **35**, 085001 (jul 2019).
- [14] Krebs, J. R., Anderson, J. E., Hinkley, D., Neelamani, R., Lee, S., Baumstein, A., and Lacasse, M.-D., “Fast full-wavefield seismic inversion using encoded sources,” *GEOPHYSICS* **74**(6), WCC177–WCC188 (2009).
- [15] Wang, K., Matthews, T., Anis, F., Li, C., Duric, N., and Anastasio, M. A., “Waveform inversion with source encoding for breast sound speed reconstruction in ultrasound computed tomography,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **62**(3), 475–493 (2015).
- [16] Boehm, C. and Ulbrich, M., “A semismooth newton-cg method for constrained parameter identification in seismic tomography,” *SIAM Journal on Scientific Computing* **37**(5), S334–S364 (2015).
- [17] Moghaddam, P. P., Keers, H., Herrmann, F. J., and Mulder, W. A., “A new optimization approach for source-encoding full-waveform inversion,” *GEOPHYSICS* **78**(3), R125–R132 (2013).
- [18] Boehm, C., Martiartu, N. K., Vinard, N., Balic, I. J., and Fichtner, A., “Time-domain spectral-element ultrasound waveform tomography using a stochastic quasi-Newton method,” in [*Medical Imaging 2018: Ultrasonic Imaging and Tomography*], Duric, N. and Byram, B. C., eds., **10580**, 92 – 100, International Society for Optics and Photonics, SPIE (2018).
- [19] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., Carey, C. J., Polat, İ., Feng, Y., Moore, E. W., VanderPlas, J., Laxalde, D., Perktold, J., Cimrman, R., Henriksen, I., Quintero, E. A., Harris, C. R., Archibald, A. M., Ribeiro, A. H., Pedregosa, F., van Mulbregt, P., and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods* **17**, 261–272 (2020).
- [20] Kuhn, H. W., “The Hungarian method for the assignment problem,” *Naval Research Logistics Quarterly* **2**(1-2), 83–97 (1955).
- [21] Li, F., Villa, U., Park, S., and Anastasio, M., “2D Acoustic Numerical Breast Phantoms and USCT Measurement Data,” (2021).
- [22] Li, F., Villa, U., Park, S., and Anastasio, M. A., “3-D Stochastic Numerical Breast Phantoms for Enabling Virtual Imaging Trials of Ultrasound Computed Tomography,” *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control* **69**(1), 135–146 (2022).
- [23] Guthe, S. and Thuerck, D., “Algorithm 1015: A fast scalable solver for the dense linear (sum) assignment problem,” *ACM Trans. Math. Softw.* **47** (apr 2021).