

The Python/Jupyter Ecosystem: Today's Problem-Solving Environment for Computational Science

Lorena A. Barba , George Washington University, Washington, DC , 20052, USA

The March/April issue of CiSE's inaugural year (1999) carried an essay by eminent computer science professor John R. Rice (who at the time was area editor for Software, together with MATLAB inventor Cleve Moler) titled "A Perspective on Computational Science in the 21st Century."¹ In it, he looked at the development directions for the future of computational science and engineering, and threaded across these was what he called "problem-solving environments." This routine-sounding term hides an ambitious vision, for the time. Rice imagined a software system for tackling problems within a science domain without all the agonizing toil of programming by hand every solution method. He and Boisvert had a previous article (1996) explaining the idea in more detail.² A problem-solving environment would include a collection of mathematical and domain-specific software libraries, offer (semi-)automatic selection of solution method for a given problem, help check the problem formulation, display or assess the correctness of solutions, allow extensibility to add new methods, and even manage the overall computational process. They envisaged an environment that could be "all things to all people," meaning: it is effective when solving simple or complex problems, it supports rapid prototyping and detailed analysis, and it can be used both in introductory teaching and in productive research at the edges of knowledge. An ideal problem-solving environment would even make decisions for the user by means of an integrated knowledge base. What fabulous ambition!

Prof. Rice led a research group at Purdue University that worked to develop early problem-solving environments. The Ellpack system for solving elliptic boundary value problems, developed in the early 1980s, included dozens of software modules

implementing solution methods and a descriptive language to formulate problems (today, we might call it a domain-specific language, DSL). For example, the line: `equation. uxx + uyy + 3 * ux - 4 * u = exp(x + y) * sin(pi * x)` would be used in an Ellpack program for defining the differential equation to be solved. Similarly expressive statements would define the boundary conditions, and the grid parameters to discretize the domain.^a Later versions of the system offered parallel solvers and a graphical user interface (screenshots of the Ellpack system from Prof. John R. Rice's website at Purdue can be found in the Internet Archive).^b

While Ellpack was licensed by Purdue University for a modest yearly fee, this project did not branch off commercially or otherwise. Perhaps a few hundred copies were distributed, mostly for use in university settings, and the project wound down by the early 2000s. By contrast, three commercial software packages for high-productivity scientific and engineering computation—Maple, Mathematica, and MATLAB—had by then become very popular.³ These systems continue to be widely used in education, industry, and government settings. Their purchase price and proprietary implementations, however, led many champions of open-source software to conceive alternatives, oftentimes closely imitating their functionality.

In March/April 2011, 12 years after Prof. Rice's Perspective essay, CiSE ran a special issue on Python for Scientific Computing, showcasing a maturing stack of tools and a highly productive environment for researchers. The issue included one of the most-widely cited articles in the history of the magazine, discussing the high-level multidimensional array structure at the core of NumPy.⁴ By this time, the scientific community had expanded Python for its purposes,

1521-9615 © 2021 IEEE
Digital Object Identifier 10.1109/MCSE.2021.3074693
Date of current version 15 June 2021.

^aA full example at <https://www.cs.purdue.edu/ellpack/example.html>
^b<https://web.archive.org/web/19990506040312/https://www.cs.purdue.edu/research/cse/index.html>

and the four keystone libraries had been put in place in the first half of the decade:

- 1) **SciPy** was consolidated as a standard collection of modules for common mathematical and statistical functions.
- 2) The first version of **IPython**, an enhanced interactive shell for Python, was created by Fernando Pérez.
- 3) **Matplotlib**, the rich 2-D visualization and now standard Python plotting library, was released by John Hunter.
- 4) Travis Oliphant created **NumPy** from a rewrite of the early Python array library Numeric, adding functionality from the competing array package called numarray.

CiSE had previously featured the developing Python support for scientific workflows in an issue organized by Paul F. Dubois, who was the project lead for Numeric from 1997 to 2002. Paul was an Editor for *Computer in Physics* (which got merged into CiSE) since 1993, and joined CiSE with its founding. He wrote and edited for the Scientific Programming department until 2006, and continued with the column "Café Dubois" until 2008. The issue he led included the Hunter piece on Matplotlib,⁵ the Pérez and Granger article about IPython,⁶ and Travis Oliphant's general overview of the Python language and its extensions with NumPy and SciPy.⁴ Other articles in the issue highlight applications in various science contexts: space observation, systems biology, robotics, nanophotonics, and more. An author team from the Simula Research Laboratory in Norway discussed new Python tooling for solving partial differential equations with finite element methods,^c in what was the early development of the Fenics project.⁸ This work heralded the compelling combination of symbolic mathematics and code generation, which Ellpack anticipated.

By the time of the 2011 special issue, the scientific Python ecosystem had gained **Sympy**, the symbolic mathematics and computer algebra system, and **Cython**,⁹ a solution to compile portions of a Python program to obtain faster execution. An important contribution for lowering the bar to adoption of scientific Python was the Enthought Python Distribution, which was free for academic users. It relieved users from the laborious installation of every library individually. The SciKits,^d or SciPy Toolkits, began appearing on the

scene: scikit-learn,^e for predictive data analysis and machine learning,¹⁰ and scikit-image,^f for image processing,¹¹ have since become essential in many scientific contexts. A new, powerful, and flexible Python library for analysis and manipulation of data in the form of labeled tables and time series, pandas, sparked a wave of adoption for statistical modeling.¹² The state of the ecosystem at that time was expertly reviewed by Fernando Pérez, Brian Granger, and John Hunter.¹³ But the last ten years have seen an explosion of innovation, beyond the wildest dreams of these leaders.

IN THE MARCH/APRIL 2021 ISSUE, CiSE IS PROUD TO FEATURE SEVERAL ARTICLES SHOWCASING JUPYTER IN COMPUTATIONAL SCIENCE.

In the March/April 2021 issue, CiSE is proud to feature several articles showcasing **Jupyter** in computational science. The showpiece is an invited article by Jupyter co-founders Brian Granger and Fernando Pérez.¹⁴ They shift the focus of our conversation about problem-solving environments to the human angle: the researcher interactively exploring a scientific question or analyzing data, and the community of people collaborating and advancing their field. Jupyter derives from the IPython Notebook, a browser-based application to compose documents that add computable content to all the other kinds of content that a browser can display: formatted text, images, video, equations, etc. It is a concept inspired by the Mathematica notebook interface, introduced in 1988, and translated to the Python world by the Sage Notebook, starting in 2006. The powerful idea of the Sage Notebook that lives on in Jupyter is connecting a web application serving as a graphical user interface—where the user's text inputs and the computational outputs are shown—to a back-end that runs the Python interpreter. While the server is running, the state is available to continue interactively computing. On shutting it down, the notebook document can be saved thereby preserving the inputs and outputs, together with any text and multimedia content added by the user. Jupyter took these ideas into the modern web era by employing open formats, protocols, and applications that work equally on a laptop and on

^c<http://www.fenics.org/>

^d<https://www.scipy.org/scikits.html>

^e<https://scikit-learn.org/>

^f<https://scikit-image.org/>

remote cloud or HPC systems. The communication protocol between the web application and the back-end is language-agnostic, which quickly led to Notebook support for other languages, like Julia and R (and the new name Jupyter). Dozens of so-called Jupyter kernels now allow computing in many different languages (including interactive use of Fortran and C++), though Python continues to be the most popular by far. Serving Jupyter to multiple-users in a corporation or university became possible with **JupyterHub**,^g which removes the need for users to install any software on their local machine while delivering a uniform environment to them. This provides an ideal solution for academic settings, where large-scale computing and data science education initiatives have always been hampered by individual students' software and hardware needs. Jupyter was rapidly adopted by tech giants like Google, Amazon, and Netflix; by financial behemoths like Bloomberg; by NASA and the LIGO collaboration, which released Jupyter notebooks with the analysis that proved the existence of gravitational waves;^h and by computing and data science educators everywhere. On being awarded the ACM Software System Award in 2017, the citation reads that the tools of Project Jupyter "have become a de facto standard for data analysis in research, education, journalism, and industry".ⁱ Like Granger and Pérez note in the previous issue of CiSE, Jupyter now has many millions of users worldwide and many thousands of organizations use Jupyter in their day-to-day operations. They explain the project's success by virtue of being a tool of thought, a new medium for communication (via computational narratives), and a community of practice.

The traction of Python in the context of scientific computing used to be ascribed to its effectiveness as a "glue language" (easy interoperability with other languages like C/C++ and Fortran), its full set of scientific libraries (NumPy, SciPy, SymPy, Matplotlib), and its notoriously shallow learning curve. But its downside was proclaimed to be performance: being an interpreted, dynamically typed language meant execution would be much slower than low-level implementations in compiled languages. In the past decade, solutions to this performance penalty have been multiplying. At first, researchers would identify performance bottlenecks in their Python applications, rewrite the relevant portions of code in C/C++, and use swig to wrap this code and interface it with the main program in Python.¹⁵ Later, they

gained Cython for numerical loops that cannot be expressed in NumPy operations. Cython compiles Python code extended with type declarations, generating code that can take advantage of the optimizations provided by the C compiler and achieve high performance on those hotspots that dominate runtime. Parallel distributed computing with Python programs using message passing became available with the `mpi4py` package, which supports communication of Python data types and definition of communicator objects according to the MPI specification. A related package, `petsc4py`, provides access to the algorithms and data structures of the PETSc library.^j It allows assembling distributed vectors and sparse matrices, solving linear systems of equations with Krylov iterative methods, and solving nonlinear equations with Newton methods—all of which are core needs in many scientific applications such as those using finite element methods.¹⁶ Access to many-core hardware from Python programs was made possible with run-time code generation via **PyCUDA** and **PyOpenCL**.¹⁷ These tools pioneered the pursuit of high-performance computing with Python, at the cost of increasingly specialized programmer effort. In the last few years, however, a new wave of innovation in scientific Python has sprung from the widespread use of Python in industry settings. Fortunately, much of this innovation has occurred under the open-source model of development and licensing pervading the Python ecosystem.

*FORTUNATELY, MUCH OF THIS
INNOVATION HAS OCCURRED UNDER
THE OPEN-SOURCE MODEL OF
DEVELOPMENT AND LICENSING
PERVADING THE PYTHON
ECOSYSTEM.*

New tools to defeat the performance penalty include **Numba**,^k which accelerates Python code by just-in-time compilation of functions to optimized machine code. The programmer only needs to add decorators, e.g., `@numba.jit(nopython = True)`, ahead of the function definition and Numba will compile the function at runtime, and it will subsequently run without involving the Python interpreter. Numba can also compile a subset of Python code to CUDA kernels for execution on Nvidia GPUs (this could possibly be the easiest way to exploit GPUs

^g<https://jupyter.org/hub>

^h<https://www.gw-openscience.org/events/GW150914/>

ⁱhttps://awards.acm.org/award_winners/perez_9039634

^j<https://www.mcs.anl.gov/petsc/>

^k<https://numba.pydata.org/>

for high-throughput computations). And the new parallel computing library for Python, **Dask**,¹ offers distributed data structures that stand in for NumPy arrays and pandas dataframes, scalable machine learning integrating with scikit-learn, and high-level tools for scheduling and distributing tasks in a cluster. This allows transitioning to parallel and distributed clusters with very little code rewrite. Numba and Dask—like NumPy, Matplotlib, SciPy, SymPy, pandas, and Jupyter—are fiscally sponsored projects of **NumFOCUS**,^m a 501(c)(3) public charity in the United States (I served on its Board of Directors from 2014 to 2020). This means that they are community governed projects developing software under a standard public license, and they both raise funds for their development and receive services (e.g., financial administration, legal support) via a nonprofit. The core developers, maintainers, and users of these projects come together at conferences, where they give technical talks and offer tutorials, participate in online conversations using discussion boards and code-repository issue trackers, and build value together tenaciously.

*PYTHON AND JUPYTER ARE ALSO
PLAYING AN INCREASINGLY
IMPORTANT ROLE IN HIGH-
PERFORMANCE COMPUTING.*

Technology companies often participate actively and benevolently in this activity. NumFOCUS receives corporate sponsorships that benefit the projects, certainly, but another impactful way companies contribute is by allowing or assigning their paid employees to work on the open-source projects of this large ecosystem. Numba and Dask were started at Anaconda, Inc., an Austin-based software and consulting company that also created the hugely popular Anaconda Python distribution. The explosion of machine learning and AI saw the tech giants developing and releasing open-source Python libraries for these applications: Google's **TensorFlow**ⁿ and Facebook's **PyTorch**^o being the most notable. Both libraries provide a Python interface while executing core operations in compiled languages and also CUDA for access to Nvidia GPUs. Google also developed a Jupyter-based cloud notebook, **Colab**,^p providing a hosted solution to run

¹<https://dask.org/>

^m<https://numfocus.org/>

ⁿ<https://www.tensorflow.org/>

^o<https://pytorch.org/>

^p[http://colab.research.google.com/](https://colab.research.google.com/)

TensorFlow with access to GPUs and Google's own TPUs (Tensor Processing Units). The Japanese company Preferred Networks led the development of **CuPy**, a NumPy-like open-source library of matrix functions for Nvidia GPUs. Nvidia embraced the PyData ecosystem creating its RAPIDS AI team^q to develop open-source libraries like **cuDF**, with a pandas-like API for manipulating dataframes on GPUs, and the machine learning library **cuML**, with a growing set of algorithms from scikit-learn. The data science community can now build high-performance workflows in Python and Jupyter, taking advantage of the latest hardware on cloud resources.

Python and Jupyter are also playing an increasingly important role in high-performance computing. Rollin Thomas and Shreyas Cholia, in the previous issue of CiSE, explained how the National Energy Research Scientific Computing Center (NERSC) began their voyage to Jupyter five years ago. They tell us that today about 25% of user interactions with the Cori supercomputer are via JupyterHub, and several scientific workflows have been made more user-friendly while at the same time enjoying parallel speed-ups with Dask. And they conclude: "Jupyter is quickly becoming the entry point to HPC for a growing class of users."¹⁸ In the next issue, CiSE will feature several scientific applications that embody the powerful idea of combining high researcher productivity via Python and high performance through code generation, just-in-time compilation, or exploiting advanced Python libraries. John Rice's perspective of a scientific problem-solving environment for the 21st century may finally be realized, as long as we continue to engage and support the thriving communities of practice of the Python/Jupyter ecosystem. Don't miss our next issue!

REFERENCES

1. J. R. Rice, "A perspective on computational science in the 21st Century," *Comput. Sci. Eng.*, vol. 1, no. 2, pp. 14–16, 1999, doi: [10.1109/5992.753042](https://doi.org/10.1109/5992.753042).
2. J. R. Rice and R. F. Boisvert, "From scientific software libraries to problem-solving environments," *IEEE Computat. Sci. Eng.*, vol. 3, no. 3, pp. 44–53, 1996, doi: [10.1109/99.537091](https://doi.org/10.1109/99.537091).
3. N. Chonacky and D. Winch, "Maple Mathematica, and MATLAB: The 3M's without the tape," *Comput. Sci. Eng.*, vol. 7, no. 1, pp. 8–16, Jan./Feb. 2005, doi: [10.1109/mcse.2005.18](https://doi.org/10.1109/mcse.2005.18).

^q<https://rapids.ai/>

4. S. van der Walt, S. C. Colbert and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Comput. Sci. Eng.*, vol. 13, no. 2, pp. 22–30, Mar./Apr. 2011, doi: [10.1109/mcse.2011.37](https://doi.org/10.1109/mcse.2011.37).
5. J. D. Hunter, "Matplotlib: A 2D graphics environment," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 90–95, May/Jun. 2007, doi: [10.1109/mcse.2007.55](https://doi.org/10.1109/mcse.2007.55).
6. F. Perez and B. E. Granger, "IPython: A system for interactive scientific computing," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 21–29, May/Jun. 2007, doi: [10.1109/mcse.2007.53](https://doi.org/10.1109/mcse.2007.53).
7. T. J. Sargent and J. Stachurski, "Python for scientific computing," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 10–20, May/Jun. 2007, doi: [10.1109/mcse.2007.58](https://doi.org/10.1109/mcse.2007.58).
8. K.-A. Mardal, O. Skavhaug, G. T. Lines, G. A. Staff, and A. Odegard, "Using python to solve partial differential equations," *Comput. Sci. Eng.*, vol. 9, no. 3, pp. 48–51, May/Jun. 2007, doi: [10.1109/mcse.2007.64](https://doi.org/10.1109/mcse.2007.64).
9. S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, "Cython: The best of both worlds," *Comput. Sci. Eng.*, vol. 13, no. 2, pp. 31–39, 2011, doi: [10.1109/mcse.2010.118](https://doi.org/10.1109/mcse.2010.118).
10. F. Pedregosa *et al.*, "Scikit-learn: Machine learning in python," *The J. Mach. Learn. Res.*, vol. 12, 2825–2830, 2011. [Online]. Available: <https://www.jmlr.org/papers/v12/pedregosa11a.html>
11. Stefan van der Walt, J. L. Schonberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors, "Scikit-image: Image processing in Python," *PeerJ*, vol. 2, 2014, Art. no. e453, doi: [10.7717/peerj.453](https://doi.org/10.7717/peerj.453).
12. W. McKinney, "Pandas: A foundational Python library for data analysis and statistics," in *Proc. Workshop Python High Perform. Sci. Comput.* [Online]. Available: https://www.dlr.de/sc/en/Portaldata/15/Resources/dokumente/pyhpc2011/submissions/pyhpc2011_submission_9.pdf
13. F. Perez, B. E. Granger, and J. D. Hunter, "Python: An ecosystem for scientific computing," *Comput. Sci. Eng.*, vol. 13, no. 2, pp. 13–21, Mar./Apr. 2011, doi: [10.1109/mcse.2010.119](https://doi.org/10.1109/mcse.2010.119).
14. B. E. Granger and F. Pérez, "Jupyter: Thinking and storytelling with code and data," *Comput. Sci. Eng.*, vol. 23, no. 2, pp. 7–14, Mar./Apr. 2021, doi: [10.1109/mcse.2021.3059263](https://doi.org/10.1109/mcse.2021.3059263).
15. D. M. Beazley, "Automated scientific software scripting with SWIG," *Future Gener. Comput. Syst.*, vol. 19, no. 5, pp. 599–609, 2003, doi: [10.1016/s0167-739x\(02\)00171-1](https://doi.org/10.1016/s0167-739x(02)00171-1).
16. L. D. Dalcin, R. R. Paz, P. A. Kler, and A. Cosimo, "Parallel distributed computing using Python," *Adv. Water Resour.*, vol. 34, no. 9, pp. 1124–1139, 2011, doi: [10.1016/j.advwatres.2011.04.013](https://doi.org/10.1016/j.advwatres.2011.04.013).
17. A. Klöckner, N. Pinto, Y. Lee, B. Catanzaro, P. Ivanov, and A. Fasih, "PyCUDA and PyOpenCL: A scripting-based approach to GPU run-time code generation," *Parallel Comput.*, vol. 38, no. 3, pp. 157–174, 2012, doi: [10.1016/j.parco.2011.09.001](https://doi.org/10.1016/j.parco.2011.09.001).
18. R. Thomas and S. Cholia, "Interactive supercomputing with Jupyter," *Comput. Sci. Eng.*, vol. 23, no. 2, pp. 93–98, Mar./Apr. 2021, doi: [10.1109/mcse.2021.3059037](https://doi.org/10.1109/mcse.2021.3059037).

LORENA A. BARBA is a professor of mechanical and aerospace engineering with The George Washington University. Her research interests include computational fluid dynamics, biophysics, and high-performance computing. She is Editor-in-Chief of *Computing in Science & Engineering* (CiSE), Co-Editor of the CiSE Reproducible Research Track, Associate Editor for *The ReScience Journal*, Associate Editor-in-Chief of the *Journal of Open Source Software*, and Editor-in-Chief of the *Journal of Open Source Education*. She received the Ph.D. degree in aeronautics from the California Institute of Technology, Pasadena, CA, USA. Contact her at labarba@gwu.edu.