

## RESEARCH ARTICLE

## gmxapi: A GROMACS-native Python interface for molecular dynamics with ensemble and plugin support

M. Eric Irrgang<sup>1</sup>, Caroline Davis<sup>1</sup>, Peter M. Kasson<sup>1,2\*</sup>

**1** Departments of Molecular Physiology and Biomedical Engineering, University of Virginia, Charlottesville, Virginia, United States of America, **2** Science for Life Laboratory, Department of Cell and Molecular Biology, Uppsala University, Uppsala, Sweden

\* [kasson@virginia.edu](mailto:kasson@virginia.edu)**OPEN ACCESS**

**Citation:** Irrgang ME, Davis C, Kasson PM (2022) gmxapi: A GROMACS-native Python interface for molecular dynamics with ensemble and plugin support. *PLoS Comput Biol* 18(2): e1009835. <https://doi.org/10.1371/journal.pcbi.1009835>

**Editor:** Dina Schneidman-Duhovny, Hebrew University of Jerusalem, ISRAEL

**Received:** July 23, 2021

**Accepted:** January 16, 2022

**Published:** February 14, 2022

**Copyright:** © 2022 Irrgang et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** The gmxapi Python package is also maintained as part of the GROMACS repository at <https://gitlab.com/gromacs/gromacs>. It can be installed from the GROMACS source ([https://gitlab.com/gromacs/gromacs/-/tree/master/python\\_packaging/src](https://gitlab.com/gromacs/gromacs/-/tree/master/python_packaging/src)) or from <https://pypi.org/project/gmxapi/> with “pip”, but the GROMACS installation must be specified. The documentation at <https://manual.gromacs.org/current/gmxapi> provides details. During installation, the gmxapi Python package builds a C++ extension module against the GROMACS installation. Gmxapi tutorials are available from

## Abstract

Gmxapi provides an integrated, native Python API for both standard and advanced molecular dynamics simulations in GROMACS. The Python interface permits multiple levels of integration with the core GROMACS libraries, and legacy support is provided via an interface that mimics the command-line syntax, so that all GROMACS commands are fully available. Gmxapi has been officially supported since the GROMACS 2019 release and is enabled by default in current versions of the software. Here we describe gmxapi 0.3 and later. Beyond simply wrapping GROMACS library operations, the API permits several advanced operations that are not feasible using the prior command-line interface. First, the API allows custom user plugin code within the molecular dynamics force calculations, so users can execute custom algorithms without modifying the GROMACS source. Second, the Python interface allows tasks to be dynamically defined, so high-level algorithms for molecular dynamics simulation and analysis can be coordinated with loop and conditional operations. Gmxapi makes GROMACS more accessible to custom Python scripting while also providing support for high-level data-flow simulation algorithms that were previously feasible only in external packages.

## Author summary

The gmxapi software provides a Python interface for molecular dynamics simulations in GROMACS. In addition to simply wrapping GROMACS commands, it supports custom user plugin code, ensemble simulation, and data-flow chaining of commands. As such, gmxapi enables the writing and execution of high-level simulation algorithms. The software ships with GROMACS and is freely available under an LGPL2 license.

## Introduction

As molecular dynamics simulations have become more complex and mature as scientific tools, typical simulation use is shifting from manual invocation of a few simulations and analysis tools to pre-defined simulation and analysis protocols, often involving many simulation

<https://github.com/kassonlab/gmxapi-tutorials>. Custom molecular dynamics extension code is illustrated in a “sample\_restraint” package: [https://gitlab.com/gromacs/gromacs/-/tree/master/python\\_packaging/sample\\_restraint](https://gitlab.com/gromacs/gromacs/-/tree/master/python_packaging/sample_restraint). BRER restraint potentials [28] are forked from this sample code and can be found at [https://github.com/kassonlab/brer\\_plugin](https://github.com/kassonlab/brer_plugin). Scripted BRER workflows are available at [https://github.com/kassonlab/run\\_brer](https://github.com/kassonlab/run_brer). Input data for the BRER example shown have been deposited at doi: [10.5281/zenodo.5122931](https://doi.org/10.5281/zenodo.5122931) Gmxapi issues are tracked with the label “gmxapi” at <https://gitlab.com/gromacs/gromacs/-/issues>. Code contributions follow the GROMACS contribution procedure (<https://manual.gromacs.org/current/dev-manual/contribute.html>). However, gmxapi is intended to allow for maximal extensibility without requiring modification to the sources. A discussion forum is available at <https://gromacs.bioexcel.eu/tag/gmxapi>.

**Funding:** This work was supported by National Institutes of Health R01 GM115790 (<https://www.nigms.nih.gov/>), National Science Foundation OAC-1835780 (<https://www.nsf.gov/div/index.jsp?div=OAC>) to P.M.K., and National Science Foundation ACI-1547580 (<https://www.nsf.gov/div/index.jsp?div=OAC>) to M.E.I. (subaward to P.M.K). Computational resources on Frontera were provided under NSF LRAAC MCB20006. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

trajectories. In addition, custom applications for advanced sampling[1–4] or molecular structure refinement[5–7] are becoming more common, where the molecular dynamics engine is used as part of a more complex data integration protocol. In these more complex use scenarios, toolchain and file-system management as well as execution of many simulations can become limiting factors. As a result, many scientists either develop custom scripts, custom modifications of the simulation source code, or adapt general-purpose workflow engines[8–10] to molecular simulation tasks. There also exist some molecular-simulation-specific workflow engines, frequently coupled to the underlying molecular simulation code via command-line interfaces[11–16]. All of the above can be brittle, particularly without robust APIs for simulation interfaces to workflow engines. Integration of parallel analysis into unified jobs can also be challenging whether this is performed via built-in tool parallelism or user-level coding, such as Python multiprocessing packages or workflow/execution managers[9,10,12,13,17–19].

We previously reported on gmxapi 0.0.4[20], which allows Python driven molecular dynamics in GROMACS[21,22] to be extended at run time with custom researcher code. Here, we describe features present in gmxapi 0.2 and beyond, which offers both more advanced ensemble simulation logic and a feature-complete interface for GROMACS tools and analysis. This framework allows ensemble methods to be implemented with less code, and without patching an official GROMACS release. Data-flow-oriented programming logic and general paradigm for integrating new software tools further enhance the utility of the package. Both the libgmxapi C++ interface and the gmxapi Python package are now maintained and distributed with GROMACS; gmxapi 0.2 is integrated with GROMACS 2021, and gmxapi 0.3 is integrated with GROMACS 2022.

## Design and implementation

We outline distinguishing features and key user functionality of gmxapi, followed by a more technical design discussion. Gmxapi offers a Python scripting interface maintained as part of the GROMACS software, and thus one basic feature is the ability to reproduce all GROMACS command-line calls. This is done through a combination of native gmxapi calls such as `gmxapi.mdrun()` and a wrapper that permits any operation from the GROMACS command-line client: `gmxapi.commandline_operation()`. These are illustrated in Fig 1A. We also highlight three key additional features not possible with a simple wrapper script: built-in ensemble parallelism, composability, and plugins.

## Ensemble parallelism

Simulations are increasingly being performed not singly but as a collection of related tasks, which we will term an ensemble. This collection could consist of a set of replicas sampling a thermodynamic ensemble, but it could also represent a parameter sweep across experimental conditions or a set of simulations with different starting conformations as part of an advanced sampling strategy. Ensembles are treated as first-class objects in gmxapi, and the python interface is designed to facilitate parallelism across such ensembles, which is in turn implemented by the gmxapi backend. This is illustrated in Fig 1B. The ability to flexibly and simply articulate high-level ensemble logic is thus a key feature of gmxapi.

## Composability

Another key design and usability feature of gmxapi is the notion of composability. Similar to Tinkertoys, different calls to GROMACS commands through the API are designed to be plugged together. This is illustrated in Fig 1C and has two advantages: a natural way to conceptualize a sequence of GROMACS calls (or calls to external programs that can use the gmxapi

```

a.
args = ['pdb2gmx', '-ff', 'amber99sb-ildn', '-water', 'tip3p']
input_files = {'-f': os.path.join(input_dir, 'start0.pdb')}
output_files = {
    '-p': os.path.join(cmd_dir, 'topol.top'),
    '-i': os.path.join(cmd_dir, 'posre.itp'),
    '-o': os.path.join(cmd_dir, 'conf.gro')
}
make_top = gmx.commandline_operation('gmx', args, input_files, output_files)

b.
grompp_input_files = {'-f': os.path.join(cmd_dir, 'grompp.mdp'),
    '-c': make_top.output.file['-o'],
    '-p': make_top.output.file['-p']}
# make array of inputs
N = 50
grompp = gmx.commandline_operation('gmx', ['grompp'], grompp_input_files * N,
    {'-o': os.path.join(cmd_dir, 'run.tpr')})
tpr_input = grompp.output.file['-o'].result()
input_list = gmx.read_tpr([tpr_input])
md = gmx.mdrun(input_list)
md.run()

c.
subgraph = gmx.subgraph(variables = {'found_native': False, 'checkpoint': '', 'min_rms': 1e6})
with subgraph:
    md = gmx.mdrun(input_list, runtime_args={'-cpi': subgraph.checkpoint, '-maxh': '2'})
    subgraph.checkpoint = md.output.checkpoint
    rmsd = gmx.commandline_operation('gmx', ['rms'],
        input_files={'-s': reference_struct,
            '-f': md.output.trajectory},
        output_files={'-o': 'rmsd.xvg'}, stdin='Backbone Backbone')
    subgraph.min_rms = numeric_min(xvg_to_array(rmsd.output.file['-o']).output.data)
    subgraph.found_native = less_than(lhs=subgraph.min_rms, rhs=0.3).output.data
    folding_loop = gmx.while_loop(operation=subgraph, condition=logical_not(subgraph.found_native))()
    folding_loop.run()

```

gmxapi command-line operation

gmxapi command on ensemble input

looping and pluggability

**Fig 1. gmxapi usage examples.** Panel (A) shows an example of the `commandline_operation` function by which `gmxapi` can reproduce any GROMACS functionality. Panel (B) shows `gmxapi` molecular dynamics calls operating on ensemble input, providing straightforward high-level parallelism in addition to parallelization within each command. Panel (C) demonstrates both a while loop and pluggability of `gmxapi` components. Together, the examples will execute an ensemble of small protein folding simulations until at least one ensemble member samples the native state.

<https://doi.org/10.1371/journal.pcbi.1009835.g001>

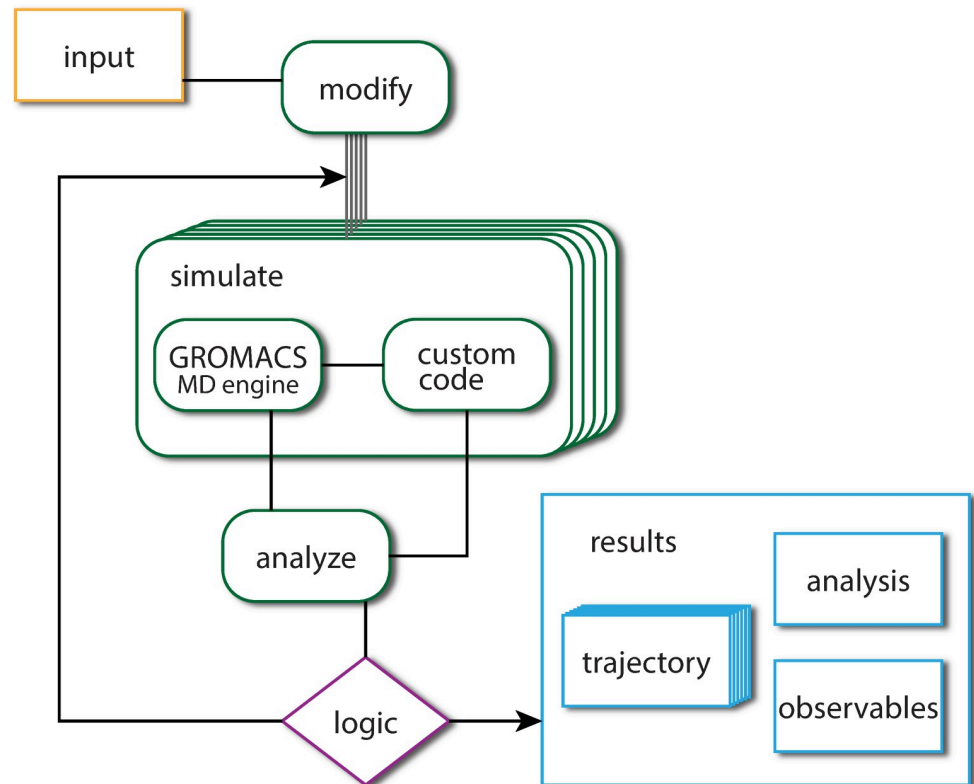
wrapper facility) and a way to parallelize using the ensemble logic and data-flow management by `gmxapi`.

## Plugins

`Gmxapi` includes the ability to add user-defined plugins that can interact with GROMACS at runtime without modifying the GROMACS source. Plugins are currently used to implement custom force routines; these run using the native GROMACS parallel decomposition and thus can benefit from parallelism as well as acceleration. More details are given in the technical design section below.

## Technical design

We used `pybind11` [23] to implement Python bindings to a C++ API for the GROMACS [21,22] molecular simulation library. The design concept of `gmxapi`, as with many modern high-level interfaces, is for the user to construct a computational graph where execution is then managed by lower layers of the software stack. Similar to `Keras` [24], we recognize that explicit data-flow programming is not natural for all users and thus provide an interface where data-flow can be specified explicitly or implicitly. [Fig 2](#) illustrates the `gmxapi` paradigm in



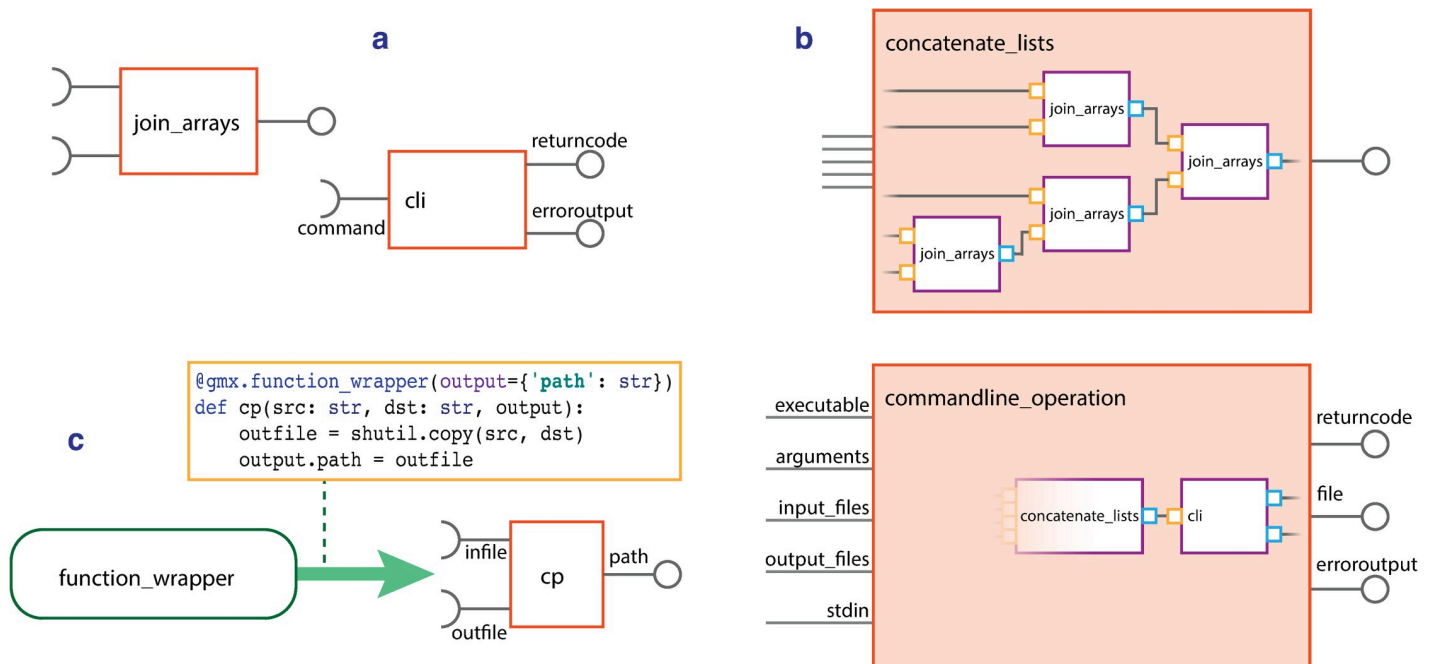
**Fig 2. Schematic of data flow and control flow for a segment of a complex simulation workflow.** Ensembles of simulations can be run (denoted by stacked rectangles) by gmxfapi merely by passing an array of inputs instead of a single input. Custom plugins can interact with running MD simulations. Finally, conditional and looping logic can create high-level simulation algorithms.

<https://doi.org/10.1371/journal.pcbi.1009835.g002>

which control signals and data are treated in the same framework. Branching and other control events require that some work is "dynamic", such that the complete workflow graph may not be determined until runtime. However, the workflow can be fully specified in terms of workflow commands and data primitives (Fig 3).

### Data flow formalism

Data-flow formalisms have become increasingly popular in contemporary workflow engines due to their ability to separate work specification and task execution, permitting more straightforward resolution of dependencies and optimization of computation and data movement. For the same reasons, execution in gmxfapi is deferred as much as possible to when and where it is required. The high-level gmxfapi Python interface creates and executes a directed acyclic graph (DAG) specifying the computational work to be done. Nodes of the work graph represent discrete *operations* that produce and consume well defined inputs and outputs. Operation and data references in the work graph are proxy objects for the computational tasks and data until the graph is run. Execution management may be optimized by running only as much of the graph as necessary to satisfy explicit data dependencies. In addition, in order to minimize unnecessary data movement, most gmxfapi operations do not transfer data from the C++ library to the Python interpreter unless explicitly requested by the user. When a gmxfapi command is used to add work to the work graph, it returns a reference to the graph node representing the operation. The operation's output attributes may be used as inputs to further



**Fig 3. Schematics for gmxfapi operations.** As shown in panel (A), gmxfapi operations have well defined inputs and outputs. These can operate on arrays of inputs and also include support for legacy GROMACS operations by wrapping the command-line toolset, as shown in panel (B). As shown in panel (C), the decorator `@function_wrapper` allows arbitrary user code to be transformed into a gmxfapi operation.

<https://doi.org/10.1371/journal.pcbi.1009835.g003>

commands, and may be used as “Future” results, as discussed below. Calling “`result()`” on a Future forces dependency resolution and data localization.

## Python interface

The gmxfapi Python interface consists of five categories of operations: 1) typing and logical operations, 2) bound GROMACS API calls, 3) legacy GROMACS command-line operations, 4) utilities for user creation of new gmxfapi operations, and 5) looping and conditional operations. These are described in sequence below.

The simplest gmxfapi operations manipulate signals, data topology, or typing. “`join_arrays`”, “`logical_not`”, and “`make_constant`” have strict definitions for inputs and outputs (Fig 3A). These operations are also composable: “`concatenate_lists`” is a helper function that reduces a number of inputs in terms of “`join_arrays`” (Fig 3B).

The preferred mechanism for interfacing to GROMACS is via bound API calls. Commands like “`read_tpr`”, “`modify_input`”, and “`mdrun`” use a binary Python extension module written in C++ using `pybind11` to interact with GROMACS operations and data through `libgmxfapi`, which is a C++ library installed by default with recent versions of GROMACS.

In order to provide legacy support for the full range of GROMACS command-line operations, “`gmxfapi.commandline.cli()`” is a simple pure-Python gmxfapi operation that wraps command-line tools. It is a thin wrapper of the Python “`subprocess`” module, so command inputs and outputs are embedded in the argument list parameter. Instances of “`gmxfapi.commandline.cli`” are not by themselves conducive to DAG representations of data flow, so a helper function creates additional gmxfapi primitives to provide a consistent set of named inputs and outputs. “`gmxfapi.commandline_operation()`” generates a graph of gmxfapi primitives around “`cli`” to translate input and output arguments into operation inputs and outputs, as specified by the user (Fig 3B).



In order to provide extensibility, arbitrary user code can be transformed into a gmxapi operation by decorating a Python function definition with “@function\_wrapper” (Fig 3C). As an example, “commandline\_operation” is written in this manner.

Looping and conditional commands are critical to adaptive ensemble simulation[25] and thus form a key part of the gmxapi repertoire. Conditional iteration takes the form of conventional gmxapi command syntax but relies on some metaprogramming under the hood. The “while\_loop” command allows a graph to be dynamically extended as a chain of repeated operations, including fused operations constructed with the “subgraph” tool. A subgraph allows inputs and outputs to be expressed in terms of other gmxapi operations. When used with “while\_loop”, internal subgraph state can be propagated from one iteration to the next to provide a consistent environment similar to a standard loop.

## The work graph

The work graph permits data flow topology to be represented independently from execution strategies or run time resource assignment details. An array of input sources provided to a command generates a corresponding set of tasks, such as for trajectory ensemble simulations. Because such ensembles may be coupled (and gmxapi currently does not have a mechanism to specify that an ensemble is uncoupled), arrays of simulations must currently be co-executed through the gmxapi mpi4py executor. Operations defined with “@function\_wrapper”, including the other built-in operations, are assumed to be uncoupled, and are launched in a sequence determined by the DAG topology during recursive resolution of data dependencies, executed sequentially by the simple built-in gmxapi 0.1 executor.

The work graph enables straightforward dependency resolution: commands are specified to operate on abstracted “handles” or references to work inputs rather than requiring the fully instantiated objects, so chains of simulations can be expressed where one command depends on the outputs of a prior command. Each individual command is then ready to execute once its inputs have been fully resolved. Because operations and data flow are represented as a DAG, arbitrarily complex topologies can be expressed unambiguously without unexpected side-effects. Trajectories can be forked or extended without re-executing simulation segments or overwriting previous results.

## Execution and control flow

The gmxapi interface enables several levels of control logic, described below. Native Python logic (if/else statements, for loops, etc.) can be used with gmxapi operations. However, since gmxapi operations construct a work graph that is then executed, expressing conditionals within this work graph requires special constructs rather than standard Python operators. Gmxapi therefore provides logic such as “while\_loop” as an efficient way to express adaptivity within a sequence of gmxapi operations. Python-level logic can be utilized by forcing work graph execution using the result() method, which causes explicit execution of all code required to produce the requested result. In contrast, gmxapi-level logic operates on Futures and is thus deferred until graph execution.

Gmxapi also provides low-level logic for simulation control and ensemble operators that can be used by third party code. For instance, gmxapi registers with the GROMACS “Stop-Signal” facility, enabling plugin code to stop a simulation based on external criteria, such as when a statistical estimator has converged. Additional examples of low-level operations include a “ReduceAll” to collect data across an in-flight ensemble. This operation can be used for adaptive updates across an ensemble of simulations, such as modifying biasing forces in restrained-ensemble simulations[26,27] or to update estimators and determine when a simulation ensemble should be terminated.

## Gmxapi as a means to extend MD code without source modification

Gmxapi establishes a facility[28] for providing MD extension code to GROMACS during launch. Binary objects are then loaded at runtime via the Python interpreter. This allows custom code to be executed during the MD integration loop with minimal overhead. Python facilitates plugin binding, but once the simulation launches, GROMACS and the extension code communicate directly via C++ interfaces, so there is no overhead added by Python. This extension code executes within the existing GROMACS parallelism framework, so it can take advantage of domain decomposition and MPI parallelism. Since this only requires a standard GROMACS installation with gmxapi, it permits flexible extension of the molecular dynamics code without modifying the main codebase or libraries.

Gmxapi does not prescribe a C/C++ bindings strategy. We used pybind11 for GROMACS bindings, but Python interaction with compiled code relies only on the Python C API (PyCapsule) and the public libgmxapi C++ interface.

## Results

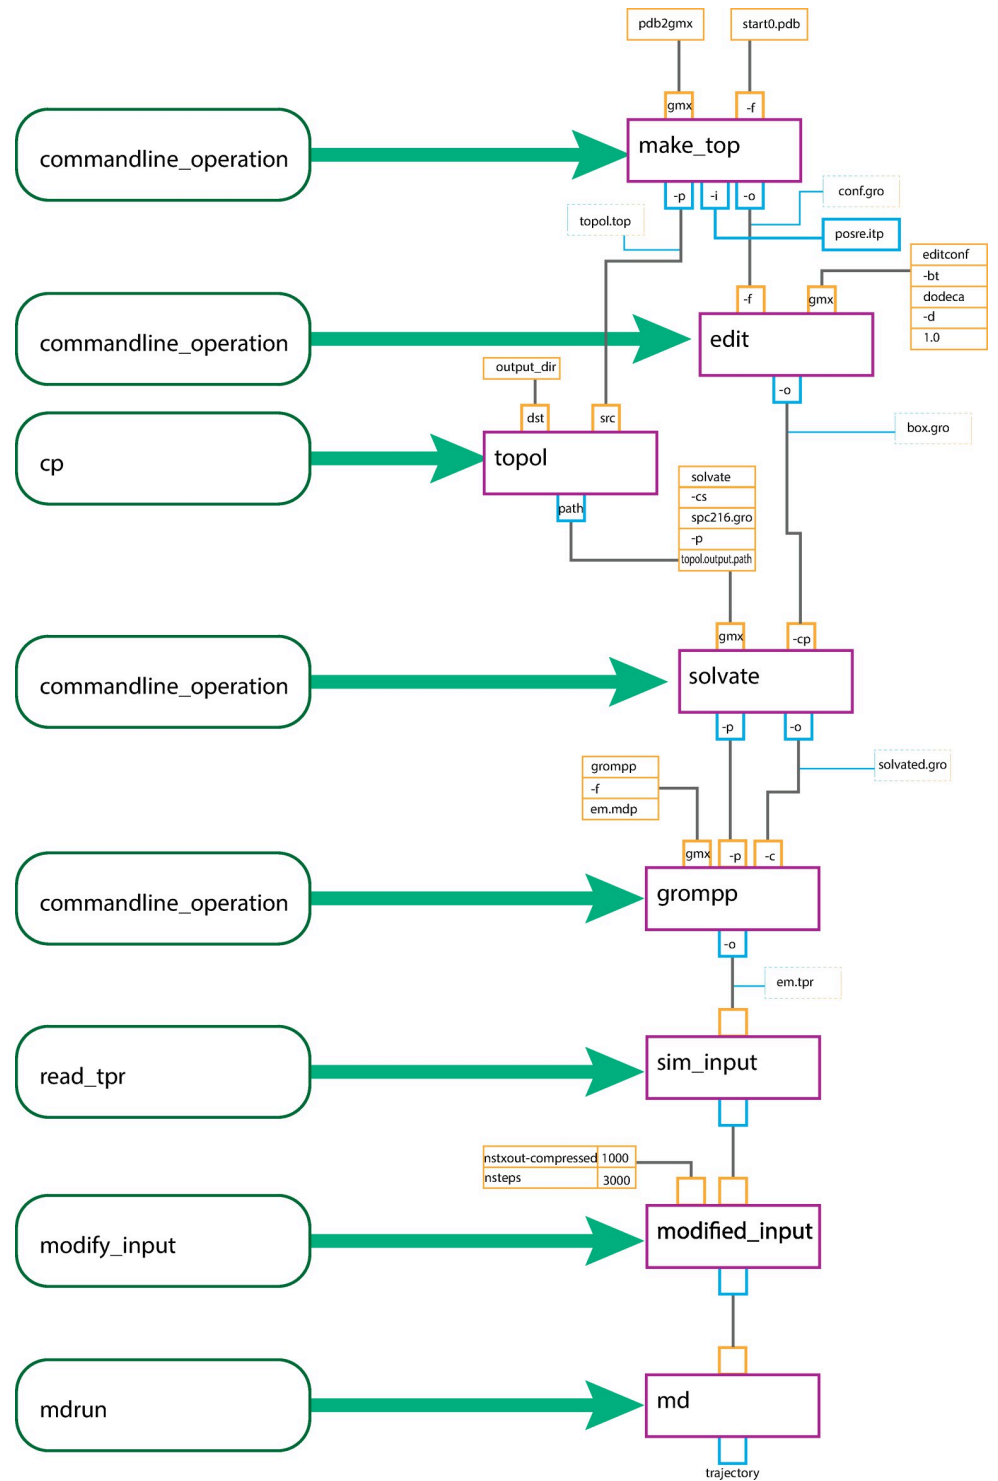
Typical molecular dynamics workflows use a chain of tools to prepare simulation input, particularly command-line programs that take file locations as input and output arguments. Gmxapi is designed to permit connecting the output of any tool to the input for another tool with consistent Python syntax. Fig 4 illustrates a chain of GROMACS tools preparing simulation input, culminating with a call to "gmxapi.mdrun" to execute an array of simulations.

To illustrate the application of gmxapi, we demonstrate refinement of an HIV gp41 conformational ensemble based on DEER spectroscopy data. We use a starting crystal structure for the BG505 SOSIP of HIV gp41[29] and previously reported DEER spectroscopy data of 5 different spin labels positions on this SOSIP[30]. The gp41 trimer was asymmetrically restrained (each of the three monomer-monomer distances was sampled separately), and the previously reported Bias-resampling ensemble refinement (BRER) algorithm for refining heterogeneous conformational ensembles[28] was applied. This involves a custom force plugin for GROMACS, and gmxapi was used to run an array of 250 refinement simulations, each randomly sampling target distance values from the experimental distribution for each measured spin-label pair. This array was executed in parallel on a supercomputing cluster, and accompanying example scripts demonstrate such a large-scale deployment. Further details are given in S1 Methods. Full convergence of the refinement would require additional simulation sampling of the 15-dimensional experimental distribution, but even after 24 wall-clock hours of simulation a reasonable sampling of the experimental distributions was obtained (Fig 5). This application illustrates the ease of simulation ensemble management and application of custom biasing algorithms within GROMACS facilitated by gmxapi.

## Discussion

Gmxapi provides a high-level Python interface for GROMACS with several key design features. It permits easy chaining of commands (GROMACS or third-party analysis tools) to create pipelines. Such pipelines can be parallel in nature, and gmxapi supports arrays or ensembles of simulations as first-class objects. Finally, gmxapi has a plugin interface that enables custom user code to be executed as part of molecular dynamics force calculations without modifying the GROMACS source but still benefitting from GROMACS native parallel decomposition. These features provide distinct advantages over shell-script-based workflows or simple Python wrappers for the GROMACS command-line interface.

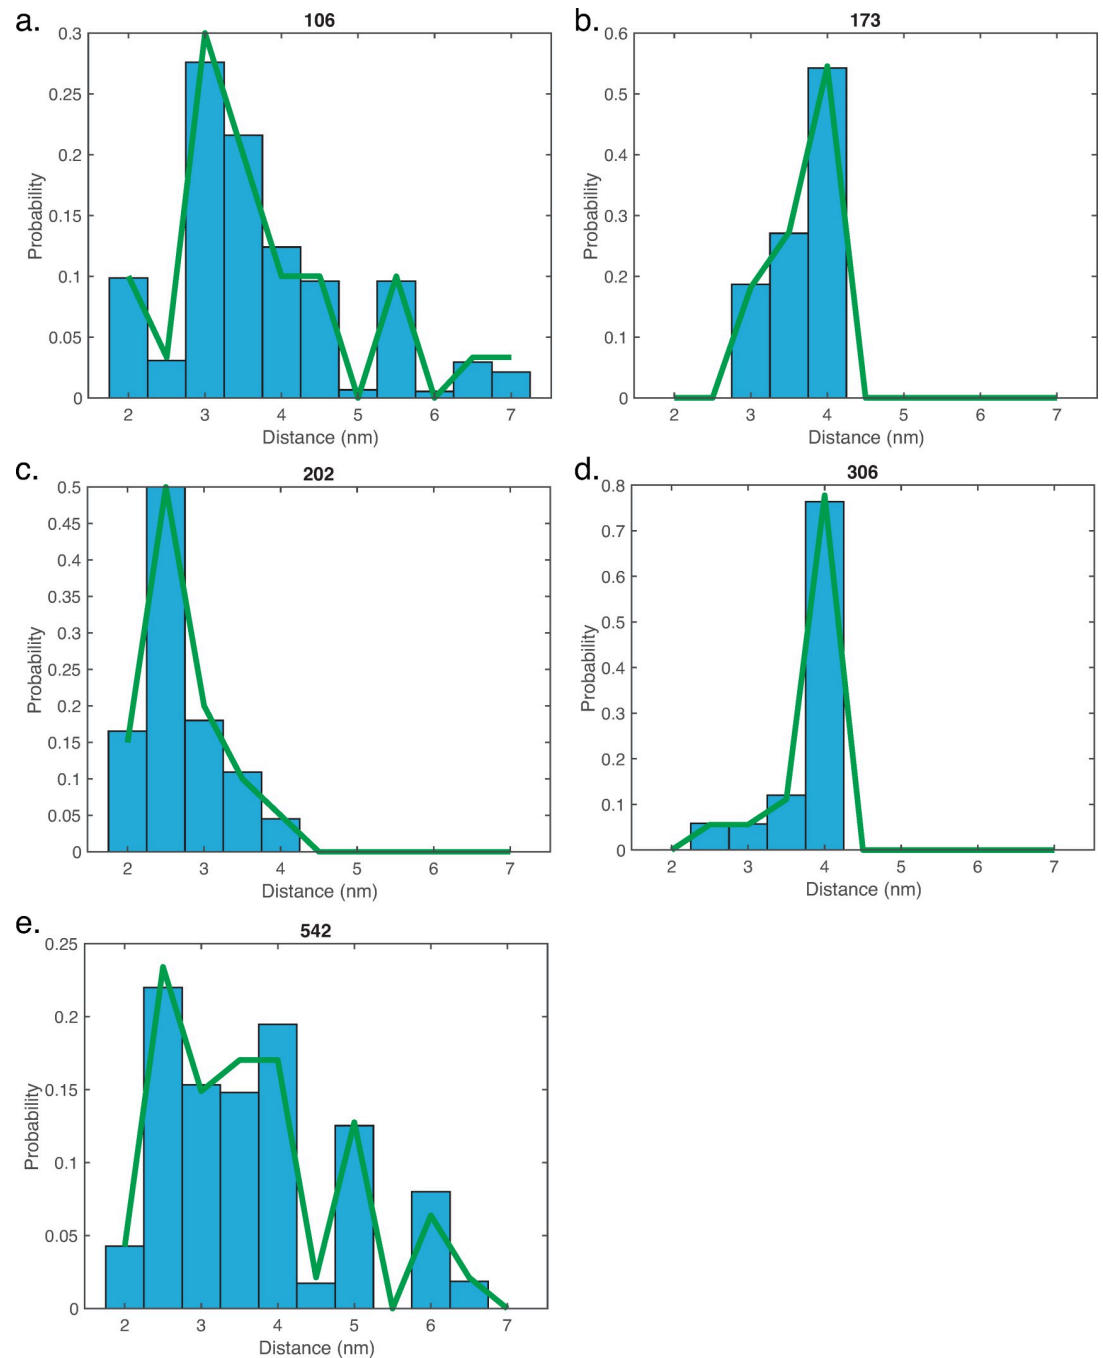
Gmxapi differs from many molecular dynamics scripting APIs in that it embraces a data-flow programming paradigm while at the same time aiming for simplicity and ease of



**Fig 4. Detailed diagram of inputs, outputs, and operations in a chain of gmxapi operations.** Operations are shown in green boxes, the corresponding nodes in the work graph in red boxes, inputs in orange boxes, and outputs in blue boxes. The diagram depicts a chain of GROMACS tools preparing simulation input, feeding into "gmxapi.mdrun" to run a batch of simulations.

<https://doi.org/10.1371/journal.pcbi.1009835.g004>





**Fig 5. Residue-residue distance distributions in a simulated ensemble of HIV gp41.** DEER spectroscopy was used to measure distance distributions between two monomers of the gp41 trimer. Thus, each residue designates a monomer-monomer residue pair. Panels A-E show plots for the 5 restrained residues in the HIV trimer: 106, 173, 202, 306, and 542 respectively. Discretized DEER distance distributions are plotted in green, and simulation results are plotted in blue bars. The simulation ensemble shows good convergence to the measured values.

<https://doi.org/10.1371/journal.pcbi.1009835.g005>

programming. This is analogous to the Keras high-level deep learning library [24]. A data-flow approach simplifies treatment of dependencies in complex parallel workflows and has been adopted by a number of high-performance computational tools.

Gmxapi scales well to high-performance clusters using MPI parallel interfaces. Current versions do not, however, provide advanced scheduler management along the lines of the prior Copernicus software [11–12] or the Radical Cybertools toolkit [16]. Interfaces that permit integration of gmxapi with more advanced schedulers is planned for future work. In comparison with Copernicus, gmxapi is a lighter-weight solution, integrating more closely with GROMACS and designed to facilitate much greater extensibility and ease of use but not including the advanced scheduler and client-server communication capabilities.

## Availability and future directions

The libgmxapi C++ interface was released as part of GROMACS 2019 release and has been part of a standard GROMACS installation since 2020. Gmxapi 0.2 shipped with GROMACS 2021; version 0.3, described here, is part of the GROMACS 2022 release and contains further enhancements for ease of use and installation. Continued development of gmxapi will expand the number of bound GROMACS API calls and reduce the need for legacy command-line support as well as generalize the plugin interface. At the programming level, one key future direction is exactly that—an increased use of the “Future” paradigm in Python. Python design patterns using Futures are becoming more widespread and standardized [14,31] as a way to refer to results that have not yet been calculated. Some additional refinement of the gmxapi Future protocol is needed to be fully compatible with native and third-party frameworks for concurrent or asynchronous program flow.

## Supporting information

### S1 Methods. Supplementary Methods.

(PDF)

### S1 Software Archive.

(TGZ)

## Acknowledgments

The authors thank Mark Abraham for many helpful discussions.

## Author Contributions

**Conceptualization:** M. Eric Irrgang, Peter M. Kasson.

**Investigation:** Caroline Davis, Peter M. Kasson.

**Methodology:** M. Eric Irrgang.

**Software:** M. Eric Irrgang, Caroline Davis, Peter M. Kasson.

**Writing – original draft:** M. Eric Irrgang, Peter M. Kasson.

**Writing – review & editing:** M. Eric Irrgang, Caroline Davis, Peter M. Kasson.

## References

1. Bonomi M, Camilloni C (2017) Integrative structural and dynamical biology with PLUMED-ISDB. *Bioinformatics* 33: 3999–4000. <https://doi.org/10.1093/bioinformatics/btx529> PMID: 28961689
2. Tribello GA, Bonomi M, Branduardi D, Camilloni C, Bussi G (2014) PLUMED 2: New feathers for an old bird. *Computer Physics Communications* 185: 604–613.

3. Olsson S, Wu H, Paul F, Clementi C, Noe F (2017) Combining experimental and simulation data of molecular processes via augmented Markov models. *Proc Natl Acad Sci U S A* 114: 8265–8270. <https://doi.org/10.1073/pnas.1704803114> PMID: 28716931
4. Chodera JD, Noé F (2014) Markov state models of biomolecular conformational dynamics. *Current opinion in structural biology* 25: 135–144. <https://doi.org/10.1016/j.sbi.2014.04.002> PMID: 24836551
5. Heo L, Feig M (2018) Experimental accuracy in protein structure refinement via molecular dynamics simulations. *Proceedings of the National Academy of Sciences* 115: 13276–13281. <https://doi.org/10.1073/pnas.1811364115> PMID: 30530696
6. Perilla JR, Zhao G, Lu M, Ning J, Hou G, et al. (2017) CryoEM structure refinement by integrating NMR chemical shifts with molecular dynamics simulations. *The Journal of Physical Chemistry B* 121: 3853–3863. <https://doi.org/10.1021/acs.jpcc.6b13105> PMID: 28181439
7. Igaev M, Kutzner C, Bock LV, Vaiana AC, Grubmuller H (2019) Automated cryo-EM structure refinement using correlation-driven molecular dynamics. *Elife* 8. <https://doi.org/10.7554/eLife.43542> PMID: 30829573
8. Goecks J, Nekrutenko A, Taylor J, Team G (2010) Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biology* 11.
9. Deelman E, Singh G, Su M-H, Blythe J, Gil Y, et al. (2005) Pegasus: A framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming* 13: 219–237.
10. Wolstencroft K, Haines R, Fellows D, Williams A, Withers D, et al. (2013) The Taverna workflow suite: designing and executing workflows of Web Services on the desktop, web or in the cloud. *Nucleic Acids Research* 41: W557–W561. <https://doi.org/10.1093/nar/gkt328> PMID: 23640334
11. Pronk S, Larsson P, Pouya I, Bowman GR, Haque IS, et al. (2011) Copernicus: A new paradigm for parallel adaptive molecular dynamics. *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*: 60.
12. Pronk S, Pouya I, Lundborg M, Rotskoff G, Wesén B, et al. (2015) Molecular Simulation Workflows as Parallel Algorithms: The Execution Engine of Copernicus, a Distributed High-Performance Computing Platform. *Journal of Chemical Theory and Computation*. <https://doi.org/10.1021/acs.jctc.5b00234> PMID: 26575558
13. Jain A, Ong SP, Chen W, Medasani B, Qu XH, et al. (2015) FireWorks: a dynamic workflow system designed for high-throughput applications. *Concurrency and Computation-Practice & Experience* 27: 5037–5059.
14. Babuji Y, Woodard A, Li Z, Katz DS, Clifford B, et al. *Parsl: Pervasive parallel programming in python 2019*. pp. 25–36.
15. Wilde M, Hategan M, Wozniak JM, Clifford B, Katz DS, et al. (2011) Swift: A language for distributed parallel scripting. *Parallel Computing* 37: 633–652.
16. Balasubramanian V, Bethune I, Shkurti A, Breitmoser E, Hruska E, et al. *Extasy: Scalable and flexible coupling of md simulations and advanced sampling techniques*; 2016. IEEE. pp. 361–370.
17. van der Aalst WMP, Jablonski S (2000) Dealing with workflow change: identification of issues and solutions. *Computer Systems Science and Engineering* 15: 267–276.
18. Ludascher B, Altintas I, Berkley C, Higgins D, Jaeger E, et al. (2006) Scientific workflow management and the Kepler system. *Concurrency and Computation-Practice & Experience* 18: 1039–1065.
19. Mattoso M, Dias J, Ocana KACS, Ogasawara E, Costa F, et al. (2015) Dynamic steering of HPC scientific workflows: A survey. *Future Generation Computer Systems-the International Journal of Esience* 46: 100–113.
20. Irrgang ME, Hays JM, Kasson PM (2018) gmxapi: a high-level interface for advanced control and extension of molecular dynamics simulations. *Bioinformatics* 34: 3945–3947. <https://doi.org/10.1093/bioinformatics/bty484> PMID: 29912282
21. Pronk S, Pall S, Schulz R, Larsson P, Bjelkmar P, et al. (2013) GROMACS 4.5: a high-throughput and highly parallel open source molecular simulation toolkit. *Bioinformatics* 29: 845–854. <https://doi.org/10.1093/bioinformatics/btt055> PMID: 23407358
22. Abraham MJ, Murtola T, Schulz R, Páll S, Smith JC, et al. (2015) GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX* 1: 19–25.
23. Wenzel J, Rhineland J, Moldovan D (2017) pybind11—Seamless operability between C++ 11 and Python.
24. Chollet F (2015) Keras. <https://keras.io>.
25. Kasson PM, Jha S (2018) Adaptive ensemble simulations of biomolecules. *Current opinion in structural biology* 52: 87–94. <https://doi.org/10.1016/j.sbi.2018.09.005> PMID: 30265901

26. Roux B, Islam SM (2013) Restrained-ensemble molecular dynamics simulations based on distance histograms from double electron-electron resonance spectroscopy. *J Phys Chem B* 117: 4733–4739. <https://doi.org/10.1021/jp3110369> PMID: 23510121
27. Hays JM, Kieber MK, Li JZ, Han JI, Moremen KW, et al. (2018) Refinement of highly flexible protein structures using simulation-guided spectroscopy. *Angew Chem Int Ed Engl* 57: 17110–17114. <https://doi.org/10.1002/anie.201810462> PMID: 30395378
28. Hays JM, Cafiso DS, Kasson PM (2019) Hybrid Refinement of Heterogeneous Conformational Ensembles Using Spectroscopic Data. *J Phys Chem Lett* 3410–3414. <https://doi.org/10.1021/acs.jpcclett.9b01407> PMID: 31181934
29. Do Kwon Y, Pancera M, Acharya P, Georgiev IS, Crooks ET, et al. (2015) Crystal structure, conformational fixation and entry-related interactions of mature ligand-free HIV-1 Env. *Nature Structural & Molecular Biology* 22: 522–531.
30. Stadtmueller BM, Bridges MD, Dam K-M, Lerch MT, Huey-Tubman KE, et al. (2018) DEER Spectroscopy Measurements Reveal Multiple Conformations of HIV-1 SOSIP Envelopes that Show Similarities with Envelopes on Native Virions. *Immunity* 49: 235–246.e234. <https://doi.org/10.1016/j.immuni.2018.06.017> PMID: 30076100
31. Hadjidoukas PE, Bartezzaghi A, Scheidegger F, Istrate R, Bekas C, et al. (2020) torcpy: Supporting task parallelism in Python. *SoftwareX* 12: 100517.