




hIPPYlib: An Extensible Software Framework for Large-Scale Inverse Problems

Umberto Villa¹, Noemi Petra², and Omar Ghattas³

¹ Institute for Computational Engineering & Sciences, The University of Texas at Austin ² Applied Mathematics, School of Natural Sciences, University of California, Merced ³ Institute for Computational Engineering & Sciences, Department of Mechanical Engineering, and Department of Geological Sciences, The University of Texas at Austin

DOI: [10.21105/joss.00940](https://doi.org/10.21105/joss.00940)

Software

- [Review](#) 
- [Repository](#) 
- [Archive](#) 

Submitted: 27 August 2018

Published: 23 October 2018

License

Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License ([CC-BY](#)).

Summary

hIPPYlib (Villa, Petra, & Ghattas, 2018) implements state-of-the-art scalable algorithms for deterministic and Bayesian inverse problems governed by partial differential equations (PDEs); see, e.g., (Bui-Thanh, Ghattas, Martin, & Stadler, 2013), (N. Petra, Martin, Stadler, & Ghattas, 2014) for an introduction to infinite dimensional PDE-constrained Bayesian inverse problems. **hIPPYlib** builds on the parallel finite element library **FEniCS**, (Logg & Wells, 2010), (Logg, Mardal, & Wells, 2012), (Langtangen & Logg, 2017), for the discretization of the PDE and on **PETSc**, (Balay et al., 2018) for scalable and efficient linear algebra operations and solvers.

Conceptually, **hIPPYlib** can be viewed as a toolbox that provides the building blocks for experimenting new ideas and developing scalable algorithms for PDE-constrained deterministic and Bayesian inverse problems.

In **hIPPYlib** the user can express the forward PDE and the likelihood in weak form using the friendly, compact, near-mathematical notation of **FEniCS**, which will then automatically generate efficient code for the discretization. Linear and nonlinear, and stationary and time-dependent PDEs are supported in **hIPPYlib**. For stationary problems, gradient and Hessian information can be automatically generated by **hIPPYlib** using **FEniCS** symbolic differentiation of the relevant weak forms. For time-dependent problems, instead, symbolic differentiation can only be used for the spatial terms, and the contribution to gradients and Hessians arising from the time dynamics needs to be provided by the user.

Noise and prior covariance operators are modeled as inverses of elliptic differential operators allowing us to build on existing fast multigrid solvers for elliptic operators without explicitly constructing the dense covariance operator.

The key property of the algorithms underlying **hIPPYlib** is that solution of the deterministic and Bayesian inverse problem is computed at a cost, measured in forward PDE solves, that is independent of the parameter dimension.

hIPPYlib provides a robust implementation of the inexact Newton-conjugate gradient algorithm to compute the maximum a posteriori (MAP) point. The gradient and Hessian actions are computed via their weak form specification in **FEniCS** by constraining the state and adjoint variables to satisfy the forward and adjoint problem. The Newton system is solved inexactly by early termination of CG iterations via Eisenstat-Walker (to prevent oversolving) and Steihaug (to avoid negative curvature) criteria. Two globalization techniques are available to the user: Armijo back-tracking line search and trust region.

In `hIPPYlib`, the posterior covariance is approximated by the inverse of the Hessian of the negative log posterior evaluated at the MAP point. This Gaussian approximation is exact when the parameter-to-observable map is linear; otherwise, its logarithm agrees to two derivatives with the log posterior at the MAP point, and thus it can serve as a proposal for Hessian-based Markov chain Monte Carlo (MCMC) methods. `hIPPYlib` makes the construction of the posterior covariance tractable by invoking a low-rank approximation of the Hessian of the log likelihood.

`hIPPYlib` also offers scalable methods for sample generation. To sample large scale spatially correlated Gaussian random fields from the prior distribution, `hIPPYlib` implements a new method that strongly relies on the structure of the covariance operator defined as the inverse of a differential operator: by exploiting the assembly procedure of finite element matrices `hIPPYlib` constructs a sparse Cholesky-like rectangular decomposition of the precision operator. To sample from a local Gaussian approximation to the posterior (such as at the MAP point) `hIPPYlib` exploits the low rank factorization of the Hessian of the log likelihood to correct samples from the prior distribution. Finally, to explore the posterior distribution, `hIPPYlib` implements dimension independent MCMC sampling methods enchanted by Hessian information.

Finally, randomized and probing algorithms are available to compute the pointwise variance of the prior/posterior distribution and the trace of the covariance operator.

The source code for `hIPPYlib` has been archived to Zenodo with the linked DOI [10.5281/zenodo.596931](https://doi.org/10.5281/zenodo.596931).

Acknowledgements

`hIPPYlib` development is partially supported by the U.S. National Science Foundation (NSF), Software Infrastructure for Sustained Innovation (SI2: SSE & SSI) Program under grants ACI-1550593, ACI-1550547; the Defense Advanced Research Projects Agency, Enabling Quantification of Uncertainty in Physical Systems Program under grant CBET-1508713; the US National Science Foundation, Division Of Chemical, Bioengineering, Environmental, & Transport Systems under grant CBET-1508713; and the Air Force Office of Scientific Research, Computational Mathematics program under grant FA9550-12-1-81243.

References

- Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., et al. (2018). PETSc Web page. <http://www.mcs.anl.gov/petsc>. Retrieved from <http://www.mcs.anl.gov/petsc>
- Bui-Thanh, T., Ghattas, O., Martin, J., & Stadler, G. (2013). A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6), A2494–A2523. doi:[10.1137/12089586X](https://doi.org/10.1137/12089586X)
- Langtangen, H. P., & Logg, A. (2017). *Solving PDEs in Python*. Springer. doi:[10.1007/978-3-319-52462-7](https://doi.org/10.1007/978-3-319-52462-7)
- Logg, A., & Wells, G. N. (2010). DOLFIN: Automated finite element computing. *ACM Transactions on Mathematical Software*, 37(2), Article 20. doi:[10.1145/1731022.1731030](https://doi.org/10.1145/1731022.1731030)
- Logg, A., Mardal, K.-A., & Wells, G. (2012). *Automated solution of differential equations by the finite element method: The FEniCS book* (Vol. 84). Springer Science & Business Media. doi:[10.1007/978-3-642-23099-8](https://doi.org/10.1007/978-3-642-23099-8)

Petra, N., Martin, J., Stadler, G., & Ghattas, O. (2014). A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet inverse problems. *SIAM Journal on Scientific Computing*, 36(4), A1525–A1555. doi:[10.1137/130934805](https://doi.org/10.1137/130934805)

Villa, U., Petra, N., & Ghattas, O. (2018). HIPPYlib: An extensible software framework for large-scale inverse problems; Part I: Deterministic inversion and linearized Bayesian inference. *In preparation*.