





















# Data publishing in mechanics and dynamics: challenges, guidelines, and examples from engineering design

Henrik Ebel<sup>1</sup> , Jan van Delden<sup>2</sup> , Timo Lüddecke<sup>2</sup> , Aditya Borse<sup>3</sup> , Rutwik Gulakala<sup>3</sup> , Marcus Stoffel<sup>3</sup> , Manish Yadav<sup>4</sup> , Merten Stender<sup>4</sup> , Leon Schindler<sup>5</sup> , Kristin Miriam de Payrebrune<sup>5</sup> , Maximilian Raff<sup>6</sup> , C. David Remy<sup>6</sup> , Benedict Röder<sup>7</sup> , Rohit Raj<sup>8</sup> , Tobias Rentschler<sup>8</sup> , Alexander Tismer<sup>8</sup> , Stefan Riedelbauch<sup>8</sup>  and Peter Eberhard<sup>7</sup> 

<sup>1</sup>Department of Mechanical Engineering, LUT University, Lappeenranta, Finland

<sup>2</sup>Institute of Computer Science, University of Göttingen, Göttingen, Germany

<sup>3</sup>Institute of General Mechanics (IAM), RWTH Aachen, Aachen, Germany

<sup>4</sup>Cyber-Physical Systems in Mechanical Engineering, Technische Universität Berlin, Berlin, Germany

<sup>5</sup>Institute for Computational Physics in Engineering, RPTU Kaiserslautern-Landau, Kaiserslautern, Germany

<sup>6</sup>Institute for Nonlinear Mechanics, University of Stuttgart, Stuttgart, Germany

<sup>7</sup>Institute of Engineering and Computational Mechanics, University of Stuttgart, Stuttgart, Germany

<sup>8</sup>Institute of Fluid Mechanics and Hydraulic Machinery, University of Stuttgart, Stuttgart, Germany

**Corresponding author:** Henrik Ebel; Email: [henrik.ebel@lut.fi](mailto:henrik.ebel@lut.fi)

**Received:** 28 November 2024; **Revised:** 21 February 2025; **Accepted:** 04 March 2025

**Keywords:** data publishing; dynamics; generative design; machine learning; mechanics

## Abstract

Data-based methods have gained increasing importance in engineering. Success stories are prevalent in areas such as data-driven modeling, control, and automation, as well as surrogate modeling for accelerated simulation. Beyond engineering, generative and large-language models are increasingly helping with tasks that, previously, were solely associated with creative human processes. Thus, it seems timely to seek artificial-intelligence-support for engineering design tasks to automate, help with, or accelerate purpose-built designs of engineering systems for instance in mechanics and dynamics, where design so far requires a lot of specialized knowledge. Compared with established, predominantly first-principles-based methods, the datasets used for training, validation, and test become an almost inherent part of the overall methodology. Thus, data publishing becomes just as important in (data-driven) engineering science as appropriate descriptions of conventional methodology in publications in the past. However, in mechanics and dynamics, quite widely, still traditional publishing practices are prevalent that largely do not yet take into account the rising role of data as much as that may already be the case in pure data-scientific research. This article analyzes the value and challenges of data publishing in mechanics and dynamics, in particular regarding engineering design tasks, showing that the latter raise also challenges and considerations not typical in fields where data-driven methods have been booming originally. Researchers currently find barely any guidance to overcome these challenges. Thus, ways to deal with these challenges are discussed and a set of examples from across different design problems shows how data publishing can be put into practice.



This research article was awarded Open Data and Open Materials badges for transparent practices. See the Data Availability Statement for details.



**Impact statement**

Data-based methods are increasingly used in mechanics and dynamics and are attractive for engineering design tasks. However, when using such methods, data and its representation become inherent parts of the proposed methodology. This has repercussions on publishing practices, for which challenges are outlined in this article, including perspectives unique to engineering design. Several solution approaches to these challenges are presented, and seven concrete examples of publishing model problems together with open data and code give practical insights. Adapting publishing practices to data-based methodologies in the proposed way can significantly increase the impact of data-driven research in mechanics, dynamics, and associated design tasks, potentially leading to automated design assistants capable of automatically producing designs conforming to specified mechanical and dynamical criteria.

**1. Introduction**

In engineering research and practice, in recent years, data-oriented thinking and the usage and research of data-based methods have become much more prevalent, especially in modeling, simulation, and control. However, engineering research in these fields relies on decades of established processes regarding research communication, publishing, and methodology that have been geared toward classical, first-principles approaches. One manifestation of this is that in areas that were pioneering the latest surge in data-based approaches, that is, mostly subfields of computer science, quite different customs regarding publications are now established. There, datasets and executable code are often openly available with articles. Value is put on open review processes and speedy dissemination of research results, where relevance is preferably ensured through convincing performance in benchmark problems and executable code that can be tested by readers. Yet, this need not mean that copying these customs to data-based research in engineering is the way to go. Engineering problems have their own character, and it may be that some of the traditional customs of disseminating engineering research have value in their own right even in new contexts. After all, engineering research nowadays often uses data-based components, but still also relies on classical methodologies, like first-principles modeling and analytical arguments, or even mathematical proofs where it is feasible and successful.

The described way of thinking has also emerged with researchers partaking in the German Research Foundation's priority program "Daring More Intelligence—Design Assistants in Mechanics and Dynamics" (SPP 2353), which unites at least thirteen senior researchers and a similar number of research associates in the common endeavor to leverage artificial intelligence to assist in designing the mechanics and dynamics of the systems of the future. During their collaborative and diverse work on problems from engineering design, it became evident that some common thought needed to be put on data publishing and sharing. Most importantly, it appeared that dealing with mechanics and dynamics and, in particular, with engineering design leads to specific challenges regarding data. This article is the result of discussions in the priority program and of the experience from nine collaborative research projects therein. It serves three key purposes.

Firstly, the article shall raise awareness to the challenges of data publishing in mechanics and dynamics and system design, in particular for engineering scientists also embarking on more data-driven research, and it shall present some possibilities and guidelines to deal with these challenges.

Secondly, the article presents a set of seven examples from the aforementioned research projects that outline how the described challenges of data publication were overcome. The authors hope that these concrete examples will further nurture ongoing discussions and change within the community and will be of help to other scientists in mechanics and dynamics who want to publish data and data-driven research.

Thirdly, while not the immediate goal, the article and its examples may show scientists from outside engineering that very challenging and worthwhile problems exist in mechanics, dynamics, and engineering design, where good data-driven contributions could lead to sustained technological progress. The authors invite data scientists to try their methods on the published data.

In general, data publishing and also scientific treatments of how and why to publish data are emerging topics. However, the authors are not aware of comparable treatises specifically for mechanics and dynamics and, therein, the topic of engineering design. Still, many existing works and guidelines are also relevant to the field looked at here, and will also partly appear in the following, such as the well-

known FAIR principles on data publishing (Wilkinson et al., 2016). Interestingly, some proposed and recognized guidelines on how to document datasets in the machine-learning community take inspiration from practices established in engineering industries (Gebru et al., 2021), used there to document technical components. Thus, they should feel natural to scientists in engineering. Moreover, whereas this article will exemplify realizations of data publishing at the examples of the authors' datasets from the studied engineering problem-fields, there is a growing number of published datasets in closely related fields for example from natural sciences. In particular, this holds true for datasets containing large libraries of solutions of (often expensive) numerical simulations, guided by the hope to use them to train ML-based solvers that can be queried to more quickly arrive at approximate solutions of expensive-to-solve differential equations. A well-known example is the PDEBENCH (Takamoto et al., 2024), but also datasets specific to computational fluid dynamics may be found (Elrefaie et al., 2024; Yagoubi et al., 2024), which is a field with notoriously costly simulations that are hence particularly worthwhile to expedite or circumvent. A rarer dataset that is directly meant for design applications is presented in Regenwetter et al. (2021), containing a large number of human-designed bicycle models. Similarly, Heyrani Nobari et al. (2022) describe a dataset consisting of 100 million planar linkage mechanisms, which can be used for data-driven kinematic design. Picard et al. (2023) provide guidelines on how to specifically set up synthetic datasets, including technical details like appropriate sampling strategies, but do not focus on challenges present in data publishing in general, independent of whether the data is generated synthetically or recorded from a natural source.

At the same time, an increasing number of methodological works exist that deal with generative design in engineering, see for instance the review paper Regenwetter et al. (2022), giving an overview of various generative methods used for engineering design. Therein, the authors identify several challenges for deep-learning based design methods, including for instance whether deep-learning based approaches can actually be creative enough to produce novel results as well as the usability and feasibility of resulting designs, for instance for production. Unlike the present work, Regenwetter et al. (2022) is focused on the learning methods themselves, not on the aspect of data publishing, its challenges, and implications for the proceedings of science on generative engineering design in mechanics and dynamics. A review of machine-learning methods commonly used in computational mechanics, not necessarily for design tasks, is given by Herrmann and Kollmannsberger (2024). In general, the breadth of work and methodology discussed in Regenwetter et al. (2022) and in Herrmann and Kollmannsberger (2024) can be seen as a motivation for this work's approach to look at data-driven design in mechanics and dynamics from a data-centric instead of a method-focused or application-centric perspective. In particular, the article focuses on what that entails for publishing practices.

The article is organized as follows. Based on a general motivation in Section 2, challenges for publishing data in mechanics, dynamics, and associated design tasks are formulated in Section 3. Ideas and guidelines on how these challenges might be addressed are discussed in Section 4. Section 5 presents examples of published datasets, program code, and model problems, showing first attempts to overcome the aforementioned challenges, and concludes with lessons learned from the examples. Finally, Section 6 draws conclusions and gives an outlook with regard to the findings specific to engineering design.

## 2. The value of publishing data

Traditionally, with classical methods of mechanics and dynamics, calls for the publication of data and code were rooted in the desire to enable a straight-forward reproducibility of results in the spirit of good science. However, if described well enough, proposed methods could also be tried and checked, sometimes quite readily, by implementing and testing them on problems with known reference solutions. It was simple to separate the proposed methodology from input data, such as the initial conditions of a simulation. In contrast to that, for data-based methods, the lines between methodology and data can blur significantly. Put bluntly, training data can become an integral part of the methodology itself (Recht et al., 2019; Sun et al., 2017). After all, there is commonly an interdependency between the makeup of the dataset, such as the selection of data and its mathematical representation, and the quality of the results

obtained with a certain data-based method. In many applications, coming up with a suitable amount of data with the required properties can be much harder than the setup of the learning method and the execution of the training process. Therefore, in such scenarios, the principles of good science require publishing datasets together with the proposed methodology or algorithm. For this, it is adequate to publish the data exactly in the way the (also to be published) code demands. To people more familiar with “traditional” first-principles methods, this may seem like a nuisance brought up by data-driven methods.

However, when the publication of data is thought of in a broader fashion, it can present significant opportunities for research. If properly published, data can serve as an abstraction layer. This enables researchers not familiar with the concrete problem behind the data to propose data-driven solutions, based on abstract properties of the data and the problem, accelerating scientific progress. The underlying problem in question, as represented by the data, can then become a benchmark or model problem used by data scientists to test and develop new methods. The term benchmark problem should be understood slightly differently than it has sometimes been used in mechanics and dynamics. In classical dynamics, it can mean a certain kind of physical problem that every simulator should be able to solve accurately to be deemed “correct” in the sense that it can reproduce an accepted ground-truth solution. Instead, subsequently, benchmark problems are to be understood as prototypes or archetypes for typical problems from engineering (design), prepared in a way suitable for data-based methods and data scientists. In particular, there need not be a ground-truth solution. At best, methods performing well with the benchmark problem may be good candidates for the whole class of problems represented by the benchmark.

It is to be expected that the availability of data and well-prepared benchmark problems can very much influence and steer in which areas progress in machine learning is made. However, an analysis of published datasets and openly accessible and executable benchmark problems shows that there are comparatively few from the field of mechanics and dynamics, especially when focusing on published material suitable for people who are not experts in the field.

Obviously, apart from the kind of data mentioned before, there have always been published datasets that are not meant to be interesting primarily because of their value as a basis for machine learning. For instance, measurements of temperatures at a certain geographic location or measurements of certain material properties are among these datasets for which the provision of the data itself is, at least classically, the purpose of the investigation and publication. This kind of data publishing is not the main purpose of this article.

### 3. Challenges in data publishing

Some challenges to be faced for the publication of data are quite common beyond the specific characteristics of the research on design-assistant systems. These include:

- **Interpretability and Domain Knowledge:** Raw data can be tough to understand for anybody unfamiliar with the generation process of the data. Furthermore, the representation of the data may heavily depend not only on the underlying system but also on the chosen analysis method and observation approach. In simulations, the deformation of a simple mechanical beam, for example, can be computed using numerical discretization schemes that involve the partitioning of the physical domain into small elements. The shape and/or numbering of those elements may however greatly vary across different modeling or simulation approaches, while displaying the same or a very similar elastic deflection of the beam. Comparability of different simulation techniques—and the resulting data—can be difficult for the very same physical system. This can happen analogously in experiments, if measurements are taken at different points of the same mechanical object or with different measurement hardware, which may come with different built-in filtering, for example for noise reduction. In light of this challenge, it is maybe not surprising that many of the first, renewed successes of machine learning in recent years happened for data interpretable without special education, like text, speech, and image data.

- **Problem Complexity:** There can be conflicting interests regarding a desirable level of the studied problem's complexity, and depending on it, it may be hard to place publications that focus on the proposition of a new model problem and accompanying datasets. In mechanics and dynamics, this type of publication is less common than in data science, and the usual conventions in the field may prefer engineering problems at the upper end of the complexity scale, deeming them publishable even outside any machine-learning context. However, overly complex scenarios significantly complicate the explanation of scenarios and data, limiting the appeal to machine-learning experts and thus the potential data-scientific impact.
- **Generalizability:** Compared with language models or computer vision, problems from mechanics and engineering design and corresponding collected data can appear (or sometimes are) overly specific in the eyes of a researcher from the machine learning community, even in cases where the domain expert may know of the potentially large impact of a solution of the underlying engineering problem or, more generally, of its generalizability to other problems. Hence, published datasets may remain unused even if they are easily usable and understandable.

Beyond these, engineering design also presents a very specific challenge, namely the challenge of **evaluation**. In realistic engineering design problems, there is typically not a “correct” reference solution. There usually is no known, globally optimal design for predefined optimality criteria, even if there is an agreed-upon single optimality criterion. In contrast, the goal is to obtain a “better” design than initially given or than achieved with other design methods. This is in stark contrast to benchmark problems widely popular in supervised machine learning, where a labeled set of test data can be used to assess the performance of the learning method using simple calculations. In engineering design, the evaluation process of a candidate design can be very intricate, involving advanced simulations that may require extensive amounts of computation, potentially on high-performance general-purpose computers or specialized architectures. Often more critically, even installing and running the software environment for these test runs may be a burden only domain experts can overcome. Hence, means to judge the quality of results need to be published in an accessible fashion.

Apart from the above-mentioned challenges, there are challenges that appear more rarely in mechanics and dynamics than in some other fields. To keep the scope of this work manageable, we do not focus on such challenges here. One important example can be ethical considerations when using and publishing human-sourced data, which is very common in medical, biological, and social-sciences research. In some situations, for instance in biomechanics or when exploring human opinions on specific results from engineering design processes, similar issues may be faced. In this case, generally, local regulations and legislation have to be followed, and advisory from and agreement by research institutions' ethics boards must be sought.

#### 4. Dealing with the challenges

The following ideas and approaches may help to overcome the identified challenges. This collection of ideas is not meant as a guide on how to overcome the challenges in general, as this still depends on the specific scenario and problem studied. However, afterwards, in [Section 5](#), we give concrete examples on how the subsequent ideas can be practically used to deal with the challenges.

**Evaluation:** If the evaluation mechanism happens to be lightweight and easy to install and use, it can be sufficient to distribute it with the dataset or model problem. In other cases, a workable, yet more involved, solution is to publish a (web-)interface to an evaluator, to which results can be uploaded. The evaluation is performed on the server, and the evaluation results are made available to the user. The advantage is that the server and evaluation toolchain are set up and configured only once by the domain experts. Luckily, the required technology is already available from other domains, albeit for slightly different reasons. In data science, it is sometimes even undesirable to publish test datasets used for evaluation since then researchers may tune their method also with the test dataset. For these reasons, there exists open-source software that greatly helps to implement such online-evaluation or benchmarking



tools. A popular example is Codabench (Xu et al., 2022), which is even used to do public competitions for defined tasks. Some challenges, like the funding and maintenance of the server, remain. However, a growing number of success stories shows that this can work and can significantly help to popularize a benchmark problem, particularly in the crucial phase when the problem is new.

To find new, interesting papers, data scientists nowadays often browse through websites such as Papers with Code (Meta AI, 2023), where papers appear jointly with their open-source code and their performance in public benchmark problems. There, papers are not ranked by traditional measures, such as the prestige of the journal they appear in, but based on benchmark results as well as on ratings of users, which may base their rating on their experience when executing the code. State-of-the-art approaches are grouped by tasks (closely related to the type of data such as time series, images, etc.). The visibility of papers is dictated by the performance in known benchmark problems rather than by their number of citations. Thus, papers without code or benchmark results will most likely remain invisible there, highlighting the culture change from the “traditional” ways of mechanics and dynamics. Another popular platform to share datasets and trained machine-learning models is Hugging Face (Hugging Face, 2023), where grouping the available datasets by tasks again reveals a focus on natural language processing and computer vision, and a severe lack of mechanics, dynamics, and engineering design examples. Also popular among different disciplines is Code Ocean (2025), with an original focus on bioinformatics, but offering a platform useful for data-driven research in many fields. It offers reproducible cloud installs of environments for scientific computing, which is important for reproducibility, and allows the joint publication of code and data, as well as the association with a scientific article.

**Interpretability and Domain Knowledge:** Some of the issues regarding interpretability and domain knowledge are appropriately dealt with by established principles for data management, such as the FAIR principles (short for findable, accessible, interoperable, reusable, Wilkinson et al. (2016)), where some of those principles can be fulfilled by including descriptive metadata with the data. There exist professionally curated data repositories implementing these principles, some operated by universities (for example the DaRUS repository, see University of Stuttgart (2023)) or by third parties. There also exist registries of research data repositories, see for instance re3data (Karlsruhe Institute of Technology, 2023). Moreover, as planning, maintaining, managing, and curating data is a complex process, professionals focusing on this task, so-called data stewards, may support the domain experts in data publishing, see Basalti et al. (2024) for more details on the role of data stewards, as well as Flemisch et al. (2024) for a practical case study from data management in simulation science. If a research unit runs its own public data repository, data stewards may also be curators of such a repository. Data stewards can directly ensure that common principles like the FAIR principles are fulfilled, and, in exchange with the domain experts that have created the data, at best can also give advice on how to deal with other challenges that first appear more domain-specific, like generalizability.

For engineering and engineering design problems, mere metadata, such as describing what is stored in which format, may still not be enough to facilitate usage, especially for researchers from other domains. Hence, at the very least, datasets should come shipped with some open-source code that, as a form of minimal example, reads the data from the data files and visualizes it in an intuitive way. The latter is very important in the field of mechanics and dynamics, as a good visualization of the data can be quite insightful—compare, for instance, an explanatory spatial animation of motion to just supplying raw data or line plots in a variety of different coordinate frames. If data and program code are published at different places, hyperlinks should point from one to the other, as well as to and from associated research articles. A further way to increase accessibility and reduce the inevitable technical hurdle to looking at published code and data is technology such as Google Colab (Google, 2023), which permits very straightforward (Python) code execution via a web browser, including access to GPUs. An explanatory worksheet, which outlines the idea behind the problem, interactively visualizes the data, and explains the basic usage of the proposed methodology, can help greatly to familiarize interested parties with a specialized problem, as the execution of the worksheet is just one hyperlink and one click away. Technology, regulations, and guidelines for research data management advance swiftly. Even binding legislation may appear in the future (Deutsche Forschungsgemeinschaft, 2023c). Many universities now have strategies for the

management of research data that provide useful information. Moreover, helpful resources include the German Research Foundation's comments on the topic (Deutsche Forschungsgemeinschaft, 2023b), including possibilities to receive funding for data management, which supplement the usual guidelines for good scientific practice (Deutsche Forschungsgemeinschaft, 2023a).

**Conflicting Interests on Complexity:** Before publishing data or a model problem, one needs to decide on the goal of the data publication, including who the target audience is (closely connected to the venue of publication), and what the benefit is of finding new, better methods to deal with the data. The complexity of the problem should then be calibrated so that it is engaging to the target audience in its difficulty, where difficulty may mean different things for different audiences. Still, domain experts may not see the value in handling a problem that they deem “simple” (e.g., already solved by “established” methods) with new methods from machine learning. This risk can be alleviated by an appropriate framing of the problem as well as an explicit explanation of the benefits of finding new methods and of the perspectives to move to more intricate problems when the simpler ones have been solved with data-driven approaches. Fortunately, engineering design processes deliver plenty of natural motivation, for instance, because traditional methods currently quickly lead to computational intractability, especially when combining multiple design aspects. In a more classic motivation for data publication, published data can also help readers check their own implementations, and thus their understanding of the published methods. In that case, the data should be composed in a way that allows an exact reproduction of the published results. The published results and their complexity should then be set up so that their correct reproduction is a good indication of a correct implementation.

**Generalizability:** To alleviate concerns regarding the specificity of tasks, once again, a proper framing of the problem is important. This is often glossed over in domain-specific publications, as, for instance, the significance of typical fluid-dynamics model problems is clear to mechanical engineers. In contrast, to have a data-scientific impact, the motivation behind and the characteristics of engineering tasks need to be explained in a way engaging to researchers usually dealing for example with language or image processing. Guiding questions may be: Why is the problem of general importance? What makes the problem engaging from a data-centric perspective? How are the mechanical and dynamical intricacies of the problem represented in the specific properties of the data? Are there problem-specific requirements, for instance on the inputs and outputs or the mathematical properties of the relationship between them? When purely looking at the data and its properties, are there similarities to data from other tasks and domains? Where are interesting differences?

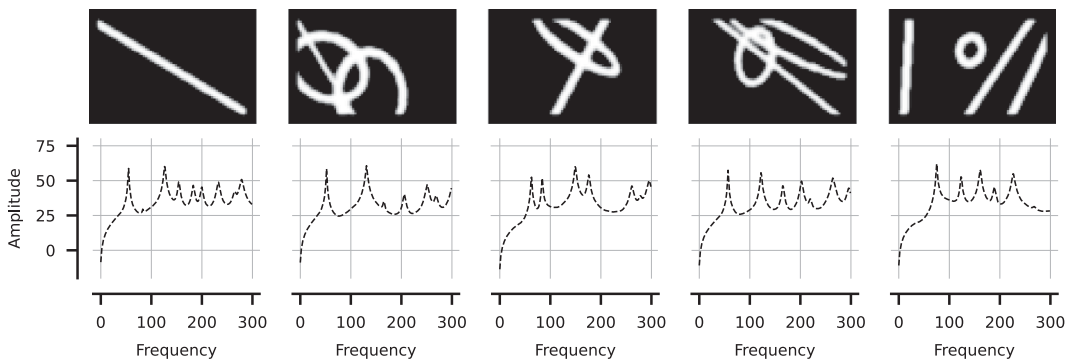
## 5. Examples: published code and data

In this section, published data and code, as they have resulted from the priority program “Daring More Intelligence—Design Assistants in Mechanics and Dynamics” (SPP 2353) of the German Research Foundation, are listed and discussed in the light of the aforementioned findings, to serve as inspiration for how some of the formulated challenges can be overcome, and to learn which types of data and problems appear in mechanics, dynamics, and engineering design.

### 5.1. Vibrating plates

Noise and vibrations of mechanical structures such as machinery and vehicles impact human well-being and health (Basner et al., 2014), which could be mitigated with the aid of design assistants aimed at reducing noise and vibrations. Motivated by this, the vibrating plates dataset addresses the design of plates with indentations to minimize vibrations. In mechanical structures, vibrations emit noise and can be transmitted from sound sources like engines to, for example, the passenger cabin in a car. Indentations in plate-like structures can help reduce vibrations if they are well-placed.

Figure 1 visualizes typical pieces of input and output data. On the input side, the dataset consists of image-type data describing the indentation patterns and scalar parameters describing the material as well as the geometric properties of the plate. On the output side, the dataset describes the dynamic response of



**Figure 1.** Example plates from the vibrating plates dataset along with the averaged mean squared velocity of the vibrations. Indentations are marked in white in the top row.

the plate at discrete excitation frequency steps. Here, different fidelity resolutions are provided: Image-type data of the velocity field represents high-fidelity data, whereas the domain averaged mean squared velocity represents an integral quantity.

#### 5.1.1. How are the challenges addressed?

Regarding **problem complexity**, this dataset exhibits many representative characteristics of a vibroacoustic problem, but the choice of indented plates instead of more complex geometries enables the straightforward application of many standard machine-learning methods. The dataset is published in a curated data repository from the University of Göttingen (van Delden et al., 2024d) and listed on Papers with Code (van Delden et al., 2023). The code is published in a repository on GitHub (van Delden et al., 2024a). The data is stored in a common format (hdf5) compressed with a publicly available compression algorithm.

The dataset can be used to train machine-learning-based regression models. In particular, the regression model should learn the mapping from beading patterns and plate properties to the velocity field. Regarding **evaluation**, metrics are defined in the paper to evaluate the performance of the regression model on a separate test dataset, for example, the classical mean squared error in the dynamic response or the error in the resonance-peak location. Results for several deep-learning-based methods are provided.

Regarding **interpretability**, the dataset (van Delden et al., 2024d), the associated paper (van Delden et al., 2024b), and the accompanying code (van Delden et al., 2024a) are designed to be accessible to non-domain experts who have a background in machine learning. To achieve this, the descriptions in the paper avoid domain-specific language where possible, and important concepts are explained within it. Also, care was taken to follow conventions for machine learning papers and code repositories. This includes a focus on reproducibility and employing standard convenience tools like conda environment definitions, jupyter notebooks, and a clear and detailed README file in the code repository. In addition, video visualizations of the dataset are displayed directly in the README and can be produced with a provided notebook. As a verification step, a researcher from a different institute with no prior knowledge was asked to download and visualize the data and was able to do this within 10 to 15 minutes. Based on this dataset, two additional publications have been accepted, dealing with extending the dataset (Schultz et al., 2024) and employing guided diffusion to targeted design (van Delden et al., 2024c).

## 5.2. Database of crashworthiness analysis of crash boxes

Every year, approximately 2.54 million people in Germany alone are involved in automotive crashes, which makes vehicle crashworthiness a crucial part of the development of automobiles (Destatis, 2023). To help reduce the seriousness of the impact and prevent serious injuries and fatalities, both active and passive crash management systems are installed in automobiles. Active safety systems are designed to




prevent accidents, and passive safety systems are usually structural components that are built to absorb impact energies and reduce potential harm to the occupants in accidents. Since these systems are static, they must be meticulously designed and optimized to perform in various scenarios. One such passive safety system component is a crash box behind the front bumper, as shown in Figure 2. The depicted model is derived from the Honda Accord model provided in the NHTSA database (National Highway Traffic Safety Administration (NHTSA), 2016). Crash boxes are thin-walled hollow structures that are installed between the subframe and the front bumper. These absorb the kinetic energy of the impact by undergoing axial deformation and local buckling, thereby reducing the forces experienced by the occupants. The optimal design of the crash box is crucial since frontal impacts are the most common type of impact in road traffic.

### 5.2.1. How are the challenges addressed?

The database is generated during the investigation to improve the methods used in the crashworthiness design and incorporate machine learning in the optimization process for the square crash-box design. This database is not just a result of an associated research project; it has been developed to provide other researchers with valuable insights and tools. The database is publicly available in a GitHub repository (Borse et al., 2024a), Papers with Code (Borse et al., 2025), and the related optimization method results are published in the article (Borse et al., 2023). The provided data encourages scientists and engineers to:

- **Compare different methods:** Researchers can utilize the data to **evaluate** their optimization and FE surrogate methods against the results from the analysis in Borse et al. (2023). This can foster constructive discussion about different approaches to improve crash safety and optimization methods. Additionally, researchers can utilize the data to investigate and perform comparative studies using classical optimization methods such as evolutionary algorithms, response surface methods (RSM), or modern random search algorithms.
- **Verify their implementations:** The detailed simulation results offer a solid foundation for testing and validating individual models or simulation techniques in crash analysis.
- **Develop new approaches:** Access to the data allows others to develop innovative methods or adapt existing techniques to meet specific requirements of their research projects better.
- **Encourage data sharing and collaboration:** Even if the collection of examples in the dataset is small, it can inspire others to share their problem-specific datasets, considering how these could be beneficial for the community.


Concerning **problem complexity**, the database maps the relation between the structural and crash-test parameters to that of various crashworthiness objectives used to evaluate the optimality of the crash-box



**Figure 2.** Crash box position in automobiles.

design. This database provides valuable data without the need to perform problem-specific and complex simulations and can be directly used to evaluate developed methods. In the provided database, crash box parameters (width, thickness, and length), crash-test parameters (impactor mass, initial impact velocity), and time required for simulations are recorded. Multiple crash-box configurations are evaluated by varying the thickness of the crash box and impact velocity. The crashworthiness objectives are peak impact force, energy absorbed, maximum deformed length, mean contact force, and mass of the crash box. They are calculated for every configuration based on the FE simulation results. The neighboring components define maximum peak impact force as they should not buckle before the crash box, and, generally, the energy absorbed should be maximized. Other parameters, such as maximum deformed length, mean contact force, and mass of the crash box, are informative parameters of the crash box's overall performance. Assuming the energy absorbed has to be maximized while reducing the maximum deformed length leads to certain non-dominated solutions, as shown in Figure 3. These form the so-called Pareto front, denoting solutions where both objectives are considered but have different importance. Similarly, different relations between the input parameters and objectives could be drawn to optimize the design parameters. Consequently, this dataset can be **generally** applicable to study relations between input parameters and crashworthiness objectives and for testing optimization methods.

Regarding **interpretability**, the database can be used for solving optimization problems, as demonstrated in the development of the multi-objective optimization method using machine learning algorithms described in Borse et al. (2023). This database can be further adapted (by varying the width and thickness of the crash boxes) for multi-parameter, multi-objective optimization of the crash box design to determine the optimal thickness and width for an ideal peak impact force and allowed maximum deformed length of the crash box (Borse et al., 2024b). These investigations provide one way to interpret the given database and can be used to develop comparable optimization methods. A list of all the dependencies required for executing the scripts is mentioned in the GitHub repository in the detailed ReadMe file, along with the Abaqus script to generate a similar database.



**Figure 3.** Pareto front for the objectives: maximum deformed length and total energy absorbed.


### 5.3. Backbone reconstruction of a non-slender soft robot

The research field of soft robotics is based on the investigation of robots made up of elastic soft materials like silicone. The use of such elastic soft materials results in a continuous deformation of the robot under load or even under its dead weight. However, compared to conventional robots made up of rigid connections, the computation of this continuous deformed state can be more involved, but is a crucial part of the control of the soft robot and for the verification of models. The so-called backbone can be seen as a one-dimensional description of the continuous deformation of the soft robot, as visualized in Figure 4 for a deformed soft robot. Especially when manufacturing soft robots by hand, variances in the geometry and the manufacturing accuracy change the behavior significantly. Hence, being able to identify the parameters of the soft robot via image processing is a fast and easy way to enhance the reliability of subsequent modeling and to validate simulation results of the soft robot with geometrically accurate beam models. To that end, first, a stereo camera system is set up and calibrated. The soft robot is then actuated with different pressures while both cameras capture images to obtain the raw data. The dataset contains the calibration data, raw images, and images annotated with the reconstructed backbones. The corresponding parameters of the reconstruction are available in Schindler and de Payrebrune (2024a). Figure 4 displays an annotated entry from this dataset. The blue, dashed line represents the reconstructed backbone, whereas the white markers indicate the estimated points used in the optimization process.

Applying the image-based method published by Hoffmann et al. (2024b), with modifications described in Schindler and de Payrebrune (2024b), we adapt the technique for reconstructing the backbone of a non-slender soft robot. The original method is provided as the Python software package *icpReconstructor*, and the modifications and extensions for a modular bending actuator, along with a short example, can be found in the associated GitHub repository (Hoffmann et al., 2024a), as detailed in Hoffmann et al. (2024b). The image dataset is available in Schindler and de Payrebrune (2024a). Code, data set, and the associated paper are listed on Papers with Code (Schindler and de Payrebrune, 2025).

#### 5.3.1. How are the challenges addressed?

In the dataset, the raw images are annotated to face the issue of **interpretability** of the results. Since the optimization for the reconstruction of the backbone operates in pixel space and can be depicted in both figures at each iteration step, the results can be interpreted and even evaluated visually, allowing the practitioner to verify the inner workings of the reconstruction method. Furthermore, with a rendering of a deformed geometry using the reconstructed backbone from the dataset, the **evaluation** can be further supported. While the reconstruction method can be implemented in software, its usefulness is only leveraged with data captured in experiments. Therefore, the dataset can be seen as part of the modified model. Even if the adaptation contained in the dataset is specific to the employed soft robot, it can be viewed as an example of how well the method introduced by Hoffmann et al. (2024b) can be **generalized**



**Figure 4.** Images of a non-slender soft robot annotated with the reconstructed backbone (blue dashed line) and estimate points (white markers). Taken from Schindler and de Payrebrune (2024a) by CC BY.

to different tasks. The application of this method is not limited to soft robotics but can be modified to fit different applications where stereo images are available to determine a backbone-like structure. For example, further possible use cases include spine reconstruction from X-ray images, analysis of animal movements, or orientation detection in conventional robots.

A single soft-robot bending actuator, as shown in Figure 4, has limited applicability on its own, as it is usually part of a larger soft robot system. However, focusing on a single module reduces the **problem complexity** that arises when considering an entire soft robot system with multiple actuation parameters and potential continuous deformations, and makes the reconstruction method more feasible. Nevertheless, the deformation resulting from the pressurization of one or more chambers of the soft robot with or without external loads can be evaluated and compared with the beam deflection. The described datasets of stereo images **offer numerous further possibilities for the analysis of soft robots**. Besides the original detection of the backbone, the dataset can be used to analyze and identify parameters of different mechanical properties of the soft robot, for example, axial stiffness, bending stiffness, and torsional stiffness, where the pressure actuation and resulting deformations have to be put in relation. In addition, the behavior of the robot can be evaluated for different material models, providing valuable insights for optimizing design and performance. These applications can significantly improve the understanding and development of efficient and adaptable soft robotic systems.

#### 5.4. Duffing oscillator response analysis (DORA)

In mechanical engineering, the state-of-the-art system design paradigms are mostly based on a number of crucially simplifying and idealizing assumptions on dynamic loads and system complexity. In particular, the non-stationarity of the system dynamics is often regarded as a perturbation to some steady-state operation. However, static load cases actually occur only rarely in real-life systems: aircraft engines, wind turbines, and vehicle components consistently operate under non-stationary, non-periodic, multi-scale, multi-physical, or in total **complex** external loads and control inputs. Hence, the transient dynamics of engineering structures must be taken into account early in the design process. At the same time, data acquisition can be expensive, rendering the available sequential **data sparse and small**: typically, only a few operating conditions and corresponding system responses are available. This situation requires models that **generalize** very well to unseen system parameters or external loads, and which are able to predict complex dynamic behavior. The **problem complexity** thus arises from four core aspects: (a) the problem involves time-dependent multivariate quantities on input and response sides, (b) many variable system parameters and external loads that affect the system response, (c) critical transitions occur in the system response due to minor changes in the system parameters, and (d) limited observability and small data in the form of time series. **Interpretability** is not at the core of this modeling task, however highly welcomed to elaborate on how a model picks up transient behavior.

The dataset presented in this section deals with the aforementioned challenges in the following ways: it provides data in the form of time series of a minimal mechanical oscillator, namely the Duffing oscillator. This model is complicated enough to show all qualitatively possible deterministic dynamics (regular to irregular) with transitions governed by parametric changes. Moreover, using a synthetic data approach, it is possible to generate training datasets of varying sizes, thereby allowing to scale the complexity stemming from data availability. The Duffing oscillator represents an externally forced nonlinear oscillator with dynamics

$$\dot{q}_1 = q_2, \quad (1)$$

$$\dot{q}_2 = -cq_2 - kq_1 - \beta q_1^3 + f \cos(\omega t) \quad (2)$$

where the parameters are chosen as  $c = 0.32$ ,  $k = -1$ ,  $\beta = 1$ ,  $\omega = 1.5$  to provoke a wide range of differing dynamics under changes to the external forcing ( $F(t) = f \cos(\omega t)$ ). Duffing-type oscillators have been studied and implemented in a wide range of physics and engineering applications such as MEMS

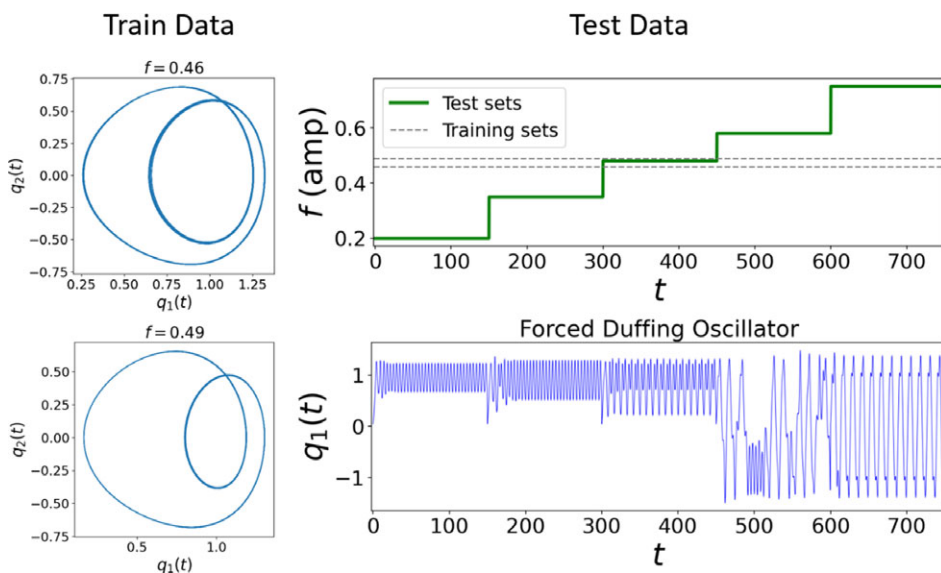
resonators (Sabarathinam and Thamilaran, 2016), in energy harvesting devices (Martens et al., 2013), and to study periodic forcing-dependent chaotic behavior of nonlinear systems (Strogatz, 2024). Various qualitative changes (bifurcations) of the system dynamics are obtained by varying the forcing amplitude ( $f$ ) while keeping the forcing frequency constant. This variation allows to generate period-1 limit cycle solutions for small forcing amplitude, but also chaotic dynamics for larger forcing amplitudes. Rapid changes in the dynamics occur under minimal changes to the forcing amplitude (period-doubling bifurcation route to chaos) as displayed in Figure 5.

The modeling challenge takes the general form of multivariate sequence prediction for parameterized dynamics and requires auto-regressive models that can advance from a given temporal context into the future. The dataset is comprised of time-series data obtained for specific values of the forcing amplitude ( $f \in \{0.46, 0.49\}$ , corresponding to period-2 cycle motions). From this dataset, the model is requested to predict vibration time series for different forcing values far from the training set, spanning values that are multiples smaller or much larger than the training regime. The modeling task hence asks for time series prediction of qualitatively strongly varying dynamics, and extreme out-of-sample **generalization** from a minimal amount of training data. Successful models will allow to predict the system's transient dynamics across a wide range of forcings, and to accurately predict critical transition points (bifurcations).

#### 5.4.1 How are the challenges addressed?

For **evaluation** purposes, the success of the prediction model will depend on the accurate capture of the overall dynamics (the *climate*) as well as the short-term fully-resolved temporal dynamics. The system response characteristics are quantified in terms of the (mean) amplitude of the system state  $q_1^2(t)$ . The mean squared error (MSE) can be used as an accuracy quantifier to obtain the deviation of predicted system response characteristics relative to the original ones for each of the external forcing amplitudes, given by

$$\varepsilon_A = \text{MSE}\left(\text{Amp}\left(q_{1,\text{prediction}}^2(t)\right), \text{Amp}\left(q_{1,\text{original}}^2(t)\right)\right), \quad (3)$$



**Figure 5.** DORA time series prediction task in the small-data limit: the training data (left) comprises two trajectories of period-2 cycle dynamics for two values of external forcing amplitudes. The modeling tasks aim at generalization to different forcing amplitudes that induce qualitatively different dynamics, among others chaotic dynamics, as displayed on the right.



$$\varepsilon_M = \text{MSE} \left( \text{Mean} \left( q_{1,\text{prediction}}^2(t) \right), \text{Mean} \left( q_{1,\text{original}}^2(t) \right) \right) \quad (4)$$

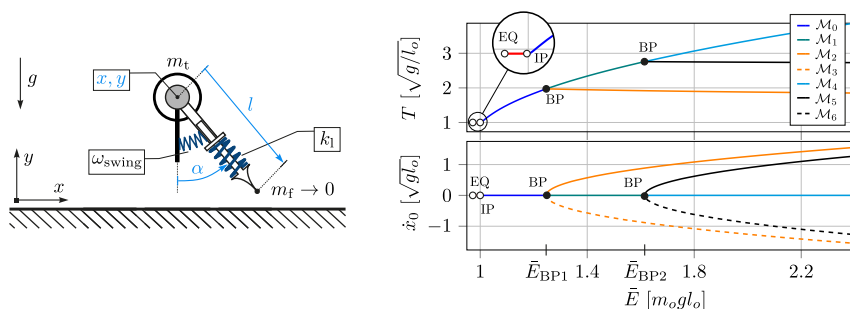
where  $\varepsilon_A$  and  $\varepsilon_M$  are the response amplitude and response mean errors, respectively. A more involved description of the DORA task, files, and data are provided in an open-access GitHub repository and published on Zenodo (Yadav and Stender, 2025).

One approach to this task is to leverage a machine-learning framework capable of sequentially generating predictions while adapting to the evolving periodic external forcing. Specifically, the ML model takes the dynamic external forcing at time  $t$ , denoted as  $F(t)$ , as input and outputs the response of the Duffing oscillator,  $[q_1(t), q_2(t)]$ . These predicted outputs, along with the next external forcing value  $F(t+1)$ , are then used to generate the subsequent predictions,  $[q_1(t+1), q_2(t+1)]$ . This iterative process continues to model the response of the Duffing oscillator for the given external forcing.

The most crucial and challenging aspect of this task is ensuring the **generalization** capability of the machine-learning model. Specifically, the model must be capable of generalizing across the system parameter, which extends over a range of external forcing amplitudes,  $f \in [0.1, 0.8]$ , despite being trained solely on time-series data from only two specific forcing amplitudes. A representative solution to this problem is demonstrated in Yadav et al. (2024a), the work is also listed on Papers with Code (Yadav et al., 2024b). In this work, the authors employ an unconventional machine learning framework—reservoir computing (Jaeger, 2001; Yadav et al., 2025)—to accurately capture the forced Duffing oscillator’s behavior using minimal training data. Reservoir computing is a machine learning paradigm that leverages a fixed, high-dimensional dynamical system (the “reservoir”) to transform inputs into a complex feature space, where a simple readout layer is trained for predictions. Notably, their approach ensures generalizability across the range of external forcing amplitudes, leading to accurate predictions of the Duffing oscillator’s response and enabling the reconstruction of its entire bifurcation diagram, showcasing the framework’s effectiveness in capturing complex dynamical responses.

### 5.5. Passive one-legged hopper

Similar to the Duffing oscillator presented in the previous section, this dataset is associated with the challenge to learn the dynamics, approximate the first-order return map, and predict the bifurcation behavior of a complex nonlinear system. It is based on the dynamic model of a one-legged hopper, see Figure 6. In comparison to the Duffing oscillator, it stands for a far more complex class of dynamical systems, which are not only nonlinear, but also hybrid and non-smooth, and which cannot be modeled using a single set of ordinary differential equations. Such dynamics greatly increase the **complexity** associated with simulation and optimization problems. They are found, for example, in robotic systems that interact with the environment such as in object manipulation or legged locomotion. This interaction



**Figure 6.** Depicted are the passive one-legged hopper and a bifurcation diagram of its periodic trajectories. Since the system conserves energy, the bifurcation parameter is the energy level  $\bar{E}$ . The projections of the periodic motions presented are the period time  $T$  and horizontal velocity  $\dot{x}_0$  at the apex of the hopper. The illustrations are taken from Figures 3 and 4 in Raff et al. (2022a).

introduces non-smooth dynamical effects (rapid changes in a system's velocities), necessitating more sophisticated simulation tools that need to detect when such non-smooth effects happen and treat their effects appropriately.

In periodic tasks, such as legged locomotion, the resulting motion is often regarded on the basis of a recurrence map, which computes how the state is updated during one full period. As a consequence, the resulting simulations are often computationally expensive. This holds in particular when gradient information must be computed within the simulation to obtain sensitivities. This derivative information is of great importance not only for the study of the bifurcation behavior but also for understanding stability and developing efficient and robust controllers for robotic systems. Despite the high complexity involved in generating time-series data for non-smooth dynamical systems, the resulting dataset can still be low dimensional, which makes it an interesting candidate for learning surrogate functions.

The one-legged hopper is a good example of such a hybrid dynamical system. It has only two parameters (swing-leg frequency  $\omega_{\text{swing}}$ , leg stiffness  $k_1$ ) and six states  $(x, y, \varphi, \dot{x}, \dot{y}, \dot{\varphi})$ . Yet, the system exhibits non-smooth behavior when the foot comes in contact with the ground, and generating time-series data is thus quite challenging. As a consequence of the nonlinear dynamics, the hopper also exhibits a rich variety of periodic motions, similar to the Duffing oscillator in the previous section. A projection of these periodic motions is depicted in the bifurcation diagram shown in Figure 6. The hopper also serves as a prominent template model within the fields of biomechanics and robotics (Blickhan, 1989; Geyer et al., 2006; Gan et al., 2018).

#### 5.5.1. How are the challenges addressed?

Data-driven predictions on the hopper's dynamics can take one of three forms, depending on how its periodic time-series data is **interpreted**:

1. The evolution of the time-series data can be used to do a sequence prediction or to train a model for the right-hand side of a discrete dynamical system. This is similar to how the data in the DORA dataset is handled.
2. Each time-series in the dataset corresponds to a periodic motion and thus represents a fixed point of a first-return map. Learning such a map significantly simplifies the system dynamics by reducing a recurrent motion to a single point. Again, this can be interpreted as a sequence prediction for parameterized dynamics.
3. The bifurcation diagram can also be learned directly, requiring a model that takes the bifurcation parameter ( $\bar{E}$ ) as its sole input.


These three approaches to training a machine-learning-based model are listed as of increasing **complexity**, as the availability of training and testing data decreases significantly for the latter approaches. The **evaluation** of the dataset depends on the chosen prediction method. For the first approach, the metrics presented in Equations (3) and (4) are suitable for data evaluation, whereas the latter two approaches can be assessed by their ability to generate the bifurcation diagram shown in Figure 6. Regarding **generalizability**, the hopper example belongs to the complex class of non-smooth dynamical systems making suitable machine-learning techniques directly applicable to other datasets from this class.

The dataset is publicly available in DaRUS, the curated data repository of the University of Stuttgart (Raff and Remy, 2024), and listed on Papers with Code (Raff et al., 2025). The model is described in detail in Raff et al. (2022a). This paper is accompanied by code on GitHub (Raff et al., 2022b). The time-series data of periodic motions is provided in comma-separated text files encoding matrices of varying sizes. The provided dataset description, as well as MATLAB and Python scripts, include the necessary information to access the time-series data and plot the bifurcation diagram in Figure 6. Additionally, the non-smooth periodic motions can be interactively visualized with the provided MATLAB live script, allowing for a detailed examination of the bifurcating trajectories. This allows for a detailed examination of the bifurcating trajectories.


### 5.6. Motion data of a four-bar mechanism

Creating virtual models of mechanical systems is a common task in engineering to gain better insights and to develop suitable control algorithms. However, this requires careful and extensive parameter identifications of the systems and the resulting models usually do not perfectly describe the behavior due to unknown dynamics, complex state-dependent forces, or other, intentionally unmodeled effects. Utilizing recorded data from hardware experiments can help to improve the virtual model and better align it with the real-world system. This motivates the development of pure data-driven methods or hybrid models that combine data-driven and physics-based approaches such as in Röder et al. (2024a).

The presented dataset contains motion and motor-current data from hardware experiments of a four-bar mechanism. Despite its simplicity, the four-bar linkage is a widely used mechanism, because when carefully designed, it can convert rotational motion into translational motion in a very flexible manner. The mechanism is driven by a Dynamixel XH430-W350-R motor that allows for current control, i.e., desired current values are supplied, which the onboard electronics strives to reach. The prescribed current trajectories serve as the input to the system, whereas the measured, actual current, and the resulting angular positions and the angular velocities of the motor shaft are recorded. Figure 7 shows the hardware prototype and Figure 8 shows one of the recorded trajectories.



**Figure 7.** Hardware prototype of a four-bar mechanism used to stir a viscous fluid, where the stirrer can be replaced. An exemplary stirrer is shown in the lower right.



**Figure 8.** (a) Prescribed current trajectory and corresponding current measurements. (b) Measured velocity for the applied currents.

By mounting a stirrer onto the mechanism and moving it through a viscous liquid, different damping effects can be introduced into the system. The complete dataset, a notebook for intuitive data handling, and a detailed description of the mechanism are published in a curated data repository of the University of Stuttgart (Röder et al., 2024b). They were used by a partner research group to compare two different hybrid modeling approaches in Wohlleben et al. (2024). Such an identified model can then be used, for example, to obtain a model-based controller, furnishing, together with the data-driven modeling process, a control-design assistant.

The measurements can be utilized by others to train an end-to-end data-driven model of the system capable of predicting the behavior without any expert knowledge. Another approach is to build a hybrid model where either an error correction model can be learned or the model parameters are fitted from the data. This can lead to a better understanding of the model and potentially allow insights into the previously unknown dynamics. Additionally, as with any data from a physical system, the presented data exhibits measurement noise and other disturbances such as coarse measurement resolutions, which are often not dealt with in purely theoretical works on data-driven methods in engineering, but which are a regular occurrence and key limiting factor on real hardware. Hence, the dataset can also be used to test different preprocessing strategies for handling noisy and quantized data.


#### 5.6.1. How are the challenges addressed?

For the **evaluation** of proposed models, a part of the data can be held back as test datasets to assess the prediction accuracy and generalization capabilities. Various metrics can be defined such as the mean error for a one-step prediction or comparisons of the complete time series data to check for slow drifts over time. Even though the system is conceptually simple, learning models from hardware data exhibits a high **problem complexity** due to unknown external influences such as manufacturing inaccuracies, the presence of measurement noise and quantization, imperfections of the motor unit, and introduced damping stemming from the fluid-structure interaction of the stirrer with a viscous liquid. Concerning the **interpretability**, the dataset itself is easy to understand and is provided in text files which can be quickly loaded into any programming environment. Additionally, a Python notebook is provided to guide the user through the data and to show how to load and visualize it. This allows also non-domain experts to develop their own data-driven models without any system knowledge of a simple mechanical system commonly used in engineering. Furthermore, the dataset can be utilized as a benchmark to compare against hybrid models developed by domain experts. Regarding **generalizability** and the scarcity of data from hardware experiments, the presented dataset also serves as a starting point for the development and verification of algorithms applicable to more complex systems.

#### 5.7. Axial turbine dataset

In recent years, the demand for renewable energies has drastically grown to reduce the dependency on fossil fuels. Among all, hydropower itself contributes 47% of total renewable energies produced in 2023 (Ritchie et al., 2024). Axial turbines are one of the popular machines for hydropower plants due to their nature to be suitable for low-head conditions. Typically, hydraulic machines are designed for an operating range with several operating points simultaneously. In their simplest form, they are designed specifically for one operating condition to be fulfilled. For hydraulic machines, the common objectives are design head ( $\Delta h$ ; energy difference between flow inlet and outlet of the turbine), efficiency ( $\eta$ ), power ( $P$ ), mechanical stresses ( $\sigma$ ), and volume of cavitation ( $V_{\text{Cav}}$ ; the volume of fluid (vapor-filled cavities) below the vapor pressure within the liquid). Hence, multi-objective optimizations are performed to get a finalized shape of the axial turbines (Fraas et al., 2022). Performing multi-objective optimization of an axial turbine is expensive because of the number of expensive computational fluid dynamics (CFD) simulations involved.

Figure 9 visualizes a typical section of an axial machine and the corresponding b-splines to parameterize the blade. The detailed descriptions of the axial turbine and the parameterization are published by Rentschler et al. (2024) and Raj et al., 2024b. Usually, hydraulic turbines are operated at different



**Figure 9.** Section of an axial turbine and the parameterization of the blade (Rentschler et al., 2024; Raj et al., 2024b).

operating conditions like part, full, and nominal loads. Hence during optimization, it is important to take care of operating conditions as well. This means that for each operating point, a separate CFD simulation has to be performed. The overarching optimization target is to obtain high-efficiency values and no cavitating regions for the desired net head and power range.

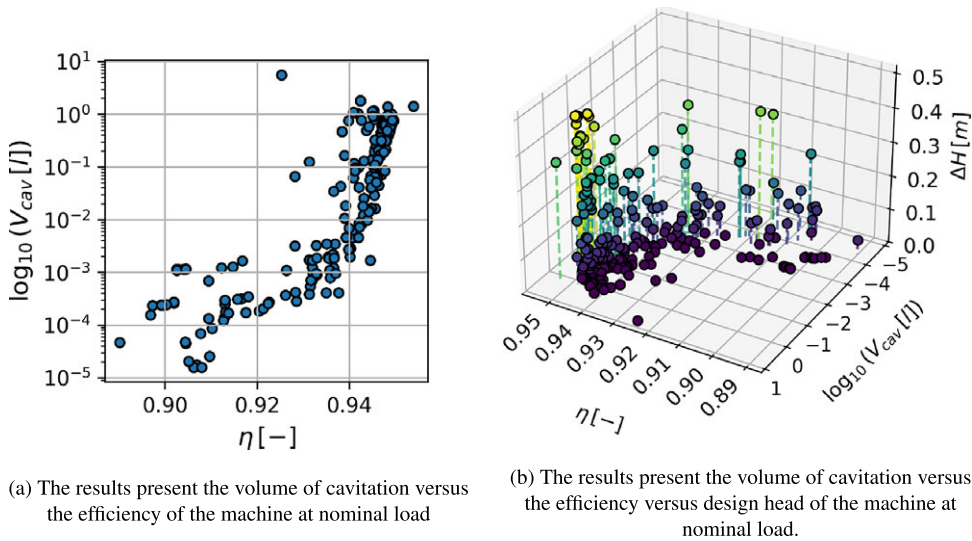
The dataset presented here is an outcome of an optimization run of an axial turbine shown in Figure 9. The inputs of the dataset consist of 30 geometry parameters, which define the blade at the hub, shroud, and mid-span. The shroud and hub diameter remain constant in the current case. The outputs of the dataset are the corresponding post-processed results of the following CFD simulations. They are power ( $P$ ), design head ( $\Delta h$ ), the efficiency of the machine ( $\eta$ ), and volume of cavitation ( $V_{Cav}$ ). The presented dataset is publicly accessible at the Github repository (Raj et al., 2024a) containing geometry and result data in tabular form. It is also listed on Papers with Code (Raj et al., 2025). Detailed descriptions of inputs are given in the open-access article by Raj et al. (2024b). It is also possible to generate and extend the dataset. The expansion of the dataset for training, validation, and testing will include CFD simulations to perform, which are computationally expensive. The docker container is available with all source code on another Github repository (Tismer, 2020a), which is based on the design tool (d)esign (t)ool (O)bject-(O)riented (dtOO) (Tismer, 2020b). Figure 10 illustrates the trade-off between different objective functions. The presented Pareto front in Figure 10a is a two-dimensional projection of the three-dimensional set of non-dominated solutions, displaying the key objectives when the volume of cavitation is minimized at the same time, maximizing the efficiency of the machine.

This comprehensive axial turbine dataset offers valuable resources for both engineers and machine learning practitioners. Engineers can utilize this data to optimize turbine designs, and benchmark performance across various operational conditions. The inclusion of cavitation data is particularly useful for improving turbine longevity and reliability. Machine learning engineers can leverage this dataset to develop predictive models for turbine performance, potentially creating algorithms that can forecast efficiency and cavitation risks based on input parameters. Such models could be invaluable for real-time monitoring and predictive maintenance in hydropower plants. Additionally, the dataset could be used to train deep learning models for automated design optimization, potentially uncovering novel turbine configurations that maximize efficiency while minimizing cavitation.

#### 5.7.1. How are the challenges addressed?

The dataset of this relatively simple turbine geometry is **interpretable** with basic fluid mechanics and turbomachinery knowledge. It is available to train machine-learning-based algorithms, specifically mapping from different design parameters to the post-processed results. Regarding **generalizability**, since hydraulic





**Figure 10.** Pareto front of the dataset after 89 generations.

machines are specifically designed for specific applications, the design or machine type cannot be generalized in the sense of large-scale production. However, this simple case may serve as a training object to optimize more complex turbine types in further work. Furthermore, the prediction of the flow field is based on the numerical simulation of the coupled nonlinear, partial differential equations (Navier–Stokes equations), which makes the generalizability in the case of hydraulic turbines impossible. In terms of the **problem complexity**, the multi-objective problem can be changed to a single-objective optimization. Neglecting other objectives can significantly reduce the complexity of the dataset. However, with the current state-of-the-art turbine design problems, the complexity will substantially increase, primarily because of an increase in the number of geometry parameters. Owing to the flow interaction between several turbine components, the optimization of just one component is not sufficient. On the contrary, it is essential to design several turbine components simultaneously in order to optimize the entire machine.

### 5.8. Lessons learned

While translating the perceived challenges and proposed solution approaches of data publishing in the studied area to the seven concrete examples, the authors learned various lessons that can be summarized as follows, focusing on aspects not already discussed at length:

- The avoidance of domain-specific wording in published data and code can, as far as possible, help to make data and code reusable for a wider range of data scientists. Domain-specific scientific details can still be discussed in associated scientific papers catering to the specific domain.
- Researchers seem to routinely underestimate the time and effort it takes for others to use their data and code.
- It can be non-trivial to manually or automatically evaluate the quality of design results in mechanics and dynamics, different from some applications of learning methods to natural language and images.
- All authors have worked with their own published datasets described here and partly with those of their co-authors to use them within engineering design. Many have experienced that the composition of the data and potential data preprocessing or alternative data representations were usually very fundamental to the performance of the developed approach. For many engineering tasks, it is still very unclear which data representation is best suited, unlike for some other fields like computer

vision where successful ways to represent images for learning have been established. Sometimes, it is even unclear which data of an engineering system to record, to begin with.

Moreover, a key lesson learned was also that some more fundamental challenges remain, that often do not directly have something to do with the scientific subject matter at hand but with conditions of the surrounding scientific process:

- Pure publications of datasets in a data repository are nowadays citable with a digital object identifier. However, for key performance indicators used to assess researchers' careers, such as the *h*-index, data publications of this kind and citations of them are often irrelevant. This is in contrast to the undisputed importance of data for data-based methods and the skill and effort it requires to obtain and prepare a useful dataset. This discourages the publication of data for which one is not publishing a paper in parallel that uses this data and neglects the worth of published data in its own right. As papers usually propose functioning solutions, in effect, the publication of data for unsolved problems is discouraged.
- A review process for data publications and associated code is usually necessary to ensure quick accessibility even for other domain experts. The reviewer needs to actually execute the software and load and work with the data to ensure its functionality on an unprepared computer. As this is not usually required for publications of any sort, one should ask fellow scientists to try it and report back how long it took them to make it work. This practice was used here where applicable. Usually, in a result-driven environment, and with research funding rarely allocating earmarked funds for data management, there is, however, currently little incentive to spend the time to polish code and data publications at the end of projects.

## 6. Conclusions and outlook

Data science, machine learning, and artificial intelligence have in recent years expanded from first success stories in computer vision and natural language processing to progress in natural sciences and engineering. Regarding the latter, artificial intelligence also starts to have a lasting impact on engineering design and will transform how humanity conceives new machines, vehicles, robots, etc. As this article has outlined, engineering design brings about characteristic perspectives and challenges of which not every aspect and challenge is present for the currently most successful applications of machine learning. In particular, this includes the challenge of **evaluation**, as for design problems, there usually is no known best or true solution relative to which the quality of results can be gauged. Also, computationally evaluating the quality of a proposed design can be very expensive, for example by performing expert-driven and computationally highly demanding simulations, whereas data from experiments is usually not yet available in early design stages or very scarce. As a consequence, the generation of large datasets becomes prohibitively expensive. Even worse, these simulations often benefit from high-performance computing architectures different than the ones most fitting for the popular deep learning as engineering simulations typically rely more significantly on general-purpose processing. In the future, this may be alleviated for instance through simulators that, themselves, are based on machine-learning approaches. Nevertheless, when proposing and evaluating radically new designs, there remains the risk that such machine-learned simulators would be evaluated in regions where they have no prior knowledge, which may yield arbitrarily unreliable results or require currently unseen generalization. Still, as the examples provided in this paper show, today's methods and state of the art already allow to publish and use data and data-based approaches in the field of engineering design, giving first answers on how to overcome associated challenges. Clearly, researchers in engineering are only at the beginning of transforming engineering design toward a data-driven future. And one of the most interesting aspects will be whether this will lead to a future full of auto-designed, highly task-specific machinery that can be extremely efficient and accurate due to designs perfected for the specific task, not limited by the amount of time, money, and knowledge available in the form of human design engineers.

However, as publication habits change to include ever more data and code, also the question arises whether reviewing processes need to adapt. Does the reviewer of the future need to rate the interpretability of supplied data and ensure that future, competitive approaches can be readily and fairly evaluated? Will there be scientifically peer-reviewed publications of data, without accompanying papers and with only the necessary data-usage information and processing code, for which citations will count toward key performance indicators for scientific careers? Can reviewers who thoroughly inspect published data be anything more than a utopia as journal editors today face a flood of submitted papers and feel a relative dearth of willing reviewers? Many more questions of this kind will accompany the future of publication habits in mechanics, dynamics, and engineering design.

Whereas this work focused on an academic perspective on publishing and working with data, it is also interesting to explore how the picture might be different from an industrial or business perspective. For instance, in academia, it is widely accepted that open, easily accessible publications of knowledge and data are desirable. In contrast, companies' commercial interests may at first glance disincentivize a publication of the data their data-driven algorithms work with, even in applications where there might be a public interest to ensure that the data is not biased in certain ways. Future work may hence focus on the industrial perspective and its differences and similarities to the academic one.

**Data availability statement.** For all examples discussed in this paper, data and/or code are openly published. Please refer to the individual references to data and code repositories in [Section 5](#), where the data creation processes and accessibility are also outlined. Below is a list of the published data and code. Each repository includes metadata and/or usage instructions, accessible via a standard web browser:

- <https://doi.org/10.5281/zenodo.11352739>
- <https://doi.org/10.25625/UWF7RB>
- [https://github.com/ecker-lab/Learning\\_Vibrating\\_Plates](https://github.com/ecker-lab/Learning_Vibrating_Plates)
- <https://github.com/maneesh51/Benchmark-Tasks/>
- [https://github.com/MrAdityaBorse/FEM\\_Data\\_Impact\\_Simulation](https://github.com/MrAdityaBorse/FEM_Data_Impact_Simulation)
- <https://github.com/MKHoffmann/icpReconstructor>
- <https://github.com/raffimax/ConnectingGaitsinEnergeticallyConservativeLeggedSystems>
- <https://doi.org/10.18419/darus-4237>
- <https://doi.org/10.18419/darus-4152>
- <https://github.com/ihs-ustutt/dtOO>
- [https://github.com/ihs-ustutt/axial\\_turbine\\_database](https://github.com/ihs-ustutt/axial_turbine_database)

**Author contribution.** All authors contributed to conceptualization, methodology, data curation, data visualization, and writing the original draft. All authors approved the final submitted draft.

**Funding statement.** All authors acknowledge the support of the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG) within the Priority Program (DFG SPP) 2353 “Daring More Intelligence: Design Assistants in Mechanics and Dynamics.” P. Eberhard acknowledges support under DFG Grant No. 501890093 “Daring More Intelligence—Design Assistants in Mechanics and Dynamics (SPP 2353)—Coordination Proposal,” and H. Ebel is grateful for the past support from the same grant. P. Eberhard and B. Röder acknowledge support under DFG Grant No. 501840485. J. van Delden and T. Lüddecke acknowledge the support by the Deutsche Forschungsgemeinschaft (German Research Foundation), project number 501927736, within the DFG Priority Program 2353. They also acknowledge the computing time made available to them on the high-performance computers HLRN-IV at GWDG at the NHR Centers NHR@Göttingen. These centers are jointly supported by the German Federal Ministry of Education and Research and the German state governments participating in the NHR ([www.nhr-verein.de/unsere-partner](http://www.nhr-verein.de/unsere-partner)). A. Borse, R. Gulakala, and M. Stoffel gratefully acknowledge the financial support provided by Deutsche Forschungsgemeinschaft Priority Program SPP 2353 (DFG Grant No. STO 469/16-1). M. Stender and M. Yadav acknowledge the support by the Deutsche Forschungsgemeinschaft under grant 501847579. L. Schindler and K. M. de Payrebrune gratefully acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant number 501861263. M. Raff and C. D. Remy acknowledge the funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—501862165. M. Raff further acknowledges the support through the International Max Planck Research School for Intelligent Systems (IMPRS-IS). The contribution of R. Raj, T. Rentschler, A. Tismer, and S. Riedelbauch was funded by the German Research Foundation (DFG) (Grant/Award Number: 501932169) and the corresponding simulations were performed on the national supercomputer HPE Apollo Hawk cluster at the High-Performance Computing Center Stuttgart (HLRS) under the grant number parDiOOMeetsCanada/44196. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interest.** The authors declare none.

**Ethical standard.** The research meets all ethical guidelines, including adherence to the legal requirements of the study countries.

## References

- Basalti C, Fazekas-Paragh J, Forni M, van Gelder C, Hasani-Mavriqi I, Janik J, Kalová T, Kuchma I, Lindroos H, Lütcke H, Pinnick J, Raga N, Thorpe D and Wildgaard L (2024) Recommendations for data stewardship skills, training and curricula with implementation examples from European countries and universities. Zenodo. DOI: [10.5281/zenodo.10573892](https://doi.org/10.5281/zenodo.10573892).
- Basner M, Babisch W, Davis A, Brink M, Clark C, Janssen S and Stansfeld S (2014) Auditory and non-auditory effects of noise on health. *The Lancet* 383(9925), 1325–1332.
- Blickhan R (1989) The spring-mass model for running and hopping. *Journal of Biomechanics* 22(11–12), 1217–1227.
- Borse A, Gulakala R and Stoffel M (2023) Machine learning enhanced optimisation of crash box design for crashworthiness analysis. *PAMM* 23(4), e202300145.
- Borse A, Gulakala R and Stoffel M (2024a) Github—mradityaborse/fem\_data\_impact\_simulation. GitHub Repository. Available at [https://github.com/MrAdityaBorse/FEM\\_Data\\_Impact\\_Simulation](https://github.com/MrAdityaBorse/FEM_Data_Impact_Simulation) (accessed 4 July 2024).
- Borse A, Gulakala R and Stoffel M (2024b) Machine learning enhanced optimisation of crash box design for crashworthiness analysis. *Archives of Mechanics* 76(1–2), 61–92.
- Borse A, Gulakala R and Stoffel M (2025) Database of axial impact simulations of the crash box dataset | Papers with Code. Available at <https://paperswithcode.com/dataset/database-of-axial-impact-simulations-of-the> (accessed 15 February 2025).
- Code Ocean (2025) Explore capsules | Code Ocean. Available at <https://codeocean.com/explore?page=1> (accessed 4 February 2025).
- Destatis (2023) Accidents registered by the police by type of damage/ location. Available at <https://www.destatis.de/EN/Themes/Society-Environment/Traffic-Accidents/Tables/accidents-registered-police.html> (accessed March 2023).
- Deutsche Forschungsgemeinschaft (2023a) DFG, German Research Foundation—Good Research Practice. Available at <https://www.dfg.de/en/principles-dfg-funding/basics-and-principles-of-funding/good-scientific-practice> (accessed 1 December 2023).
- Deutsche Forschungsgemeinschaft (2023b) DFG, German Research Foundation—Handling of Research Data. Available at <https://www.dfg.de/en/principles-dfg-funding/basics-and-principles-of-funding/research-data>, (accessed 1 December 2023).
- Deutsche Forschungsgemeinschaft (2023c) DFG supports plans for legislation on research data. Press Release No. 10. Available at <https://www.dfg.de/en/service/press/press-releases/2023/press-release-no-10> (accessed 1 December 2023).
- Elrefaie M, Morar F, Dai A and Ahmed F (2024) DrivAerNet++: A large-scale multimodal car dataset with computational fluid dynamics simulations and deep learning benchmarks. *Preprint arXiv:2406.09624*.
- Flemisch B, Hermann S, Herschel M, Pflüger D, Pleiss J, Range J, Roy S, Takamoto M and Uekermann B (2024) Research data management in simulation science: Infrastructure, tools, and applications. *Datenbank-Spektrum* 24, 97–105.
- Fraas S, Tismer A and Riedelbauch S (2022) Sensitivity study of numerical and geometrical parameters for structural mechanical analyses in the automatic design process of hydraulic machines. *IOP Conference Series: Earth and Environmental Science* 1079, 012084.
- Gan Z, Yesilevskiy Y, Zaytsev P and Remy CD (2018) All common bipedal gaits emerge from a single passive model. *Journal of the Royal Society Interface* 15(146), 20180455.
- Gebru T, Morgenstern J, Vecchione B, Wortman Vaughan J, Wallach H, Daumé III H and Crawford K (2021) Datasheets for datasets. *Preprint arXiv:1803.09010v8*.
- Geyer H, Seyfarth A and Blickhan R (2006) Compliant leg behaviour explains basic dynamics of walking and running. *Proceedings of the Royal Society B: Biological Sciences* 273(1603), 2861–2867.
- Google (2023) Welcome to colabory—colabory. Available at <https://colab.research.google.com/> (accessed 1 December 2023).
- Herrmann L and Kollmannsberger S (2024) Deep learning in computational mechanics: A review. *Computational Mechanics* 74, 281–331.
- Heyrani Nobari A, Srivastava A, Gutfreund D and Ahmed F (2022) LINKS: a dataset of a hundred million planar linkage mechanisms for data-driven kinematic design. In *Proceedings of the 48th Design Automation Conference (DAC)*. St. Louis, MO, pp. V03AT03A013.
- Hoffmann MK., Mühlenhoff J, Ding Z, Sattel T and Flaßkamp K (2024a) Github—mkhoffmann/icpreconstructor. GitHub Repository. Available at <https://github.com/MKHoffmann/icpreconstructor> (accessed 4 July 2024).
- Hoffmann MK, Mühlenhoff J, Ding Z, Sattel T and Flaßkamp K (2024b) An iterative closest point algorithm for marker-free 3d shape registration of continuum robots. *Preprint arXiv:2405.15336*.
- Hugging Face (2023) Hugging face—the AI community building the future. Available at <https://huggingface.co/> (accessed 1 December 2023).
- Jaeger H (2001) The “echo state” approach to analysing and training recurrent neural networks—with an erratum note. Technical Report 148, German National Research Center for Information Technology GMD, Bonn, Germany.
- Karlsruhe Institute of Technology (2023) re3data.org. Available at <https://www.re3data.org/> (accessed 1 December 2023).
- Martens W, von Wagner U and Litak G (2013) Stationary response of nonlinear magneto-piezoelectric energy harvester systems under stochastic excitation. *The European Physical Journal Special Topics* 222, 1665–1673.
- Meta AI (2023) The latest in machine learning | Papers with Code. Available at <https://paperswithcode.com/> (accessed 1 December 2023).
- National Highway Traffic Safety Administration (NHTSA) (2016) Crash simulation vehicle models. Model: Honda Accord (2017). Available at <https://www.nhtsa.gov/crash-simulation-vehicle-models> (accessed 16 July 2024).

- Picard C, Schiffmann J and Ahmed F** (2023) Dated: Guidelines for creating synthetic datasets for engineering design applications. In *Proceedings of the International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*. Boston, MA, pp. V03AT03A015.
- Raff M and Remy CD** (2024) Periodic trajectories of a passive one-legged hopper. DaRUS, V2, <https://doi.org/10.18419/DARUS-4237>.
- Raff M, Rosa N and Remy CD** (2022a) Connecting gaits in energetically conservative legged systems. *IEEE Robotics and Automation Letters* 7(3), 8407–8414.
- Raff M, Rosa N and Remy CD** (2022b) Connecting gaits in energetically conservative legged systems. github repository. Available at <https://github.com/raffmax/ConnectingGaitsinEnergeticallyConservativeLeggedSystems>.
- Raff M, Rosa N and Remy CD** (2025) Monopedal gaits dataset | Papers with Code. Available at <https://paperswithcode.com/dataset/hopper-gaits> (accessed 15 February 2025).
- Raj R, Rentschler T and Tismer A** (2024a) Axial turbine database. Available at [https://github.com/ihs-ustutt/axial\\_turbine\\_database](https://github.com/ihs-ustutt/axial_turbine_database) (accessed 25 October 2024).
- Raj R, Rentschler T and Tismer A** (2025) Axial turbine dataset | Papers with Code. Available at <https://paperswithcode.com/dataset/axial-turbine-dataset> (accessed 15 February 2025).
- Raj R, Tismer A, Gaisser L and Riedelbauch S** (2024b) A deep learning approach to calculate elementary effects of Morris sensitivity analysis. *PAMM* 24(4), e202400104.
- Recht B, Roelofs R, Schmidt L and Shankar V** (2019) Do ImageNet classifiers generalize to ImageNet? In Chaudhuri K and Salakhutdinov R (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*. Long Beach, CA, pp. 5389–5400.
- Regenwetter L, Curry B and Ahmed F** (2021) BIKED: A dataset for computational bicycle design with machine learning benchmarks. *Preprint arXiv:2103.05844v3*.
- Regenwetter L, Nobari AH and Ahmed F** (2022) Deep generative models in engineering design: A review. *Journal of Mechanical Design* 144(7), 071704.
- Rentschler T, Berkemeier MB, Fraas S, Tismer A, Raj R, Peitz S and Riedelbauch S** (2024) Multi-criteria hydraulic turbine optimization using a genetic algorithm and trust-region postprocessing. *PAMM* 24(4), e202400126. <https://doi.org/10.1002/pamm.202400126>
- Ritchie H, Roser M and Rosado P** (2024) Renewable Energy. Available at <https://ourworldindata.org/renewable-energy> (accessed 24 October 2024).
- Röder B, Ebel H and Eberhard P** (2024a) Motion and Motor-Current Data of a Four-Bar Linkage. *DaRUS, VI*, <https://doi.org/10.18419/darus-4152>.
- Röder B, Hajipour S, Ebel H, Eberhard P and Bestle D** (2024b) Design of a four-bar mechanism from hand drawings of desired coupler trajectories and velocity profiles (submitted) *Mechanics Based Design of Structures and Machines*.
- Sabarathinam S and Thamilmaran K** (2016) Implementation of analog circuit and study of chaotic dynamics in a generalized duffing-type mems resonator. *Nonlinear Dynamics* 87, 2345–2356.
- Schindler L and de Payrebrune KM** (2024a) Data for “image-based backbone reconstruction for non-slender soft robots”. Zenodo, v1.0.0, <https://doi.org/10.5281/zenodo.11352739>.
- Schindler L and de Payrebrune KM** (2024b) Image-based backbone reconstruction for non-slender soft robots. *PAMM* 24(4), e202400130. <https://doi.org/10.1002/pamm.202400130>
- Schindler L and de Payrebrune KM** (2025) Image-based backbone reconstruction for non-slender soft robots | Papers with Code. Available at <https://paperswithcode.com/paper/image-based-backbone-reconstruction-for-non> (accessed 18 February 2025).
- Schultz J, van Delden J, Blech C, Langer SC and Lüddecke T** (2024) Deep learning based surrogate modeling for frequency response prediction of structural dynamic models. *Fortschritte der Akustik–DAGA 2024*.
- Strogatz SH** (2024) *Nonlinear Dynamics and Chaos*, 3 Edn. Boca Raton, FL: CRC Press.
- Sun C, Shrivastava A, Singh S and Gupta A** (2017) Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision (ICCV)*. Venice, pp. 843–852.
- Takamoto M, Praditia T, Leiteritz R, MacKinlay D, Alesiani F, Pflüger D and Niepert M** (2024) PDEBENCH: An extensive benchmark for scientific machine learning. *Preprint arXiv:2210.07182*.
- Tismer A** (2020a) Design tool object oriented. Available at <https://github.com/ihs-ustutt/dtOO> (accessed 24 October 2024).
- Tismer A** (2020b) *Entwicklung einer Softwareumgebung zur automatischen Auslegung von hydraulischen Maschinen mit dem Inselmodell*. PhD thesis, Universität Stuttgart.
- University of Stuttgart** (2023) Darus. Available at <https://darus.uni-stuttgart.de/> (accessed 1 December 2023).
- van Delden J, Schultz J, Blech C, Langer SC and Lüddecke T** (2023) Vibrating plates dataset | Papers with Code. Available at <https://paperswithcode.com/dataset/vibrating-plates> (accessed 15 February 2025).
- van Delden J, Schultz J, Blech C, Langer SC and Lüddecke T** (2024a) Github—ecker-lab/learning\_vibrating\_plates. GitHub Repository, Available at [https://github.com/ecker-lab/Learning\\_Vibrating\\_Plates](https://github.com/ecker-lab/Learning_Vibrating_Plates) (accessed 27 June 2024).
- van Delden J, Schultz J, Blech C, Langer SC and Lüddecke T** (2024b) Learning to predict structural vibrations. In *The thirty-eighth annual conference on neural information processing systems*. Available at <https://openreview.net/forum?id=i4jZ6fCDdy>, Vancouver.



- van Delden J, Schultz J, Blech C, Langer SC and Lüddecke T (2024c) Minimizing structural vibrations via guided diffusion design optimization. In *ICLR 2024 workshop on AI4DifferentialEquations in science*, Available at <https://openreview.net/forum?id=z4dcQodnoo>, Vienna.
- van Delden J, Schultz J, Blech C, Langer SC and Lüddecke T (2024d) Vibrating plates dataset. GRO.data, V2.1, <https://doi.org/10.25625/UWF7RB>.
- Wilkinson MD, Dumontier M, Aalbersberg IJ, Appleton G, Axton M, Baak A, Blomberg N, Boiten J-W, da Silva Santos LB, Bourne PE, Bouwman J, Brookes AJ, Clark T, Crosas M, Dillo I, Dumon O, Edmunds S, Evelo CT, Finkers R, Gonzalez-Beltran A, Gray AJ, Groth P, Goble C, Grethe JS, Heringa J, 't Hoen PA, Hooft R, Kuhn T, Kok R, Kok J, Lusher SJ, Martone ME, Mons A, Packer AL, Persson B, Rocca-Serra P, Roos M, van Schaik R, Sansone S-A, Schultes E, Sengstag T, Slater T, Strawn G, Swertz MA, Thompson M, van der Lei J, van Mulligen E, Velterop J, Waagmeester A, Wittenburg P, Wolstencroft K, Zhao J and Mons B (2016) The fair guiding principles for scientific data management and stewardship. *Scientific Data* 3, 160018.
- Wohlleben M, Röder B, Ebel H, Muth L, Sextro W and Eberhard P (2024) Hybrid modelling of multibody systems: Application of two discrepancy models for trajectory prediction. *PAMM* 24 (2), e202400027. <https://doi.org/10.1002/pamm.202400027>
- Xu Z, Escalera S, Pavão A, Richard M, Tu W-W, Yao Q, Zhao H and Guyon I (2022) Codabench: Flexible, easy-to-use, and reproducible meta-benchmark platform. *Patterns* 3(7), 100543.
- Yadav M, Chauhan S, Shrimali MD and Stender M (2024a) Predicting multi-parametric dynamics of an externally forced oscillator using reservoir computing and minimal data. *Nonlinear Dynamics* 113, 5977–5990.
- Yadav M, Chauhan S, Shrimali MD and Stender M (2024b) Predicting multi-parametric dynamics of an externally forced oscillator using reservoir computing and minimal data | Papers with Code. Available at <https://paperswithcode.com/paper/predicting-multi-parametric-dynamics-of-an>, accessed on February 15, 2025.
- Yadav M, Sinha S and Stender M (2025) Evolution beats random chance: Performance-dependent network evolution for enhanced computational capacity. *Physical Review E* 111, 014320.
- Yadav M and Stender M (2025) Benchmark-tasks: Duffing oscillator response analysis (DORA). Zenodo, <https://doi.org/10.5281/zenodo.14851014>.
- Yagoubi M, Danan D, Leyli-Abadi M, Brunet J-P, Mazari JA, Bonnet F, Gmati M, Farjallah A, Cinnella P, Gallinari P and Schoenauer M (2024) NeurIPS 2024 ML4CFD competition: Harnessing machine learning for computational fluid dynamics in airfoil design. *Preprint arXiv:2407.01641*.