

# Towards a benchmark framework for model order reduction in the Mathematical Research Data Initiative (MaRDI)

Peter Benner  | Kathryn Lund  | Jens Saak 

Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, Germany

## Correspondence

Kathryn Lund, Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstr. 1, 39106 Magdeburg, Germany. Email: [lund@mpi-magdeburg.mpg.de](mailto:lund@mpi-magdeburg.mpg.de)

<https://www.mardi4nfdi.de/>

## Abstract

The race for the most efficient, accurate, and universal algorithm in scientific computing drives innovation. At the same time, this healthy competition is only beneficial if the research output is actually comparable to prior results. Fairly comparing algorithms can be a complex endeavor, as the implementation, configuration, compute environment, and test problems need to be well-defined. Due to the increase in computer-based experiments, new infrastructure for facilitating the exchange and comparison of new algorithms is also needed. To this end, we propose a benchmark framework as a set of generic specifications for comparing implementations of algorithms using test cases native to a community. Its value lies in its ability to fairly compare and validate existing methods for new applications, as well as compare newly developed methods with existing ones. As a prototype for a more general framework, we have begun building a benchmark tool for the model order reduction (MOR) community. The data basis of the tool is the collection of the Model Order Reduction Wiki (MORWiki). The wiki features three main categories: benchmarks, methods, and software. An editorial board curates submissions and patrols edited entries. Data sets for linear and parametric-linear models are already well represented in the existing collection. Data sets for non-linear or procedural models, for which only evaluation data, or codes/algorithmic descriptions, rather than equations, are available, are being added and extended. Properties and interesting characteristics used for benchmark selection and later assessments are recorded in the model metadata. Our tool, the Model Order Reduction Benchmark (MORB), is under active development for linear time-invariant systems and solvers. An ontology (MORBO) and knowledge graph are being developed in parallel. They catalog benchmark problem sets and their metadata and will also be integrated into the Mathematical Research Data Initiative (MaRDI) Portal to help improve the findability of such data sets. MORB faces a number of technical and field-specific challenges, and we hope to recruit community input and feedback while presenting some initial results.

This is an open access article under the terms of the [Creative Commons Attribution-NonCommercial](https://creativecommons.org/licenses/by-nc/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2023 The Authors. *Proceedings in Applied Mathematics & Mechanics* published by Wiley-VCH GmbH.

## 1 | INTRODUCTION

The Mathematical Research Data Initiative (MaRDI)<sup>1</sup> is a consortium of the National Research Data Initiative (NFDI)<sup>2</sup>, whose overarching goal is to improve and promote responsible research data management practices in the German scientific landscape and beyond. MaRDI itself concentrates on several fields of mathematics featured as content-specific task areas (TA):

- (TA1) Computer Algebra, (exact data)
- (TA2) Scientific Computing, (inexact data)
- (TA3) Statistics and Machine Learning, and (uncertain data)
- (TA4) Cooperation with Other Disciplines. (mixed data and interoperability)

Within TA2, we focus primarily on numerical algorithms and their software implementations, which are used to compute approximate solutions and simulations of scientific models. Via collaboration between the Max Planck Institute for Dynamics of Complex Technical Systems and the University of Münster, TA2 addresses the following measures (M):

- (M1) Knowledge graph of numerical algorithms;
- (M2) Open interfaces for scientific computing;
- (M3) Benchmark framework; and
- (M4) Description and design of Findable, Accessible, Interoperable, and Reproducible (FAIR)<sup>3</sup> workflows for computational science and engineering (CSE).

For this report, we will concentrate on progress in (M3).

## 2 | MARDIMARK

We begin with a domain-independent specification of a benchmark workflow, MaRDIMark. MaRDIMark is broken into five modules, as displayed in Figure 1. The content and function of each module are detailed in the following sections.

### 2.1 | Problems

This module consists of the benchmark problems themselves, both as mathematical abstractions and as datasets stored on a computer. Note that datasets might take the form of both assembly scripts and static floating-point numbers. A good benchmark collection should minimally comprise a database of unique problem IDs, download links, mathematical descriptions, fields referring to domain-specific mathematical and numerical properties, file formats, license data, creator and editor identifiers (ideally ORCID<sup>4</sup>), and any other information that helps with data provenance, findability, and accessibility. Ontologies and knowledge graphs like AlgoData<sup>5</sup> can add value to such a database, as they allow users to query complex relations, instead of just look up values in a table. For example, a user may want to know which methods can be applied to a given benchmark, which brings us to the next module.

### 2.2 | Methods

The Methods module is similar to the Problems module, except that the objects of interest are instead mathematical procedures and algorithms for solving a class of problems. The definition of a method for a particular field can be nontrivial, especially when all implementation details are taken into account. For example, most methods in numerical mathematics are initially presented as pseudocode, which a practitioner must translate into actual code to use on a computer. Minimally,

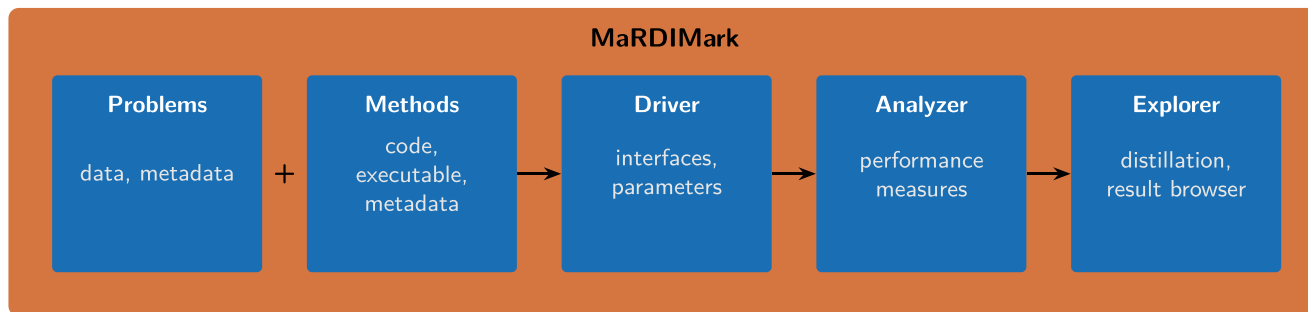
<sup>1</sup> [www.mardi4nfdi.de](http://www.mardi4nfdi.de)

<sup>2</sup> [www.nfdi.de](http://www.nfdi.de)

<sup>3</sup> [www.go-fair.org/fair-principles](http://www.go-fair.org/fair-principles)

<sup>4</sup> [www.orcid.org](http://www.orcid.org)

<sup>5</sup> A deliverable of (M1). See <https://algodata.mardi4nfdi.de/>.



**FIGURE 1** Overview of MaRDIMark.

the practitioner must choose hardware (i.e., a physical machine), an operating system (and version), and a programming language (and version). Within that language, there are multiple ways to execute any mathematical operation, and each may have slightly different floating-point behavior, which further depends on what libraries the language and operating system are running on the backend. We loosely refer to a method, its pseudocode, and its implementation choices as an *algorithm isotope*, in analogy to elemental isotopes. Each community must decide on its own how fine-grained their algorithm isotopes should be, and these definitions will necessarily vary from application to application, depending on what aspects of a class of methods are being studied.

### 2.3 | Driver

The linchpin of a benchmark workflow is a routine (or set of routines) that executes specified methods on specified problems, i.e., the Driver module. In practice, a benchmark workflow may be written in a general-purpose language but needs to call routines written in another language, without introducing unnecessary overhead. Here open interfaces developed by (M2) can be applied, which facilitate communication between different software packages and programming languages.

### 2.4 | Analyzer

Benchmarking by definition is concerned with the comparison of methods with respect to different measures, such as, time to solution, accuracy, quality, optimized cost function, and so forth. We refer to the collection and calculation of these measures as the Analysis module. A community must determine on its own which measures are relevant and how to implement them in an unbiased manner, so that no method is artificially penalized by the procedure used to calculate a measure.

### 2.5 | Explorer

Finally, it is important that practitioners are able to visualize the results and share them with colleagues. The Explorer module thus comprises graphs, plots, tables, and reports that MaRDIMark synthesizes. Convenient formats for many mathematical fields might be encoded in TeX, HTML, or Markdown, or as a Jupyter notebook.

## 3 | MODEL ORDER REDUCTION BENCHMARKER (MORB)

To demonstrate MaRDIMark in action, we have begun developing a benchmark tool for the field of model order reduction (MOR), relying on the curated benchmark database hosted in the MORWiki [1].

### 3.1 | MORWiki

MORWiki was launched in 2012 as a platform for the MOR community. It hosts established benchmark collections like the Oberwolfach Collection [2] and SLICOT [3, 4] and new datasets. Currently, approximately 100 distinct datasets have

been documented for a variety of linear, nonlinear, and parametric real-world and test examples. Various methods have also been documented, along with a comparison table for MOR software and a BibTeX file for MOR literature. Since the start of MaRDI, a searchable database for problem metadata has been compiled, along with an initial ontology that will be integrated into AlgoData and the MaRDIPortal<sup>6</sup>. Eventually, a search tool will be published that allows users to look up benchmarks with desired properties.

A number of challenging tasks remain open for MORWiki. In particular, licenses are missing for older datasets like SLICOT and the Oberwolfach, despite the fact that they have been treated as if they were part of the public domain for nearly 20 years. This poses a unique legal situation, especially if we upload updated versions of the data, to the MORWiki Zenodo community<sup>7</sup> that conform to the file format standards of MORB (cf. Section 3.2). Furthermore, it is important that such a database is built by the larger community, even if the overall editorial work is maintained by a small group. Incentivizing researchers to not only publish their data, but also prepare it in a specific format to improve interoperability is difficult. A key selling point is that papers with public code and data accumulate higher citations over time [5]. Finally, much work remains regarding the mathematical metadata for MORWiki benchmarks. Systems theory enjoys a plethora of problem varieties, which must often be determined by hand or via numerical eigensolvers, which can be prohibitively expensive for large systems.

### 3.2 | Design specifications

For the initial prototype of MORB, we have aimed to keep things as simple as possible. We have focused on linear, time-invariant, first-order systems (LTI-FOS) of the form

$$\begin{aligned} E\dot{x}(t) &= Ax(t) + Bu(t), \\ y(t) &= Cx(t) + Du(t), \end{aligned}$$

where  $A, E \in \mathbb{R}^{n \times n}$ ,  $B \in \mathbb{R}^{n \times m}$ ,  $C \in \mathbb{R}^{q \times n}$ , and  $D \in \mathbb{R}^{q \times m}$ . When  $E$  and  $C$  are not explicitly specified, the identity matrix is assumed for both; when  $D$  is not specified, it is assumed to be zero. For interoperability, we require that all the matrices for a given benchmark are stored as .mat files, version 7, which is supported by MATLAB<sup>®</sup><sup>8</sup> and SciPy<sup>9</sup>, as well as GNU Octave<sup>10</sup>, Julia<sup>11</sup>, and can be linked into most compiled languages via libmatio<sup>12</sup>. We have ensured that all LTI-FOS benchmarks in MORWiki are converted to such a format, and will be made public<sup>7</sup>, as soon as the open licensing questions have been resolved.

A key component of MORB is the configuration file. It is written in JSON<sup>13</sup>, a format compatible with most widely used programming languages, and is relatively straightforward to set up and modify. The highest level must correspond to the filename (i.e., benchmark ID). The secondary level should minimally have the “alg\_iso” field and optional fields

- “meas\_opt” (options for Analyzer module);
- “bode\_opt” (options for Bode plots; part of Explorer module);
- “plot\_opt” (general options for plots; part of Explorer module); and
- “report\_opt” (options for final report; part of Explorer module).

See Figure 2, for an example configuration file. In this example, the solver bt (Balanced Truncation) is being configured for the packages cst (Control Systems Toolbox)<sup>14</sup>, emgr (Empirical Grammian Framework) [6], mess (Matrix Equation Sparse Solver) [7], morlab (Model Order Reduction Laboratory) [8], and pymor (Model Order Reduction with Python) [9]. Depending on the package, different parameters like tol (Hankel singular value tolerance) and max\_order

<sup>6</sup> <https://portal.mardi4nfdi.de/>

<sup>7</sup> <https://zenodo.org/communities/morwiki/>

<sup>8</sup> [https://de.mathworks.com/help/matlab/import\\_export/mat-file-versions.html](https://de.mathworks.com/help/matlab/import_export/mat-file-versions.html)

<sup>9</sup> <https://docs.scipy.org/doc/scipy/reference/generated/scipy.io.loadmat.html>

<sup>10</sup> <https://octave.org/>

<sup>11</sup> <https://julialang.org>

<sup>12</sup> <https://github.com/telehan/libmatio>

<sup>13</sup> <https://www.json.org/>


<sup>14</sup> [www.mathworks.com/products/control.html](http://www.mathworks.com/products/control.html)

```

{
  "newEngland_n66m1q1":
  {
    "alg_iso":
    {
      "bt":
      {
        "cst":
        [
          [
            {
              "tol": 1E-6
            },
            {
              "tol": 1E-12
            }
          ],
          "emgr": null,
          "mess":
          [
            [
              {
                "max_order": 100,
                "tol": 1E-6
              },
              {
                "tol": 1E-12
              }
            ],
            "morlab":
            {
              "tol": 1E-6
            },
            "pymor":
            {
              "tol": 1E-6
            }
          ],
          "bg":
          {
            "emgr": null,
            "mess": null
          }
        },
        "meas_opt":
        {
          "norm_id": ["l0", "l1", "l2", "linf", "h2"],
          "time_points": 250,
          "h2_method": "lyap",
          "ml_bodemag":
          {
            "FreqRange": [-8, 8],
            "ShowPlot": 0,
            "MaxPoints": 500
          },
          "ml_sigmaplot":
          {
            "FreqRange": [-8, 8],
            "ShowPlot": 0,
            "MaxPoints": 500
          },
          "ml_frobeniusplot":
          {
            "FreqRange": [-8, 8],
            "ShowPlot": 0,
            "MaxPoints": 500
          }
        },
        "bode_opt":
        {
          "FreqRange": [-8, 8],
          "ShowPlot": 0,
          "MaxPoints": 500
        },
        "plot_opt":
        {
          "save_eps": true,
          "save_fig": true
        }
      }
    }
  }
}

```

**FIGURE 2** Example of a configuration file for the benchmark problem newEngland\_n66m1q1.



**FIGURE 3** Sample run time bar chart for steelProfile\_n1357m7q6. Note that timings are scaled by the maximum in the series.

**TABLE 1** Sample error table for steelProfile\_n1357m7q6.

Algorithm Isotope	$L_2$	$L_\infty$	$H_2$
bt-mess	$2.23 \times 10^{-6}$	$2.83 \times 10^{-6}$	$4.66 \times 10^{-6}$
bt-pymor	$1.56 \times 10^{-5}$	$1.78 \times 10^{-5}$	$2.92 \times 10^{-5}$
bt-cst	$2.45 \times 10^{-4}$	$1.95 \times 10^{-4}$	$2.90 \times 10^{-3}$
bt-morlab	$2.86 \times 10^{-6}$	$3.00 \times 10^{-6}$	$5.55 \times 10^{-6}$

(maximum order of the reduced model) may be set. Note that actually two algorithm isotopes are specified for both the `cst` and `mess` packages, because two different `tol` options have been provided. Because `emgr` is set to `null`, six total algorithm isotopes will be run for `bt`.

For this version of MORB, the encasing workflow has been written in MATLAB. This choice has allowed us to minimize interface issues, as most MOR software for this class of models is written in MATLAB<sup>15</sup>. Furthermore, the MATLAB package MORLAB contains well-documented routines for computing standard errors in the  $L_0$ ,  $L_1$ ,  $L_2$ ,  $L_\infty$ , and  $H^2$  norms, as well as Bode, Frobenius, and Sigma plots, which we have relied on for the Analyzer module.

### 3.3 | Early prototype and sample reports

A prototype of MORB that works on the steel profile benchmark [10] can be found on Zenodo [11]. Furthermore, sample reports generated from other datasets are included there that demonstrate what the Explorer module looks like. Details about the run itself—a timestamp, the benchmark ID, operating system, software version, and problem size—are reported, along with a table and graph of the run times and a table of error norms, as well as error, Bode, Sigma, and Frobenius plots. See Figure 3 and Table 1 for a preview of the full report. In this version of MORB, reports are auto-generated as TeX files, which can then be compiled as PDFs. Additional details about the run are stored in the comments of the TeX file, generated by the MATLAB command `system('systeminfo')`<sup>16</sup>.

## 4 | CONCLUSIONS AND OUTLOOK

We have demonstrated the versatility of MaRDIMark on a curated database of MOR benchmarks. The resulting tool, MORB, is an initial attempt to compare solvers written in different languages with different syntax and parameter-parsing standards for a large class of benchmark examples. There is still much room for improvement. First of all, we plan to refac-

<sup>15</sup> [https://morwiki.mpi-magdeburg.mpg.de/morwiki/index.php/Comparison\\_of\\_Software](https://morwiki.mpi-magdeburg.mpg.de/morwiki/index.php/Comparison_of_Software)

<sup>16</sup> This command only works on Windows. <https://www.mathworks.com/help/matlab/ref/system.html>

tor it in Python, a free, open-source language with a large, active user-base. This will ensure that we comply with FAIR principles, especially accessibility and reproducibility, which the MATLAB version currently violates due to the requirement for paid licenses. Second, the configuration file as described in Section 3.2 will be updated to remove redundancies in the hierarchy, to reflect better the modular design of the tool and to enable parallelization for large numbers of algorithm isotopes. In addition, we plan to integrate more algorithms from the field to allow for benchmarks on a wider variety of problems, particularly nonlinear and parametric examples. Finally, the report itself will be updated to allow for a wider variety of formats (e.g., HTML and Markdown) and to incorporate progress from (M4) and MaRDIFlow [12]

Principles from MaRDIMark can already be seen in other benchmark tools. One of the co-authors has actively incorporated the modular design in the stability analysis of communication-avoiding Gram–Schmidt and Krylov subspace methods [13, 14]. Within MaRDI, TA3 has developed `mlr3`, which provides a unified interface for machine learning in the R language [15]. The toolkit includes a benchmark feature<sup>17</sup>, which generates a grid of learner configurations, runs them on a set of tasks, and returns performance results that can be explored as a table or box plot. Another well-developed open-source tool is `FitBenchmarking` [16], which compares fitting and minimizer frameworks from software that is callable from Python. The primary output is an interactive table, allowing users to quickly navigate results for interesting solver and benchmark combinations. A key feature of both `mlr3` and `FitBenchmarking` is that they have been developed long-term by large teams of domain experts and software engineers. Furthermore, both rely on curated benchmark databases established and maintained by their respective academic communities, and they both provide pathways for users to contribute new software to be benchmarked.

## ACKNOWLEDGMENTS

We would like to thank Christian Himpe and Tim Mitchell for conversations that laid the foundation of the benchmark tool, as well as Alexander Stage and Malte Speidel, for work on the MORB ontology and search tool. We further thank Martin Köhler for helping debug MORB. The authors are supported by MaRDI, funded by the Deutsche Forschungsgemeinschaft (DFG), project number 460135501, NFDI 29/1 “MaRDI – Mathematische Forschungsdateninitiative”.

Open access funding enabled and organized by Projekt DEAL.

## ORCID

Peter Benner  <https://orcid.org/0000-0003-3362-4103>

Kathryn Lund  <https://orcid.org/0000-0001-9851-6061>

Jens Saak  <https://orcid.org/0000-0001-5567-9637>

## REFERENCES

1. The MORwiki Community. MORwiki—Model Order Reduction Wiki. <http://modelreduction.org>. Accessed 30 June 2023.
2. Korvink, J. G., & Rudnyi, E. B. (2005). Oberwolfach Benchmark Collection. In P. Benner, V. Mehrmann, & D. C. Sorensen (Eds.), *Dimension Reduction of Large-Scale Systems. Lecture Notes in Computational Science and Engineering* (pp. 311–315). Berlin, Heidelberg: Springer.
3. Varga, A. (1999, December). Task II. B.1 – selection of software for controller reduction. SLICOT Working Note 1999–18, The Working Group on Software (WGS). [www.slicot.org](http://www.slicot.org)
4. Varga, A. (2001). Model reduction software in the SLICOT library. In B. N. Datta (Ed.), *The Kluwer International Series in Engineering and Computer Science: Vol. 629. Applied and computational control, signals, and circuits* (pp. 239–282). Kluwer Academic Publishers.
5. Colavizza, G., Hrynaszkiewicz, I., Staden, I., Whitaker, K., & McGillivray, B. (2020). The citation advantage of linking publications to research data. *PLOS ONE*, *15*(4), e0230416.
6. Himpe, C. (2022). `emgr` – EMpirical GRamian framework (version 5.99). <https://gramian.de>.
7. Benner, P., Köhler, M., & Saak, J. (2021). Matrix equations, sparse solvers: M-M.E.S.S.-2.0.1 – philosophy, features and application for (parametric) model order reduction. In P. Benner, T. Breiten, H. Faßbender, M. Hinze, T. Stykel, & R. Zimmermann (Eds.), *International Series of Numerical Mathematics: Vol. 171. Model reduction of complex dynamical Systems* (pp. 369–392). Birkhäuser.
8. Benner, P., & Werner, S. W. R. (2021). Frequency- and time-limited balanced truncation for large-scale second-order systems. *Linear Algebra and its Applications*, *623*, 68–103.
9. Milk, R., Rave, S., & Schindler, F. (2016). `pyMOR`—Generic algorithms and interfaces for model order reduction. *SIAM Journal on Scientific Computing*, *38*(5), S194–S216.
10. Saak, J., & Behr, M. (2021). The Oberwolfach Steel-Profile benchmark revisited. <https://zenodo.org/record/5113560>
11. Lund, K., Saak, J., & Benner, P. (2023). mini-MORB (Model Order Reduction Benchmark). <https://zenodo.org/record/8093833>.
12. Veluvali, P., Heiland, J., & Benner, P. (2023). MaRDIFlow: A workflow framework for documentation and integration of FAIR Computational Experiments. <https://zenodo.org/record/7863520>.

<sup>17</sup> [https://mlr3book.mlr-org.com/chapters/chapter3/evaluation\\_and\\_benchmarking.html](https://mlr3book.mlr-org.com/chapters/chapter3/evaluation_and_benchmarking.html)



13. Lund, K. (2022). LowSyncBlockArnoldi. <https://gitlab.mpi-magdeburg.mpg.de/lund/low-sync-block-arnoldi/-/releases/v1.0>.
14. Lund, K., Carson, E., & Oktay, E. (2023). BlockStab. <https://github.com/katlund/BlockStab/releases/tag/v2.2023-beta>.
15. Bischl, B., Sonabend, R., Kotthoff, L., & Lang, M. (2023). Mlr3: Machine learning in R. <https://mlr3.mlr-org.com/>.
16. Markvardsen, A., Rees, T., Wathen, M., Lister, A., Farmer, J., Farooq, R., Kitenge, J., Wilkins, J., Hanrahan, S., Peters, B., Applin, R., Fowkes, J., Odagiu, P., Anuchitanukul, A., Farmer, T., Lim, A., Pouzols, F. M., Snow, T., & McCluskey, A. (2022). FitBenchmarking: An open source Python package comparing data fitting software. <https://github.com/fitbenchmarking/fitbenchmarking>.

**How to cite this article:** Benner, P., Lund, K., & Saak, J. (2023). Towards a benchmark framework for model order reduction in the Mathematical Research Data Initiative (MaRDI). *Proceedings in Applied Mathematics and Mechanics*, 23, e202300147. <https://doi.org/10.1002/pamm.202300147>