

MDAKits: A Framework for FAIR-Compliant Molecular Simulation Analysis

Irfan Alibay^{‡†*}, Lily Wang^{‡†}, Fiona Naughton^{§†}, Ian Kenney^{¶†}, Jonathan Barnoud^{||}, Richard J Gowers[‡], Oliver Beckstein[¶]



Abstract—The reproducibility and transparency of scientific findings are widely recognized as crucial for promoting scientific progress. However, when it comes to scientific software, researchers face many barriers and few incentives to ensure that their software is open to the community, thoroughly tested, and easily accessible. To address this issue, the MDAKit framework has been developed, which simplifies the process of creating toolkits for the MDAnalysis simulation analysis package (<https://www.mdanalysis.org/>) that follow the basic principles of FAIR (findability, accessibility, interoperability, and reusability). The MDAKit framework provides a cookiecutter template, best practices documentation, and a continually validated registry. Registered kits are continually tested against the latest release and development version of the MDAnalysis library and their code health is indicated with badges. Users can browse the registry frontend (<https://mdakits.mdanalysis.org/>) to find new packages, learn about associated publications, and assess the package health in order to make informed decisions about using a MDAKit in their own research. The criteria for registering an MDAKit (open source, version control, documentation, tests) are similar to the criteria required for publishing a paper in a software journal, so we encourage and support publication in, e.g., the Journal of Open Source Software, creating further academic incentive for researchers to publish code. Through the MDAKits framework, we aim to foster the creation of a diverse ecosystem of sustainable community-driven downstream tools for MDAnalysis and hope to provide a blueprint for a model for growing communities around other scientific packages.

Index Terms—Molecular Dynamics Simulations, Python, MDAnalysis, ecosystem

Introduction

Software has become increasingly essential to research. In many areas, it underlies fundamental tasks such as generating, processing, analyzing, storing, visualizing, and communicating the key results and insights ultimately published.

Scientific code frequently fails to meet FAIR tenets, impeding scientific progress

Despite the importance of software, it is typically not central to the publication peer review process in many scientific fields.

† These authors contributed equally.

* Corresponding author: ialibay@mdanalysis.org

‡ Open Molecular Software Foundation, Irvine, CA, USA

§ Cardiovascular Research Institute, University of California, San Francisco, San Francisco, CA, USA

¶ Arizona State University, Tempe, AZ, USA

|| Centro Singular de Investigación en Tecnoloxías Intelixentes, Santiago de Compostela, Spain

Copyright © 2023 Irfan Alibay et al. This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Consequently, scientific code frequently fails to meet the basic tenets of FAIR: findability, accessibility, interoperability, and reusability [1], [2]. With the publication of “The FAIR Guiding Principles for scientific data management and stewardship” in 2016 and the follow-up “FAIR Principles for Research Software (FAIR4RS Principles)” in 2022, it has become increasingly acknowledged that abiding by the principles of FAIR is crucial to promoting robust, reproducible, and efficient scientific discovery and innovation [1], [2]. We believe that extending FAIR principles to include open-source software not only significantly advances that goal, but furthermore is necessary for transparent research. Open sharing of code brings a number of substantial benefits to the scientific community. For example, scientists can accurately replicate a given methodology or re-use previous code, reducing duplication of effort and reducing the risk of implementation errors. Indeed, the molecular simulation community in particular has made a concerted effort over recent years to encourage the open sharing of scientific codes [3]. For example, as of July 2022, over 4700 GitHub repositories containing Python code that makes use of MDAnalysis [4], [5] have been made publicly available¹.

However, simply sharing code is not sufficient to fulfill FAIR guidelines. In fact, making software FAIR compliant requires significant investment and often expert knowledge on the part of the developers, especially if the code was written specifically for a particular research project. For example, the Python ecosystem is so dynamic that it is common for research code to rapidly become obsolete or unusable if a new version of a key library is released. To fulfill the Reusability tenet of FAIR alone, code should include documentation, version control, and dependency management. Ideally, it would also include unit tests, examples, and packaging. Even when code is released in reference to a publication, it often falls short of ideal FAIR standards. A short survey of publications in Scopus [6] and the Journal of Open Source Software [7] over 2017–2021 identified that out of a total 720 papers citing MDAnalysis [4], [5], only 43 linked to code available on a version control platform such as GitHub, GitLab, or Bitbucket. Of these, only 18 met the requirements of best practices: they implemented unit tests, comprehensive documentation, and some means of installation.

Two major factors contribute to the lack of open-source FAIR compliant code. Firstly, code is typically written by scientists with no formal training or support in programming, for whom implementing FAIR principles can pose an intimidating and tedious barrier. Secondly, despite the substantial investment of effort and time required to implement best practices, publishing FAIR

software is not typically appreciated with academic recognition or reward. Fostering a culture of open-source FAIR software requires addressing both.

Centralized open-source packages such as MDAnalysis offer a limited solution

One solution is to consolidate scientific code around a small number of large, central packages. MDAnalysis [4], [5] is a widely-used open-source Python library for molecular simulation data. With over 18 years of development by more than 180 developers, MDAnalysis has refined its code base to offer a mature, robust, flexible API that offers a range of high-performance tools to extract, manipulate, and analyze data from the majority of common simulation formats. MDAnalysis tools have been used for a variety of scientific applications ranging from exploring protein-ligand interactions [8], [9], [10], to understanding lipid behavior [11], [12], to assessing the behavior of novel materials [13], [14].

Initially, MDAnalysis focused on growing the developer and user community by encouraging users to contribute their code directly to the MDAnalysis library. Notable examples of this include the waterdynamics [15] and ENCORE [16] analysis modules. This approach of encouraging code to be contributed to a central package has also been successfully taken by packages such as cpptraj [17] and the GROMACS tools [18]. It has a number of key advantages for users and the original developers:

- MDAnalysis can ensure that the code follows best practices (including documentation and tests).
- Code is promoted and made freely accessible to all MDAnalysis users.
- Maintenance, support, and potential updates are performed by the experienced MDAnalysis developer team, ensuring that the contributed code remains functional even while the other parts of the library change. The original developers can thus focus on other work.

However, the many costs of this approach can, under some conditions, result in unsustainable, untenable disadvantages:

- Ensuring that the code follows best practices often requires long review periods and strict code-style adherence, thus slowing down the availability of the new code in a released version of the package.
- The necessity of keeping the API stable between major releases precludes quick releases of breaking changes. In general, a mature package such as MDAnalysis has a slow release cycle, so new features and bug fixes can take months to become available in new releases.
- As MDAnalysis implicitly agrees to maintain any code that we release, a certain level of understanding and expertise is required from the maintainers. If the core developer team lacks expertise in a specific discipline or subdiscipline, adding new code in these areas introduces a substantial maintenance burden should the original code contributors not be available to help with maintenance. Consequently, it is impractical to include recently released or cutting-edge techniques in the core library.
- Introducing new package dependencies incurs software stack maintenance costs for many users who may not require this additional code.

- Code contributors lose complete control of their code.

The many disadvantages listed above can severely limit the usefulness of centralizing code around one monolithic package. Indeed, encountering these issues when attempting to expand the core MDAnalysis library attests that this approach is not the most suited for the MDAnalysis community.

Implementing an ecosystem of downstream packages for more sustainable progress

We believe that a sustainable alternative solution is for communities such as MDAnalysis to encourage, educate, and foster researchers in their efforts towards developing individual software. We have developed a program of structured technical assistance to help researchers implement best practices and publish their code within a growing ecosystem of toolkits that we have called MDAKITS (MDAnalysis Toolkits). We have also developed a platform called the "MDAKit registry" (<https://mdakits.mdanalysis.org/mdakits.html>) where packages that meet certain standards are advertised to the community. The MDAKit ecosystem builds on the success of other community packages such as PLUMED's PLUMED-NEST [19], AiiDA's plugin registry [20], or the napari-hub [21] of plugins for the napari image viewer [22], all of which list available tools that are known to work in their respective user communities.

Our technical assistance begins with cookiecutter templates and example repositories. Here we model best practices, promote the use of helpful tools, e.g., for checking code coverage, and reduce the work required to set up processes such as continuous integration, versioned documentation, packaging and deployment. Developers can also reach out to the MDAnalysis community for feedback, technical assistance, or even make connections with new co-developers and potential users. Decoupled from MDAnalysis's release cycle, developers are able to introduce new changes as required, keeping complete control over their code-base. Joining an MDAnalysis registry allows for frequent and streamlined communication between MDAnalysis and downstream developers, allowing developers to be efficiently forewarned about potential breaking changes.

Although establishing such an ecosystem of MDAnalysis-supported packages requires substantial investment from MDAnalysis developers, this approach is nonetheless likely to be far more sustainable than centralizing around a super-package. Offering technical assistance to individual developers in implementing best practices constitutes a large part of the effort; however, this level has thus far proven much lower than the effort associated with adding additional functionality to the core MDAnalysis library, and we believe that it will continue to remain so. Furthermore, as the ecosystem grows, we hope that an increasing portion of the community will participate in taking care of the packages and registry, and that the culture of following best practices and publishing code will gain momentum in itself.

In part, we hope that this momentum will be driven by users and user expectations. Users of the MDAnalysis ecosystem gain huge benefit from the MDAKit registry. They are able to see new software as it gets added, rather than having to comb through literature or rely on developers advertising the code themselves. They are also able to easily verify the current development status of a package and whether it is being actively maintained and passing tests with both released and in-development versions of MDAnalysis. In the future, the registry could contain information about the health of a given codebase, such as its activity

1. Based on a search for repositories containing `import MDAnalysis`.

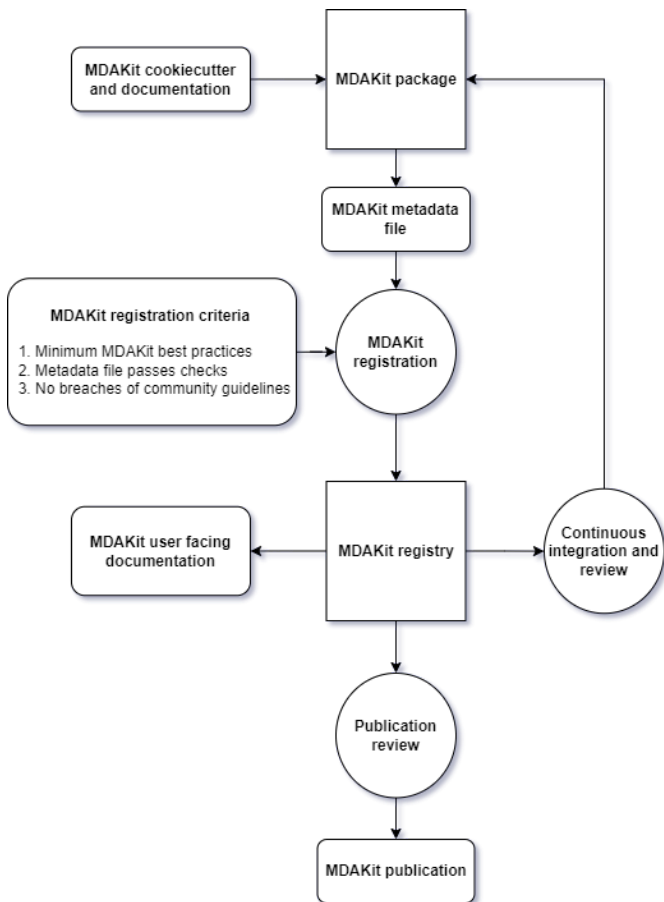


Fig. 1: Workflow diagram of the MDAKit framework. Starting from the creation of an MDAKit package, with the help of documentation and the MDAKit cookiecutter, the package then goes through the process of being added to the MDAKit registry, undergoing continuous validation and review and eventually reaching the stage of publication.

status and if it optimally leverages high performance MDAnalysis components (e.g. highly optimised PBC-aware distance routines). Packages on the registry also come with easy-to-find instructions on how to easily install and run a given package, significantly lowering the technical barrier to use and experimentation. As the maintenance remains the burden of the package owners, unfortunately the risk remains that packages on the registry may eventually become out-of-date, which is indeed one of the major disadvantages of this approach. However, the registry significantly increases the likelihood that packages will reach users who will become sufficiently motivated to contribute or take over their maintenance and development.

In the rest of this document we outline our expectations for MDAKits in terms of best practices and how we implement their registration and continuous validation.

The MDAKit framework

The MDAKit framework (Fig. 1) is designed to be a complete workflow to help and incentivize developers to go from the initial stages of package development all the way through to the long-term maintenance of a mature codebase, while adhering to best practices.

Main goals

As such, the main goals of the proposed MDAKit framework are:

- 1) To help as many packages as possible implement best practices and develop user communities.
- 2) To ensure that members of the MDAnalysis community can easily identify new packages of interest and know to what extent they are suitable for production use.
- 3) To improve contacts between MDAnalysis core library developers and those developing packages using MDAnalysis.
- 4) To encourage participation from the community at all steps of the process.

We wish to state three main points that the framework is *not* designed for:

- 1) The MDAKit framework is not intended to restrict the packages which can participate. It is our view that all packages at any stage of their development are of value to the community. As such, we aim for framework components to be as non-blocking as possible.
- 2) It is not the intention of any parts of this framework to take control or ownership of the packages that participate within it. The original code developers retain full ownership, control, and responsibility for their packages and may optionally participate in any part of this framework.
- 3) We also do not want to block future contributions to the core library. If new code in MDAKits prove particularly popular, and the MDAKit developers are amenable to contributing these back into the core library, the MDAnalysis team will work with them to integrate additional functionality into MDAnalysis itself.

Overview of the framework

The MDAKit framework (Fig. 1) is a multi-step process. In the first step of the MDAKit framework, developers create an initial package which is intended to achieve a set purpose of their choice. To help with this process, MDAnalysis provides a cookiecutter template specifically for MDAKits [23], alongside documentation on best practices and how to optimally use the MDAnalysis API. An overview of what we consider to be best practices for the contents of MDAKit packages is included in Section [Defining MDAKits: best practice package features](#). We note that at this point MDAKits are not expected to fully adhere to best practices, but should at least meet the minimum requirements defined in Section [Defining MDAKits: best practice package features](#) before moving to the next step along this process.

Once a package is suitably developed, code owners are encouraged to add the details of their code to the “MDAKit registry” which advertises their package to the MDAnalysis community and offers continual validation and review tools to help with package maintenance. Section [The MDAKit registry](#) contains more information about the MDAKit registry, including the registration process (Section [Registering MDAKits](#)). Briefly, the registration process involves submitting a metadata file to the registry that contains essential information about the MDAKit, such as where the source code is provided, who the code authors are, and how to install the MDAKit. The contents of this metadata file is reviewed both by automatic code checks and the MDAnalysis developer team before being added to the registry. We want to highlight that this process does not include checks on scientific validity or code health. In fact, none of the processes in this framework account for the scientific validity of the MDAKits. While members of the

community are free to offer help, scientific or technical validity is beyond the scope of what is feasible with the MDAnalysis registry.

Upon registration, the MDAKit is automatically advertised to the MDAnalysis community (see Section [Advertising MDAKits](#)). In the first instance this amounts to a set of auto-generated pages that expose the details in the metadata file provided in the registration step. Additional tags and badges are also included that reflect the current status and health of the package. Examples include:

- whether or not it is compatible with the latest versions of MDAnalysis
- what percentage of the codebase is covered by unit tests
- what type or extent of documentation is provided
- what Python versions are currently supported.

This status information is provided as part of checks done during the continual validation and review steps (see Sections [Continual validation](#) and [Continual review](#)) of the framework. These steps involve a mix of regularly scheduled automatic (e.g., linters and unit test execution) checks and more infrequent manual (e.g., code reviews) processes. It is our intention that code health analysis will help developers maintain and improve their codes, as well as suitably warn potential users about issues they may encounter when using a given codebase.

Where possible, the framework encourages a code review process to be carried out by members of the MDAnalysis community. The aim here is to work with developers in identifying potential areas of improvements for both MDAKits and the core MDAnalysis library (see Sections [Continual review](#) and [Feeding back into the MDAnalysis library](#)). We aim to tie this process closely to the review processes of journals such as the Journal of Open Source Software [7], which would help lower the barrier towards and encourage an eventual publication (Section [Towards publication](#)).

Defining MDAKits: best practice package features

Here we list requirements that we believe MDAKits should strive to fulfill in order to meet best practices in Python package usability and maintenance. To help with implementing these, a cookiecutter is provided which offers a template for potential MDAKits to follow [23]. We want to emphasize again that the aim of the MDAKit project is to encourage best practices whilst also minimizing barriers to sharing code where possible. Therefore, only a minimal set of requirements listed here as *required* are necessary for MDAKits to be included in the MDAKit registry. Similarly, we do not mean to enforce the label of MDAKit on any package; the process is fully optional and the code owners may choose whether to associate themselves with it.

All MDAKits must implement the features on the list of **required features** in order to become registered:

- Code in the package *uses MDAnalysis* ([Code using MDAnalysis \(required\)](#)).
- Open source code is published under an *OSI approved license* ([Open source code under an OSI approved license \(required\)](#)).
- Code is *versioned* and provided in an *accessible version-controlled repository* ([Versioning and provision under an accessible version-controlled repository \(required\)](#)).
- Code *authors and maintainers are clearly designated* ([Designated code authors and maintainers \(required\)](#)).

- *Documentation* is provided ([Documentation \(required\)](#)).
- *Tests and continuous integration* are present ([Tests and continuous integration \(required\)](#)).

The following are **highly recommended features**:

- Code is *installable as a standard package* ([Packaging](#)).
- Information on *bug reporting, user discussions, and community guidelines* is made available ([Bug reporting, user discussions, and community guidelines](#)).

Code using MDAnalysis (required): This is the base requirement of all MDAKits. The intent of the MDAKit framework is to support packages existing downstream from the MDAnalysis core library. MDAKits should therefore contain code using MDAnalysis components which are intended by the package authors to address the MDAKit's given purpose.

Open source code under an OSI approved license (required): The core aim of MDAKits is to encourage the open sharing of codes to potential users within the MDAnalysis community and beyond. To achieve this, we require that codes under this framework be released as open source. Here we define open source as being under an Open Source Initiative (OSI) approved license [24].

As of writing, the MDAnalysis library is currently licensed under GPLv2+ [25]. Due to limitations with this license type, we cannot currently recommend other licenses than GPLv2+ for codes importing MDAnalysis. However, we hope to relicense to a less restrictive license. In this event, MDAKits will be able to adopt a wider range of OSI approved licenses.

Versioning and provision under an accessible version-controlled repository (required): The ability to clearly identify changes in a codebase is crucial to enabling reproducible science. By referencing a specific release version, it is possible to trace back any bug fixes or major changes which could lead to a difference in results obtained with a later version of the same codebase. Whilst we encourage the use of Semantic Versioning ("semver") [26], any PEP440 [27] compliant versioning specification would be suitable for MDAKits.

Beyond versioning releases, it is also crucial to be able to develop code in a sustainable and collaborative manner. The most popular way of achieving this is through the use of version control through Git [28]. We require all MDAKits to be held in a publicly facing version controlled repository such as GitHub [29], GitLab [30], or Bitbucket [31].

Designated code authors and maintainers (required): In order for users to be able to contact the code owners and maintainers, all MDAKits should clearly list their authors and a means of contacting the persons responsible for maintaining the codebase. To incentivize and recognize contributors throughout the life of a project, we recommend the use of a version controlled "authors" file which lists the authors to a codebase over time.

Documentation (required): Describing what a given code does and how to use it is a key component of open sharing. Ideally a package would include a complete description of the entire codebase, including both API documentation and some kind of user guide with worked examples on how the code could be used in certain scenarios. Whilst this is recommended as best practices for an MDAKit, we recognize that this is not always feasible, especially in the early stages of development. Therefore, the minimum requirement for MDAKits is to have a readme file which details the key aspects of the MDAKit, such as what it is intended to do, how to install it, and a basic usage example.

For best practices, we strongly recommend using docstrings (see PEP 257 [32]) to document code components and using a tool such as ReadTheDocs [33] to build, version and host documentation in a user-friendly manner. We also recommend using duecredit [34] to provide the correct attributions to a given method if it has been published previously.

Tests and continuous integration (required): Testing is a critical component to ensure that code behaves as intended. Not only does it prevent erroneous coding, but it also assures users that the code they rely on is working as intended. We require at least a single regression test for major functionality to qualify for the registry (i.e. if a toolkit implements a new analysis method, at least one test that checks to see if the analysis code yields the expected value on provided data; regression tests can often double as example documentation).

Ideally one should do full unit testing of the contents of a code, ensuring that not only a specific outcome is reached, but also that each smaller component works. As part of best practices, we highly recommend implementing tests using a framework such as pytest [35] for executing tests and codecov [36] to capture which lines are covered by the tests. We strongly encourage that a minimum of at least 80% of the code lines be covered by tests.

To ensure that tests are run regularly, the recommended best practice is to implement a continuous integration pipeline that performs the tests every time new code is introduced. We encourage the use of free pipelines such as GitHub Actions [37] to implement continuous integration.

Packaging: Providing a standard means of installing code as a package is important to ensure that other code can correctly link to (i.e., `import` in the case of Python) and use its contents. Whilst it can be easy to expect users to simply read a Python script, look at its required dependencies, and install them manually, this can quickly become unreasonable should the code grow beyond a single file. Additionally, the lack of clearly defined versions, including the intended Python versions, can lead to inoperable code.

As best practices we heavily encourage the use of `setuptools` [38] or an alternative such as `poetry` [39] for package installation. We also encourage that packages be available on common package repositories such as PyPi [40] and `conda-forge` [41]. The use of such repositories and their respective package managers can significantly lower the barrier to installing a package, enabling new users to rapidly get started using it.

Bug reporting, user discussions, and community guidelines: To help maintain and grow the project, it is important to specify where users can raise any issues they might have about the project or simply ask questions about its operation. To achieve this, we recommend at the very least adding documentation that points users to an issue tracker.

Key to successfully building a user community is ensuring that there are proper guidelines in place for how users will interact with a project [42]. As best practices we recommend making a code of conduct available that defines how users should interact with developers and each other within a project. It is also advised to provide information on how users can contribute to the project as part of its documentation.

The MDAKit registry

As defined in Section [The MDAKit framework](#), once MDAKits are created, we encourage that they be added to the MDAKit

registry. The registry not only provides a platform to advertise MDAKits to the MDAAnalysis user community at the web page <https://mdakits.mdanalysis.org/>, but also offers tools and workflows to help packages improve and continue to be maintained. Here we describe the various processes that occur within the registry. We note that we expect the exact details of how these processes are implemented to evolve over time based on feedback from MDAKit developers and other members of the MDAAnalysis community.

MDAKit registry contents

The main aim of the registry is to hold information about MDAKits. The contents of the registry therefore center around a list of packages and the metadata associated with each MDAKit. This metadata has the form of two files: one containing user-provided information on the package contents (see Section [Registering MDAKits](#)), and the other a set of mostly auto-generated details indicating the code health of the package (see Section [Advertising MDAKits](#)).

This metadata is used for two purposes: continuous integration testing and documentation. Continuous testing, helper methods and workflows are used to regularly install MDAKits and run their test suite (if available) to check if they still work as intended. Should the tests fail, package maintainers are automatically contacted and failure information is recorded in the code health metadata to inform users. For the registry documentation, the metadata is used to provide user-facing information about the various MDAKits in the registry, their contents, how to install them, and their current status as highlighted by continuous integration tests. The registry also includes further information such as user guides and tutorials on the MDAKit framework, helping developers to implement their own MDAKits.

Registering MDAKits

A key feature of the MDAKit framework is the process of adding MDAKits to the registry. As previously defined, our intent is to offer a low barrier to entry and have packages be registered early in their development cycles. This allows developers to benefit from the MDAKit registry validation and review processes early on, hopefully lowering the barrier to further improvements and encouraging early user interactions and feedback.

From an MDAKit developer standpoint, the registration process involves opening a pull request against the MDAKit registry that adds a new YAML file with metadata about the project. The metadata, as detailed in Fig. 2, contains information such as the MDAKit description, source code location, installation instructions, how to run tests, and where to find documentation. Complete details about the metadata file specification are provided in the MDAKit registry documentation.

After a pull request is opened, the MDAAnalysis developers review the contents of the submission based on the following criteria:

- 1) If the required features for MDAKits are met (Section [Defining MDAKits: best practice package features](#)), that is:
 1. Does the MDAKit contain code using MDAAnalysis?
 2. Is the MDAKit license appropriate?
 3. Is the MDAKit code offered through a suitable version-controlled platform?

```

1  ## Required entries
2  project_name: propkatraj
3  authors:
4    - https://github.com/Becksteinlab/propkatraj/blob/main/AUTHORS
5  maintainers:
6    - ianmkenney
7    - IALibay
8    - orbeckst
9  description:
10   pKa estimates for proteins using an ensemble approach
11 keywords:
12   - pKa
13   - protein
14 license: GPL-2.0-or-later
15 project_home: https://github.com/Becksteinlab/propkatraj/
16 documentation_home: https://becksteinlab.github.io/propkatraj/
17 documentation_type: UserGuide + API
18
19 ## Optional entries
20 install:
21   - pip install propkatraj
22 src_install:
23   - pip install git+https://github.com/Becksteinlab/propkatraj@main
24 python_requires: ">=3.8"
25 mdanalysis_requires: ">=2.0.0"
26 run_tests:
27   - pytest --pyargs propkatraj.tests
28 test_dependencies:
29   - mamba install pytest MDAnalysisTests
30 project_org: Becksteinlab
31 development_status: Mature
32 publications:
33   - https://zenodo.org/record/7647010
34   - https://doi.org/10.1021/ct200133y
35   - https://doi.org/10.1085/jgp.201411219

```

Fig. 2: YAML metadata file for an MDAKit entry of the propkatraj package, stored as `mdakits/propkatraj/metadata.yaml` in the registry repository.

4. Are the MDAKit authors and maintainers clearly designated in the metadata file?
 5. Is there at least minimal documentation in place detailing the MDAKit and its functionality?
 6. Are there at least minimal regression tests available within the MDAKit code?
- 2) If the metadata file passes linting and integration checks
 - 3) That there are no potential breaches of community guidelines

Once the criteria are fulfilled the metadata is merged and the MDAKit is considered registered. Updates to the MDAKit metadata can be carried out at any time after registration by opening pull requests to change the metadata file contents.

Advertising MDAKits

Registered MDAKits are automatically added to the registry's public facing documentation at <https://mdakits.mdanalysis.org/mdakits.html>. This involves an indexable list of entries for all registered MDAKits. Each entry displays available information from the provided metadata, e.g., what the MDAKit does, any relevant keywords, how to obtain the source code, how to install the package, and where to find relevant documentation. Alongside this information is also a set of badges which describe the current health of the codebase, allowing users to rapidly identify which packages are currently active, and their level of code maturity. This includes information such as which MDAnalysis library versions

the package is compatible with. We further plan to add more information, such as how much test coverage the package has, and what type of MDAnalysis API extensions are provided (e.g., using base classes such as `AnalysisBase` or `ReaderBase`).

Information about MDAKits is continually updated, either through automatic checks or manual additions provided by package owners updating the metadata files. We aim for the MDAKit registry to be immutable (aside from special cases covered by Section [Raising issues, concerns, and paths to registry removal](#)). Therefore, should an MDAKit stop being maintained, it will not be removed from the index but instead labeled as abandoned.

Continual validation

The MDAKit registry implements workflows to validate the code health of registered packages. This mostly centers around a test matrix that regularly runs to check if the latest MDAKit release can be installed and if unit tests pass with both the latest release of MDAnalysis and the development version. Should tests fail regularly, an issue is automatically raised on the MDAKit registry issue tracker contacting the package maintainers and letting them know of the failure. The auto-generated code health metadata for the MDAKit is also updated to reflect whether or not the tests are currently failing or passing.

In the future we hope to expand these tests to include more historical releases of the MDAKits and the MDAnalysis library, checks for different architectures (non-x86), and operating systems. We may also expand the checks to consider the cross-compatibility of MDAKits with each other, offering insights on which packages can be safely used together.

Continual review

To help package growth and improvements, it is our goal for the registry to become a platform that allows members of the MDAnalysis community to offer feedback on MDAKits over the lifetime of their inclusion on the registry. Unfortunately, as MDAnalysis developers can only devote limited time towards the registry, offering regularly scheduled comprehensive reviews of packages is too large an undertaking to be practical.

Instead, we aim to use a system of badges and achievements to push packages towards gradual improvements. For example, we may offer an achievement that encourages MDAKits to use high performance PBC-aware distance routines defined in `MDAnalysis.lib.distances` instead of relying on NumPy's `linalg` method to find the distance between two points. Once MDAKit owners believe that they have suitably updated their code to fulfill the relevant badge criteria, they can open a pull request highlighting these changes and have developers review these smaller, more focused updates.

MDAKit users are also encouraged to provide feedback, request improvements, and report bug fixes. However, this should happen outside the scope of the registry; instead, we ask users to use the MDAKit's own issue tracker for these.

Feeding back into the MDAnalysis library

The existence of the MDAKits framework does not preclude the addition of new codes and methods to the core MDAnalysis library. The MDAKit registry, and especially the ongoing review process, provides a platform for MDAnalysis and MDAKit developers to interact and work together to identify common goals and areas of improvements for both upstream and downstream packages. In particular, MDAnalysis developers will work with

MDAKit developers to see if any popular MDAKit methods, components or other means to improve core method performance and lower the barrier to downstream package development can and should be implemented back into the core MDAnalysis library.

Towards publication

We have laid out a number of best practices here that we encourage MDAKits to fulfill. These essentially amount to the majority of the contribution criteria for submissions to software-focused journals such as the Journal Open Source Software (JOSS) [7]. In order to incentivize developers, we heavily encourage MDAKits to consider submission to a journal such as JOSS once they meet the required levels of best practices. To aid in this process, the MDAnalysis developers will in the first instance work with journal editors at JOSS to create a streamlined process to submit MDAKits as JOSS entries [43]. The details of this process are still under development.

Raising issues, concerns, and paths to registry removal

If community members (users, developers or otherwise) have concerns about an MDAKit, we primarily encourage them to raise issues on the MDAKit's own issue tracker. However, in situations where the MDAKit maintainers cannot respond, or if the concern relates to code of conduct breaches, MDAnalysis developers may step in. If an MDAKit has systemic issues with its correctness, the MDAKit may be given special annotations warning users about the issues before using the code. We generally view the MDAKit registry as a permanent record, and avoid removing packages after registration even if they become fully obsolete. However, we reserve the right to remove packages at our discretion in specific cases, notably code of conduct breaches and violation of the GitHub terms of service [44].

Long term registry maintenance and support

As with most MDAnalysis projects, long-term support for the MDAKit framework and especially the registry is expected to be carried out by contributors from the MDAnalysis community. Members of the MDAnalysis core development team lead the maintenance of the registry and are also responsible for passing judgment on serious events such as code of conduct breaches. In the long term, we hope that any gains in popularity of the MDAKits framework are accompanied by an increase in community involvement in reviews and other maintenance tasks.

Examples of MDAKits

The web frontend of the registry (Fig. 3) provides a searchable database of packages. At the moment, seven MDAKits are registered that already showcase the breadth of specialized tools for the analysis of biomolecular simulations. For example, *mdacl* provides a commandline interface to analysis tools in MDAnalysis itself. *openmm-mdanalysis-reporter* enhances the interoperability with the popular OpenMM MD engine [45]. *hole2-mdakit* interfaces with the legacy HOLE2 program for the analysis of pores and tunnels in proteins such as ion channels [46], [47]. The *lipyds* package provides a suite of tools for the analysis of biological membranes in simulations [11]. *ProLIF* quantitatively analyzes the interactions between small molecules such as drugs and biomolecules (protein, nucleic acids) [10].

MDAKit	Keywords	Authors	CI badges
hole2-mdakit	pores, ion channels, transporters, HOLE	hole2-mdakit authors	latest: passed develop: passed
lipyds	lipids, membranes	lilyminium	latest: passed develop: unavailable
mdachecker	structure analysis, conformation checks	IAlibay	latest: passed develop: passed
mdacl	command line interface, molecular-dynamics	mdacl authors	latest: passed develop: passed
openmm-mdanalysis-reporter	OpenMM, Reporters	openmm-mdanalysis-reporter authors	latest: passed develop: passed
ProLIF	drug-design, cheminformatics, molecular-dynamics	Cédric Bouysset	latest: passed develop: passed
propkatraj	pKa, protein	propkatraj authors	latest: passed develop: passed

Showing 1 to 7 of 7 entries Previous 1 Next

Fig. 3: Web front end of the searchable MDAKit registry with registered MDAKits. Badges indicate code health based on continuous validation against the latest release and development version of the MDAnalysis library.

Conclusions

We introduce the MDAnalysis MDAKits framework for scientific software packages. This framework is designed to assist and incentivize the creation of FAIR-compliant (findable, accessible, interoperable, and reusable) packages that use and extend MDAnalysis. We describe the current state of scientific code, which is typically published either in independent repositories of varying quality, or as additions to a large, monolithic package. We summarize the limitations of each approach that result in code that falls short of FAIR principles, or may end up impractical to sustain as a long-term strategy. We propose the MDAKits framework as an alternative solution to support developers in creating new packages, guiding them through the process of achieving best practices and FAIR compliance.

In Section [The MDAKit framework](#) we lay out the aims and structure of an MDAKit, summarizing the minimal and optimal requirements that we think necessary to build sustainable, reusable software. These include publishing code under a suitable open-source license, the use of version control, comprehensive documentation, thorough unit tests, and packaging the software following modern best practices. In Section [Defining MDAKits: best practice package features](#) we outline our vision and implementation of the MDAKit registry, a public facing repository that promotes MDAKits to the MDAnalysis community. The MDAKit registry offers regular checks and reviews in order to help improve and maintain the listed MDAKits. We describe a structured workflow that begins from the initial registration of MDAKits and reaches as far as eventual publication in software-focused journals such as JOSS.

This document is just the first step and broad guide to our vision of developing a rich, diverse software ecosystem, and we are still in the early stages of implementing MDAKits. While we expect that we may need to revisit and refine our strategy to best serve the needs of the community, we believe that the fundamental framework outlined here will bring great benefit to the software written and used by scientists, and thereby empower transparent

and reproducible research.

Acknowledgments

We gratefully acknowledge the 184 developers and countless community members who have contributed to the MDAnalysis project since its inception and NumFOCUS for its support as our fiscal sponsor.

The work on the MDAKits project and this publication have been made possible in part by CZI grant DAF2021-237663 and grant DOI <https://doi.org/10.3792/1426590wiobus> from the Chan Zuckerberg Initiative DAF, an advised fund of Silicon Valley Community Foundation (funder DOI 10.13039/100014989).

Jonathan Barnoud has received financial support from the Agencia Estatal de Investigación (Spain) (REFERENCIA DEL PROYECTO / AEI / CÓDIGO AXUDA), the Xunta de Galicia - Consellería de Cultura, Educación e Universidade (Centro de investigación de Galicia accreditation 2019-2022 ED431G-2019/04 and Reference Competitive Group accreditation 2021-2024, CÓDIGO AXUDA) and the European Union (European Regional Development Fund - ERDF).

REFERENCES

- [1] N. P. Chue Hong, D. S. Katz, M. Barker, A.-L. Lamprecht, C. Martinez, F. E. Psomopoulos, J. Harrow, L. J. Castro, M. Gruenpeter, P. A. Martinez, and T. Honeyman, "FAIR Principles for Research Software (FAIR4RS Principles)," *Research Data Alliance*, 2022, <https://doi.org/10.15497/RDA00068>. [Online]. Available: <https://zenodo.org/record/6623556#.YqCJTJNBwlv>
- [2] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, J. Bouwman, A. J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C. T. Evelo, R. Finkers, A. Gonzalez-Beltran, A. J. G. Gray, P. Groth, C. Goble, J. S. Grethe, J. Heringa, P. A. C. 't Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S. J. Lusher, M. E. Martone, A. Mons, A. L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S.-A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M. A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao, and B. Mons, "The FAIR Guiding Principles for scientific data management and stewardship," *Scientific Data*, vol. 3, no. 1, p. 160018, Mar. 2016, <https://doi.org/10.1038/sdata.2016.18>. [Online]. Available: <https://www.nature.com/articles/sdata201618>
- [3] W. P. Walters, "Code Sharing in the Open Science Era," *Journal of Chemical Information and Modeling*, vol. 60, no. 10, pp. 4417–4420, Oct. 2020, <https://doi.org/10.1021/acs.jcim.0c01000>. [Online]. Available: <https://doi.org/10.1021/acs.jcim.0c01000>
- [4] N. Michaud-Agrawal, E. J. Denning, T. B. Woolf, and O. Beckstein, "MDAnalysis: A Toolkit for the Analysis of Molecular Dynamics Simulations," *J Comp Chem*, vol. 32, pp. 2319–2327, 2011, <https://doi.org/10.1002/jcc.21787>. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3144279/>
- [5] R. J. Gowers, M. Linke, J. Barnoud, T. J. E. Reddy, M. N. Melo, S. L. Seyler, D. L. Dotson, J. Domanski, S. Buchoux, I. M. Kenney, and O. Beckstein, "MDAnalysis: A Python package for the rapid analysis of molecular dynamics simulations." in *Proceedings of the 15th Python in Science Conference*, S. Benthall and S. Rostrup, Eds., Austin, TX, 2016, pp. 102–109, <https://doi.org/10.25080/Majora-629e541a-00e>.
- [6] "Scopus," <https://www.scopus.com/>. [Online]. Available: <https://www.scopus.com/>
- [7] "Journal of Open Source Software," <https://joss.theoj.org>. [Online]. Available: <https://joss.theoj.org>
- [8] I. Alibay, "IALibay/MDRestrainsGenerator: MDRestrainsGenerator 0.1.0," Mar. 2021, <https://doi.org/10.5281/zenodo.4570556>. [Online]. Available: <https://zenodo.org/record/4570556>
- [9] D. B. Kokh, B. Doser, S. Richter, F. Ormersbach, X. Cheng, and R. C. Wade, "A workflow for exploring ligand dissociation from a macromolecule: Efficient random acceleration molecular dynamics simulation and interaction fingerprint analysis of ligand trajectories," *The Journal of Chemical Physics*, vol. 153, no. 12, p. 125102, Sep. 2020, <https://doi.org/10.1063/5.0019088>. [Online]. Available: <https://aip.scitation.org/doi/10.1063/5.0019088>
- [10] C. Bouysset and S. Fiorucci, "ProLIF: a library to encode molecular interactions as fingerprints," *Journal of Cheminformatics*, vol. 13, no. 1, p. 72, Sep. 2021, <https://doi.org/10.1186/s13321-021-00548-6>. [Online]. Available: <https://doi.org/10.1186/s13321-021-00548-6>
- [11] K. A. Wilson, L. Wang, Y. C. Lin, and M. L. O'Mara, "Investigating the lipid fingerprint of SLC6 neurotransmitter transporters: a comparison of dDAT, hDAT, hSERT, and GlyT2," *BBA Advances*, vol. 1, p. 100010, Jan. 2021, <https://doi.org/10.1016/j.bbadv.2021.100010>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2667160321000090>
- [12] P. Smith and C. D. Lorenz, "LiPyphilic: A Python Toolkit for the Analysis of Lipid Membrane Simulations," *Journal of Chemical Theory and Computation*, vol. 17, no. 9, pp. 5907–5919, Sep. 2021, <https://doi.org/10.1021/acs.jctc.1c00447>. [Online]. Available: <https://doi.org/10.1021/acs.jctc.1c00447>
- [13] R. Gowers, M. Matta, and H. Nguyen, "kugupu/kugupu: v0.1.2," Feb. 2021, <https://doi.org/10.5281/zenodo.4545322>. [Online]. Available: <https://zenodo.org/record/4545322>
- [14] P. Loche, H. Jaeger, A. Schlaich, M. Becker, S. Gravelle, P. Stärk, and S. Velpuri, "MAiCoS," Feb. 2022. [Online]. Available: <https://gitlab.com/maicos-devel/maicos>
- [15] R. Araya-Secchi, T. Perez-Acle, S.-g. Kang, T. Huynh, A. Bernardin, Y. Escalona, J.-A. Garate, A. D. Martínez, I. E. García, J. C. Sáez, and R. Zhou, "Characterization of a Novel Water Pocket Inside the Human Cx26 Hemichannel Structure," *Biophysical Journal*, vol. 107, no. 3, pp. 599–612, Aug. 2014, <https://doi.org/10.1016/j.bpj.2014.05.037>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0006349514006018>
- [16] M. Tiberti, E. Papaleo, T. Bengtsen, W. Boomsma, and K. Lindorff-Larsen, "ENCORE: Software for Quantitative Ensemble Comparison," *PLOS Computational Biology*, vol. 11, no. 10, p. e1004415, Oct. 2015, <https://doi.org/10.1371/journal.pcbi.1004415>. [Online]. Available: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004415>
- [17] D. R. Roe and T. E. Cheatham, "PTRAJ and CPPTRAJ: Software for Processing and Analysis of Molecular Dynamics Trajectory Data," *Journal of Chemical Theory and Computation*, vol. 9, no. 7, pp. 3084–3095, Jul. 2013, <https://doi.org/10.1021/ct400341p>. [Online]. Available: <https://doi.org/10.1021/ct400341p>
- [18] M. J. Abraham, T. Murtola, R. Schulz, S. Páll, J. C. Smith, B. Hess, and E. Lindahl, "GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers," *SoftwareX*, vol. 1-2, pp. 19–25, Sep. 2015, <https://doi.org/10.1016/j.softx.2015.06.001>. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352711015000059>
- [19] M. Bonomi, G. Bussi, C. Camilloni, G. A. Tribello, P. Banáš, A. Barducci, M. Bernetti, P. G. Bolhuis, S. Bottaro, D. Branduardi, R. Capelli, P. Carloni, M. Ceriotti, A. Cesari, H. Chen, W. Chen, F. Colizzi, S. De, M. De La Pierre, D. Donadio, V. Drobot, B. Ensing, A. L. Ferguson, M. Filizola, J. S. Fraser, H. Fu, P. Gasparotto, F. L. Gervasio, F. Giberti, A. Gil-Ley, T. Giorgino, G. T. Heller, G. M. Hocky, M. Iannuzzi, M. Invernizzi, K. E. Jelfs, A. Jussupow, E. Kirilin, A. Laio, V. Limongelli, K. Lindorff-Larsen, T. Löhner, F. Marinelli, L. Martin-Samos, M. Masetti, R. Meyer, A. Michaelides, C. Molteni, T. Morishita, M. Nava, C. Pissoni, E. Papaleo, M. Parrinello, J. Pfandner, P. Piaggi, G. Piccini, A. Pietropaolo, F. Pietrucci, S. Pipolo, D. Provasi, D. Quigley, P. Raiteri, S. Raniolo, J. Rydzewski, M. Salvalaglio, G. C. Sosso, V. Spiwok, J. Šponer, D. W. H. Swenson, P. Tiwary, O. Valskov, M. Vendruscolo, G. A. Voth, A. White, and The PLUMED consortium, "Promoting transparency and reproducibility in enhanced molecular simulations," *Nature Methods*, vol. 16, no. 8, pp. 670–673, Aug. 2019, <https://doi.org/10.1038/s41592-019-0506-8>. [Online]. Available: <https://www.nature.com/articles/s41592-019-0506-8>
- [20] "AiiDA plugin registry," <https://aiidateam.github.io/aiida-registry/>. [Online]. Available: <https://aiidateam.github.io/aiida-registry/>
- [21] Chan Zuckerberg Initiative, "napari hub," <https://www.napari-hub.org/about>, last accessed 2022-08-05. [Online]. Available: <https://www.napari-hub.org/about>
- [22] N. Sofroniew, T. Lambert, K. Evans, J. Nunez-Iglesias, G. Bokota, P. Winston, G. Peña-Castellanos, K. Yamauchi, M. Bussonnier, D. Doncila Pop, A. Can Solak, Z. Liu, P. Wadhwa, A. Burt, G. Buckley, A. Sweet, L. Migas, V. Hilsenstein, L. Gaifas, J. Bragantini, J. Rodríguez-Guerra, H. Muñoz, J. Freeman, P. Boone, A. Lowe, C. Gohlke, L. Royer, A. PIERRÉ, H. Har-Gil, and A. McGovern, "napari: a multi-dimensional image viewer for Python," May 2022, <https://doi.org/10.5281/zenodo.3555620>. [Online]. Available: <https://doi.org/10.5281/zenodo.3555620>
- [23] L. Wang, I. Alibay, and F. Naughton, "Cookiecutter for MDAnalysis-

- based packages,” <https://github.com/MDAnalysis/cookiecutter-mdakit>. [Online]. Available: <https://github.com/MDAnalysis/cookiecutter-mdakit>
- [24] Open Source Initiative, “Licenses and Standards,” <https://opensource.org/licenses>, last accessed 2022-08-05. [Online]. Available: <https://opensource.org/licenses>
- [25] “GNU General Public License v2.0 - GNU Project - Free Software Foundation,” <https://www.gnu.org/licenses/old-licenses/gpl-2.0.en.html>, last accessed 2022-08-04. [Online]. Available: <https://www.gnu.org/licenses/old-licenses/gpl-2.0.en.html>
- [26] T. Preston-Werner, “Semantic Versioning 2.0.0,” <https://semver.org/>, last accessed 2022-08-04. [Online]. Available: <https://semver.org/>
- [27] “PEP 257 – Docstring Conventions | peps.python.org,” <https://peps.python.org/pep-0257/>. [Online]. Available: <https://peps.python.org/pep-0257/>
- [28] “Git,” <https://git-scm.com/>, last accessed 2022-08-04. [Online]. Available: <https://git-scm.com/>
- [29] GitHub, Inc, “GitHub,” <https://github.com>, last accessed 2022-08-04. [Online]. Available: <https://github.com>
- [30] GitLab Inc., “GitLab,” <https://about.gitlab.com/>, last accessed 2022-08-04. [Online]. Available: <https://about.gitlab.com/>
- [31] Atlassian, “Bitbucket,” <https://bitbucket.org/product>, last accessed 2022-08-04. [Online]. Available: <https://bitbucket.org/product>
- [32] “PEP 440 – Version Identification and Dependency Specification | peps.python.org,” <https://peps.python.org/pep-0440/>, last accessed 2022-08-04. [Online]. Available: <https://peps.python.org/pep-0440/>
- [33] Read the Docs, Inc, “Read the Docs,” <https://readthedocs.org/>, 2022. [Online]. Available: <https://readthedocs.org/>
- [34] Y. O. Halchenko, M. Visconti di Oleggio Castello, M. Hanke, J. Gors, M. Szczepanik, C. Barnes, E. Irvine, P. R. Raamana, C. J. Markiewicz, J. Wilk, D. Volgyes, K. Leinweber, L. Estève, O. Beckstein, and O. F. Gulban, “duccredit/duccredit: 0.9.1,” Apr. 2021, <https://doi.org/10.5281/zenodo.4685131>. [Online]. Available: <https://zenodo.org/record/4685131>
- [35] H. Krekel, B. Oliveira, R. Pfannschmidt, F. Bruynooghe, B. Laughner, and F. Bruhin, “pytest-dev/pytest,” <https://github.com/pytest-dev/pytest>, 2004. [Online]. Available: <https://github.com/pytest-dev/pytest>
- [36] Codecov LLC, “Codecov,” <https://about.codecov.io/>, 2022. [Online]. Available: <https://about.codecov.io/>
- [37] GitHub, Inc, “GitHub Terms of Service,” <https://docs.github.com/en/site-policy/github-terms/github-terms-of-service>, last accessed 2022-08-04. [Online]. Available: <https://docs.github.com/en/site-policy/github-terms/github-terms-of-service>
- [38] “pypa/setuptools,” <https://github.com/pypa/setuptools>, Aug. 2022, original-date: 2016-03-29T14:02:33Z. [Online]. Available: <https://github.com/pypa/setuptools>
- [39] “Poetry - Python dependency management and packaging made easy,” <https://python-poetry.org/>. [Online]. Available: <https://python-poetry.org/>
- [40] “PyPI · The Python Package Index,” <https://pypi.org/>. [Online]. Available: <https://pypi.org/>
- [41] Conda-Forge Community, “The conda-forge Project: Community-based Software Distribution Built on the conda Package Format and Ecosystem,” *Zenodo*, Jul. 2015, <https://doi.org/10.5281/ZENODO.4774216>. [Online]. Available: <https://zenodo.org/record/4774216>
- [42] A. Grossfield, “How to be a Good Member of a Scientific Software Community [Article v1.0],” *Living Journal of Computational Molecular Science*, vol. 3, no. 1, pp. 1473–1473, 2021, <https://doi.org/10.33011/livecoms.3.1.1473>. [Online]. Available: <https://livecomsjournal.org/index.php/livecoms/article/view/v3i1e1473>
- [43] “Submitting a paper to JOSS,” <https://joss.readthedocs.io/en/latest/submitting.html>, 2018, last accessed 2022-08-03. [Online]. Available: <https://joss.readthedocs.io/en/latest/submitting.html>
- [44] GitHub, Inc, “GitHub Actions,” <https://github.com/features/actions>, 2022. [Online]. Available: <https://github.com/features/actions>
- [45] P. Eastman, J. Swails, J. D. Chodera, R. T. McGibbon, Y. Zhao, K. A. Beauchamp, L.-P. Wang, A. C. Simmonett, M. P. Harrigan, C. D. Stern, R. P. Wiewiora, B. R. Brooks, and V. S. Pande, “OpenMM 7: Rapid development of high performance algorithms for molecular dynamics,” *PLOS Computational Biology*, vol. 13, no. 7, pp. 1–17, 07 2017, <https://doi.org/10.1371/journal.pcbi.1005659>.
- [46] O. S. Smart, J. G. Neduvilil, X. Wang, B. A. Wallace, and M. S. P. Sansom, “HOLE: A program for the analysis of the pore dimensions of ion channel structural models,” *J Molecular Graphics*, vol. 14, pp. 354–360, 1996, [https://doi.org/10.1016/s0263-7855\(97\)00009-x](https://doi.org/10.1016/s0263-7855(97)00009-x). [Online]. Available: <http://www.holeprogram.org/>
- [47] L. S. Stelzl, P. W. Fowler, M. S. P. Sansom, and O. Beckstein, “Flexible gates generate occluded intermediates in the transport cycle of LacY,” *J Mol Biol*, vol. 426, pp. 735–751, 2014,
- <https://doi.org/10.1016/j.jmb.2013.10.024>. [Online]. Available: <http://doi.org/10.1016/j.jmb.2013.10.024>