Cross-gradient regularization for multiparameter FWI of physically modeled data

Scott Keating, Kris Innanen and Joe Wong

ABSTRACT

In 2022 an in-plane, surround acquisition ultrasonic survey was acquired in the physical modelling lab. Buzzer sources, operating at the low end of the ultrasonic spectrum were employed. These data were inverted using a variable density acoustic, frequency domain full waveform inversion. Results demonstrated relatively robust capability to characterize the longer wavelength features of target cross-sections. Shorter wavelength features were more difficult to recover. This effect is likely linked to the lack of complete reflection measurements in the data, which may have impeded accurate characterization of reflection behaviour. A cross-gradient regularization was able to enforce geometric similarity between the density and velocity models recovered by the inversion.

INTRODUCTION

Physically modelled data at the ultrasonic scale can be useful for developing inversion approaches for use on real-data problems in relatively controlled environments, and offers the potential for connections to ultrasonic inversion problems to be drawn. Recently, CREWES has performed several circular, in-plane ultrasonic modelling experiments in the physical modelling lab, in part to investigate the potential for seismic inversion methodologies in medical-like acquisition settings (Wong, 2022; Henley, 2022). In this spirit, the application of full-waveform inversion to these data sets is a current topic of interest.

Water-tank ultrasonic datasets should be dominated primarily by acoustic behaviour, due to the inability of fluids to support shear waves. Tomographic inversion of these datasets, then, can be accomplished by inverting only for a P-wave velocity term, as this is the only significant factor in determining the measured traveltimes. In a full-waveform inversion, however, the amplitudes of measurements also play a key role, and so density becomes a second necessary parameter. Unfortunately, while density contrasts play a key role in determining measured reflection amplitudes, density is not typically well constrained by ultrasonic (or seismic) measurements. In addition, cross-talk with P-wave velocity can introduce errors in density estimation. These considerations generally make density estimation from acoustic measurements a fraught procedure.

Through a regularization penalty term in FWI, we can inform the inversion about prior information independent of the data alone. This type of prior information may help to better constrain density. While case-specific prior information may be useful, in general we do not know exactly how density will vary given a change in velocity: an increase in one may correspond to an increase or decrease of greater or lesser magnitude in the other. In this sense, we do not usually have a strong prior estimate of density as a function of velocity. We do, however, generally know that contrasts in density should occur at the same locations as contrasts in density: regardless of the correlation between the values of the different properties, the correlation between the positions of the contrasts should be strong. For

this reason, we investigate the use of a cross-gradient regularization here (Gallardo and Meju, 2004). This type of regularization is widely used in joint inversion applications, as it promotes geometrical similarity between different inverted quantities. Here, we expect that this regularization will help to constrain the density to reasonable values based on our higher-confidence velocity estimates.

In the remainder of this report, we outline the physically modelled data-set we consider and describe the processing used. We then investigate the effects of the cross-gradient regularization term on our inversion results.

DATA

The data we consider in this report were gathered using low ultrasonic frequency buzzer sources and receivers, as described in Wong (2022). Data were gathered for 72 shots into 72 receivers in the physical modelling tank, using the constant-depth acquisition geometry shown in Figure 1. The data acquisition setup in the lab is fundamentally limited to one receiver per source, so the effective 72 receivers actually result from 72 single-receiver experiments with the same shot location. Regardless, for computational efficiency, we treat the data in the inversion as though a single source was measured by 72 receivers. Provided that the source characteristics are sufficiently repeatable, this should be a reasonable approximation. Each shot was measured several times and stacked in order to boost signal to noise ratio. In our plotting of the data-sets, the leftmost shot is considered shot 1, followed by increasing numbers a clockwise progression through the shots is taken. The same convention is applied for the receivers.

Because of physical limitations in the measurement apparatus, not all receivers were active for every source. Receiver locations to the left (in our plotting convention), or negative x-direction, of the sources were not acquired in the survey. Only the leftmost shot had all receivers active, whereas the rightmost shots only had a few active receivers.

Several sets of surveys were obtained with the same acquisition. The first, which we will refer to as the 'no target' survey, was acquired with only water in the physical modelling tank. The other surveys were acquired with a set of PVC targets in the tank. Two targets were present in the tank for the 'target' dataset. Both targets were parts of PVC cylinders with a radius of 48 mm and a height spanning most of the physical modelling tank. The first target was a cylindrical ring with an inner radius of 39 mm, this was placed in approximately the center of the acquisition. The second target was a semi-cylinder, offset in the +x and +y directions from the center with the direction normal to the flat edge pointing in the +y direction. Surveys of each of these targets individually were also obtained, which we will refer to as the 'Ring' and 'Semi-cylinder' surveys.

PROCESSING

An example shot of the acquired data before processing is shown in Figures 2 and 3 for the leftmost shot in the no-target and target cases, respectively. In these figures, the rightmost receiver is in the center, and receiver numbering increases as we progress clockwise about the receiver array. The shortest sourcer-receiver offsets are then at the left and right edges of the figures, while the longest source-receiver offset is at the center of the figures.

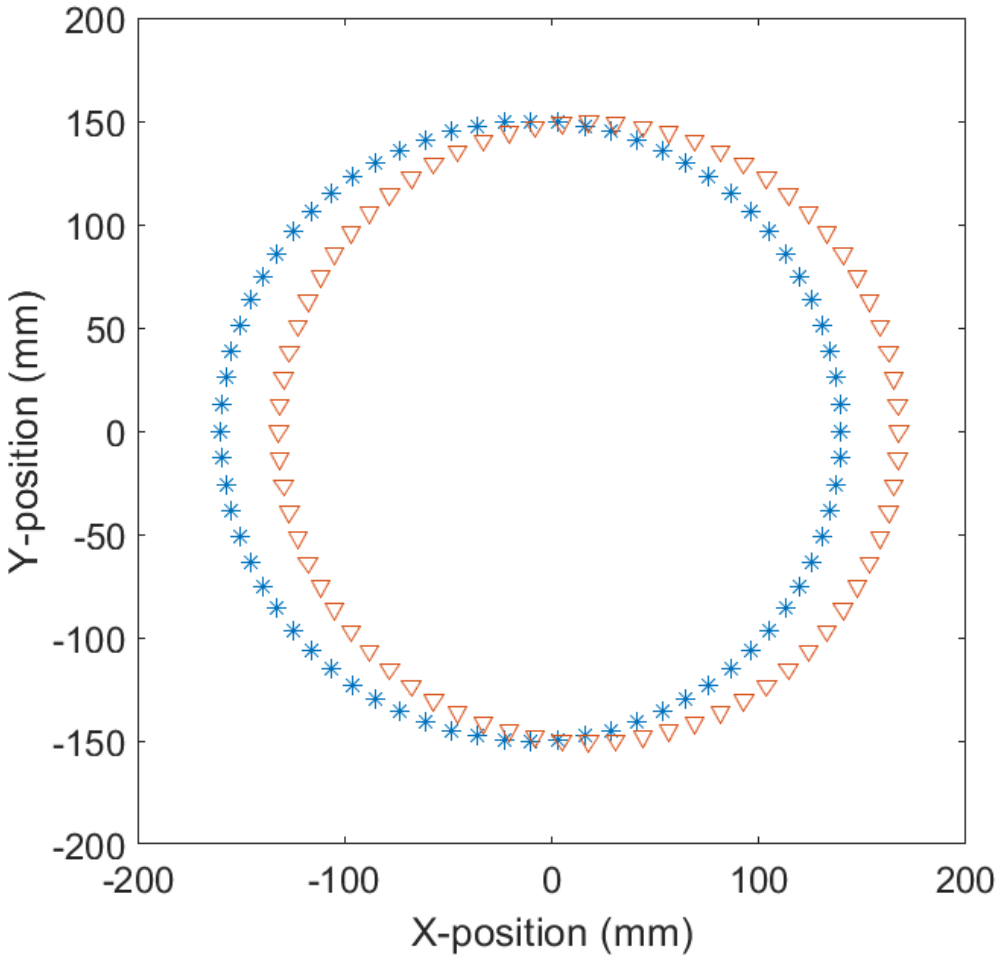

FIG. 1. Locations of sources (stars) and receivers (triangles) relative to the center of the tank.

Two essential processing steps are required before inversion can be performed: amplitude correction and muting. Amplitude correction is necessary for this problem because the physical modelling was performed in three dimensions while the inversion we use makes use of 2D wave propagation, and muting is required to remove unwanted reflections in the data from tank boundaries.

We expect two major contributions to the amplitude differences between our modelling and the data measurements: differences in geometrical spreading (which should be more pronounced in 3D), and source orientation signature. While the former is easily predictable, the latter is not, and in our testing, we found that geometrical spreading alone did not account for amplitude differences. To compensate, we took the amplitude of the envelope of both the physically modelled data in the case with no target (water only), and the corresponding envelope for the synthetic data in a water-only medium, calculated a correction factor as a function of traveltime, and applied it to the measured data to produce pseudo-2D amplitudes.

Muting is necessary in physical modelling experiments due to the finite size of the physical modelling tank: eventually reflections from the tank boundary enter the measurements.

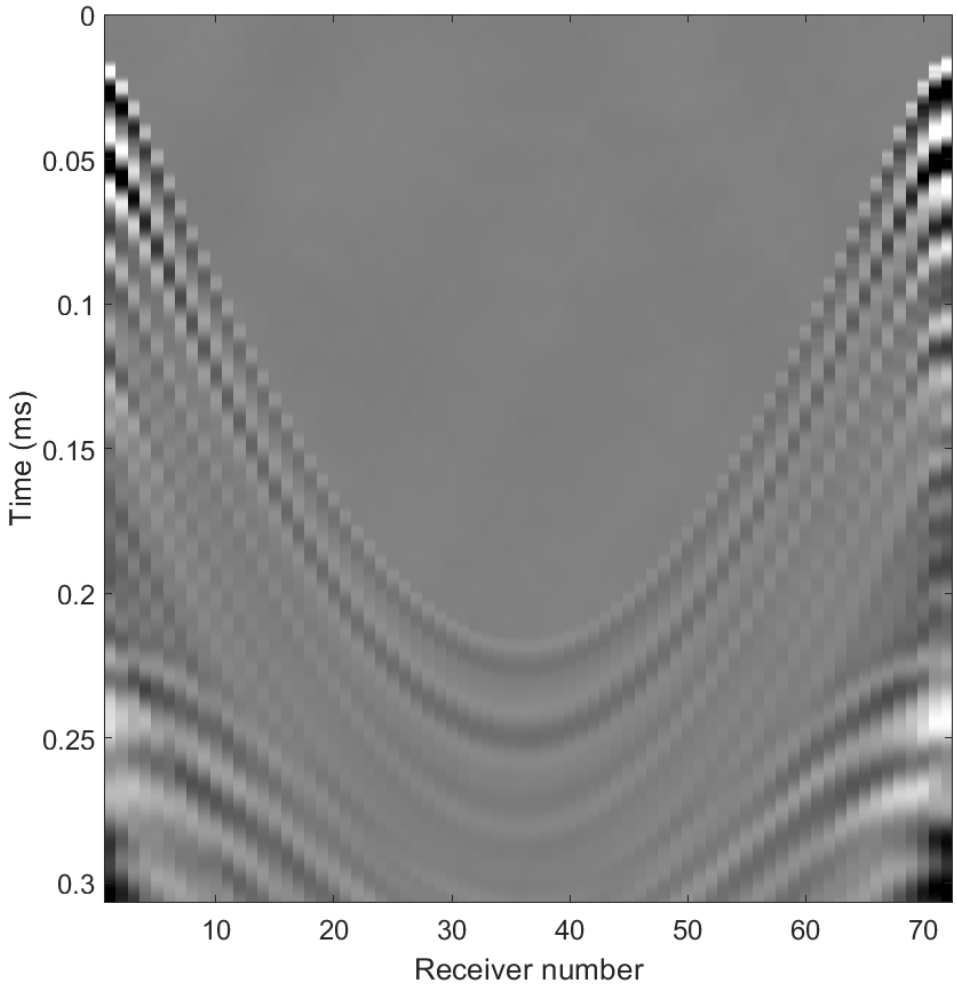

FIG. 2. Shot record for leftmost shot, no target survey.

While this effect could, with difficulty, be accounted for in a 3D inversion scheme, in two dimensions there is no way to accurately account for this unwanted signal. In consequence, the later times of each shot record need to be muted in order to prevent the boundary reflections from contaminating the signal. These reflections arrive sooner at receivers near the source (which are on the same side of the tank) than at those opposite the source (which are on the other side of the tank). The timing of the reflections is such that most of the reflections at near and mid offsets are obscured and so removed by the mute, while those at far offsets arrive before the boundary reflections and are preserved by the mute. This likely has significant implications for features of the inversion results which depend strongly on reflection data.

Figures 4 and Figure 5 show the data after our preliminary processing, both with and without the targets present for the leftmost shot. The difference between the target and no target data (the data residual we hope to invert for) for the both the unprocessed and the processed data sets (again, for the leftmost shot) are shown in Figures 6 and 7. Notably,

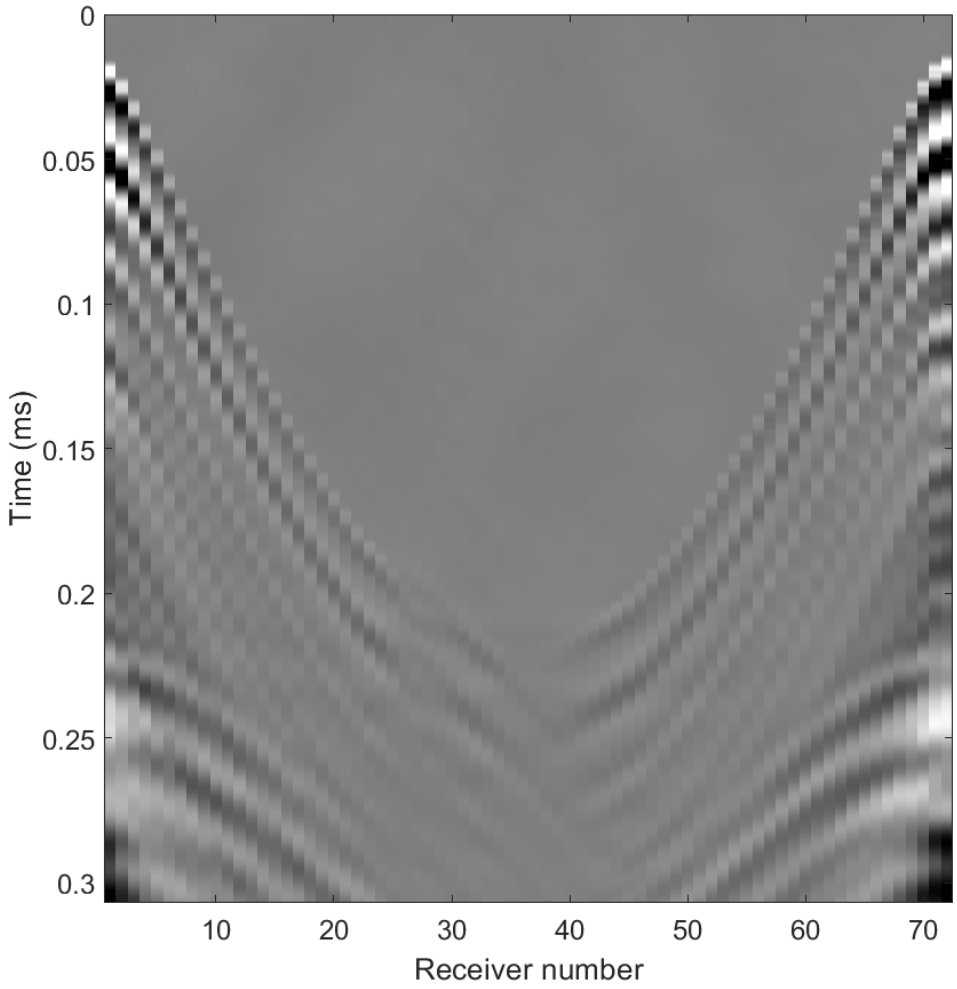

FIG. 3. Shot record for leftmost shot, target survey.

both the tank boundary reflections (about 0.25-0.3 ms on the closest traces) and the direct arrivals (about 0.05-0.1 ms on the closest races) do not subtract perfectly in the unprocessed data. This highlights the need for a bottom mute, and also demonstrates that these measurements are not entirely reproducible: otherwise the direct arrivals at short offset (which are not affected by the targets) should match perfectly. This non-repeatability suggests that it will not be possible, or even desirable, to achieve a very low data misfit with the inversion result, as a substantial part of the measurement is subject to non-target related variation.

INVERSION METHOD

We consider 2D, frequency-domain, variable-density acoustic full-waveform inversion in this report. While synthetic data for a homogeneous water medium are very similar to the measurements for the no-target case, some differences persist. These could arise from incomplete conversion from 3D to 2D amplitudes in the preprocessing, from small errors in the recorded positions and directionality of the sources and receivers, from measurement noise, or from changes in the wavelet between different source-receiver experiments in the

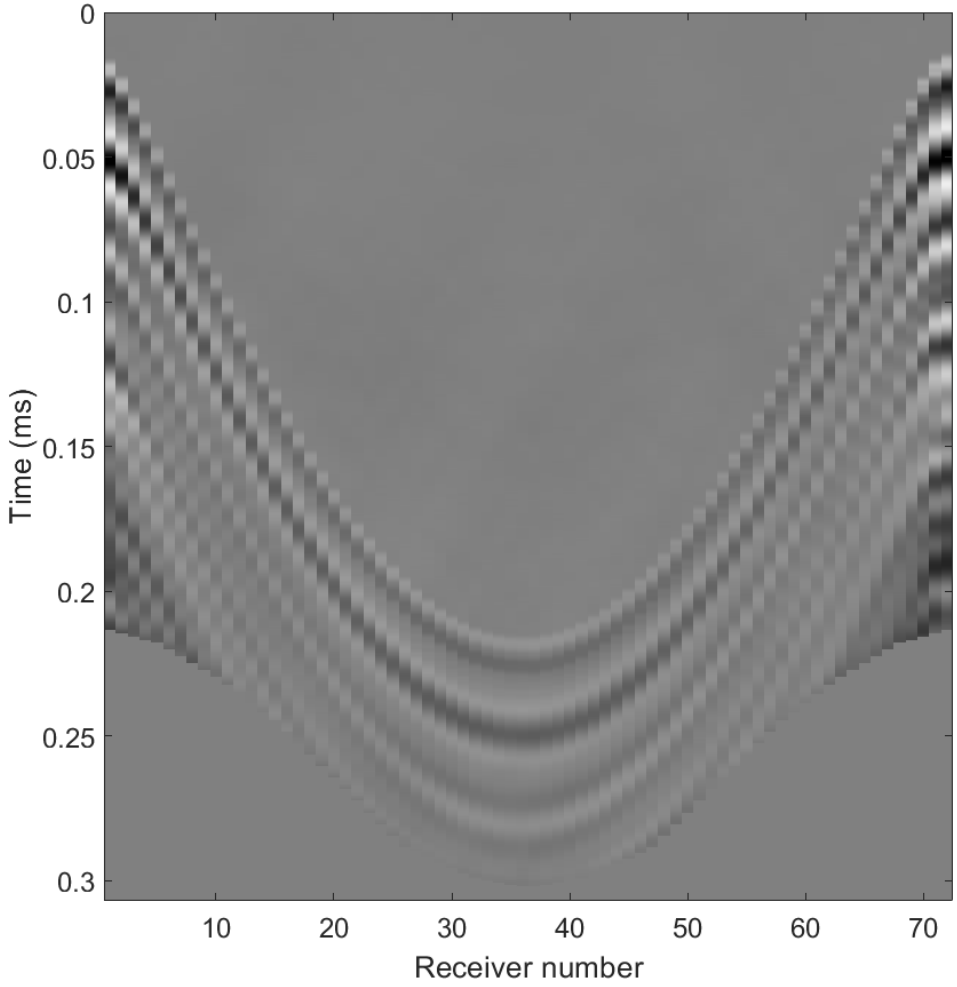

FIG. 4. Shot record for leftmost shot, no target survey after processing.

same "shot". To help compensate for the differences that cannot be reproduced even in the water-only case, we treat the "observed data" in the inversion as the difference between the target and no-target measurements, plus the no-target synthetic. This ensures that all data residuals for the initial model arise from measured data differences, and not from synthetic vs modelled differences. This methodology helps reduce modelling error effects, but some spurious differences remain due to challenges in reproducibility: the target and no-target data differ in places that cannot be attributed to the target alone, as demonstrated in Figure 6.

Cross-gradient regularization

Given the substantial density contrasts between water and the materials usable as objects in the physical modelling tank, density is likely important for accurately inverting the measured data. Unfortunately accurate estimates of density can be difficult to obtain in FWI, given the relatively low sensitivity of the inversion to this parameter and the potential

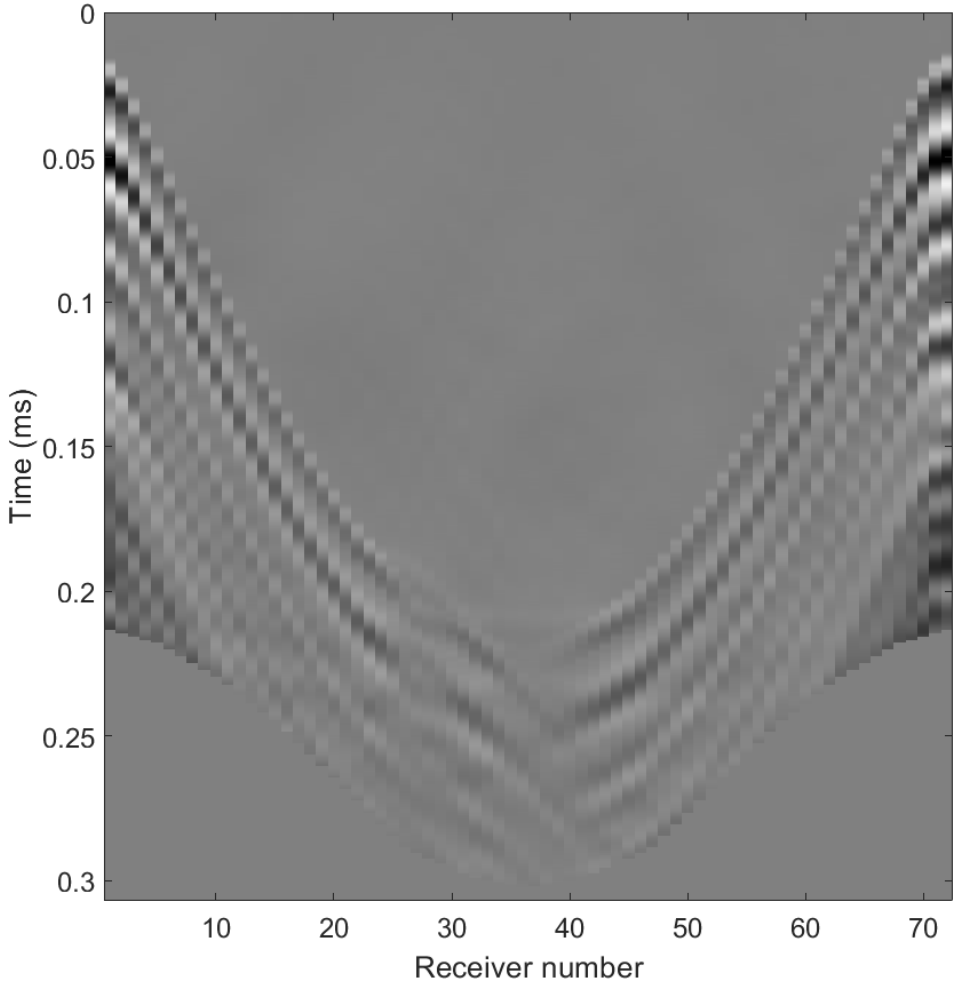

FIG. 5. Shot record for leftmost shot, target survey after processing.

for cross-talk with P-wave velocity. For this reason, we introduce a regularization term intended to limit density inversion results to those consistent with the P-wave velocity model. In many inversion problems we have high confidence that contrasts in different parameters will occur at the same locations (as these will be the locations of contrasts between materials), but little confidence in the relative sign of the parameter changes. For instance, if velocity increases across a given interface, we can expect that density may change there as well, but not have a strong expectation of whether a density increase or decrease is expected at the interface (both may be plausible). In this case, a cross-gradient regularization (Gallardo and Meju, 2004) may be appropriate. This type of regularization penalizes the cross-product of the spatial gradients of different parameters:

$$\phi_X = \sum |\nabla_x P_1 \times \nabla_x P_2|, \tag{1}$$

where P_1 and P_2 are the inversion parameters and ∇_x is the spatial gradient operator. This penalty term is small only when the spatial gradients of the parameters are either small for one parameter or parallel to one another, as these are the only ways to make the cross-

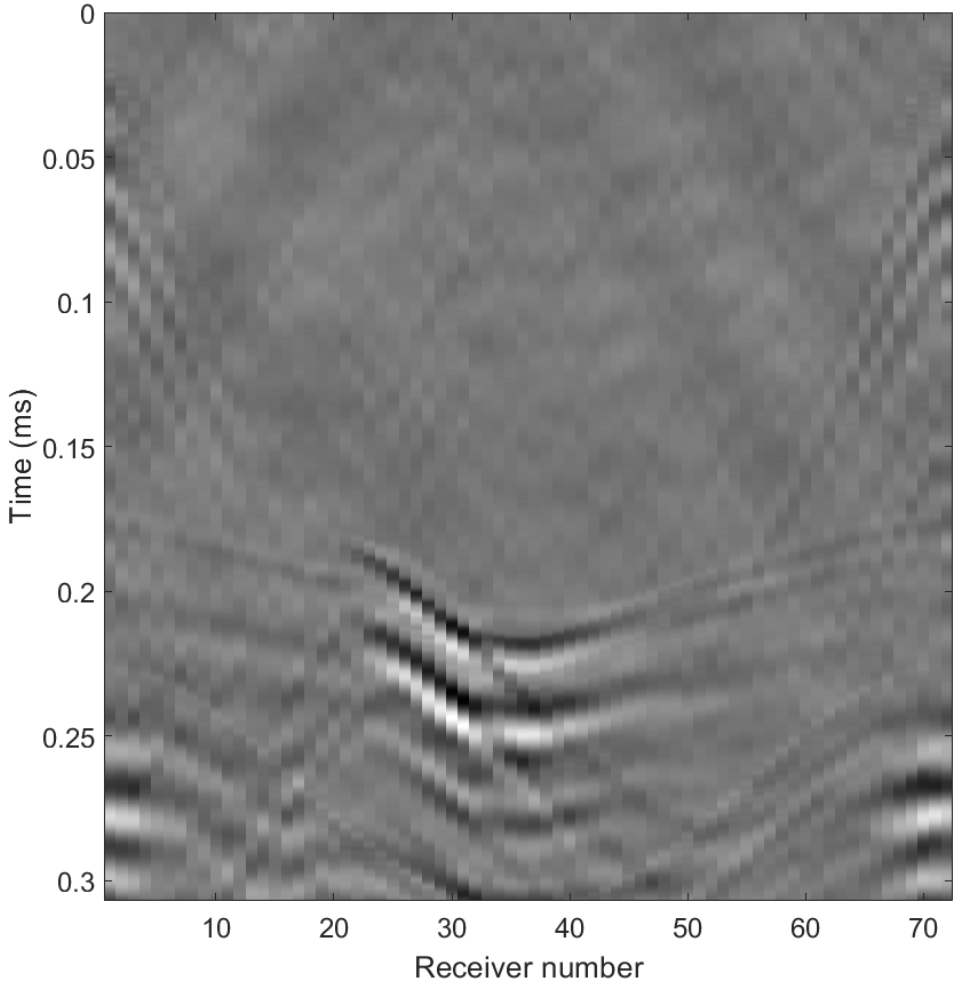

FIG. 6. Difference between target and no target datasets before processing.

product term small. By using this regularization term, we promote inversion results in which both parameters have features with similar shapes, as boundaries are penalized if they are not either parallel or small for one parameter.

INVERSION RESULTS

In our inversion tests, we consider eight bands of four frequencies each. Each band has a low frequency of 10 kHz, and a high frequency which begins at 55 kHz, and increases with each band, ending at 130 kHz. At each band, we perform 10 iterations of L-BFGS optimization. The finite-difference grid size used for wave modelling was 1.5 mm, and the inversion was performed on a x by y grid, spanning about 30 cm in both directions.

Figure 8 shows the inversion result for v_P (top) and density (bottom). The v_P result successfully identifies the locations of both anomalies, provides an estimate of the semi-cylinder target shape, and some suggestion of the ring outline. The specific geometry of the ring, and especially the interior, however, is not well recovered. The PVC regions are

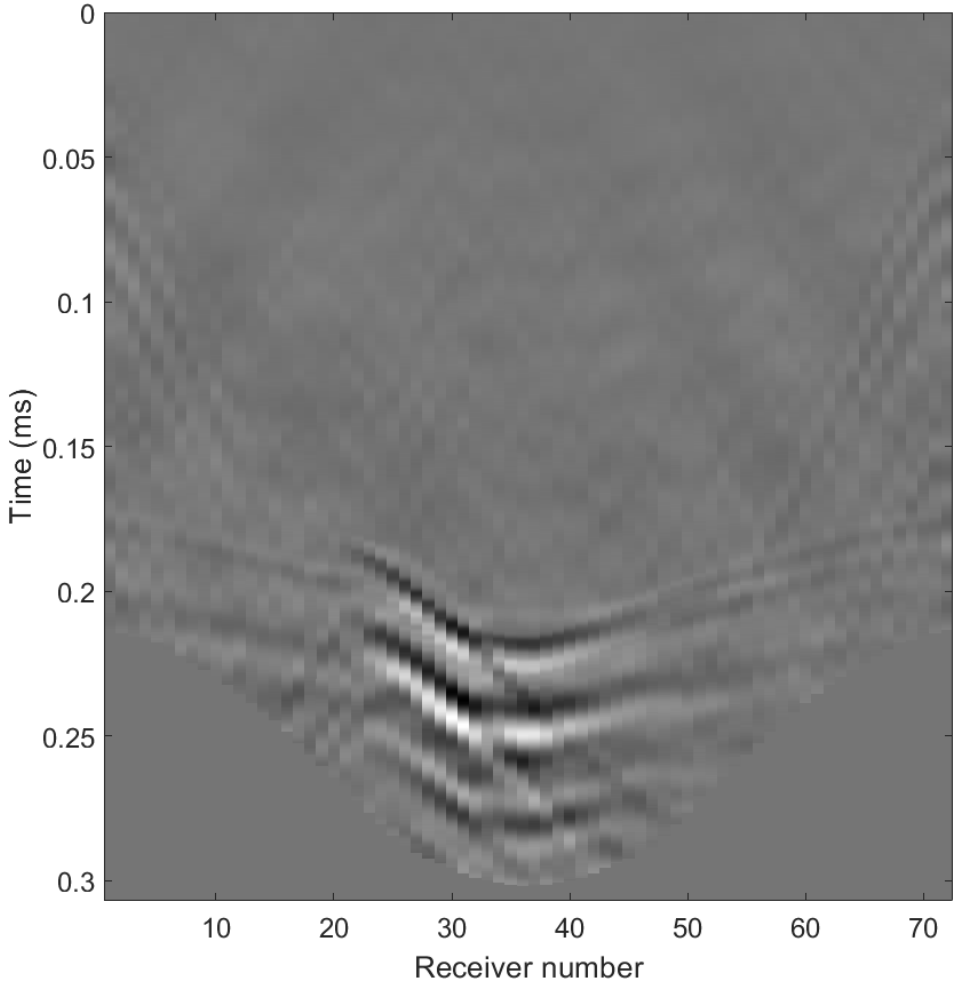

FIG. 7. Difference between target and no target datasets after processing.

correctly identified with a velocity increase in the inversion result, and minimal changes outside of the target regions are recovered. The density inversion result is largely geometrically consistent with the v_P result, suggesting that the cross-gradient penalty term is performing as desired. The density result also contains apparent ringing artifacts centered on both targets; these likely arise because the data sensitivity is primarily to density contrasts, meaning the ringing may marginally improve data fit, and because this effect is not penalized by the cross-gradient term, as there are little or no v_P variations in these locations. While the ringing is undesirable, the most problematic feature of the density result is the very low densities recovered, especially in the interior of the ring. These features probably arise because of the insensitivity of the inversion to longer scale density changes: the data suggest reflections caused by density contrasts, but not an average density, leading to unrealistically low density inside the ring.

The data-fit of the inversion result is relatively low after inversion. Figure 9 shows the difference between the modelled data from the inversion result and the modelled data for

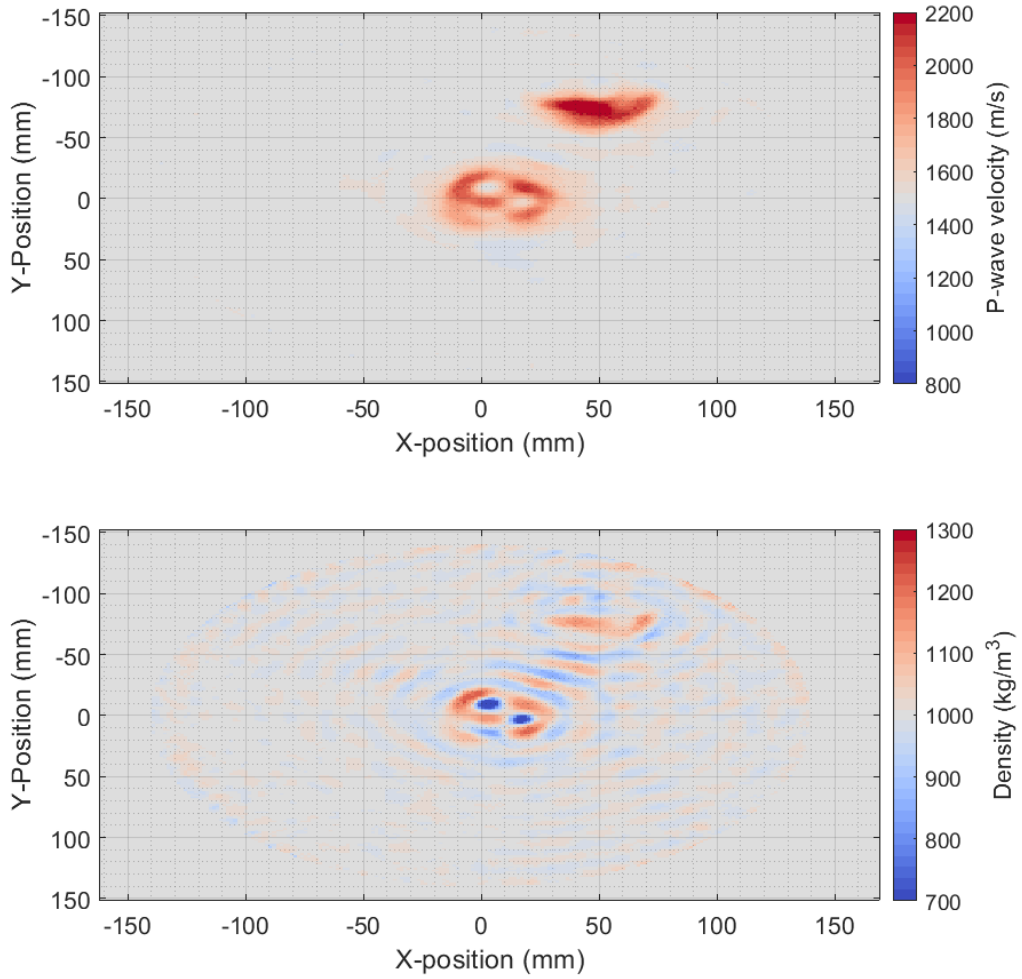

FIG. 8. P-wave velocity (top) and density (bottom) inversion results using the cross-gradient penalty term.

the no target case for the leftmost shot. Comparing to Figure 7, we can see that most of the prominent features of the data are reproduced by the inversion. What is not reproduced, however, is the mute region, which is likely a key factor limiting the quality of the inversion.

To highlight the effect of the cross-gradient penalty term, we repeat the inversion without it, generating the result shown in Figure 10. The v_P result is largely similar, though a halo region of compensation for density is introduced around each anomaly, and the noise background increases. The density result, however, is made much worse by the removal of the cross-gradient term: noise amplitudes increase, and the geometry of the targets is largely lost in the density result.

DISCUSSION

The ring structure is not well recovered in these inversions. This is likely caused, at least in part, by the current limitations of the datasets we consider. In particular, much of

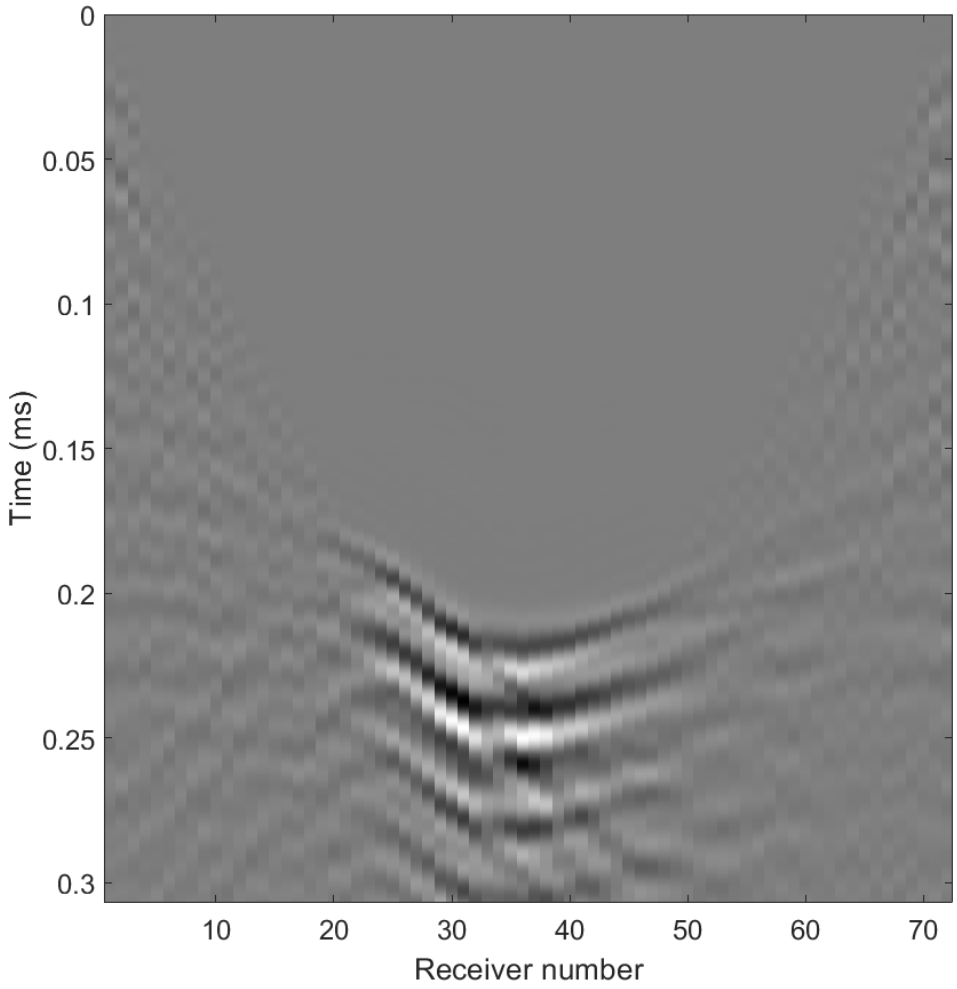

FIG. 9. Difference between data modeled from inversion result and data modelled for the no target case.

the reflection data recorded is overlapped by tank boundary reflections, and consequently removed by the mute. This leaves the inversion with relatively little reflection data to base the short-scale structure of the inversion result on. For this reason, tomography-scale features, like the semi-cylinder are relatively well recovered, while the details of the ring are not. This issue is being addressed by use of a larger physical modelling tank for future datasets, which should allow for more of the reflected wavefield to be measured.

CONCLUSIONS

Physical modelling at the ultrasonic scale can help inform both seismic and medical imaging inversion approaches. Here, we considered a circular acquisition geometry with ultrasonic buzzer sources and receivers. We found that the long-scale features of the model were well recovered, and that the modelled data fit well to the measurements, though very fine-scale characterization remained difficult. The use of a cross-gradient regularization term was key factor in allowing for the amplitude variations caused by density to be accu-

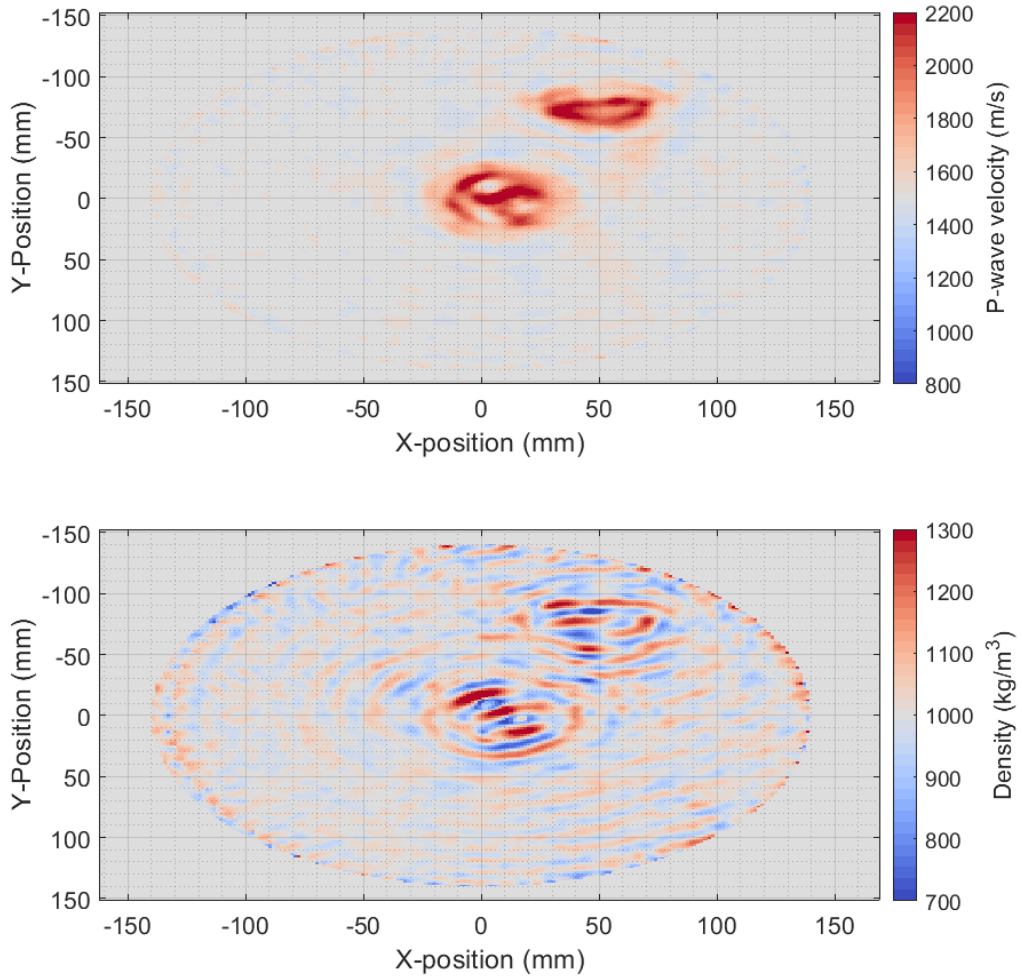

FIG. 10. P-wave velocity (top) and density (bottom) inversion results without the cross-gradient penalty term.

rately modelled, without introducing the highly unrealistic model features that an unconstrained density model would cause.

ACKNOWLEDGEMENTS

The authors thank the sponsors of CREWES for continued support. This work was funded by CREWES industrial sponsors and NSERC (Natural Science and Engineering Research Council of Canada) through the grant CRDPJ 543578-19. Scott Keating was also supported by the Canada First Research Excellence Fund, through the Global Research Initiative at the University of Calgary.

REFERENCES

Gallardo, L. A., and Meju, M. A., 2004, Joint two-dimensional dc resistivity and seismic travel time inversion with cross-gradients constraints: Journal of Geophysical Research: Solid Earth, **109**, No. B3.

Henley, D. C., 2022, Shadow imaging: attenuation projection of acoustic waves from a circular array without arrival pickinga: CREWES Annual Report, **34**.

Wong, J., 2022, Transmission traveltime tomography of physically modelled data: CREWES Annual Report, **34**.