Accelerating innovation with software abstractions for scalable computational geophysics

Mathias Louboutin[*,1], *Philipp Witte*[2], *Ali Siahkoohi*[1], *Gabrio Rizzuti*[3], *Ziyi Yin*[1], *Rafael Orozco*[1] *and Felix J. Herrmann*[1]
[1] *Georgia Institute of Technology,* [2] *Microsoft Research,* [3] *Utrecht University*

SUMMARY

We present the SLIM open-source software framework for computational geophysics, and more generally, inverse problems based on the wave-equation (e.g., medical ultrasound). We developed a software environment aimed at scalable research and development by designing multiple layers of abstractions. This environment allows the researchers to easily formulate their problem in an abstract fashion, while still being able to exploit the latest developments in high-performance computing. We illustrate and demonstrate the benefits of our software design on many geophysical applications, including seismic inversion and physics-informed machine learning for geophysics(e.g., loop unrolled imaging, uncertainty quantification), all while facilitating the integration of external software.

INTRODUCTION

Software development for exploration geophysics has been traditionally driven by performance. Though this has resulted in very efficient code, the usability and portability of these codes has been compromised as a result of the absence of user-level design. Abstractions and user interfaces at a high level have been developed in recent decades to facilitate research and development. A number of such interfaces are available, ranging from programming languages, e.g., Python (van Rossum and Drake, 2009), Julia (Bezanson et al., 2017), and Matlab, to domain-specific languages (DSLs), including RVL (Padula et al., 2009), Firedrake (Rathgeber et al., 2016), and Devito (Louboutin et al., 2019; Luporini et al., 2020). Additionally, there has been a community initiative towards open-source codes and reproducibility, led by Madagascar (Madagascar, 2017). These efforts have brought attention to the need for user abstraction to easily translate mathematics into code without having to manually refactor thousands of lines of code.

Motivated by this background, we introduce our fully open-source software SLIM framework based on high-level abstractions and separation of concerns. Our software design relies on three principles: (1) A high-level abstraction that represents the mathematical problem at hand; (2) Scalability with vertical integration of high-performance software and DSLs; (3) Interoperability through the adoption of language standards (object oriented, multiple dispatch, inheritance, ...). In the following, we will detail these three points, highlighting their importance and implementation. We will follow with some illustrations of the integration with external software to demonstrate how research and development can be eased and potentially accelerated by these design principles.

SOFTWARE DESIGN

The fundamental objective of a scientific software framework is to provide users with an interface that allows them to express their own scientific problems. In particular, the interface should provide abstractions to define the problem as closely as possible to its mathematical formulation. With this design philosophy, research and development turnaround time can be drastically reduced with quick prototyping and testing of new ideas and algorithms. Our software framework is grounded in this idea and led to the development of legacy software (Lin and Herrmann, 2015; Silva and Herrmann, 2019) adopted by industry for real-world application. Based on these ideas and modern programming paradigms, we designed a high-level framework that encapsulates the mathematical definition of geoscientific problems at every level. These different levels are handled through the vertical integration of domain-specific languages (DSLs). At the lowest level, Devito provides a symbolic DSL for the definition of wave-equation and a just-in-time compiler to target the available hardware with near-optimal performance. On top of Devito, we developed JUDI, a linear algebra DSL for wave-equation based modeling and inversion. Finally, we created a machine learning framework that allows us to integrate Devito propagators and JUDI Linear Operators into machine learning (ML) frameworks (PyTorch and Flux) opening the door to ML-augmented geophysical inverse problems and uncertainty quantification (UQ). Finally, throughout these different layers, we committed to follow language standards allowing easy interfacing and integration with external software and frameworks.

WAVE-EQUATION BASED INVERSION

Linear operators are at the core of applied geophysics. Some relevant examples include data filtering tools such as Fourier-based f-k filters, the Radon transform, NMO corrections to traditional imaging operator such as Kirchhoff migration, post-stack migration and wave-equation based inversion where the discrete wave-equation is a linear operator. One of the first frameworks for abstract matrix-free linear algebra was sPOT(van den Berg and Friedlander, 2009), later extended to wave-equation based inversion and distributed computing for wavefield reconstruction (Kumar, 2010). While its adoption was limited by its implementation in Matlab, such user oriented abstraction laid the ground for modern abstracted frameworks such as `JOLI`(Modzelewski et al., 2022) in Julia and `pyLops` in Python (Ravasi and Vasconcelos, 2020).

Similarly, wave-equation based inversion can be trivially formulated as simple least-squares optimization problem for the non-linear wave-equation represented as an abstract matrix $\mathbf{A}(\mathbf{m})$ parametrized by the model parameter $\mathbf{m}$:

$$\underset{\mathbf{m}}{\text{minimize}} \quad \Phi(\mathbf{m}) = \sum_{i=1}^{n_s} ||\mathbf{P}_r \mathbf{A}(\mathbf{m})^{-1} \mathbf{P}_s^\top \mathbf{q}_i - \mathbf{d}_i||_2^2,$$

Building on modern software solutions JUDI (Witte et al., 2019a; Louboutin et al., 2022a) provides a high-level interface allowing to define a solver for this problem with a few lines of code. For example, FWI using Gauss-Newton updates at each iteration can be summarized in as little as five lines (Listing 1).

```
1   # Gauss-Newton method
2   F = judiModeling(model, srcGeom, recGeom)
3   J = judiJacobain(F, q)
4   for j=1:maxiter
5       d_pred = F*q
6       fhistory_GN[j] = .5f0*norm(d_pred - d_obs)^2
7       # Gauss-Newton update
8       p = lsqr(J, d_pred - d_obs)
9       model0.m .= proj(model0.m .- reshape(p, model0.n))
10  end
```

Listing 1: FWI with Gauss-Newton updates using JUDI. The complete example is available and reproducible in the JUDI github repository.

In addition to offering a high-level, mathematical interface to the wave-equation and related linear operators, JUDI builds on multiple layers of abstractions while still providing computational performance that is comparable to the state-of-the-art(Luporini et al., 2020). JUDI integrates Devito as a backend for the actual wave-equation solves. Devito is a stencil DSL (Louboutin et al., 2019; Luporini et al., 2020) that offers a symbolic user-level interface to the underlying numerical solver. This symbolic interface allows the user to easily and mathematically define generic partial differential equations (PDEs), such as the acoustic, tilted-transverse-isotropic (TTI), or elastic wave-equations. The underlying software then provides a code generation framework than can generate highly-performant code for numerous architectures, (CPU, GPUs, ARM, POWER). Its ease of use and high-level interface have contributed to its adoption in the oil-and-gas industry for production and research at scale (Washbourne et al., 2021). We summarize the vertical integration of the different technologies (compiler, Devito, JUDI, . . .) in Figure 1, which highlights the different levels of abstraction. The 2007 BP TTI dataset, for example, can be setup and processed with RTM within days. The results are shown in Figure 2.
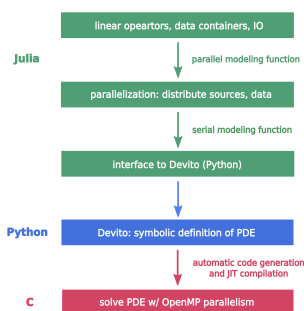
Figure 1: Schematics of the vertical integration in JUDI that enables at scale inversion through layers of abstractions.

Interoperability

While a proper separation of concerns can lead to highly usable, portable, and performant software, another important aspect to consider is interoperability, namely the ability to integrate software from different sources and organizations. We now describe how the design of our software framework allows for easy interfacing with external frameworks. One major drawback of proprietary and low-level software is the limited capability to interface and interact with external software. Through open-source code development, and committing to language-
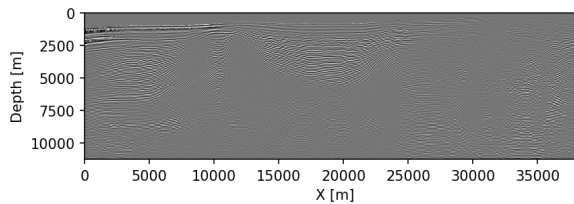
Figure 2: 2D RTM on the 2007 BP TTI model with a marine acquisition. This RTM was run on a GPU without moving off to the CPU at any time using randomized trace estimation as an extension of JUDI (Louboutin and Herrmann, 2021b; Louboutin, 2021).

specific standards (such as proper usage of types, hierarchy, inheritance, etc.) in Julia and Python, we enable seamless interoperability with external software and frameworks. This greatly increases the ability to perform research, which involves being able to easily test and combine different ideas. For instance, we combine core Julia packages, our numerical optimization toolbox SlimOptim.jl (Louboutin and Herrmann, 2021a), a constrained optimization software SetIntersectionProjection.jl (Peters and Herrmann, 2019), and COFII's wave-equation propagators (Washbourne et al., 2021), in order to setup a full waveform inversion (FWI) exercise. We were able to perform FWI, in parallel, on the Marmousi-II model by taking advantage of the best technology available. Additionally, this demonstrative example, and in general JUDI (or COFII), can trivially be deployed in the cloud using once again dedicated software abstractions (Washbourne et al., 2021; Louboutin et al., 2022b).

```
1   function objective(F, velocity, d_obs)
2       J = jacobian(F, velocity)
3       d_pred = F * velocity # Forward modeling with COFFI
4       G = J' * ({d_pred.- d_obs) # Gradient with COFII
5       f = .5f0*norm(d_pred .- d_obs)^2 # Misfit
6       return f, G
7   end
8   # Projection on TV+bounds, SetIntersectionProjection.jl
9   prj = setup_projection(...)
10  # Projected Quasi-Newton (l-BFGS) with SlimOptim.jl
11  sol = pqn(x->objective(F, x, d_obs), vec(v2), prj);
```

Listing 2: Projected Quasi-Newton FWI with multiple software packages. The complete example can be found here

Since different software modules can be pieced together with minimal effort, it greatly aids the development and testing of new research. This flexibility allows, for instance, the integration of machine learning into conventional algorithm pipelines in a plug-and-play fashion. In the following section, we provide some notable examples of this approach.

MACHINE LEARNING

Thanks to recent theoretical and practical advances, deep learning has seen a wide adoption in various geophysical applications ranging from processing (e.g., denoising, multiple removal) to inverse problems and seismic interpretation (e.g., Zhang et al., 2018; Wang et al., 2019; Yang and Ma, 2019). One of the core challenges in the adoption of deep learning for seismic applications is the integration of existing codes (e.g. seismic propagators) into deep-learning frameworks based on automatic differentiation (AD). To be able to differentiate through networks that contain both standard PyTorch/Tensorflow layers, as

well as third party functions, such as a forward modeling kernel, the third party code must have manually implemented gradients and must be properly integrated into the deep learning framework. Our high-level linear algebra abstraction framework, which sits on top of our Devito-based propagators, facilitates this integration, as deep learning frameworks like PyTorch or Flux are already able to backpropagate through a linear operator (namely by applying its adjoint to the data residual). This allows us to implement deep neural networks in Flux, in which we can combine standard deep learning layers with external third party functions, e.g. convolutional layers with migration/demigration operators. This capability has enabled various research projects in our group, including surface-related multiple elimination (Siahkoohi et al., 2019c), dispersion attenuation (Siahkoohi et al., 2019a), and ML-augmented imaging and uncertainty quantification (Siahkoohi et al., 2019b, 2020). We show in the following two examples of machine learning for exploration geophysics that take advantage of these high-level abstractions to interface PDE solvers (JUDI, Devito) with deep learning frameworks.

Seismic imaging with deep priors and uncertainty quantification

Since we can integrate Devito's highly optimized wave-equation solvers with the PyTorch deep learning library, we propose to use deep priors (Lempitsky et al., 2018) to regularize seismic imaging, where we reparameterize the seismic image as the output of an untrained convolutional neural network (CNN). This approach acts as a regularization in the image space (Lempitsky et al., 2018; Cheng et al., 2019; Dittmer et al., 2020), which exploits the inductive bias of the CNN in representing images without noisy artifacts (Lempitsky et al., 2018). We perform Bayesian inference via a gradient-based Markov chain Monte Carlo (MCMC) algorithm (Welling and Teh, 2011), where each iteration requires differentiating the action of the linearized Born scattering operator on the CNN output with respect to the CNN's weights. In order to have access to the automatic differentiation utilities of PyTorch, we expose Devito's matrix-free implementations for the migration operator and its adjoint to PyTorch via Devito4PyTorch (Siahkoohi and Louboutin, 2021). This is exemplified in Listing 3. Figure 3 summarizes the results of seismic imaging and uncertainty quantification with this approach, which are borrowed from Siahkoohi et al. (2021).

Loop-unrolled seismic imaging

Another set of applications that are enabled by the proper integration of abstract seismic modeling operators into deep learning frameworks are loop-unrolled optimization algorithms such as the learned primal-dual reconstruction (Andrychowicz et al., 2016; Adler and Öktem, 2018). These are special types of networks that follows the general structure of gradient-based optimization algorithms, in which conventional gradients are augmented by additional neural network layers. Using our integration of JUDI operators into Flux via the JUDI4Flux package (Witte et al., 2019b), we apply the loop-unrolled network architecture from (Adler and Öktem, 2017) to seismic imaging. Every iteration of the loop-unrolled algorithm consists of computing the (conventional) gradient of the LS-RTM objective

```
1   def deep_prior_imaging(d_obs, n, maxiter=1000):
2       # PyTorch Integrated Devito Jacobian.
3       J = ForwardBornLayer(model, geometry)
4       # Initialize deep prior CNN with random input.
5       G = deep_prior_net()
6       z = torch.randn([1, 1, n[1], n[2]])
7       # Setup pSGLD MCMC sampler.
8       optim, samples = pSGLD(G.parameters(), lr=0.001), []
9       for i in range(maxitr):
10          # Compute predicted perturbation model and data.
11          x = G(z)
12          d_pred = J(x)
13          # Compute the negative log-posterior.
14          nlp = n_log_posterior(d_pred, d_obs, G.↩
                parameters())
15          # Compute the gradient w.r.t. the CNN weights,
16          nlp.backward()
17          # Sample with pSGLD.
18          optim.step()
19          samples.append(x.detach().numpy())
20      return samples
```

Listing 3: Seismic imaging and uncertainty quantification with PyTorch and Devito as a wave-equation solver.

function g, using the forward and adjoint linearized Born scattering operator, followed by the application of a shallow CNN (Listing 4). The network takes a single simultaneous (super-) shot record as the input and predicts an LS-RTM (or true) image. We train the network in a supervised fashion, in which we minimize the misfit between the predicted image and the true image (whereas in reality, one would train with available LS-RTM images). Figure 4 shows the output of the loop-unrolled gradient descent algorithm, Listing 4, after training the network for 2000 iterations on a training dataset of 2,000 data-image pairs. The example is not meant to represent a realistic scenario (as the true images which were used in the training process are obviously unknown), but serve as a proof of concept on how to augment physical models with data-driven approaches. The complete example and additional variations are available at LoopUnrolledSeismicImaging.

```
1   function loop_unrolled_lsrtm(d_obs, n; maxiter=10)
2       x = randn(Float32, n[1], n[2], 1, 1)
3       nx, ny, nc, nb = size(x)
4       s = randn(Float32, nx, ny, 5, nb) # memory term
5       for j=1:maxiter
6           g = adjoint(J)*(J*vec(x) - d_obs)
7           g = reshape(Flux.normalise(g), nx, ny, 1, 1)
8           u = cat(x, g, s, dims=3)
9           u = bn1(conv1(u)); u = bn2(conv2(u)); u = bn3(↩
                conv3(u));
10          s = relu.(u[:, :, 2:6, :])
11          dx = u[:, :, 1:1, :]
12          x += dx
13      end
14      return vec(x)
15  end
```

Listing 4: Example of a physics-augmented neural network for seismic imaging. The network consists of 10 iterations of a loop-unrolled gradient descent algorithm, in which the conventional LS-RTM gradient g is augmented through convolution layers. The input into the network is the observed seismic data and the output is the predicted image.

End-to-end inversion for geological carbon storage monitoring

Finally, we show that we can easily build a framework for seismic monitoring of geological carbon storage that integrates PDE solvers, deep learning models and physical constraints. These monitoring problems rely on three different types of
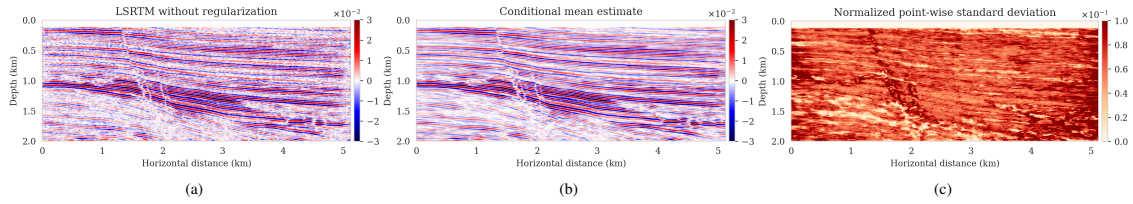
Figure 3: Imaging a 2D subset of the Parihaka (WesternGeco., 2012) dataset with deep priors (Siahkoohi et al., 2021). (a) LSRTM without any regularization. (b) The conditional (posterior) mean estimate, using the deep prior as regularization. (c) Normalized pointwise standard deviation.
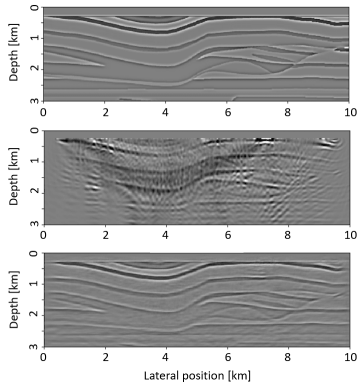

Figure 4: Loop-unrolled LSRTM on a test 2D slice of the 2D overthrust model. (a) True image, (b) RTM of simultaneous shot and (c) loop-unrolling of.

physics (Li et al., 2020a): fluid-flow physics, rock physics and wave physics. Given time-lapse seismic data collected over multiple years, we jointly invert for the rocks intrinsic permeability which can be used to recover the CO_2 concentration. By leveraging our high-level abstractions and AD, we can directly differentiate through all three physics solvers which map permeability to seismic data and minimize a fully coupled data misfit objective function. At each iteration, given an estimate of the permeability K, we model seismic data where the subsurface velocity is translated from the solution of the fluid-flow simulation. We can then compute the standard data misfit and use automatic differentiation to compute the permeability update. Because each step is abstracted, we can easily swap each physical solver with a different one, such as replacing the fluid-flow solver by a trained Fourier Neural Operator (FNO) for computational efficiency (Li et al., 2020b). We show in Figure 5 that we can recover the permeability from seismic measurements using a pre-trained FNO in place of a fluid-flow solver.

CONCLUSIONS

Through this paper, we have introduced a software design philosophy aimed at enabling research and development at scale in geoscience, with the potential to generalize to any inverse problem such as medical imaging. We demonstrated through carefully chosen example that high-level abstractions allow to express complex problem in a clear and representative way without incurring additional computational costs. Our soft-

```
1  opt = RMSprop()
2  for j=1:maxiter
3      theta = params(K)
4      grads = gradient(theta) do
5          c = S(K); v = R(c); d_pred = F(v)
6          return 0.5f0 * norm(d_pred-d_obs)^2f0
7      end
8      for p in theta
9          update!(opt, p, grads[p])
10     end
11 end
```

Listing 5: Example of an end-to-end coupled inversion for seismic monitoring of geological carbon storage. A pre-trained FNO S is used as a surrogate for fluid-flow simulation. In each iteration, we calculating the seismic data misfit, compute the gradient with respect to input permeability, K, via AD, and update it according to an optimizer opt.
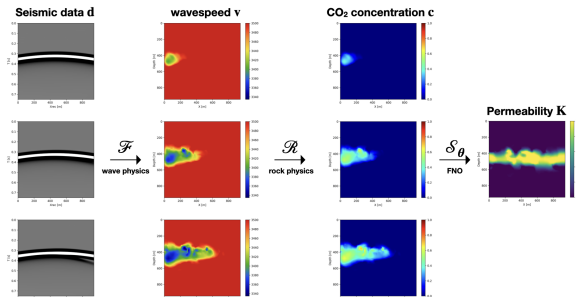

Figure 5: End-to-end inversion of CO_2 concentration from seismic measurements. In order to obtain this end-to-end result, the three operators representing the fluid-flow modeling, rock property modeling, wave modeling, and their derivatives are integrated by combining AD and manually defined gradients.

ware framework already allows for a wide range of application, including industry scale inversion and cutting-edge physics-informed deep learning for geophysics. We intend to build additional capabilities to tackle modern computing environment such as Cloud computing, task dedicated accelerators (i.e TPUs) and non-convex optimization.

ACKNOWLEDGMENT

References

Adler, J., and Öktem, O., 2017, Solving ill-posed inverse problems using iterative deep neural networks: Inverse Problems, **33**, 124007, doi: https://doi.org/10.48550/arXiv.1704.04058.

Adler, J., and Öktem, O., 2018, Learned primal-dual reconstruction: IEEE Transactions on Medical Imaging, **37**, 1322–1332, doi: https://doi.org/10.1109/TMI.2018.2799231.

Andrychowicz, M., M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, and N. De Freitas, 2016, Learning to learn by gradient descent by gradient descent: Advances in Neural Information Processing Systems, **29**, doi: https://doi.org/10.48550/arXiv.1606.04474.

Cheng, Z., M. Gadelha, S. Maji, and D. Sheldon, 2019, A Bayesian perspective on the deep image prior: Presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 5443–5451.

Dittmer, S., T. Kluth, P. Maass, and D. Otero Baguer, 2020, Regularization by architecture: A deep prior approach for inverse problems: Journal of Mathematical Imaging and Vision, **62**, 456–470, doi: https://doi.org/10.48550/arXiv.1812.03889.

Kumar, N., 2010, Parallelizing operations with ease using parallel sPOT: SINBAD. Retrieved from https://slim.gatech.edu/Publications/Public/Conferences/SINBAD/2010/Fall/kumar2010SINBADpoe/kumar2010SINBADpoe_pres.pdf.

Lempitsky, V., A. Vedaldi, and D. Ulyanov, 2018, Deep image prior: 2018 iEEE/CVF Conference on Computer Vision and Pattern Recognition, 9446–9454, doi: https://doi.org/10.1109/CVPR.2018.00984.

Li, D., K. Xu, J. M. Harris, and E. Darve, 2020, Coupled time-lapse full-waveform inversion for subsurface flow problems using intrusive automatic differentiation: Water Resources Research, **56**, e2019WR027032, doi: https://doi.org/10.1029/2019WR027032.

Li, Z., N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, 2020, Fourier neural operator for parametric partial differential equations: Machine Learning, doi: https://doi.org/10.48550/arXiv.2010.08895.

Lin, T. T., and F. J. Herrmann, 2015, The student-driven hPC environment at sLIM: Inaugural Full-Waveform Inversion Workshop.

Louboutin, M., 2021, Slimgroup/TimeProbeSeismic.jl: Zenodo, doi: https://doi.org/10.5281/zenodo.5605055.

Louboutin, M., and F. Herrmann, 2021a, Slimgroup/SlimOptim.jl: V0.1.6: Zenodo, doi: https://doi.org/10.5281/zenodo.5196691.

Louboutin, M., and F. Herrmann, 2021b, Ultra-low memory seismic inversion with randomized trace estimation: 91st Annual International Meeting, SEG, Expanded Abstracts, doi: https://doi.org/10.1190/segam2021-3584072.1.

Louboutin, M., M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, and G. J. Gorman, 2019, Devito (v3.1.0): An embedded domain-specific language for finite differences and geophysical exploration: Geoscientific Model Development, **12**, 1165–1187, doi: https://doi.org/10.5194/gmd-12-1165-2019.

Louboutin, M., P. Witte, Z. Yin, H. Modzelewski, and C. da Costa, 2022a, Slimgroup/JUDI.jl: V2.6.4: Zenodo, doi: https://doi.org/10.5281/zenodo.5893940.

Louboutin, M., Z. Yin, and F. Herrmann, 2022b, Slimgroup/JUDI4Cloud.jl: FIrst public release: Zenodo, doi: https://doi.org/10.5281/zenodo.6386831.

Luporini, F., M. Louboutin, M. Lange, N. Kukreja, P. Witte, J. Hückelheim, and G. J. Gorman, 2020, Architecture and performance of devito, a system for automated stencil computation: ACM Transactions on Mathematical Software, **46**, 1–28, doi: https://doi.org/10.1145/3374916.

Madagascar, 2017, Retrieved from http://www.ahay.org/wiki/Main_Page.

Modzelewski, H., M. Louboutin, Z. Yin, and R. Orozco, 2022, Slimgroup/JOLI.jl: V0.7.16: Zenodo, doi: https://doi.org/10.5281/zenodo.5933224.

Padula, A. D., S. D. Scott, and W. W. Symes, 2009, A software framework for abstract expression of coordinate-free linear algebra and optimization algorithms: ACM Transactions on Mathematical Software, **36**, 1–36, doi: https://doi.org/10.1145/1499096.1499097.

Peters, B., and F. J. Herrmann, 2019, Algorithms and software for projections onto intersections of convex and non- convex sets with applications to inverse problems: arXiv:1902.09699.

Rathgeber, F., D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. T. Mcrae, and P. H. J. Kelly, 2016, Firedrake: Automating the finite element method by composing abstractions: ACM Transactions on Mathematical Software, **43**, 1–27, doi: https://doi.org/10.1145/2998441.

Ravasi, M., and I. Vasconcelos, 2020, PyLops — A linear-operator python library for scalable algebra and optimization: SoftwareX, **11**, 100361, doi: https://doi.org/10.1016/j.softx.2019.100361.

van Rossum, G., and F. L. Drake, 2009, Python 3 reference manual: CreateSpace.

Siahkoohi, A., and M. Louboutin, 2021, Slimgroup/Devito4PyTorch: Initial release: Zenodo, doi: https://doi.org/10.5281/zenodo.4610087.

Siahkoohi, A., M. Louboutin, and F. J. Herrmann, 2019a, Neural network augmented wave-equation simulation: No. TR-CSE-2019-1: Georgia Institute of Technology. Retrieved from https://slim.gatech.edu/Publications/Public/TechReport/2019/siahkoohi2019TRnna/siahkoohi2019TRnna.html.

Siahkoohi, A., M. Louboutin, and F. J. Herrmann, 2019b, The importance of transfer learning in seismic modeling and imaging: Geophysics, **84**, doi: https://doi.org/10.1190/geo2019-0056.1.

Siahkoohi, A., G. Rizzuti, and F. J. Herrmann, 2020, A deep-learning based Bayesian approach to seismic imaging and uncertainty quantification: EAGE Annual Conference Proceedings. Retrieved from https://slim.gatech.edu/Publications/Public/Conferences/EAGE/2020/siahkoohi2020EAGEdlb/siahkoohi2020EAGEdlb.html.

Siahkoohi, A., G. Rizzuti, and F. J. Herrmann, 2021, Deep Bayesian inference for seismic imaging with tasks: arXiv:2110.04825.

Siahkoohi, A., D. J. Verschuur, and F. J. Herrmann, 2019c, Surface-related multiple elimination with deep learning: 89th Annual International Meeting, SEG, Expanded Abstracts, doi: https://doi.org/10.1190/segam2019-3216723.1.

Silva, C. D., and F. J. Herrmann, 2019, A unified 2D/3D large scale software environment for nonlinear inverse problems: ACM Transactions on Mathematical Software. Retrieved from https://slim.gatech.edu/Publications/Public/Journals/ACMTOMS/2019/dasilva2017uls/dasilva2017uls.html.

van den Berg, E., and M. P. Friedlander, 2009, Spot: A linear-operator toolbox for Matlab: SCAIM.

Wang, B., N. Zhang, W. Lu, and J. Wang, 2019, Deep-learning-based seismic data interpolation: A preliminary result: Geophysics, **84**, no. 1, V11–V20, doi: https://doi.org/10.1190/geo2017-0495.1.

Washbourne, J., S. Kaplan, M. Merino, U. Albertin, A. Sekar, C. Manuel, and A. Loddoch, 2021, Chevron optimization framework for imaging and inversion (cOFII) – An open source and cloud friendly Julia language framework for seismic modeling and inversion: First International Meeting for Applied Geoscience & Energy, SEG/AAPG, Expanded Abstracts, 792–796, doi: https://doi.org/10.1190/segam2021-3594362.1.

Welling, M., and Y. W. Teh, 2011, Bayesian learning via stochastic gradient Langevin dynamics: Presented at the Proceedings of the 28th International Conference on Machine Learning, 681–688, doi: https://doi.org/10.5555/3104482.3104568.

WesternGeco, 2012, Parihaka 3D PSTM Final Processing Report: No. New Zealand Petroleum Report 4582. New Zealand Petroleum & Minerals.

Witte, P. A., M. Louboutin, N. Kukreja, P. Witte, M. Lange, G. J. Gorman, and F. J. Herrmann, 2019a, A large- scale framework for symbolic implementations of seismic inversion algorithms in Julia: Geophysics, **84**, no. 3, F57–F71, doi: https://doi.org/10.1190/geo2018-0174.1.

Witte, P. A., M. Louboutin, and F. J. Herrmann, 2019b, JUDI4Flux: Seismic modeling for deep learning: Zenodo.

Yang, F., and J. Ma, 2019, Deep-learning inversion: A next-generation seismic velocity model building method: Geophysics, **84**, no. 4, R583–R599, doi: https://doi.org/10.1190/geo2018-0249.1.

Zhang, G., Z. Wang, and Y. Chen, 2018, Deep learning for seismic lithology prediction: Geophysical Journal International, **215**, 1368–1387, doi: https://doi.org/10.1093/gji/ggy344.