POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Boundary-aware Poisson-based P–S separation of elastic wave-fields in isotropic media

Laurea Magistrale in Mathematical Engineering - Ingegneria Matematica

Author: Antonio Andretta

Advisor: Prof. Anna Scotti

Academic year: 2025-2026

1. Introduction

Elastic seismic imaging relies on the simulation and analysis of elastic wave-fields, in which compressional (P) and shear (S) contributions are generally superposed. In isotropic linear elastic media, P-waves are associated with compressional deformation, whereas S-waves are associated with shear deformation. In heterogeneous configurations, and especially in the presence of interfaces and free surfaces, these modes interact and convert into one another.

P–S wave-field separation is therefore useful both for analyzing elastic wave propagation and for improving seismic imaging workflows. In elastic reverse time migration, separated wave-fields can reduce artifacts due to undesired modal correlations, while in Elastic Full Waveform Inversion (E–FWI) they provide a way to mitigate parameter cross-talk, in particular the coupling between V_P and V_S, a well-known difficulty in multiparameter elastic inversion [1].

This thesis focuses on Poisson-based P–S separation methods, following the line of Helmholtz-type approaches proposed in [3, 4]. In these formulations, the P component is identified with the irrotational part of the elastic wave-field and reconstructed as the gradient of a scalar potential obtained by solving an auxiliary Poisson problem, while the S component is recovered by subtraction. This strategy is computationally attractive, since it reduces vector-field separation to scalar elliptic problems that can be efficiently treated, in rectangular domains, by fast transform-based solvers.

However, in bounded computational domains, the Poisson equation alone does not uniquely define the separation: the auxiliary problem must be closed with boundary conditions. In many existing implementations, these conditions are not prescribed explicitly, but are implicitly determined by the numerical solver, for instance, through sine transforms (see [4]) or FFT-based periodic extensions (see [3]).

The main contribution of this thesis is a boundary-aware reinterpretation of Poisson-based P–S separation, based on the Helmholtz–Weyl decomposition. This viewpoint makes explicit the role of the boundary conditions used to close the auxiliary Poisson problem and allows their mathematical, physical and numerical consequences to be analyzed. Based on this framework, the thesis studies Neumann, Dirichlet and mixed formulations, develops efficient staggered-grid-consistent solvers, and investigates the use of separated fields in adjoint-state gradient computations for Elastic Full Waveform Inversion.

2. Methodology

2.1. Helmholtz–Weyl formulation

The starting point of the thesis is the classical irrotational/solenoidal criterion for elastic wave-field separation. In homogeneous, unbounded and isotropic elastic media, P-waves correspond to the irrotational component of the elastic field, whereas S-waves, and, in two dimensions, SV-waves, correspond to its solenoidal component. In more general configurations, involving heterogeneities or domains that are not unbounded in every direction, this identification is no longer exact in a strict modal sense, but it is adopted as a physically meaningful operational criterion. In bounded computational domains, this criterion is formulated through the *Helmholtz–Weyl decomposition theorem*, which provides the mathematical framework for interpreting P–S separation as a decomposition of a vector-valued function into its irrotational, solenoidal and *harmonic-gradient* contributions, with the latter being simultaneously curl-free and divergence-free.

Within this framework, the P component is reconstructed as the gradient of a scalar potential obtained by solving an auxiliary Poisson problem driven by the divergence of the total field. The S component is then recovered as the residual between the original field and the reconstructed P field. The separation workflow can therefore be summarized as

$$\underline{u} \longrightarrow \begin{cases} -\Delta f = -\operatorname{div}(\underline{u}) & \text{in } \Omega \\ \text{BCs} & \text{on } \partial\Omega \end{cases} \longrightarrow \\ \longrightarrow \underline{u}_P = \nabla f \longrightarrow \underline{u}_S = \underline{u} - \underline{u}_P \tag{1}$$

Consequently, in bounded domains, the choice of boundary conditions for the auxiliary Poisson problem becomes part of the mathematical definition of the separated fields.

2.2. Boundary closures and numerical solvers

Two boundary closures are explicitly analyzed: the Poisson–Neumann formulation and the Poisson–Dirichlet formulation, respectively defined by

$$\nabla f \cdot \hat{\nu} = \underline{u} \cdot \hat{\nu} \quad \text{on } \partial\Omega$$

$$f = 0 \quad \text{on } \partial\Omega$$

These choices correspond to different assignments of the *harmonic-gradient* contribution: in the Poisson–Neumann formulation, this contribution is included in the irrotational P component, whereas in the Poisson–Dirichlet formulation it is assigned to the solenoidal S component. Therefore, the two boundary closures define two mathematically distinct P–S separations. The explicit treatment of the non-homogeneous Neumann case, less common in standard Poisson-based implementations, is one of the distinctive aspects of the work.

The theoretical formulation holds in both two and three space dimensions, while the numerical developments and experiments of the thesis are specialized to the two-dimensional P–SV setting. The numerical workflow is carried out snapshot by snapshot: for a given snapshot of the elastic wave-field, the separation is performed according to the procedure (1). The numerical Poisson solvers are designed to be consistent with the staggered-grid arrangement inherited from the elastic finite-difference solver and are implemented through fast trigonometric transforms, so as to keep the computational cost compatible with large-scale seismic simulations.

For the Dirichlet formulation, the thesis adopts a different numerical strategy from the explicit finite-difference discretization of differential operators employed in the Poisson–Neumann case. Instead, the solver exploits the analytical spectral representation of the continuous Poisson solution, which is explicitly computable in rectangular domains. The corresponding coefficients are approximated from the information available on the staggered grid. This yields a continuous-like spectral approach based on numerical quadrature, while retaining a log-linear asymptotic complexity compatible with fast-transform methods. This strategy has the advantage of being exact for wave-fields that are band-limited with respect to the trigonometric spectral expansion underlying the solver, and the spectral-content analysis performed on the numerically simulated wave-fields confirms that, in the tested configurations, the modes carrying the relevant energy lie within the spectral resolution of the grid.

The comparison between the Neumann and Dirichlet closures motivates a Poisson–mixed formulation specifically designed for fluid free-surface configurations. In this case, the boundary is split as $\partial\Omega = \Gamma_N \cup \Gamma_D$, where Γ_N denotes the physical free surface and Γ_D the artificial boundaries equipped with absorbing layers. Non-homogeneous Neumann conditions are imposed on Γ_N, whereas homogeneous Dirichlet conditions are imposed on Γ_D, namely

$$\begin{cases} \nabla f \cdot \hat{\nu} = \underline{u} \cdot \hat{\nu} & \text{on } \Gamma_N \\ f = 0 & \text{on } \Gamma_D \end{cases}$$

The P scalar potential is decomposed into a harmonic lifting of the Neumann datum and a residual contribution satisfying homogeneous mixed boundary conditions. Both contributions admit analytical representations in rectangular domains, leading to a solver that follows the same continuous-like spectral strategy used for the Dirichlet formulation.

2.3. Application to Elastic Full Waveform Inversion

Finally, the thesis studies the incorporation of P–S separation into the computation of adjoint-state gradients for Elastic Full Waveform Inversion, in the spirit of mode-decomposition-based strategies such as [2], with the aim of mitigating cross-talk between the V_P and V_S updates. The *forward* and *adjoint wave-fields* are first separated into their P and S components, and the gradients are then constructed by selecting specific modal time cross-correlations between the separated fields. In particular, only P–P correlations are retained for the modified V_P gradient, whereas P–S, S–P and S–S interactions are retained for the modified V_S gradient.

3. Results

3.1. Boundary-induced artifacts and mitigation

A first result of the thesis is the physical interpretation of the boundary conditions and of the associated boundary artifacts. The Poisson–Neumann formulation implicitly enforces a vanishing normal trace for the residual S component at the boundary:

$$\underline{u}_S \cdot \hat{\nu} = 0 \quad \text{on } \partial\Omega \tag{2}$$

whereas the Poisson–Dirichlet formulation enforces the tangential component of the reconstructed P field to vanish at the boundary:

$$\underline{u}_P \cdot \hat{\tau} = 0 \quad \text{on } \partial\Omega \tag{3}$$

As a consequence, the two closures may generate dual leakage mechanisms. The Neumann closure may induce S-to-P leakage when an S wave reaches the boundary with a non-vanishing normal component, while the Dirichlet closure may induce P-to-S leakage when a P wave reaches

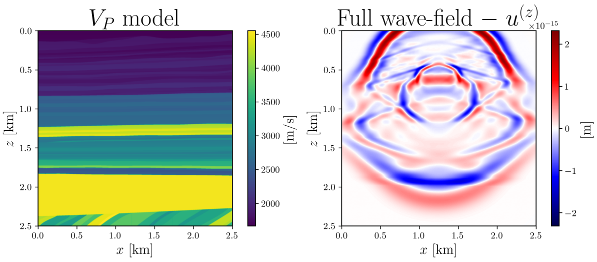

Figure 1: Land Marmousi-II configuration used to assess Poisson-based P–SV separation in an effectively unbounded setting. Right: P-wave velocity model V_P. Left: vertical component of the full wave-field at $t \approx 0.50\,\mathrm{s}$, generated by an explosive Ricker wavelet source with central frequency $f_0 = 10\,\mathrm{Hz}$, injected at $(x_s, z_s) = (1.25\,\mathrm{km},\ 0.50\,\mathrm{km})$.

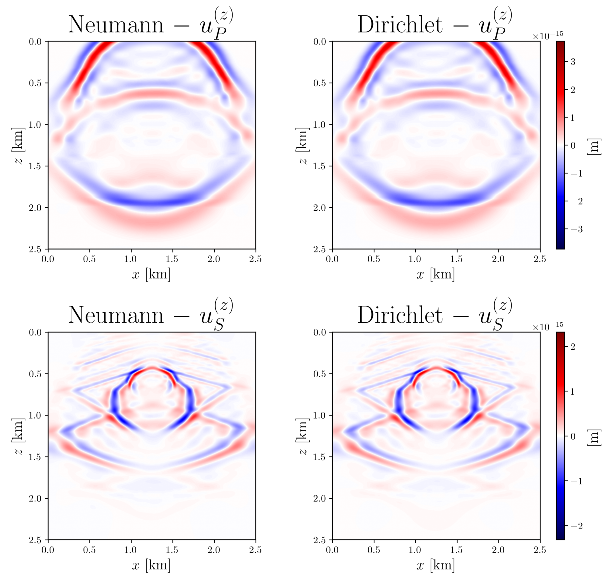

Figure 2: Comparison between Poisson–Neumann (left column) and Poisson–Dirichlet (right column) separations in the land Marmousi-II configuration with absorbing layers on all boundaries. The vertical components of the reconstructed P field (top row) and SV field (bottom row) at $t \approx 0.50\,\mathrm{s}$ are shown.

the boundary with a non-vanishing tangential component. The artifacts observed in the two-dimensional numerical tests can therefore be interpreted as consequences of the mathematical closure used to define the auxiliary Poisson problem in a bounded domain.

Based on this interpretation, the thesis adopts a mitigation strategy consisting of extending the separation domain so as to include the absorbing layers. Numerical tests in effectively unbounded configurations, i.e., with all boundaries equipped with absorbing layers, show that, owing to this strategy, the boundary-induced artifacts are strongly mitigated. As a result, the Poisson–Neumann and Poisson–Dirichlet formulations produce practically equivalent separated wave-fields within the inner physical region, even in realistic heterogeneous models. This confirms that the artifacts associated with artificial boundaries can be strongly reduced by making the interaction between the wave-field and the boundary of the separation domain negligible. An example in a realistic heterogeneous model is shown in figures 1 and 2. Only the vertical component is displayed, both for the full wave-field and for the separated wave-fields obtained with the Poisson–Neumann and Poisson–Dirichlet separation strategies.

3.2. Fluid free surface configurations

This picture changes in the presence of a free surface. In this case, the boundary is physical and the choice of the Poisson closure becomes essential, especially in fluid free surface configurations. Since fluids do not support shear-wave propagation, any SV component reconstructed inside the fluid region can be interpreted as numerical leakage. The numerical results show that the Dirichlet closure may generate spurious P-to-SV conversions near the fluid free surface. This observed behavior agrees with the implicit constraint (3) imposed by the Poisson–Dirichlet formulation, which is not compatible with the purely compressional nature of the wave-field in a fluid medium. By contrast, the Poisson–Neumann formulation implicitly imposes the boundary constraint (2), which is consistent with the absence of shear waves in the fluid, and has been shown to suppress this boundary-induced leakage mechanism. This suggests that explicitly accounting for the physical nature of the boundary can mitigate free-surface-induced artifacts that remain problematic in the recent FFT-based Poisson separation approach proposed in [3].

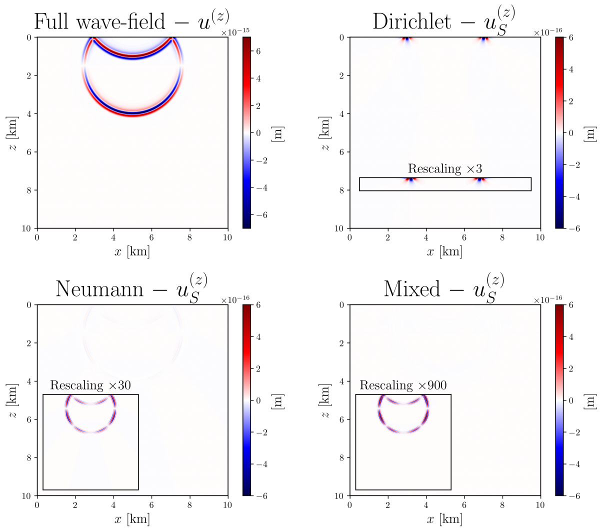

Figure 3: Homogeneous fluid free surface test at $t \approx 2.0\,\mathrm{s}$, generated by an explosive Ricker wavelet source with central frequency $f_0 = 5\,\mathrm{Hz}$, injected at $(x_s, z_s) = (5.0\,\mathrm{km}, 1.5\,\mathrm{km})$. The vertical component of the full wave-field (top-left) is shown as a reference, while the other panels compare the vertical SV component reconstructed with the Poisson–Dirichlet (top-right), Poisson–Neumann (bottom-left) and Poisson–mixed (bottom-right) formulations. The insets show rescaled views of the residual SV signal.

The numerical behavior of the Poisson–mixed formulation confirms the twofold motivation behind its construction. The non-homogeneous Neumann boundary condition imposed at the fluid free surface suppresses the structural boundary-induced P-to-SV leakage observed with the Dirichlet closure. At the same time, the residual SV signal inside the fluid, interpreted as numerical noise rather than as a structural artifact associated with the Poisson closure, is reduced by approximately one order of magnitude compared with the finite-difference Poisson–Neumann solver. This improvement can be reasonably attributed to the continuous-like spectral approach employed for the numerical solution. This behavior is illustrated in figure 3, which compares the vertical SV components obtained with the Poisson–Neumann, Poisson–Dirichlet and Poisson–mixed formula-

tions in a homogeneous fluid configuration.
In more realistic configurations, the free surface behavior remains coherent, although additional artifacts appear near the fluid–solid interface and are observed for all three separation strategies.

3.3. Application to Elastic Full Waveform Inversion

In the E–FWI application, the numerical results show an asymmetric effect of the separated-gradient strategy on the $V_P - V_S$ coupling. On the one hand, the imprint of compressional structures in the V_S updates is reduced. This behavior is interpreted as the removal of the direct volumetric pathway of the $V_P \to V_S$ coupling from the modified V_S gradient. After the modal selection, the residual coupling can only persist through deviatoric strain correlations involving at least one S component. On the other hand, the coupling in the opposite direction remains essentially unchanged, consistently with the fact that the modified V_P gradient coincides with the standard one at the continuous level. The cross-talk mitigation is therefore targeted and asymmetric, rather than a complete decoupling of the multiparameter inverse problem.
This effect is illustrated in figures 4 and 5 for a two-inclusion model, characterized by a homogeneous background in which two circular anomalies are embedded: one affecting only V_P, in the upper-left region, and the other affecting only V_S, in the bottom-right region. The two anomalies are spatially separated in order to better assess the effects of the $V_P - V_S$ coupling during the inversion. Figure 5 displays the residual model errors, computed as the difference between the final reconstructed model and the true model, both for the inversion with the separated-gradient strategy and for the standard inversion. The comparison shows that the V_P residual remains essentially unchanged, whereas the separated-gradient strategy substantially reduces the imprint of the compressional structure in the V_S residual.
In the numerical tests performed, this modification of the gradient did not lead to any observed deterioration of the convergence behavior of the optimization algorithm.

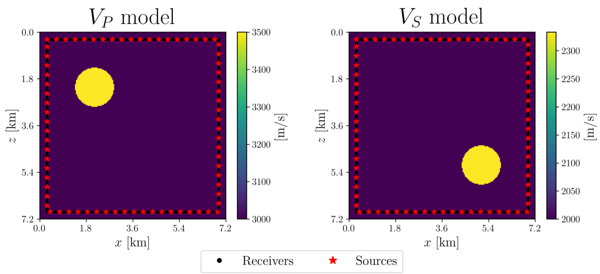

Figure 4: Two-inclusion E–FWI test case. The true V_P (right) and V_S (left) models are shown together with the acquisition geometry.

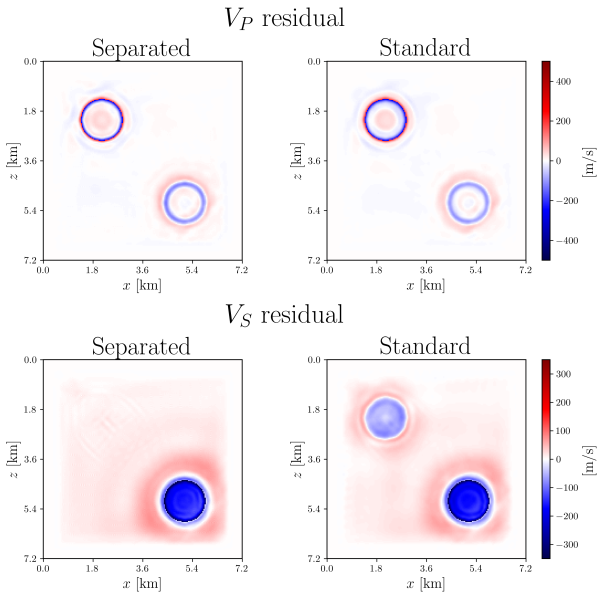

Figure 5: Residual V_P (top row) and V_S (bottom row) model errors in the two-inclusion E–FWI test. Left column: inversion with separated-gradient strategy; right column: standard inversion.

4. Conclusions

The central message of this thesis is that boundary conditions in Poisson-based P–S separation are not a secondary numerical detail, but an intrinsic part of the mathematical definition of the separation in bounded computational domains. If this role is not made explicit, the boundary closure is effectively delegated to the numerical solver and becomes embedded in the chosen computational procedure.
The Helmholtz–Weyl decomposition makes this ambiguity explicit. It shows that Neumann and Dirichlet closures assign the *harmonic-gradient* contribution differently between the reconstructed P field and the residual S field. As a

result, the separated fields have different boundary properties and physical meanings, leading to dual leakage mechanisms that can be interpreted, predicted and mitigated.

In practice, when the boundaries are artificial, these artifacts can be strongly reduced by extending the separation domain into the absorbing layers, making the Neumann and Dirichlet formulations practically equivalent in the physical bulk. The distinction becomes essential at physical boundaries: near a fluid free surface, the Neumann closure is consistent with the absence of shear-wave propagation in the fluid, whereas the Dirichlet closure may induce spurious P-to-SV contamination. The proposed mixed formulation addresses this issue by combining Neumann conditions on the physical free surface with Dirichlet conditions on the artificial boundaries. It preserves the physically coherent behavior of the Neumann closure at the fluid free surface, while reducing residual SV numerical noise through a spectral solver analogous to the one developed for the Dirichlet formulation and retaining the same log-linear complexity.

The application to Elastic Full Waveform Inversion shows that the same modal-separation framework can also inform the construction of adjoint-state gradients. Selecting modal correlations yields a targeted and structurally asymmetric mitigation of $V_P \to V_S$ cross-talk, contributing to more reliable V_S updates in the tested configurations, but without producing a complete decoupling of the multiparameter inverse problem.

Several directions may be considered for future developments of this work. From the numerical point of view, the proposed two-dimensional algorithms are conceptually extendable to three-dimensional rectangular domains, where the same staggered-grid structure and fast-transform strategy can be generalized. However, since the present framework strongly relies on structured grids and rectangular geometries, an important perspective is the development of efficient boundary-aware Poisson solvers and discrete projectors compatible with discretizations on unstructured meshes and general geometries. From the application point of view, future work should investigate whether the separated-gradient strategy introduced in this work can effectively improve parameter cross-talk reduction and reconstruction quality within E–FWI or imaging workflows in more realistic seismic configurations. Finally, a major perspective is the extension of the boundary-aware separation framework to anisotropic media, where the classical irrotational/solenoidal criterion should be replaced by projection operators based on the anisotropic polarization structure.

References

[1] S. Operto, Y. Gholami, V. Prieux, A. Ribodetti, R. Brossier, L. Métivier, and J. Virieux. A guided tour of multiparameter full-waveform inversion with multicomponent data: From theory to practice. *The leading edge*, 32:1040–1054, 2013.

[2] T. F. Wang and J. B. Chen. Elastic full waveform inversion based on mode decomposition: the approach and mechanism. *Geophysical Journal International*, 209:606–622, 2017.

[3] Q. Zheng and G. Yao. An improved elastic wavefield separation method based on the Helmholtz decomposition. *Geophysics*, 88:202–210, 2023.

[4] H. Zhu. Elastic wavefield separation based on the Helmholtz decomposition. *Geophysics*, 82:173–183, 2017.