ML4Seismic Open Source Software environment

Mathias Louboutin et. al.

October 22, 2021

Georgia Institute of Technology



Georgia Tech College of Computing School of Computational Science and Engineering



Georgia Tech College of Sciences School of Earth and **Atmospheric Sciences**

Released to public domain under Creative Commons license type BY (https://creativecommons.org/licenses/by/4.0) **Copyright (c) 2021, Mathias Louboutin (Georgia Tech)**





Georgia Tech College of Engineering School of Electrical and Computer Engineering



Some history

SINBAD 2006-2017, UBC

- inversion, data processing, acquisition design
- (parallel) Matlab code
- commercially adopted & in production (e.g. seismic data acquisition w/ Compressive Sensing)

Closed proprietary software framework stifled innovation due to

- restricted access to commercial (Matlab) code
- issues w/ scalability on parallel clusters
- Imited uptake by sponsors

Moved to an open-source software (OSS) model in the Cloud.





ML4Seismic – Software

- To improve our software stack, we
- improve access via Open Source (MIT license)
- build on years of experience
- Ieverage hyperscale serverless Cloud computing
- employs high-level abstractions to combat complexity
- are engaged in several (international) collaborations



Open source -

$\leftarrow \rightarrow C \triangle$ \triangleq github.com/slimgroup		★ 🕸 🖻 G 🗯 M 🗄		
Search or jump to / Pull requests Is	ssues Marketplace Explore	Ļ + → 🧶 •		
Image: Solution of the second state	M group JSA 🔗 https://slim.gatech.edu 🗹 Felix.herrmann@ga ople 13 २२ Teams 2 ाण Projects 🔯	tech.edu Settings		
Pinned repositories Customize pinned repositories				
☐ JUDI.jI ∷ Julia Devito inversion. ● Julia ☆ 35 ♀12	ାଲି JOLI.ji ። Julia Operators LIbrary ● Julia ☆ 13 ୫ 4	InvertibleNetworks.jl \vdots A Julia framework for invertible neural networksJulia \bigstar 11 $\%$ 5		
SetIntersectionProjection.jl III Constrained optimization IIII	ExtendedConv.jl :: Julia implementation for extended convolutional layers	XConv III Memory efficient convolution networks		
Julia な 5 % 2	Julia	Julia		
Q Find a repository	Type - Language - Sort -	Rew New		

https://github.com/slimgroup



Demos

Afternoon breakout session:

- ML4Seismic software environment preinstalled
- showcase code development on simple problems
- allows for evaluation of software
- platform to interact w/ users

https://ml4shub.eastus.cloudapp.azure.com



https://github.com/slimgroup

Active & adopted by the community



30

25



[1] P. A. Witte, M. Louboutin, N. Kukreja, F. Luporini, M. Lange, G. G. Gorman and Felix J. Herrmann, 2019, A large-scale framework for symbolic implementations of seismic inversion algorithms in Julia, Geophysics, 84, 3. [2] M. Louboutin, M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, P. Velesko and G. J. Gorman, 2019, Devito 3.1.0: an embedded domain-specific language for finite differences and geophysical exploration: Geoscientific model development, 12, 3.

[3] F. Luporini, M. Lange, M. Louboutin, N. Kukreja, J. H<u>ückelheim</u>, C. Yount, P. A. Witte, P. H. J. Kelly, F. J. Herrmann and G. J. Gorman, 2019, Architecture and performance of Devito, a system for automated stencil computation, arXiv preprints.

High-level Abstractions for Seismic Inversion & ML

JUDI.jl & JUDI4Cloud.jl:^[1]

- matrix-free linear operators and abstract data (SEGY) containers
- parallel I/O based on look-up tables, parallelism (OMP + soon MPI)
- interface to Julia's machine learning package Flux.jl
- <u>https://github.com/slimgroup/JUDI.jl/</u> (MIT license)

Devito*^{[2][3]} *Devito is a joint project between Imperial & Georgia Tech with Mathias one of the two main developers.

- a domain-specific language for finite-differences in Python
- code generation with automated performance optimizations
- model parallelism (MPI, multithreading)
- <u>https://github.com/devitocodes/devito</u> (MIT license)
- Adopted by industry (CVX, BP, DUG, ...)

InvertibleNetworks.jl:

- building blocks for invertible neural networks/normalizing flows
- scalable & gradients derived by hand
- https://github.com/slimgroup/InvertibleNetworks.jl (MIT license)



State-of-the-art performance

Modeling performance on AMD 7VI2 2.5 Ghz





Devito – seismic

Louboutin, Mathias, et al. "Devito (v3. 1.0): an embedded domain-specific language for finite differences and geophysical exploration." *Geoscientific Model Development* 12.3 (2019): 1165-1187. Lange, Michael, et al. "Devito: Towards a generic finite difference dsl using symbolic python." 2016 6th Workshop on Python for High-Performance and Scientific Computing (PyHPC). IEEE, 2016.

GPU support roadmap

- Support for multiple target languages OpenMP, OpenACC
 - potentially: CUDA, HIP, SYCL, ... Unreliability of the target languages' software stack
- Multi-GPU support:
 - Make it possible to run different shots on different GPUs Single-node multi-GPU via domain decomposition Multi-node multi-GPU via domain decomposition
 - Data movement (optimized) Data streaming (optimized)





Done Nearly done In progress **Potentially later this year**

Kernel performance (best so far: 27 GPOINTS on iso-acoustic O(2, 8))



JUDI – true vertical integration



parallel modeling function

serial modeling function

automatic code generation and JIT compilation

students

math/optimizers/cs/ seismic practitioners

students

CS/math/physics people

polyhydral compiler people



[1] P. A. Witte, M. Louboutin, F. Luporini, G. G. Gorman and Felix J. Herrmann, 2019, *Compressive least-squares migration with on-the-fly Fourier transforms*, Geophysics, 84, 5.

Example : Compressive seismic imaging

Least-squares migration as an elastic net:^[1]

$$\begin{array}{ll} \text{minimize} & \lambda \left\| \mathbf{C} \,\delta \mathbf{m} \right\|_{1} + \frac{1}{2} \left\| \mathbf{C} \,\delta \mathbf{m} \right\|_{1} \\ \text{subject to:} & \sum_{i=1}^{n_{s}} \sum_{j=1}^{n_{f}} \left\| \mathbf{M}_{l}^{-1} \mathbf{J} \mathbf{M}_{r}^{-1} \right\|_{1} \end{array}$$

Solve via the linearized Bregman method:

- at each iteration: random subsets of sources + frequencies
- memory per source: $\mathcal{O}(ar{n}_f)$
- compressive sensing: no. of samples \sim no. of grid points, non-zero entries

$\delta \mathbf{m} \|_2^2$

 $|\delta \mathbf{m} - \mathbf{M}_l^{-1} \bar{\mathbf{d}}_{ij}|| \leq \sigma$



Compressive seismic imaging

Linearized Bregman method with JUDI:

for j=1:maxiter

```
# Compute residual and gradient
i = randperm(d_obs.nsrc)[1:batchsize_source]
select_frequencies!(J, batchsize_freq)
r = Ml*J[i]*Mr*x - Ml*d_obs[i]
g = Mr'*J[i]'*Ml'*proj_l2(r)
```

```
# Residual and function value
res[j] = norm(r, 2)
fval[j] = \lambda * norm(C*z, 1) + .5f0 * norm(C*z, 2)^2
```

```
# Update variables
    global z -= \alpha * g
    global x = C*soft_thresholding(C*z, \lambda)
end
```





Compressive seismic imaging

Linearized Bregman method with JUDI:

for j=1:maxiter

```
# Compute residual and gradient
i = randperm(d_obs.nsrc)[1:batchsize_source]
select_frequencies!(J, batchsize_freq)
r = Ml*J[i]*Mr*x - Ml*d_obs[i]
g = Mr'*J[i]'*Ml'*proj_l2(r)
```

```
# Residual and function value
res[j] = norm(r, 2)
fval[j] = \lambda * norm(C*z, 1) + .5f0 * norm(C*z, 2)^2
```

```
# Update variables
    global z -= \alpha * g
    global x = C*soft_thresholding(C*z, \lambda)
end
```





Ultra-long offset SEG workshop - 5 iterations 4 shots each





x, iter 5



Witte, Philipp A., et al. "An event-driven approach to serverless seismic imaging in the cloud." IEEE Transactions on Parallel and Distributed Systems 31.9 (2020): 2032-2049.

Serverless seismic imaging on Azure "Small" 3D Imaging case study w/ Devito DSL + JUDI

- Data set: 1,500 source locations (~2.1 TB data)
- Model: 10 x 10 x 3.325 km (**270 million** unknowns)
- PDE: tilted transversely isotropic (TTI) wave equation, 3,500 time steps
- **Cost:** < 10,000\$ on 100 E64/E64s instances (2 VMs per gradient with MPI)
- Peak performance: 140 TFLOPs
- cost single image is comparable to training a large NN







JUDI4Cloud



students

math/optimizers/cs/
seismic practitioners

students

CS/math/physics people

polyhydral compiler people



Clusterless implementation on Azure

JUDI operators in Azure Batch w/ JUDI4Cloud.jl

JUDI4Cloud.init_clusterless(n_instances; n_julia_per_instance=2) # Setup operators Pr = judiProjection(info, recGeometry) F = judiModeling(info, model; options=opt) F0 = judiModeling(info, model0; options=opt) Ps = judiProjection(info, srcGeometry) J = judiJacobian(Pr*F0*adjoint(Ps), q) # Nonlinear modeling

dobs = Pr*F*adjoint(Ps)*q

Once again, abstractions pay off:

- no need to refactor code
- readable codes

Setup Azure Batch (see previous talk)



SLIM 🔶 ML4Seismic

Compressive seismic imaging

Linearized Bregman method with JUDI:

for j=1:maxiter

```
# Compute residual and gradient
i = randperm(d_obs.nsrc)[1:batchsize_source]
select_frequencies!(J, batchsize_freq)
r = Ml*J[i]*Mr*x - Ml*d_obs[i]
g = Mr'*J[i]'*Ml'*proj_l2(r)
```

```
# Residual and function value
res[j] = norm(r, 2)
fval[j] = \lambda * norm(C*z, 1) + .5f0 * norm(C*z, 2)^2
```

```
# Update variables
    global z -= \alpha * g
    global x = C*soft_thresholding(C*z, \lambda)
end
```











JUDI the Julia Devito Inversion framework

Additional advantages:

- Devito propagators usable as standalone python
- SEGY I/O
- Extended source imaging (rank 1 source q[t]w[x])
- GPU offloading, simple ENV-variable change
- Low-memory gradients with randomized trace estimation (See Tue)

Future work: expose in JUDI

- Devito's MPI domain decomposition
- AD support



Designed for Interoperability

Combine JOLI matrix-free linear operators w/ COFII's Jets

using JOLI, JetPack, Jets
1D real DFT with 1000 samples
J = joDFT(1000; DDT=Float64, RDT=Float64)
Jet operator for circular shift by 25 samples
jop = JopCircShift(JetSpace(Float64, 1000), 25);

a = randn(1000)
Parenthesis needed to prevent julia to compute J*jop first
a = randn(1000);b = J*(jop*a); c = jop'*(J'*b);
dot test
@show dot(b, b), dot(c, a), dot(b, b) - dot(c, a), dot(b, b) / dot(c, a)
490.00330785, 490.003307859, 1.1368683772161603e-13, 1.000000000000002)

Once again, abstractions pay off:

- no need to refactor code
- readable codes

Seamless integration JUDI & COFII \implies super fast developments cycle = innovation.



https://github.com/slimgroup/ConstrainedFWIExamples/blob/master/notebooks/02_constr_fwi_jetpack.ipynb https://github.com/ChevronETC/Examples/blob/main/50 fwi/11 constrained fwi pqn.ipynb

Mored advanced integration synergetic



SetIntersectionProjection.jl

Step Length Function Val Opt Cond 0.00000e+00 5.47293e+06 1.00000e+01

5.28328e+06 9.59644e+00 1.00000e+00



https://github.com/ChevronETC/Examples/blob/main/50_fwi/11_constrained_fwi_pqn.ipynb

Mored advanced integration constrains from SLIM & wave propagators from CVX's COFII





COFII vs SLIM

	SLIM	COFII	Compatible
Language	julia, python	julia, C++	yes
Software philosophy	Linear Algebra	Jets	yes (see above)
Propagators	Devito	Legacy C++, potential migration to Devito	NA
IO	SEGY (via SegylO), no cloud support yet	JavaSeis CloudSeis	not yet
Cloud	Azure Batch (via AzureClusterlessHPC) fully remote (local desktop/laptop)	Azure Virtual Machine ScaleSet Custom cloud cluster manager Azure Network only	no
Cloud storage	OpenVDS (future plan) Cloud native OSDU compliant	Blob manager	yes
OSS	Yes, all	Yes, transitioning	Yes



ML Software

InvertibleNetworks.jl: A Julia package for invertible CNNs

- Manually implemented derivatives and Jacobians that make use of invertibility
- Unit testing (adjoints, gradients, invertibility)
- <u>https://github.com/slimgroup/InvertibleNetworks.jl</u> (MIT license)
- GPU support through CUDA.jl

Main features:

- Additive and affine invertible coupling layers (e.g. as in NICE, Glow)
- Invertible recurrent inference machines (i-RIM)
- Hyperbolic networks
- Hierarchical invertible neural transport (HINT) and conditional HINT
- 1 x 1 convolutions, wavelet squeezing, etc.
- Integration of Flux.jl layers
- AD with ChainRules.jl



InvertibleNetworks.jl

Class structure:

struct CouplingLayerGlow <: NeuralNetLayer</pre> C::Conv1x1 **RB::**ResidualBlock logdet::Bool forward::Function inverse::Function backward::Function

end

Usage:

```
# Create Layer
LayerGlow = CouplingLayerGlow(nx, ny, nchannel, nhidden, batchsize)
# Operations
Y = LayerGlow.forward(X)
X_ = LayerGlow.inverse(Y)
```

```
\Delta X, X_{-} = LayerGlow.backward(\Delta Y, Y)
```



Loop unrolling revisited

Invertible version of loop-unrolled network:

- Inspired by i-RIM and Real NVP
- Invertible + tractable logdet







Uncertainty Quantification

Examples:



Siahkoohi, Ali, and Felix J. Herrmann. "Learning by example: fast reliability-aware seismic imaging with normalizing flows." arXiv preprint arXiv:2104.06255 (2021).

Orozco, Rafael, et al. "Photoacoustic imaging with conditional priors from normalizing flows." NeurIPS 2021 Workshop on Deep Learning and Inverse Problems. 2021.

Uncertainty Quantification

Posterior sampling implemented with conditional layers in InvertibleNetworks.jl

struct ConditionalLayerHINT <: NeuralNetLayer</pre> CL_X::CouplingLayerHINT CL_Y::CouplingLayerHINT CL_YX::CouplingLayerBasic C_X::Union{Conv1x1, Nothing} C_Y::Union{Conv1x1, Nothing} logdet::Bool is_reversed::Bool end

First, train with joint distribution target then sample from posterior distribution:

```
Zy_fixed = H.forward_Y(y_obs);
post_sampling_size = 50
Zx = randn(Float32, nx, ny, n_in, post_sampling_size)
X_post = H.inverse(Zx, Zy_fixed.*ones(Float32, nx, ny, n_in, post_sampling_size))[1];
#statistics over pixels in batch
mean(X_post; dims=4)
std(X_post; dims=4)
```

Kruse, Jakob, et al. "HINT: Hierarchical invertible neural transport for density estimation and Bayesian inference." arXiv preprint arXiv:1905.10687 (2019)

Next steps

Scaling of ML in 3D

Integration w/ OLIVES & generative classifiers

Further integration AD tools to combine wave-equation solvers w/ ML

Integration of Julia's new AD tool when available

- Ieverages JIT compiler
- \blacktriangleright offers more control \rightarrow scales well

This research was carried out with the support of Georgia Research Alliance and partners of the ML4Seismic consortium.

