





Numerics: A .NET Library for Numerical Computing, Statistical Analysis, and Risk Assessment

C. Haden Smith ¹, Woodrow L. Fields ¹, Julian Gonzalez ¹, Sadie Niblett ¹, Brennan Beam ², and Brian Skahill ³

¹ U.S. Army Corps of Engineers, Risk Management Center, Lakewood, Colorado, USA ² U.S. Army Corps of Engineers, Hydrologic Engineering Center, Davis, California, USA ³ Fariborz Maseeh Department of Mathematics and Statistics, Portland State University, Portland, Oregon, USA  Corresponding author

DOI: [10.21105/joss.10540](https://doi.org/10.21105/joss.10540)

Software

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

Editor: [Owen Lockwood](#)  

Reviewers:

- [@KaranamLokesh](#)
- [@martinmodrak](#)

Submitted: 06 April 2026

Published: 24 June 2026

License

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

Summary

Numerics is a free and open-source numerical computing library for .NET, developed by the U.S. Army Corps of Engineers Risk Management Center (USACE-RMC). The library provides 43 univariate probability distributions with L-moment and maximum likelihood parameter estimation, eight Markov Chain Monte Carlo (MCMC) samplers for Bayesian inference, bootstrap uncertainty quantification, bivariate copulas, optimization algorithms, and numerical methods for integration, differentiation, root finding, and linear algebra. Numerics targets engineers, scientists, and analysts working in infrastructure risk assessment, flood frequency analysis, and Monte Carlo simulation within .NET-based enterprise systems. The library supports .NET 8.0 through 10.0 and .NET Framework 4.8.1 and is validated by over 1,000 unit tests against published reference values.

Statement of Need

Infrastructure risk assessment for dams and levees requires the integration of extreme value statistics, uncertainty quantification, and Monte Carlo simulation. Practitioners face challenges including small sample sizes from historical records, regulatory requirements for specific statistical methods such as the Log-Pearson Type III distribution mandated by USGS Bulletin 17C ([England et al., 2018](#)), robust outlier handling via the Multiple Grubbs-Beck Test ([Cohn et al., 2013](#)), and the need to incorporate expert judgment through Bayesian inference.

The .NET framework is a widely used platform for enterprise and government desktop applications in the United States, yet the ecosystem lacks comprehensive numerical libraries tailored for engineering risk analysis. Python offers mature tools such as NumPy ([Harris et al., 2020](#)) and SciPy ([Virtanen et al., 2020](#)), and R provides specialized hydrology packages such as `lmom` ([Hosking, 2019](#)), `nsRFA` ([Viglione, 2024](#)), and `SpatialExtremes` ([Ribatet, 2022](#)), but neither ecosystem integrates with .NET-based engineering workflows without introducing language interoperability complexity, performance overhead, and deployment challenges for regulated government systems. In addition, Numerics benefits from .NET's ahead-of-time compilation, absence of a global interpreter lock, and native `Parallel.For` support, delivering significantly higher throughput for computationally intensive methods such as bootstrap resampling and MCMC sampling compared to interpreted alternatives.

Numerics fills this gap by providing domain-specific capabilities within .NET:

- **Hydrology-specific distributions:** Log-Pearson Type III, Generalized Extreme Value, Pearson Type III, Generalized Pareto, and Kappa-Four, with L-moment parameter

estimation ([Hosking, 1990](#); [Hosking & Wallis, 1997](#)) that outperforms conventional moments for the small samples typical of flood records.

- **Mixed-population and competing-risk models:** Mixture distributions ([Alila & Mtiraoui, 2002](#); [Waylen & Woo, 1982](#)) for sites where floods arise from distinct causal mechanisms such as rainfall, snowmelt, or dam-regulated releases, and competing-risks models ([Bedford & Cooke, 2001](#); [Crowder, 2001](#)) for reliability analysis with multiple failure modes.
- **Bayesian inference:** Eight MCMC samplers, including Random Walk Metropolis-Hastings, Adaptive RWMH ([Haario et al., 2001](#)), Differential Evolution MCMC (DE-MCz and DE-MCzs) ([Braak & Vrugt, 2008](#)), Hamiltonian Monte Carlo ([Neal, 2011](#)), No-U-Turn Sampler (NUTS) ([Hoffman & Gelman, 2014](#)), Gibbs, and Self-Normalizing Importance Sampling (SNIS), with improved Gelman-Rubin convergence diagnostics ([Gelman & Rubin, 1992](#); [Vehtari et al., 2021](#)).
- **Uncertainty quantification:** Bootstrap resampling methods ([Efron & Tibshirani, 1993](#)) for confidence intervals on design estimates, a requirement in dam and levee safety risk assessments.
- **Global optimization:** Differential Evolution ([Storn & Price, 1997](#)), Shuffled Complex Evolution ([Duan et al., 1994](#)), Particle Swarm Optimization ([Kennedy & Eberhart, 1995](#)), Multi-Level Single-Linkage (MLSL) ([Rinnooy Kan & Timmer, 1987](#)), and Nelder-Mead ([Nelder & Mead, 1965](#)) for calibrating complex, multi-modal objective functions.
- **Adaptive numerical integration:** Gauss-Kronrod quadrature ([Piessens et al., 1983](#)) and VEGAS adaptive Monte Carlo integration ([Lepage, 1978](#)) for efficiently evaluating complex risk integrals involving multiple system components and failure modes, a core computation in quantitative risk assessment.
- **Machine learning:** Supervised and unsupervised algorithms, including generalized linear models, decision trees, random forests, k-nearest neighbors, k-means clustering, and Gaussian mixture models for regression, classification, and clustering tasks in risk assessment workflows.

State of the Field

General-purpose .NET numerical libraries exist, most notably Math.NET Numerics ([Ruegg et al., 2024](#)), which provides linear algebra, probability distributions, and basic statistics. However, Math.NET Numerics does not include L-moment estimation, hydrology-specific distributions (Log-Pearson Type III, Kappa-Four), mixture and competing-risk models, adaptive integration, bootstrap resampling, copulas, or global optimization algorithms. Math.NET Numerics includes basic MCMC samplers (Metropolis-Hastings, Hybrid Monte Carlo, Slice), but lacks the adaptive and ensemble methods (Adaptive RWMH, DE-MCzs, NUTS) commonly applied for Bayesian inference of hydrologic models. Contributing these features to Math.NET Numerics was not pursued because the scope of domain-specific functionality, the distinct API design requirements for risk assessment workflows, and the need for long-term maintenance by a domain-expert organization warranted an independent library.

In other language ecosystems, Python's SciPy ([Virtanen et al., 2020](#)) provides broad numerical capabilities but lacks specialized hydrological distributions and L-moment estimation. R packages such as `lmom` ([Hosking, 2019](#)) and `evd` ([Stephenson, 2002](#)) offer these features individually, but integrating R into .NET production applications introduces runtime dependencies and complicates deployment in regulated government environments. No single package in any ecosystem consolidates mixture distributions, competing-risk models, adaptive integration, extreme-value statistics, global optimization, bootstrapping, machine learning, and MCMC into a unified library for engineering risk analysis.

Numerics consolidates these capabilities into a single, self-contained .NET library purpose-built for quantitative risk assessment in water resources engineering. Computationally intensive operations, including bootstrap resampling and MCMC chain evaluation, are parallelized using

Parallel.For for high-throughput execution on modern multi-core hardware. Backed by long-term USACE-RMC maintenance, Numerics provides the ecosystem, performance, and institutional support needed for quantitative risk analysis in dam and levee safety.

Research Impact Statement

Numerics serves as the computational engine for six USACE-RMC production applications used for infrastructure safety decisions:

- **RMC-BestFit**: Bayesian estimation and fitting software that supports flood frequency analysis for flood risk management and dam and levee safety studies.
- **RMC-RFA**: Reservoir frequency analysis software to facilitate flood hazard assessments for dam safety studies.
- **RMC-TotalRisk**: Quantitative risk analysis software designed to support dam and levee safety investment decisions.
- **LifeSim**: A consequence estimation engine used to simulate population redistribution during an evacuation.
- **Levee Screening Tool**: A web-based portfolio-level levee risk screening tool that has been used to evaluate more than 7,500 levee segments across the national inventory.
- **Dam Screening Tool**: A web-based portfolio-level dam risk screening tool that will be used to evaluate more than 90,000 dams across the national inventory.

These applications support risk-informed decisions for thousands of dams and levees managed by USACE, other federal agencies, and private dam and levee owners, directly affecting public safety for millions of Americans. The RMC software suite is also used internationally by organizations in Australia, the Netherlands, and Canada. By open-sourcing Numerics, USACE-RMC enables independent verification of the algorithms underpinning these critical assessments and invites contributions from the broader water resources engineering community.

Software Design

The library is organized around abstract base classes that define consistent interfaces for each computational domain. `UnivariateDistributionBase` provides PDF, CDF, inverse CDF, log-likelihood evaluation, and random variate generation for all 43 univariate distributions. Parameter estimation is decoupled from distribution classes through segregated interfaces (`IMaximumLikelihoodEstimation`, `ILinearMomentEstimation`, `IMomentEstimation`), so each distribution declares which estimation methods it supports and analysts can swap candidate distributions without changing workflow code. `Optimizer` provides a common minimization interface shared by all local and global optimization algorithms, with built-in function evaluation tracking, convergence testing, and numerical Hessian computation. `MCMCSampler` manages chain initialization, parallel execution via `Parallel.For`, thinning, and posterior output collection for all eight samplers, while convergence diagnostics are handled by a separate `MCMCDiagnostics` utility class. Bivariate copulas follow the same pattern through `BivariateCopula` and its `ArchimedeanCopula` specialization.

The library has zero external runtime dependencies, relying only on .NET Base Class Libraries, which eliminates NuGet dependency conflicts in large government applications and simplifies deployment in regulated environments. Multi-targeting (.NET 8.0, 9.0, 10.0, and .NET Framework 4.8.1) ensures compatibility with both modern and legacy enterprise systems. Because .NET 8 and later are cross-platform, Numerics runs on Windows, Linux, and macOS, supporting both desktop engineering tools and cloud-hosted web applications. Python users can also access Numerics through PythonNet, making the library's performance and capabilities available to the broader scientific and academic community.

Continuous integration via GitHub Actions runs the test suite (over 1,000 tests validated

against published references (England et al., 2018; Press et al., 2007; Rao & Hamed, 2000)) on every pull request, preventing regressions in numerical accuracy.

Example

The following example illustrates a bootstrap uncertainty analysis (Efron & Tibshirani, 1993) for flood frequency quantiles, a common workflow in dam and levee safety studies. A Log-Pearson Type III distribution is fit to a synthetic sample using L-moments, and bias-corrected and accelerated (BCa) confidence intervals are computed for a range of non-exceedance probabilities:

```
using Numerics.Distributions;

// Generate a synthetic flood record
int n = 80;
int B = 10000;
var lp3 = new LogPearsonTypeIII(
    meanOfLog: 3,
    standardDeviationOfLog: 0.5,
    skewOfLog: 0.25);
double[] sample = lp3.GenerateRandomValues(n, 12345);

// Estimate parameters using L-moments
lp3.Estimate(sample, ParameterEstimationMethod.MethodOfLinearMoments);

// Configure the bootstrap analysis
var boot = new BootstrapAnalysis(
    distribution: lp3,
    estimationMethod: ParameterEstimationMethod.MethodOfLinearMoments,
    sampleSize: n,
    replications: B,
    seed: 12345);

// Non-exceedance probabilities for quantile estimation
var probabilities = new double[] { 0.999, 0.99, 0.9, 0.8, 0.7, 0.5, 0.3,
    0.2, 0.1, 0.05, 0.02, 0.01 };

// Compute 90% BCa confidence intervals
var CIs = boot.BCaQuantileCI(sample, probabilities, alpha: 0.1);
```

AI Usage Disclosure

Anthropic Claude Opus (versions 4.5, 4.6, and 4.7) was used to assist with the following: drafting markdown documentation content, reviewing and editing XML documentation within the source code, and reviewing code for bugs. All core library code, class architecture, numerical methods, and algorithms were designed and implemented by the authors. The authors reviewed, edited, and validated all AI-assisted output, made all core design decisions, and accept full responsibility for the work.

Acknowledgements

The authors acknowledge the U.S. Army Corps of Engineers Risk Management Center for supporting the development and open-source release of this software. The Numerics software

would not exist without the support of Risk Management Center leadership, in particular the RMC Director Bryant A. Robbins and RMC Lead Engineers David A. Margo (retired) and John F. England. The authors are grateful to all who have contributed to the development of Numerics and the content of this paper.

References

- Alila, Y., & Mtiraoui, A. (2002). Implications of heterogeneous flood-frequency distributions on traditional stream-discharge prediction techniques. *Hydrological Processes*, *16*(5), 1065–1084. <https://doi.org/10.1002/hyp.346>
- Bedford, T., & Cooke, R. (2001). *Probabilistic risk analysis: Foundations and methods*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511813597>
- Braak, C. J. F. ter, & Vrugt, J. A. (2008). Differential evolution Markov Chain with snooker updater and fewer chains. *Statistics and Computing*, *18*(4), 435–446. <https://doi.org/10.1007/s11222-008-9104-9>
- Cohn, T. A., England, J. F., Berenbrock, C. E., Mason, R. R., Stedinger, J. R., & Lamontagne, J. R. (2013). A generalized Grubbs-Beck test statistic for detecting multiple potentially influential low outliers in flood series. *Water Resources Research*, *49*(8), 5047–5058. <https://doi.org/10.1002/wrcr.20392>
- Crowder, M. J. (2001). *Classical competing risks*. Chapman; Hall/CRC. <https://doi.org/10.1201/9781420035902>
- Duan, Q., Sorooshian, S., & Gupta, V. K. (1994). Optimal use of the SCE-UA global optimization method for calibrating watershed models. *Journal of Hydrology*, *158*(3–4), 265–284. [https://doi.org/10.1016/0022-1694\(94\)90057-4](https://doi.org/10.1016/0022-1694(94)90057-4)
- Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. Chapman & Hall. <https://doi.org/10.1201/9780429246593>
- England, J. F., Cohn, T. A., Faber, B. A., Stedinger, J. R., Thomas, W. O., Veilleux, A. G., Kiang, J. E., & Mason, R. R. (2018). *Guidelines for determining flood flow frequency—Bulletin 17C* (Techniques and Methods Book 4, Chapter B5; p. 148). U.S. Geological Survey. <https://doi.org/10.3133/tm4B5>
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, *7*(4), 457–472. <https://doi.org/10.1214/ss/1177011136>
- Haario, H., Saksman, E., & Tamminen, J. (2001). An adaptive Metropolis algorithm. *Bernoulli*, *7*(2), 223–242. <https://doi.org/10.2307/3318737>
- Harris, C. R., Millman, K. J., Walt, S. J. van der, Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., & others. (2020). Array programming with NumPy. *Nature*, *585*, 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, *15*(47), 1593–1623. <https://jmlr.org/papers/v15/hoffman14a.html>
- Hosking, J. R. M. (1990). L-moments: Analysis and estimation of distributions using linear combinations of order statistics. *Journal of the Royal Statistical Society: Series B (Methodological)*, *52*(1), 105–124. <https://doi.org/10.1111/j.2517-6161.1990.tb01775.x>
- Hosking, J. R. M. (2019). *Lmom: L-moments*. <https://doi.org/10.32614/cran.package.lmom>
- Hosking, J. R. M., & Wallis, J. R. (1997). *Regional frequency analysis: An approach based on l-moments*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511529443>
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. *Proceedings of ICNN'95*

- *International Conference on Neural Networks*, 4, 1942–1948. <https://doi.org/10.1109/ICNN.1995.488968>
- Lepage, G. P. (1978). A new algorithm for adaptive multidimensional integration. *Journal of Computational Physics*, 27(2), 192–203. [https://doi.org/10.1016/0021-9991\(78\)90004-9](https://doi.org/10.1016/0021-9991(78)90004-9)
- Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. L. Jones, & X.-L. Meng (Eds.), *Handbook of markov chain monte carlo* (pp. 113–162). CRC Press. <https://doi.org/10.1201/b10905-6>
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer Journal*, 7(4), 308–313. <https://doi.org/10.1093/comjnl/7.4.308>
- Piessens, R., Doncker-Kapenga, E. de, Überhuber, C. W., & Kahaner, D. K. (1983). *QUADPACK: A subroutine package for automatic integration*. Springer-Verlag. <https://doi.org/10.1007/978-3-642-61786-7>
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (2007). *Numerical recipes: The art of scientific computing* (3rd ed.). Cambridge University Press. ISBN: 978-0521880688
- Rao, A. R., & Hamed, K. H. (2000). *Flood frequency analysis*. CRC Press. <https://doi.org/10.1201/9780429128813>
- Ribatet, M. (2022). *SpatialExtremes: Modelling spatial extremes*. <https://doi.org/10.32614/cran.package.SpatialExtremes>
- Rinnooy Kan, A. H. G., & Timmer, G. T. (1987). Stochastic global optimization methods part II: Multi level methods. *Mathematical Programming*, 39(1), 57–78. <https://doi.org/10.1007/BF02592071>
- Ruegg, C., Cuda, M., & Van Gael, J. (2024). *Math.NET numerics*. <https://numerics.mathdotnet.com/>
- Stephenson, A. G. (2002). evd: Extreme value distributions. *R News*, 2(2), 31–32. <https://doi.org/10.32614/cran.package.evd>
- Storn, R., & Price, K. (1997). Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4), 341–359. <https://doi.org/10.1023/A:1008202821328>
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., & Bürkner, P.-C. (2021). Rank-normalization, folding, and localization: An improved R-hat for assessing convergence of MCMC. *Bayesian Analysis*, 16(2), 667–718. <https://doi.org/10.1214/20-BA1221>
- Viglione, A. (2024). *nsRFA: Non-supervised regional frequency analysis*. <https://doi.org/10.32614/cran.package.nsrfa>
- Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., & others. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods*, 17, 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Waylen, P., & Woo, M.-K. (1982). Prediction of annual floods generated by mixed processes. *Water Resources Research*, 18(4), 1283–1286. <https://doi.org/10.1029/WR018i004p01283>