

mlpack 4: a fast, header-only C++ machine learning library

Ryan R. Curtin, Marcus Edel, Omar Shrit, Shubham Agrawal, Suryoday Basak, James J. Balamuta, Ryan Birmingham, Kartik Dutt, Dirk Eddelbuettel, Rishabh Garg, Shikhar Jaiswal, Aakash Kaushik, Sangyeon Kim, Anjishnu Mukherjee, Nanubala Gnana Sai, Nippun Sharma, Yashwant Singh Parihar, Roshan Swain, Conrad Sanderson

Abstract

For over 15 years, the mlpack machine learning library has served as a “swiss army knife” for C++-based machine learning. Its efficient implementations of common and cutting-edge machine learning algorithms have been used in a wide variety of scientific and industrial applications. This paper overviews mlpack 4, a significant upgrade over its predecessor. The library has been significantly refactored and redesigned to facilitate an easier prototyping-to-deployment pipeline, including bindings to other languages (Python, Julia, R, Go, and the command line) that allow prototyping to be seamlessly performed in environments other than C++. mlpack is open-source software, distributed under the permissive 3-clause BSD license; it can be obtained at <https://mlpack.org>

Background

The use of machine learning has become ubiquitous in almost every scientific discipline and countless commercial applications [3, 14]. There is one important commonality to virtually all of these applications: machine learning is often computationally intensive, due to the large number of parameters and large amounts of training data. This was the main motivator for the original development of mlpack in the C++ language, which allows for efficient close-to-the-metal implementations [5, 6].

But speed is not everything: development and deployment of applications that use machine learning can also be significantly hampered if the overall process is too difficult or unwieldy [17, 21]. Furthermore, deployment environments often have computational or engineering constraints that make a full-stack Python solution infeasible [12]. As such, it is important that lightweight and easy-to-deploy machine learning solutions are available. This has motivated our refactoring and redesign of mlpack 4: we pair efficient implementations with easy and lightweight deployment, making mlpack suitable for a wide range of deployment environments. A more complete set of motivations can be found in the mlpack vision document [19].

mlpack is a general-purpose machine learning library, targeting both academic and commercial use; for instance, data scientists who need efficiency and ease of deployment, or, e.g., by researchers who need flexibility and extensibility. While there are other machine learning libraries intended to be used from C++, many, such as FAISS [13] and FLANN [20], are limited to a few specific algorithms, instead of a full range of machine learning algorithms, like mlpack provides. dlib-ml [15], on the other hand, does provide a broad toolkit of machine learning algorithms, but its extensibility is somewhat limited as it does not use policy-based design [1] to provide arbitrary user-defined behavior, and the range of machine learning algorithms provided is smaller than mlpack’s.

Functionality

The library contains a wide variety of machine learning algorithms, some of which are new to mlpack 4. The list of algorithms includes linear regression, logistic regression, random forests, furthest-neighbor search [8], accelerated k-means variants [4], kernel density estimation [18], and fast max-kernel search [9]. There is also a module for deep neural networks, which has implementations of numerous layer types, activation functions, and reinforcement learning applications. Details of the available functionality are provided in the online mlpack documentation (<https://www.mlpack.org/docs.html>). The efficiency of these implementations has been shown in various works [5, 11] using mlpack’s benchmarking system [10].

The algorithms are available via automatically-generated bindings to Python, R, Go, Julia, and the command line. Each of these bindings has a unified interface across the languages; for example, a model trained in Python can be used from Julia or C++ (or any other language with mlpack bindings). The bindings are available in each language's package manager, as well as system-level package managers such as *apt* and *dnf*. Furthermore, ready-to-use Docker containers with the environment fully configured are available on DockerHub, and an interactive C++ notebook interface via the *xeus-cling* project (<https://github.com/QuantStack/xeus-cling>) is available on BinderHub.

Once a user has developed a machine learning workflow in the language of their choice, deployment is straightforward. The mlpack library is now header-only, and directly depends only on three libraries: *Armadillo* [24], *ensmallen* [7], and *cereal* (<https://github.com/USCILab/cereal>). When using C++, the only linking requirement is to an efficient implementation of BLAS and LAPACK (required via Armadillo). This significantly eases deployment; a standalone C++ application with only a BLAS/LAPACK dependency is easily deployable to many environments, including standard Linux-based Docker containers, Windows environments, and resource-constrained embedded environments. To this end, mlpack's build system now also contains a number of tools for cross-compilation support, including the ability to easily statically link compiled programs (important for some deployment environments).

Major Changes

Below we detail a few of the major changes present in mlpack 4. For a complete and exhaustive list (including numerous bug fixes and new techniques), the *HISTORY.md* file (distributed with mlpack) can be consulted.

- **Removed dependencies.** In accordance with the vision document [19], the majority of the refactoring and redesign work focused on reducing dependencies and compilation overhead. This has motivated the replacement of the *Boost C++ libraries* (<https://www.boost.org>), upon which mlpack previously depended, with lightweight alternatives including *cereal* (<https://github.com/USCILab/cereal>) for serialization. The entire neural network module was refactored to avoid the use of *Boost* (amounting to an almost complete rewrite). This effort was rewarded handsomely: with mlpack 3, a simple program would often require several gigabytes of memory just for compilation. After refactoring and removing dependencies, compilation generally requires just a few hundred megabytes of memory, and is often an order of magnitude faster.
- **Interactive notebook environments.** mlpack can be used in a Jupyter notebook environment [16] via the *xeus-cling* project (<https://github.com/QuantStack/xeus-cling>). This is demonstrated interactively on the mlpack homepage (<https://www.mlpack.org>). Examples of C++ notebooks can be found in the mlpack examples repository (<https://github.com/mlpack/examples>), and these can easily be run on BinderHub.
- **New bindings and enhanced availability.** Support for the Julia [2], Go [22], and R languages [23, 25] has been added via mlpack's automatic binding system. These bindings can be used by installing mlpack from the language's package manager (*Pkg.jl*, *go get*, *install.packages('mlpack')*). Furthermore, since mlpack's reduced dependency footprint has significantly simplified the deployment process, mlpack's Python dependencies are now available for numerous architectures both on PyPI and in conda-forge.
- **Cross-compilation support and build system improvements.** mlpack's build configuration now supports easy cross-compilation, for instance via toolchains such as *buildroot* (<https://buildroot.org>). By specifying a few flags, a user may produce a working mlpack setup for a variety of embedded systems. This required the implementation of a dependency auto-downloader, which is capable of downloading *OpenBLAS* (<https://github.com/xianyi/OpenBLAS>) and compiling (if necessary) for the target architecture. The auto-downloader can also be enabled and used for any situation, thus easing installation and deployment.

Acknowledgements

Development of mlpack is community-led. It is the product of hard work by over 220 individuals (at the time of writing). We are also indebted to people that have provided bug reports over the years. The development has been supported by Google, via a decade-long participation the Google Summer of Code program, and also by NumFOCUS, which fiscally sponsors mlpack.

References

- [1] A. Alexandrescu. *Modern C++ Design: generic programming and design patterns applied*. Addison-Wesley, 2001.
- [2] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017.
- [3] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4):045002, 2019.
- [4] R. R. Curtin. A dual-tree algorithm for fast k-means clustering with large k. In *SIAM International Conference on Data Mining (SDM '17)*, pages 300–308. SIAM, 2017.
- [5] R. R. Curtin, J. R. Cline, N. P. Slagle, W. B. March, P. Ram, N. A. Mehta, and A. G. Gray. MLPACK: A scalable C++ machine learning library. *The Journal of Machine Learning Research*, 14(1):801–805, 2013.
- [6] R. R. Curtin, M. Edel, M. Lozhnikov, Y. Mentekidis, S. Ghaisas, and S. Zhang. mlpack 3: a fast, flexible machine learning library. *Journal of Open Source Software*, 3(26):726, 2018.
- [7] R. R. Curtin, M. Edel, R. G. Prabhu, S. Basak, Z. Lou, and C. Sanderson. The ensmallen library for flexible numerical optimization. *Journal of Machine Learning Research*, 22:1–6, 2021.
- [8] R. R. Curtin and A. B. Gardner. Fast approximate furthest neighbors with data-dependent candidate selection. In *International Conference on Similarity Search and Applications (SISAP 2016)*, pages 221–235. Springer, 2016.
- [9] R. R. Curtin and P. Ram. Dual-tree fast exact max-kernel search. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 7(4):229–253, 2014.
- [10] M. Edel, A. Soni, and R. R. Curtin. An automatic benchmarking system. In *NIPS Workshop on Software Engineering for Machine Learning*, 2014.
- [11] D. Fang and D. H. Chau. M3: Scaling up machine learning via memory mapping. In *International Conference on Management of Data*, pages 2249–2250, 2016.
- [12] L. Fischer, L. Ehrlinger, V. Geist, R. Ramler, F. Sobieszky, W. Zellinger, D. Brunner, M. Kumar, and B. Moser. AI System Engineering—Key Challenges and Lessons Learned. *Machine Learning and Knowledge Extraction*, 3(1):56–83, 2020.
- [13] J. Johnson, M. Douze, and H. Jégou. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.
- [14] M. I. Jordan and T. M. Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.
- [15] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [16] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. E. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. B. Hamrick, J. Grout, S. Corlay, et al. Jupyter Notebooks—a publishing format for reproducible computational workflows. In *Proceedings of the 20th International Conference on Electronic Publishing*, pages 87–90, 2016.
- [17] A. Lavin, C. M. Gilligan-Lee, A. Visnjic, S. Ganju, D. Newman, S. Ganguly, D. Lange, A. G. Baydin, A. Sharma, A. Gibson, et al. Technology readiness levels for machine learning systems. *Nature Communications*, 13(1):1–19, 2022.
- [18] D. Lee and A. G. Gray. Fast high-dimensional kernel summations using the Monte Carlo Multipole Method. *Advances in Neural Information Processing Systems*, 21:929–936, 2008.
- [19] mlpack community. mlpack: A vision for an efficient prototype-to-deployment machine learning library. <https://www.mlpack.org/papers/vision.pdf>.
- [20] M. Muja and D. G. Lowe. Fast approximate nearest neighbors with automatic algorithm configuration. In *Proceedings of the 2009 International Conference on Computer Vision Theory and Applications (VISAPP '09)*, pages 331–340, 2009.
- [21] A. Paleyes, R.-G. Urma, and N. D. Lawrence. Challenges in deploying machine learning: a survey of case studies. *ACM Computing Surveys (CSUR)*, 2020.
- [22] R. Pike. Go at Google: Language design in the service of software engineering. In *Proceedings of the 3rd Annual Conference on Systems, Programming, and Applications: Software For Humanity*, pages 5–6, 2012.
- [23] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2022.
- [24] C. Sanderson and R. R. Curtin. Armadillo: a template-based C++ library for linear algebra. *Journal of Open Source Software*, 1(2):26, 2016.
- [25] Y. Singh Parihar, R. R. Curtin, D. Eddelbuettel, J. Balamuta, et al. *mlpack: 'Rcpp' Integration for the 'mlpack' Library*, 2022. R package version 4.0.0.