

FULL PAPER

Open Chemistry, JupyterLab, REST, and quantum chemistry

Marcus D. Hanwell^{1,2}  | Chris Harris¹  | Alessandro Genova¹  |
Mojtaba Haghaghatlari³  | Muammar El Khatib⁴  | Patrick Avery¹  |
Johannes Hachmann^{3,5,6}  | Wibe Albert de Jong⁴ 

¹Kitware Inc, Clifton Park, New York

²National Synchrotron Light Source II (NSLS-II), Brookhaven National Laboratory, Upton, New York

³Department of Chemical and Biological Engineering, University at Buffalo, SUNY, Buffalo, New York

⁴Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California

⁵Computational and Data-Enabled Science and Engineering Graduate Program, University at Buffalo, SUNY, Buffalo, New York

⁶New York State Center of Excellence in Materials Informatics, Buffalo, New York

Correspondence

Marcus D. Hanwell, National Synchrotron Light Source II (NSLS-II), Brookhaven National Laboratory, Upton, NY.
Email: mhanwell@bnl.gov

Funding information

National Science Foundation CAREER Program, Grant/Award Number: OAC-1751161; New York State Center of Excellence in Materials Informatics, Grant/Award Numbers: CMI-1140384, CMI-1148092; NSF Molecular Sciences Software Institute, Grant/Award Numbers: ACI-1547580, ACI-1547580-479590; Office of Science, Grant/Award Numbers: DE-SC0012704, DE-SC0017193

Abstract

Quantum chemistry must evolve if it wants to fully leverage the benefits of the internet age, where the worldwide web offers a vast tapestry of tools that enable users to communicate and interact with complex data at the speed and convenience of a button press. The Open Chemistry project has developed an open-source framework that offers an end-to-end solution for producing, sharing, and visualizing quantum chemical data interactively on the web using an array of modern tools and approaches. These tools build on some of the best open-source community projects such as Jupyter for interactive online notebooks, coupled with 3D accelerated visualization, state-of-the-art computational chemistry codes including NWChem and Psi4, and emerging machine learning and data mining tools such as ChemML and ANI. They offer flexible formats to import and export data, along with approaches to compare computational and experimental data.

KEY WORDS

automation, computational chemistry, open infrastructure, reproducibility, visualization

1 | INTRODUCTION

We live in a golden age where data can be transmitted at light speed to almost anywhere on the planet. Yet modern scientific publishing, and indeed much of the research process, languishes in a time where printing presses rule. It is long past time that these processes and publication practices were modernized—there is no need to primarily rely on PDF tables and static images in an age where advanced JavaScript applications

Referee reports and author responses are made public under through the "Peer-Review Details" section.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *International Journal of Quantum Chemistry* published by Wiley Periodicals LLC.

can run on cell phone browsers using hardware-accelerated rendering in 2D and 3D and where augmented and virtual reality is emerging as a powerful, immersive alternative to more traditional approaches.

It is time for a paradigm shift that embraces fully digital workflows where all steps of research projects are documented, one in which modern tools can be used to “show all steps of our research process.” To that end, the authors of this paper have been developing an ambitious open-source platform to realize the dream of open science in a fully connected digital age. PDFs and flat images are relics of the past, attempts to map the limitations of the printing press to the digital world. It is time to push the scientific research enterprise forward and drive innovation in software to a point where interactive data outputs can be published for the whole world to see. We must look at the technologies available in modern web browsers to offer seamless publication of data, code, and written results by leveraging the best open-source software, extending this software where necessary to offer specialized views, entry points, and direct access to raw, partly processed data, to tools that extract the key information from the data, and even notes, linked from figures in published works.

The modern scientific workflow takes place increasingly on a computer, or distributed set of computers/nodes/images. This needs to be reflected in, and embraced throughout, the dissemination of research. Computational chemistry by its very nature is a prime example of this, and we have the tools to document all steps in producing, processing/analyzing data, visualizing data, and linking that with the final publication. More and more research is multidisciplinary in nature, with experimentalists, theoreticians, and computational scientists working closely together in their pursuit of scientific discoveries that advance society. Integration of diverse data resources and analysis tools into complex research workflows can be made much more accessible with modern technologies. In this paper, we will discuss our approach to the handling of data and metadata, storing them on a flexible data server platform, representing rich data that can be exported and integrating computational chemistry data generators (quantum chemistry and machine learning codes) with experimental data and approaches to generate and access the data through the web using browser-based and desktop-based tools.

The project described has been developed in the open on GitHub using Open Source Initiative (OSI) approved open source licenses,^[1] primarily the three-clause Berkeley Software Distribution (BSD) license.^[2] Issues were used to track new features to be developed and bugs in existing functionality. Pull requests offer a number of commits in a logical branch to add new features of fixing bugs; where possible, these feature automated testing and line-by-line review of code additions/removals using the web interface. Focused software repositories were used for each component, such as the server code in Python,^[3] the web client code in JavaScript/TypeScript, and the deployment code/image build instructions. A coordinating repository was also used to discuss high-level development of the platform within the “OpenChemistry”^[4] organization on GitHub.

The GitHub platform offers access control for registered users, including individual and team-level access control. Integrations with a number of other platforms create a rich ecosystem where continuous integration jobs can be run for pull requests, when the main development branches are merged and when release tags are created. Commits and release tags can be cryptographically signed, and releases offer snapshots of the software at a particular point in time. Pull requests can be made by anyone, code can be reviewed by anyone, but only approved developers can merge code to the main development branches. This offers a level of transparency and accessibility that is often the norm in open-source software development but is not so common when developing code for scientific research. All code developed, along with all previous versions of the code, remain available using these standard technologies and platforms. More recently, this has been used to document the data/metadata formats, the build process for images, and the deployment of services.

2 | HANDLING DATA AND METADATA

If things work well, the data and metadata standards used will rarely be seen by end users, but they will enable the storage of data within server instances and the seamless transfer of data into or out of these systems. Flexible standards must have formats that can be used from a number of programming languages ranging from compiled languages such as C, C++, and Fortran on supercomputers/high-performance computing resources to perform quantum chemistry calculations to interpreted languages such as Python for data analysis and JavaScript/TypeScript in web frontends or C/C++ in desktop applications. They should also be suited to the needs of persistent data publication independent of any particular database technology or programming language.

These considerations led to the choice to using JavaScript Object Notation (JSON)^[5] as a core standard for data and metadata, with a view of using related technologies such as JSON-LD (JSON for Linked Data)^[6] where appropriate. Large data are preferably stored in binary formats, which is where HDF5^[7] was seen as one strong contender, and more recently, MessagePack^[8] has gained traction due to its JSON-like structure and wide language support thanks to its simple binary specification. JSON also lends itself for use in BSON^[9] and jsonb^[10] two binary JSON specifications used in MongoDB^[11] and PostgreSQL^[12] respectively.

In order to effectively share chemical data, we must establish data and metadata standards capable of representing everything we wish to communicate. Furthermore, it must offer routes to extending the standards without causing breakage and churn in existing data. Ideally, communities should form to establish best practices and propagate this to a number of codes to prove viability and offer a body of work that demonstrates the advantages of the approaches shown. Several existing formats have been used, such as XYZ,^[13] SDF, XML-based formats such as

CML,^[14–17] and—more recently—JSON-based formats such as Chemical JSON.^[18] Open Babel,^[19] RDKit,^[20] cclib^[21], and ASE^[22] offer conversion between these and many other formats, aiding in the normalization of data from different sources when ingesting data.

The JSON specification was chosen for its wide language support, with fast simple parsers available in every language considered, including those outlined previously: Python, C, C++, Fortran, JavaScript, and TypeScript. It is not a part of the core language, but it has become established as a simple container format that can easily map to data structures in all of these languages, for example, Python dictionaries, C++ maps, and—with a little work—the array and some shape information easily maps to NumPy arrays^[23] in Python or one of a number of structures in C++ such as the Eigen^[24] matrix structures or Vector3 for position information. The text format is easy to inspect for casual users, and there are clear routes to migrate to very similar binary containers in the future as outlined earlier. Due to the simplicity of the container format, fast parsers can be used, and the loss of precision is of a secondary concern for visualization. It will, of course, be very important if exchanging data between quantum chemistry codes and would warrant further consideration there.

The block below shows the Chemical JSON^[25] representation of a water molecule. It contains a unique key of “chemicalJson” that can be used to identify the convention being used, along with a numeric version identifier of “1.” The first iteration used “chemical json” with a version of “0” before the decision was taken to switch to camelCase to simplify client code. The “atoms” object contains “elements” identifying the atoms by their atomic number. The “coords” object contains a “3d” array that has 3N numbers representing the Cartesian coordinates in Angstroms. There is also a “bonds” object containing connectivity of atoms (using a zero-based index) and the bond orders.

```
{
  "chemicalJson": 1,
  "atoms": {
    "elements": { "number": [ 8, 1, 1 ] },
    "coords": {
      "3d": [ 0.0000, 0.0000, -0.0685,
              0.0000, -0.7890, 0.5437,
              0.0000, 0.7890, 0.5437 ]
    }
  },
  "bonds": {
    "connections": {
      "index": [ 0, 1,
                 0, 2 ]
    },
    "order": [ 1, 1 ]
  }
}
```

The equivalent molecule can be seen using the QC JSON Schema,^[26] where the molecule geometry is expressed in Bohr, and the atom type uses the string symbols rather than atomic numbers. The QC JSON Schema is at an earlier stage of development, with some pull requests under review adding support for bonding, basis sets, and more. It also has a fairly comprehensive set of possibly calculated properties expressing the numeric results of calculations such as the energy, number of iterations, etc. The authors have been collaborating with MolSSI developers to develop the QC JSON Schema and help drive its standardization in the hopes of wider adoption in the community. Recent work has added support for basis sets into the standard.

```
{
  "schema_name": "qc_schema_output",
  "schema_version": 1,
  "molecule": {
    "geometry": [ 0.0, 0.0, -0.1294,
                  0.0, -1.4941, 1.0274,
                  0.0, 1.4941, 1.0274],
    "symbols": [ "O", "H", "H" ],
    "connectivity": [ [ 0, 1, 1 ], [ 0, 2, 1 ] ]
  }
}
```

The Chemical JSON schema has been in development and use since 2011, focusing on the needs of the Avogadro chemical editor^[27] and the associated Open Chemistry open-source projects. It has, as such, been developed to enable the visualization and analysis of chemical systems, supporting elements such as the basis set, molecular orbitals, and other electronic structure calculation outputs of the system to offer a flexible format. This has been extended^[18] to support, for example, multiple steps in a reaction coordinate calculation (one basis set, multiple 3D geometries, and molecular orbital matrices). There is also support for vibrational modes and 3D cubes representing calculated molecular orbitals. An extended example for water, available in the Chemical JSON repository, shows a basis set, molecular orbitals, and related data from a Psi4^[28] energy calculation.

The layout is designed to be compact and follow the convention of using linear arrays, where feasible, to store sets of values. It also continues the index-based extension of atom properties with the `shellToAtomMap`. Chemical JSON is not meant to be used directly but read into an appropriate data structure in the programming language of choice; this means that simple index-based look ups can be performed, but it may not be as human-readable as desired. It will map to more efficient use of binary analogs such as BSON, jsonb, and MessagePack and slightly less analogous HDF5 storage formats. The project has validated the approach and uses this format in C++, Python, and JavaScript using in-memory representations, translation to other tools, and storage in MongoDB on the data server.

The data standards developed for the platform reuse community best practices, such as the work carried out when developing the IUPAC International Chemical Identifier,^[29] InChI key (a hashing standard built on InChI), and Simplified Molecular-Input Line-Entry System (SMILES).^[30] These standards provide simplified, unambiguous string representations of chemical structures (describing both elemental makeup and connectivity) that can be used as common identifiers. Generating these and augmenting other data enable simple linking of data to existing databases such as PubChem,^[31] ChemSpider,^[32] Wikipedia, and more.

The Chemical JSON format forms the core of the data model in the platform described. It has an informal standardization process that would benefit from more discipline in the future. The development team worked with MolSSI from an early point in their founding to encourage the standardization of a flexible JSON-based format and has actively participated in discussions. The Chemical JSON format is a useful internal representation, the documentation of which should be improved, but the ultimate goal is to offer seamless export to other formats such as those supported by current libraries in existing formats and new ones such as QCSchema. It is clear that, for large data, binary formats will be essential, and a number of them can be mapped from JSON, such as MessagePack and HDF5, as they offer a familiar container structure.

3 | FLEXIBLE DATA SERVER PLATFORM

At the core of the project sits a RESTful data server, with a well-defined set of programming language-agnostic interfaces. These endpoints take the form of web URLs using the HTTP protocol underpinning the web. Simple queries are encoded in the address, with responses often returned as JSON. The approach is standardized and widely used, with software that can automatically generate documentation of the interfaces. The data server provides persistence of data, metadata, and workflow execution to clients, such as the Jupyter notebooks,^[33,34] desktop tools, command line clients, and other consumers of the platform.

The Girder project^[35] was used as the central component because it already offered a large amount of the more general functionality, such as user authentication, file/directory storage, upload, download, and access control. It also offers a structured way to augment its functionality using plugins developed in Python. An early version of the data server was described in a previous publication along with the Python plugins developed.^[18] This has since evolved into a more ambitious project, but the data server concept described there is still at the core of the design. Existing capabilities in the framework, such as authentication, access control, computational job control, and integration with high-performance computing (HPC) queuing systems, offered clear advantages in this project.

The data server provides the central point of integration; it is data-centric and avoids direct generation of any HTML/visualization on the server. The RESTful application programming interfaces (APIs) are consumed by different clients, principally the single-page web application, JupyterLab^[36] web widgets, and Python-based Jupyter kernels. The RESTful APIs are completely programming language-agnostic and can be consumed by others in the community either directly or by adapting reference examples developed as part of the project. We have developed a Python command-line client, Jupyter Python kernel extension, web widgets, and an Avogadro plugin to search and display data from the server.

The server platform forms the central component of the open software platform, providing persistence, authentication, metadata generation, coordination with microservices, and generation of new data. It has been developed to use a plugin to extend the server component with specific functionality, and critical services have been moved into specialized microservices such as an Open Babel microservice for file translation, 3D coordinate generation, chemical utility functions, and an Avogadro microservice offering some file services and molecular orbital cube generation. They provide reference interfaces that can be improved or replaced as the need arises in future projects. In addition, the use of microservices enables the decoupling of services so that more compute-intensive tasks can be scaled separately, for example, 3D coordinate generation or cube generation (Figure 1).

The high-level software architecture was developed to enable researchers to pick and choose the most appropriate components or to use the entire platform as a turnkey solution. The complete solution requires administrative privileges in order to run JupyterHub, but it can be run in a cutdown form where the Jupyter/JupyterHub integration is disabled, but prompts are offered so that users can run their own Jupyter notebook and connect to the data server. This works for most things but removes the ability to store notebooks and to link from calculations to the notebooks that generated them.

As the project has evolved, it has become clear that, while a complete solution provides a broader overview of the data generation, analysis, and publication process, it can be too onerous. It may also be preferable to export as much of the process as possible to services that are able to provide greater guarantees of longevity, that is, months, years, decades, centuries depending upon need. This led to adapting the project to serve in different roles where active data generation benefits greatly from a running data server. While it may not be feasible to keep a project server up and running once projects are complete, there are a number of community services that can be employed to store static assets that clearly document the path from data generation to final figures such as Zenodo,^[37] Binder^[38] coupled with GitHub,^[39] and Figshare.^[40] To this end, the data

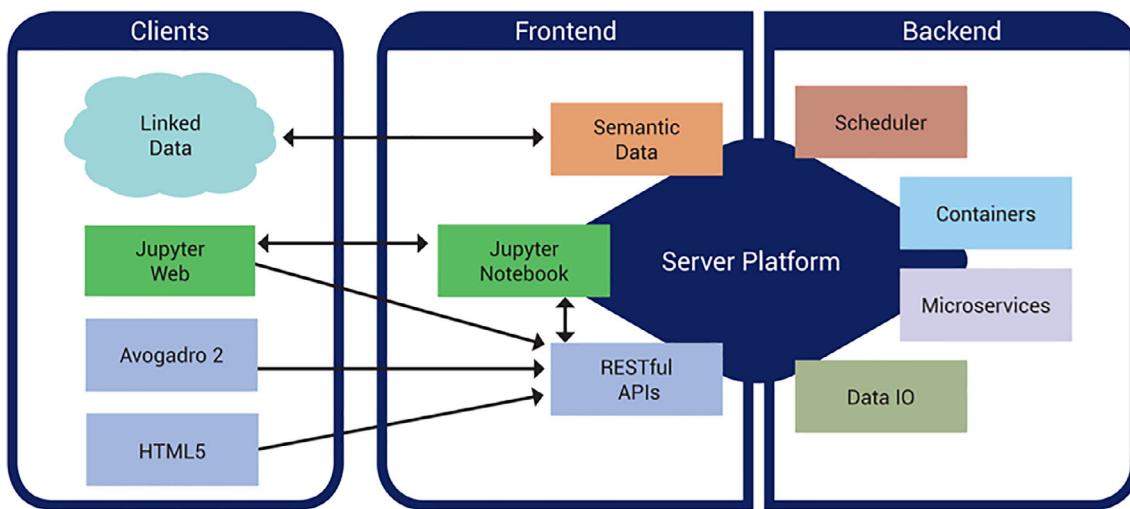


FIGURE 1 Overview of the platform's architecture, with the data server platform providing core services

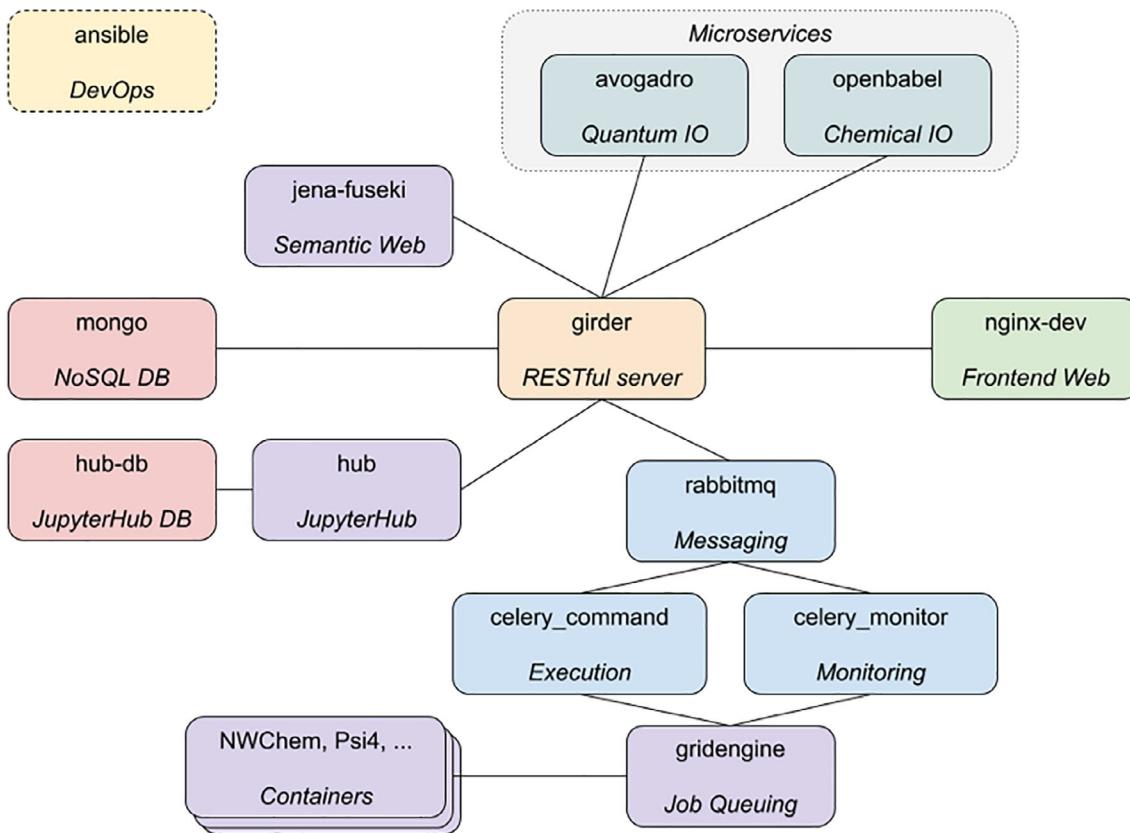


FIGURE 2 The containers for the complete platform; the “girder” container is the data server

and metadata standards enable the export of individual calculations, and the notebooks can be archived. The Binder export discussed later shows one example of moving to a static repository containing data and a notebook; the process is currently manual, and we are working on improved automation (Figure 2).

The bulk of the code for the extensions developed for the Girder server are in the server repository,^[41] with deployment code in a separate deployment repository.^[42] They define the chemistry-specific endpoints and hook into the documentation capabilities provided by Girder and use MongoDB to persist data/metadata. Much more time could be spent on other aspects of the project by using a framework developed for scientific data, which also reuses many open-source projects to provide the RESTful server, database access, schema validation, etc. Figure 2 shows the containers that make up the platform, as deployed using Ansible, with the “girder” container hosting the data server and providing integration of other components.

```

class Molecule(Resource):
    def __init__(self):
        super(Molecule, self).__init__()
        self.resourceName = 'molecules'
        self.route('GET', (), self.find)
        self.route('GET', ('inchikey', ':inchikey'), self.find_inchikey)

    @access.public
    def find(self, params):
        return MoleculeModel().find_molecule(params)

    @access.public
    def find_inchikey(self, inchikey, params):
        mol = MoleculeModel().find_inchikey(inchikey)
        if not mol:
            raise RestException('Molecule not found.', code=404)
        return self._clean(mol)

```

The code snippet above is taken from “girder/molecules/molecules/molecule.py” in the server repository and shows an abbreviated version of the Molecule endpoint. The server can be extended with new endpoints such as this, where routes are added, and the capabilities of the underlying model can be exposed. Decorators define the access level required, for example, “public” in this case for unauthenticated access or “user” to indicate that a user must be logged in. The model interacts with the database to find, update, delete, and amend records, providing an object-relational mapping (ORM) to map from the RESTful API to the underlying data representation.

```

class Molecule(AccessControlledModel):
    def __init__(self):
        super(Molecule, self).__init__()
        self.ensureIndex('properties.formula')
        self.ensureIndex('inchikey')

    def initialize(self):
        self.name = 'molecules'

    def find_molecule(self, search = None):
        limit, offset, sort = parse_pagination_params(search)

        if search is None:
            search = {}

        query = {}
        if 'queryString' in search:
            # queryString takes precedence over all other search params
            query_string = search['queryString']
            try:
                query = mol_query.to_mongo_query(query_string)
            except mol_query.InvalidQuery:
                raise RestException('Invalid query', 400)
        elif search:
            # If the search dict is not empty, perform a search
            if 'name' in search:
                query['name'] = { '$regex': '^' + search['name'],
                                  '$options': 'i' }
            if 'inchi' in search:
                query['inchi'] = search['inchi']
            if 'inchikey' in search:
                query['inchikey'] = search['inchikey']
            if 'smiles' in search:
                # Make sure it is canonical before searching
                query['smiles'] = openbabel.to_smiles(search['smiles'], 'smi')
            if 'formula' in search:
                formula_regex = re.compile('^%s$' % search['formula'],
                                           re.IGNORECASE)
                query['properties.formula'] = formula_regex
            if 'creatorId' in search:
                query['creatorId'] = ObjectId(search['creatorId'])

        fields = [
            'inchikey',
            'smiles',
            'properties',
            'name'
        ]

        cursor = self.find(query, fields=fields, limit=limit, offset=offset,
                           sort=sort)

        num_matches = cursor.collection.count_documents(query)

        mols = [x for x in cursor]
        return search_results_dict(mols, num_matches, limit, offset, sort)

    def find_inchikey(self, inchikey):
        query = { 'inchikey': inchikey }
        mol = self.findOne(query)
        return mol

```

A truncated example of the corresponding model is shown above. The `find_molecule` function is quite complex, supporting multiple fields, pagination, etc. These are necessary to support larger result sets where returning all results would take too much space and would not be desirable. It enables clients to display pages and request offsets into the results returned by a search query. The `find_inchikey` is one of the simplest functions where a query is created, and the query is executed with the result returned directly. There is much freedom within the model to add chemistry-specific logic or perform very simple database queries depending on the need. This also introduces a separation of concerns where the model code interacts with the database, whereas the endpoint code deals with issuing the correct response, documenting endpoints, and dispatching to the appropriate parts of the data model.

Having an open-source data server not only offers useful capabilities now, but through the use of Python and by documenting the code using modern approaches, it offers a starting point for future projects that want to extend an already capable data server. The use of microservices works towards making this even easier by providing containers that perform specific tasks that could be reused within other platforms or where their capabilities can be extended and improved within this platform without the need to work within the data server itself. Both microservices were migrated from the server component to simplify the data server and offer greater modularity. The Open Babel and Avogadro microservices were able to separate out longer running tasks from the data server thread, ensuring that the data server is less likely to stall, and the microservices can be scaled more easily.

4 | CONTAINERS FOR CHEMISTRY CODES

The use of software containers has revolutionized how services are deployed on the internet, and science is beginning to reap some of these rewards. At their core, containers such as Docker,^[43] Singularity,^[44] and Shifter^[45] are self-contained Linux-based binary environments that contain everything from Linux system calls to the full call stack of software libraries. This enables the development of a complete build environment from a known base image to packages that should be installed and software repositories to clone, build, and install.

The software images have an entry point that can be called from the host operating system and a flexible set of capabilities to mount file systems, get data in and out, etc. The use of a Dockerfile enables a full description of the base image, build steps, and configuration. The binary images can be uploaded to DockerHub and similar services and downloaded using standard tools developed by the community. The container specification developed for the project offers a standard method of invoking a code, with input directives delivered in a simple JSON format, geometry using Chemical JSON input, and output translated to a simple JSON output format or Chemical JSON. All relevant files are in the "docker" folder in the deployment repository,^[42] with a folder for each of the containers.

Docker was the first widely adopted container standard, and it acts as a base for the images used in this project. It provides network isolation, which is not necessary for most computational chemistry tasks or HPC tasks in general. The Singularity project was founded to address this issue, concentrating more on the execution of binary programs without the need for privileged system access that is often not available on HPC clusters/supercomputers. Shifter offers a solution developed at NERSC for executing containers within their supercomputing environment. If Singularity was available on all environments, that would probably be the only container format used, but at the present time, this is not the case, and so, multiple container types are supported depending upon deployment requirements. The Docker images are used to generate the Singularity and Shifter containers so that it is only necessary to maintain a single image for each code.

The containers developed to support the project go beyond providing a useful delivery mechanism for executable code. They reduce the barriers to collaboration and enable binary containers to be used as reference implementations that can easily move between host systems. Private container registries can host images containing proprietary software, and automated builds can enable sharing codes that have complex dependencies with build systems that can hamper seamless testing of changes. As supercomputers, HPC, and cloud environments evolve, the utility of containers will only make these advantages clearer to the wider community.

4.1 | Quantum chemistry codes

Thanks to the container infrastructure described above, it is easy to extend the platform to perform calculations using new computational chemistry packages. Assuming the general flow follows that of an existing code, the Dockerfile and Python driver can be used as a starting basis. If the code is packaged using a Linux distribution or a Python packaging repository, they can simply be installed as part of the image build process. The Python driver script provides a basis for adapting input data and directives to input files for the code and translating output data to a suitable form that can be consumed by the platform. Once these pieces have been implemented, the container can be tested on the command line before attempting to use it on the server side.

The NWChem^[46] and Psi4^[28] packages were used as our first computational chemistry codes due to their popularity, open-source licensing, and the fact that they offer similar features at a high level with very different implementations. These images serve as an example for others interested in adding new codes to the platform. The NWChem and Psi4 images support a similar set of features:

- Single-point energy, geometry optimization, and frequency calculations
- Basis set selection
- Theory-level (Hartree Fock, DFT, MP2) and exchange-correlation functional selection
- Charge and multiplicity

This only represents a select subset of what these quantum chemistry codes are capable of. The features exposed by the images will be extended in the future, prioritizing things that are relevant to the needs of our platform's users. The input mechanism through the notebooks and API accepts arguments that are passed through to the driver script, making it easy to add custom arguments to new containers. Container names consist of an organization, name, and version currently pulled from DockerHub, making it easy to add custom versions and fetch them.

```
FROM conda/miniconda3:latest

RUN apt-get update && \
    apt-get install -y gcc git && \
    rm -rf /var/lib/apt/lists/*

RUN conda update --all && \
    conda install psi4=1.2 psi4-rt=1.2 python=3.6 -c psi4

COPY requirements.txt /app/
RUN pip install -r /app/requirements.txt

ENV LC_ALL=C.UTF-8
ENV LANG=C.UTF-8

COPY src/* /app/

ENTRYPOINT ["python", "/app/main.py"]
```

The Psi4 container is one of the simplest; it uses the Python 3 Conda-based image as a base and then installs Psi4 from the Conda package provided by the Psi4 project. The Python driver script is then set as the entry point and acts as the gatekeeper for getting data into and out of the container. When new releases are made, updated images can be made. The NWChem container is a little more complex as it clones the git repository of a fork of the NWChem codebase that supports JSON output; it then uses Avogadro to translate that output to Chemical JSON. An example of running one of the containers is shown below.

```
docker run -v $(pwd):/data openchemistry/psi4:latest -g /data/geometry.xyz -p /data/parameters.json -o
/data/out.cjson -s /data/scratch
```

The JSON used as input follows a simple format that is translated by the driver script into the appropriate input parameters. One of the simplest examples is shown below, which can be used with an input geometry to drive the container on the command line for testing purposes (or even use within other platforms).

```
{
  "theory": "hf",
  "task": "energy",
  "basis": "6-31g",
  "charge": 0,
  "multiplicity": 1
}
```

These input parameters have been added to the output Chemical JSON in order to make it easier to see what calculation parameters were used. The keys are common across the codes already present, but it is up to the driver script to decide how to interpret them. The "psi4" in the image name can be substituted with "nwchem," keeping all other arguments/JSON input the same, and an NWChem calculation will be executed. Existing containers aim for consistency, but casting an eye to the future, the platform will support the development of new methods or the extension of existing containers beyond their current capabilities.

4.2 | Machine learning

From the point of view of the platform, there is no difference between running a quantum chemistry or a machine learning code. They both take as input a molecule and some parameters and return a conforming output. The main difference is that machine learning codes may not require

the existence of a 3D structure and can still operate when only string-based formats, such as InChI or SMILES, are available. The server was extended to support the addition of molecules that have no 3D structure so that machine learning can be used more effectively within the platform as 3D coordinate generation can be quite expensive when considering high-throughput models.

If many small calculations (especially machine learning calculations on line format structures) are to be performed on a large set of molecules, it can be very inefficient to have a task monitor and a Docker container for each individual calculation. For this reason, the user is able to submit a batch of calculations as a single task. In this task, there is only one task monitor and one Docker container used, and the Docker container is given the freedom to choose how to run the calculations (one at a time, one per core, etc.). In addition, each calculation is checked to see if it has been performed before. If so, it will skip that calculation and use the output already present in the database.

Images were developed for ChemML^[47,48] and ANI.^[49] These represent two very different classes of machine learning models. The illustrated model developed by the ChemML package takes in a SMILES string and predicts corresponding properties that are output as key-value pairs. ANI behaves more like a typical quantum chemistry code that takes in the 3D chemical structure; can perform an energy calculation, geometry optimization, or frequency calculation; and outputs (potentially updated) a 3D chemical structure. It does not require as much input as is typical for other codes, and it does not output electronic structure. Both required further generalization of our treatment of tasks, in order to adapt to their unique requirements, and added examples of quite distinct types of the model, paving the way to easily add more using similar patterns.

ChemML is a general-purpose machine learning package developed in Python for the analysis and modeling of chemical and materials data. Although ChemML provides essential tools to develop and train machine learning models in general, we aim to deploy trained models and broaden their utilization in the chemistry community. Our long-term goal is to compile a collection of predictive models for certain properties that can be used as alternatives to corresponding physics-based modeling or simulation approaches. To start, we developed a deep learning model for prediction of the refractive index values of organic compounds and made it available as a Docker container.

The model takes as input the SMILES string of an organic molecule and returns as output a few quantities of interest on which the model has been trained. The predictions are available almost instantaneously (unlike a physics-based model via the Lorentz-Lorenz equation parametrized by inputs from quantum chemistry and molecular dynamics calculations). The results of the model are comparable with those of other data-derived prediction models in terms of diversity of molecular candidates and the accuracy of predictions. The current implementation enables the user to retrain each model for a better or more generalized prediction power. Moreover, a trained model can leverage other relevant ML models through the concept of transfer learning design methodologies. Our near-term plans for future development include the addition of models trained on other datasets and for other material properties.

The ANI image used the PyTorch^[50] implementation of the ANI potentials—TorchANI.^[51] We make use of the algorithms in atomic simulation environment (ASE) to drive tasks such as geometry optimization and normal mode calculations. The TorchANI image features the *ani-1x* and *ani-1ccx* optimized potentials that can be used to generate single-point energies, perform geometry optimizations, and compute Hessians for frequency calculations. This makes the ANI image look quite similar to the quantum chemistry codes, but no basis set selection is necessary, and there is only support for a limited set of chemical elements.

For the elements that are supported, ANI offers superior performance to geometry optimizations performed using conventional quantum chemistry codes, often trained on a higher level of theory than would typically be used. It also offers performance comparable with molecular mechanics codes without the need for atom typing and other parameterizations. This often makes a geometry optimization desirable before using other codes assuming the chemical structure only uses the elemental set supported. It is also possible to use other techniques and compare the output in order to validate its use for particular chemical systems.

5 | GENERATING AND ACCESSING DATA ON THE WEB

One of the primary goals of this project is to make computational chemistry tools more accessible and to enhance their use within a collaborative online environment where data can be generated within the web browser. This does not preclude the use of existing tools, but it provides an alternative way of accessing tools that have seen significant investment using one of the most pervasive mediums in the modern world. The use of web-based tools makes the platform accessible on low-powered phones to the most powerful desktop workstations. The platform itself can provide access to calculations already run almost instantaneously while enabling authenticated users to generate new data.

The previous section describes the nongraphical server-side capabilities developed for the software platform. These are necessary tools, but they do not offer an approachable user interface for generating, interacting with, analyzing, and visualizing data. This section discusses a number of capabilities developed to offer web-based interfaces to look at data, search it, and categorize it, and in many cases, these components can be used independent of the complete platform. The high-level web technologies used include JavaScript and TypeScript for code development, WebGL for 3D visualization, and React as a framework for exposing elements in web applications.

5.1 | Web widgets

A set of reusable widgets have been developed that can be embedded in any web environment, from a React-,^[52] Angular-,^[53] or Vue^[54]-based single-page web application to a JupyterLab extension or a static HTML page. They were developed in the TypeScript language using the Stencil library.^[55] Once they have been transpiled into JavaScript, they become standard web components that can be used just like any other HTML tag, such as a `<div/>` or a ``.

The core widgets are exposed using the `<oc-molecule-moljs/>` and `<oc-molecule-vtkjs/>` tags. The widgets have a common interface and can be used interchangeably. They take Chemical JSON as input, and a set of parameters can be used to tweak the visualization (such as the isosurface value, the active normal mode, or the molecule representation type). They render a three-dimensional representation of the molecule that the user can interact with using the mouse or touch-based interaction, such as the 3Dmol.js-based widget shown in Figure 3. The difference between them is that `<oc-molecule-moljs/>` uses 3Dmol.js^[56] to render the 3D scene, while `<oc-molecule-vtkjs/>` uses vtk.js.^[57]

The major features of the widgets are the display of the molecular structure and connectivity, animation of normal modes, and the display of isosurface data from a cube of data (typically molecular orbital data from a quantum chemical calculation). The vtk.js-based widget can display the crystal unit cell and offers volume rendering of cube data, shown in Figure 4, in addition to isosurface rendering.

The `<oc-molecule/>` tag wraps the core `<oc-molecule-moljs/>` and `<oc-molecule-vtkjs/>` and adds a menu so that the user can easily change the visualization parameters. The menu (click on the three dots in the right top corner of the figure) offers:

- Structure representation (ball and stick, Van der Waals, or custom)
- Selection of the rendering library to use (3Dmol.js or vtk.js)
- Selection of the active normal mode for vibrational mode visualization
- Selection of the isosurface value for molecular orbital visualization
- Selection of the color map used by the volume renderer
- Modification of the opacity function used by the volume renderer

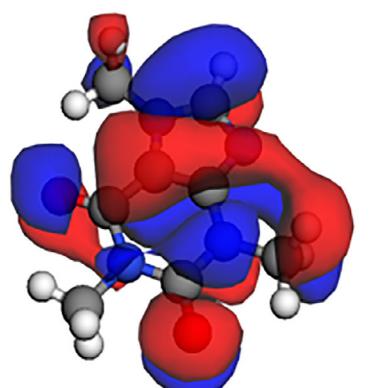


FIGURE 3 An interactively rendered caffeine molecule saved in a Chemical JSON file with a molecular orbital rendered using the 3Dmol.js-based backend. JSON, JavaScript Object Notation

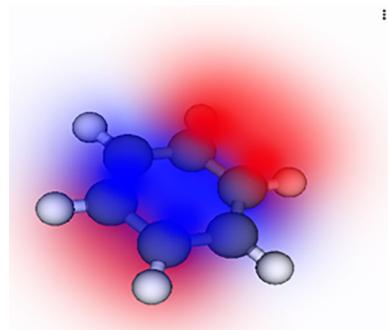


FIGURE 4 An interactively rendered benzene molecule saved in a Chemical JSON file with a molecular orbital rendered using the vtk.js-based backend. JSON, JavaScript Object Notation

5.2 | Jupyter extensions

JupyterLab is the new web frontend for the Jupyter project. It can be extended using a mime renderer extension; this has been used to serve two purposes: displaying interactive visualizations of molecules in the output of a notebook cell and monitoring computational jobs being executed through the platform. The decision was taken at an early point in the project to support JupyterLab extensions exclusively, although it would certainly be possible to add support for the older Jupyter web frontend if there was a big call for it.

The visualization of the molecular structure (see Figures 3 and 4) is accomplished by using the `<koc-molecule/>` component described earlier. The widget is displayed when the user opens a `cjson` file in JupyterLab or when selecting the result of a calculation as the output of a notebook cell. Because of some utility functions in the `openchemistry` Python package, the molecular visualization can be used for arbitrary Avogadro molecule objects. So, no matter the provenance of the molecular structure, if the user is able to convert it into an Avogadro molecule, they will be able to take advantage of this JupyterLab extension to obtain interactive 3D visualizations in the notebook.

A JupyterLab notebook is the preferred environment to dispatch and analyze calculations performed through the platform. For this reason, part of the extension offers a widget to monitor running jobs. The monitor widget (see Figure 5) displays a table with a row per job and information such as the Docker image running the job, the version, and the current status. If a job is selected, it is possible to see log messages and error messages in real time. The job-monitoring widget will also instruct the user to execute the cell again to see the result once the job has completed, or it will display a failure condition if things did not go as planned.

In addition to the JupyterLab web extensions described, an `openchemistry` Python package was developed to extend the Jupyter kernel running on the server. It uses the same RESTful API endpoints that were described in the previous section to query the data server, initiate new calculations, and monitor progress. The module also has the utility functions described so that it can work independently using local data or can take environment variables to run the notebook kernel locally and communicate with a remote data server. This results in a flexible deployment strategy where a fully integrated server can invoke a notebook with all the right environment variables set up, a local (or centralized) notebook server can be invoked with the right environment variables set, or a static notebook/local notebook can avoid using the data server and refer to local data files.

The execution of notebooks and the execution of containers are two of the most demanding aspects of the project in terms of infrastructure and computational resources required. They can both be disabled, resulting in a data server that offers data query, retrieval, and visualization capabilities. They both allow a great deal of freedom to execute code on a shared resource and so must be considered carefully in terms of offering access due to the security considerations and resource constraints.

There is a balance to strike between what is executed in the notebook's kernel, what is executed by the data server (including its micro-services), and what is sent to the queuing system for batch execution. Some things may be duplicated, but we have tried to create a balance based

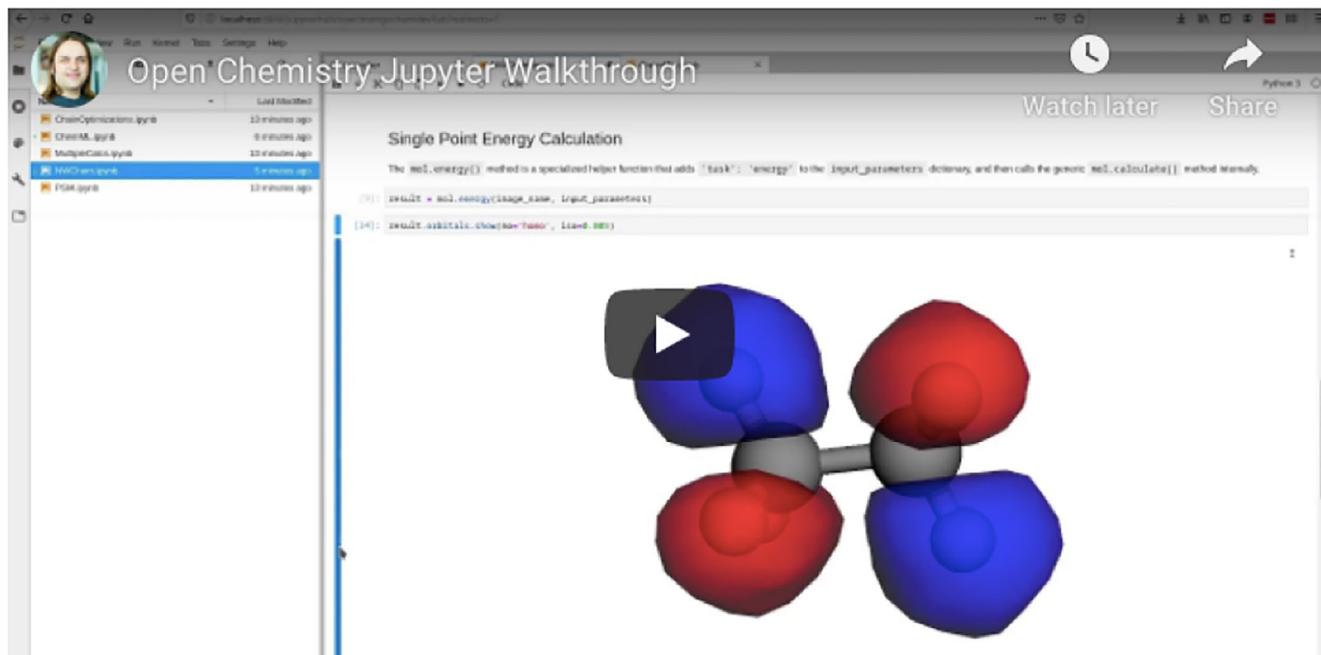


FIGURE 5 A walk-through of the Open Chemistry Jupyter interface generating data through remote calculation with job monitoring and visualizing it. Video also available at <https://youtu.be/7otfejIAoqg>

on the task. When things are executed in the Python kernel they tend to be relatively short-lived tasks, or calls to the data server. The data server is able to cache results so that they can be retrieved in other notebooks, over the web interface, or using other methods without executing each time, such as 3D coordinate generation using Open Babel or molecular orbital calculation. Longer-lived geometry optimizations, executing quantum codes, and running machine learning kernels are usually performed using the queuing system so that they can use HPC resources and images.

5.3 | Export to binder

The Binder project has led to the development of reproducible Jupyter notebooks that can be executed, changed, and examined for free on the web. A GitHub repository can be used to specify the environment used, including what Python modules are needed for the Jupyter kernel and what extensions should be installed for JupyterLab before the notebook opens in the browser. A Docker container is generated on the fly in response to the specification in the repository, and once ready, the repository is mounted inside the container and the notebook opened.

This is clearly a powerful tool that is ready for use by the wider community and offers a path to providing full access to the data and notebooks for anyone without the need to create accounts. The Python module that had been developed to extend the Jupyter kernel was extended to support execution without a data server and packaged using pip so that it could be installed via the mechanism offered by the Binder project. A blog post was written in December, 2018, describing the approach^[58] and showing the Binder repository with example links to data for calculations^[59] (Figure 6).

5.4 | HTML5 web application

The main entry point to the platform is a state-of-the-art web application developed. It offers a landing page, with anonymous access to public data; authentication; and access to notebooks, molecules, calculations, and profile services once authenticated. Both the AWS instance using Google's OAuth provider and the NERSC instance integrate with their authentication provider, negating the need to create new accounts. The profile page can be used to associate Twitter and/or ORCID identifiers with an account and to access a summary of the notebooks, molecules, and calculations authored/initiated by the account's owner.

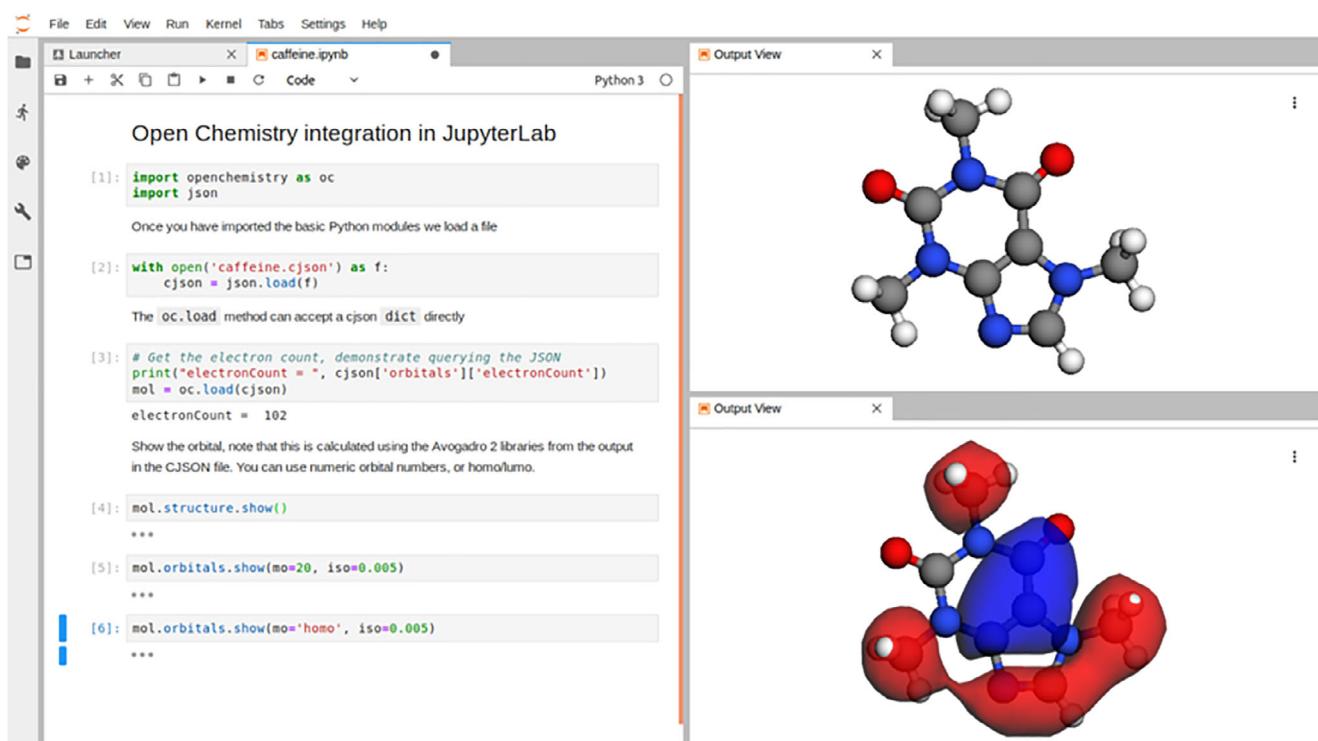


FIGURE 6 The JupyterLab extension allows the visualization of molecular structures in a notebook, with Binder used here to offer a static repository to offer a more interactive exploration of data

Basic search is offered to filter structures and calculations based on the molecule name, InChI, or SMILES. Internal navigation is possible through the cross linking of all calculations associated with a specific molecule and vice versa. In the detailed page of a structure or calculation, there are links to download the data in various formats, such as Chemical JSON, XYZ, SDF, or the raw output of the simulation code.

The application is written using the React + Redux framework in a mixture of JavaScript and TypeScript. The server component is completely decoupled from the application, and communication between the two only happens via the REST API. For the 3D visualization elements, we rely on the reusable `<oc-molecule/>` component that was discussed earlier, with the same look and feel as in the Jupyter notebooks and the static web page examples. They are coupled with the search, linking, and other facilities to provide an interactive view of the data present on the instance being explored.

Figure 7 shows a demonstration of the main web application, demonstrating how anonymous users can access public data, and authenticated users can do more once they supply their credentials. In both cases, it is possible to see the notebook that generated the data, but for anonymous users, that is a read-only view, whereas authenticated users will be able to launch the notebook. The molecular data and calculation results can be downloaded from the web application using buttons on the relevant pages, with permanent URLs identifying specific molecules, calculations, and even the orbital/contour being visualized. These views of the data are designed to be simpler and more focused and use responsive elements so that they can be viewed on mobile telephones up to full desktop systems.

5.5 | Extensions to Python and Avogadro

The server platform has an extensive RESTful API that is used by the web widgets, Jupyter components, and web application. These same interfaces have been reused in a Python client that can interact with the server using an API key, offering simple job monitoring, structure upload, etc. This gives a simple path to upload a number of existing files from the command line in an automated fashion or to augment existing workflows with upload to a data server in order to benefit from the capabilities provided by this platform.

The Avogadro application has made use of web databases for many years. These have been read-only and offered an intuitive way to search and import existing structures. This includes searching the Protein Data Bank or the Cactus structure resolution services,^[60] where users can simply type the name of a chemical to search for and retrieve a structure if found. These capabilities were very popular with users but offered no ability to share results or generate new data that were made available through a chemical data platform.

A plugin has been developed for the Avogadro application that can perform basic searches of the data available. Far more interesting is the ability to use an API key for an authenticated user to upload chemical structures and request calculations. This makes a host of capabilities available to the desktop tool and readily enhances all data passing through the platform by adding additional metadata, search, and discoverability.

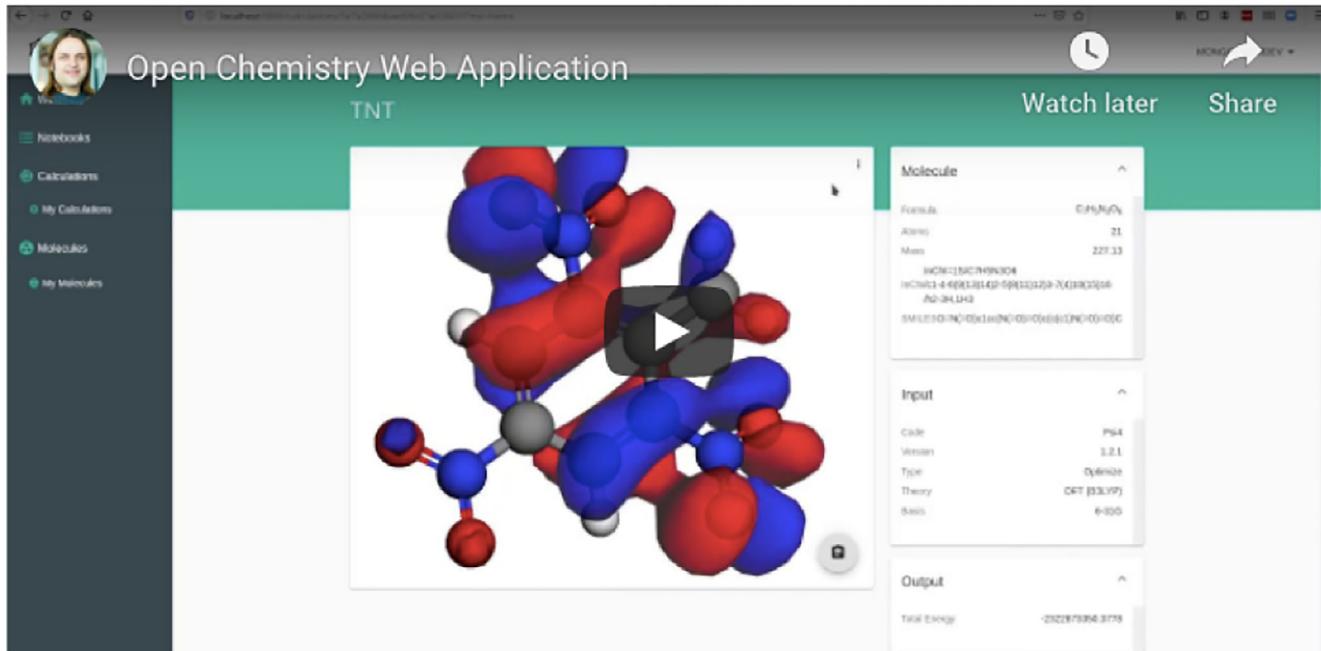


FIGURE 7 The main web application serving the platform showing access to molecule and calculation data generated by notebooks (with links back to the notebooks). Video also available at https://youtu.be/5nUff6O_Gf0

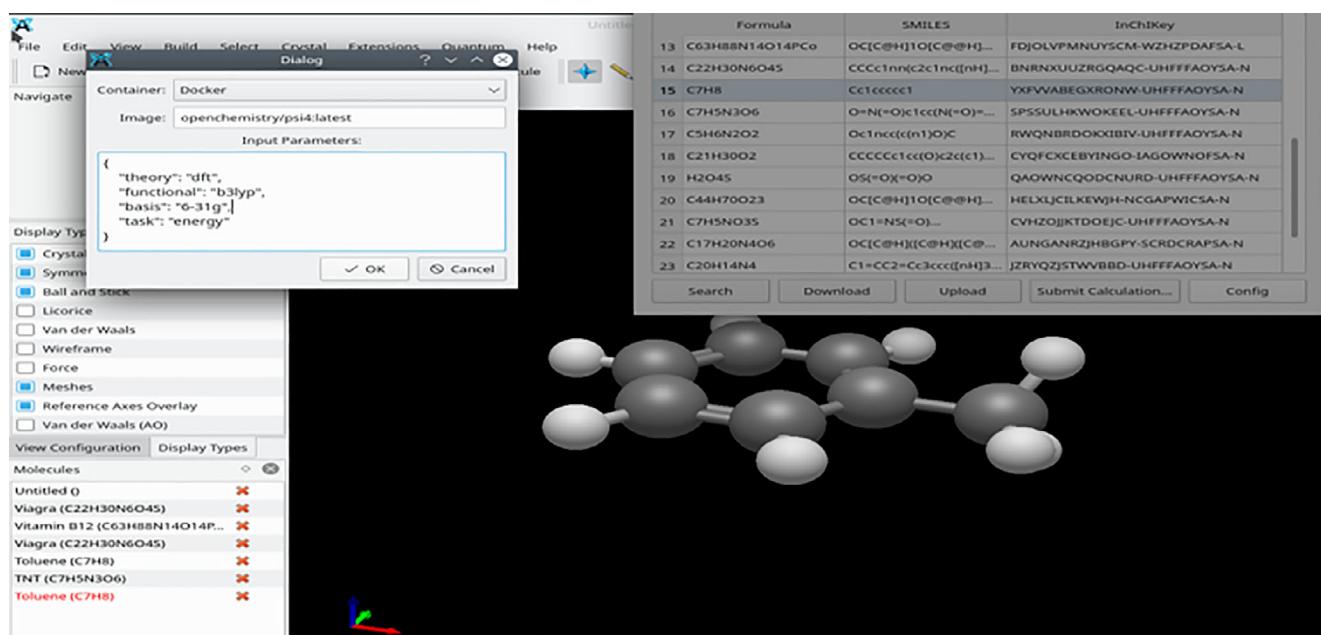
The Avogadro tool can be used to download a structure, or draw one freehand, with basic geometry optimization available locally using Open Babel. The user can then upload the structure to the server, making it available on the web frontend, and they can request a calculation using the plugin interface (Figure 8). The result can be downloaded and visualized when ready, or the user can share a link to the record created and share it more widely on their local intranet or the wider internet if using a publicly hosted instance.

The plugin can switch between instances of the data server, with future enhancements planned to offer a selection of one or more servers, and basic interaction capabilities. This could then search across accessible instances to find existing data and accelerate workflows by making data more easily available. The addition of unique identifiers enables users to retrieve data based on an unambiguous identifier so that coordinates can be edited efficiently, conformations examined, and more before uploading a performing calculation on the selected geometry. These can then be linked to directly using identifiers or searched across sites using InChI, InChI key, SMILES, etc.

6 | INTEGRATION OF EXPERIMENTAL DATA

Over the last decade, the electronic structure community has refined its methods to make them more computationally efficient and accurate. As a result, computational chemistry calculations are beginning to be on par with experimental results.^[61,62] This revolution has opened the path for experimentalists and theoreticians to integrate their data sources and establish more consistent experiments and theories. Providing a platform to access scientific data from both worlds is therefore important and useful. Computational chemistry calculations can be used to examine a hypothesis before carrying out the experiments, saving time and money, and they can serve to interpret the experimental observations. On the other side, access to experimental results enables the improvement of theories and methods.^[63]

The Open Chemistry platform allows the calculation of theoretical spectra using the provided quantum chemistry containers. For instance, the prediction of theoretical infrared spectra (IR) can be routinely computed by quantum chemistry packages. Users can request the calculation of a theoretical IR with a chosen level of theory and visualize it within the notebook. In addition, a keyword argument can be passed to overlap experimental data queried from the NIST Chemistry WebBook database.^[64] In Figure 9, we show a demonstrative example of an overlap between theoretical and experimental IR within a Jupyter Notebook in Open Chemistry. When visually inspecting the spectra to carry out the vibrational band assignment, we identify an intense IR active peak at about 750 cm⁻¹, which corresponds to the characteristic absorption of C—H bending in aromatic compounds. A broad active IR peak is also present in the range of 2800 to 3200 cm⁻¹, which corresponds to C—H stretching. Finally, the typical signals in the range 1500 to 1700 cm⁻¹ are characteristic of the absorption of C=C bending. Having the ability to create band assignments together with experimental IR spectrum and theoretical animated vibrational modes is a feature that eases this task.



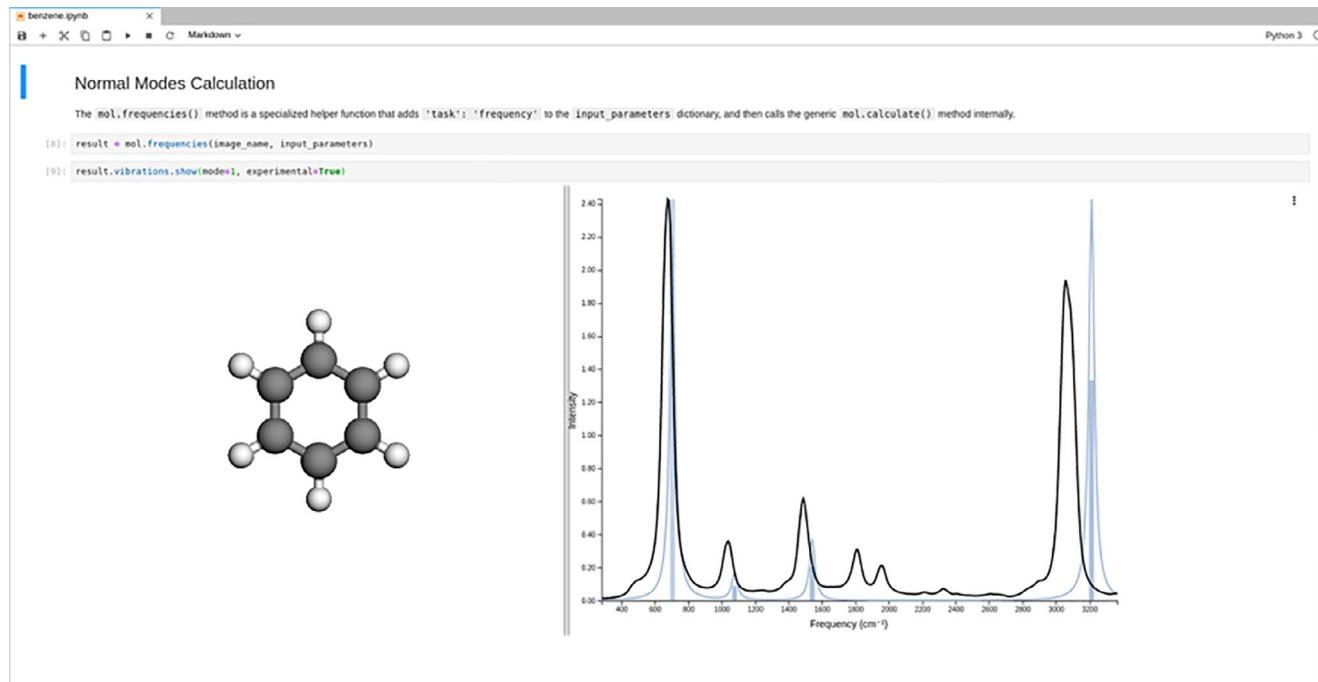


FIGURE 9 Overlap of experimental infrared spectrum of benzene over theoretical results

```
[1]: import openchemistry as oc
import matplotlib.pyplot as plt

[2]: structure_smiles = ['CN1CNCN(C1)C1CCC(C2C1CCCC2)C1NCNCN1', 'SC1CSCC(S1)(S)C1CCCN1', 'C1CCC(NC1)C1SCCC(C1)C1CCNC(C1)C1CCC2C(C1)CCCC2']
mols = []
for smiles in structure_smiles:
    mols.append(oc.import_structure(smiles, gen3d=False))

[3]: image_name = 'openchemistry/chemml/latest'
input_parameters = {}

[4]: results = oc.run_calculations(mols, image_name, input_parameters)

[5]: x = []
y = []
for result in results:
    x.append(result.data()['properties']['polarizability'])
    y.append(result.data()['properties']['refractiveIndex'])

[6]: fig, ax = plt.subplots()
ax.plot(x, y, 'ro')
ax.set(xlabel='polarizability', ylabel='refractive index', title='Machine learning...')
ax.grid()
```

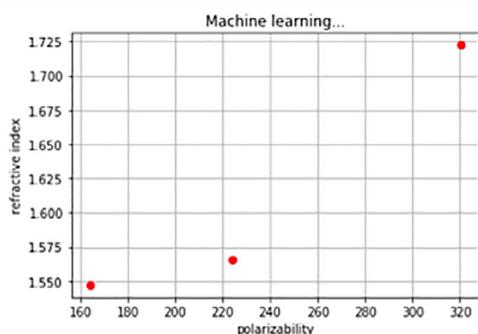


FIGURE 10 A batch of results may be obtained by running multiple calculations in one task, as is done here with ChemML. 3D coordinates are not generated in this example

7 | DATA ANALYTICS AND MINING

Coupling a data server to a computational notebook like Jupyter offers the possibility of performing analytics on the data generated and compiled within the platform. The platform allows us to import classes of compound, run various tools in the batch queuing system, and parse out important information from the outputs. Once a set of results is available, the notebook environment of the platform and its binding to ChemML allows users to analyze and mine the data in its entirety, rather than only referring to individual results. This statistical workup allows for benchmarking insights into the performance (accuracy and timing) of different simulation techniques, algorithms, and codes. More importantly, the datasets can also be mined to generate data-derived machine learning prediction models that can serve as surrogates for the underlying physics-based models as discussed above. The notebook platform also provides plotting functionality and access to other graphic representations. An example of this strategy is shown in the Figure 10, where ChemML results for a list of SMILES are compared in a plot.

8 | REPRODUCIBILITY IN QUANTUM CHEMISTRY

Scientific teams and organizations should embrace reproducible workflows where all data can be exported, shared, and peer reviewed. For computational chemistry, data should include the full flow of data from initial structures to final coordinates, energies, and software versions/binary environments used. The Jupyter and JupyterLab projects have many of the components needed, such as an extensible programming interface, visualization, and analysis of data in a common format, but it lacks specific workflows for quantum chemistry. This project adds those, and couples notebooks with a data server, and uses an extensible data format definition for static export suitable for long-term archiving of results.

The target audience ranges from a quantum chemistry code developer publishing new methods, where this platform enables them to show the input, execution, and output of a development snapshot, to end users running calculations on production code. The ability to specify the organization, container name, and version offers the ability to use known versions of codes and even rerun when fixes are made. It also enhances the peer-review and publication process by offering a full record of what was done computationally, along with the results obtained and a recipe to replicate. As a community, we must move toward the routine publication of all of these steps and consider data standards along with software platforms to reduce the upfront costs of doing so.

The Space Telescope^[65] selected Jupyter as the primary analysis platform for many of the same reasons it has been used in this project. Fields of scientific research must converge on shared platforms where reproducibility is built in and share the cost of improving them, with customization for each field where it makes sense. The platform described can interface with public databases such as PubChem and QCArchive^[66] to import existing data and produce new data with structures suitable for wider dissemination. The data and metadata standards discussed seek to embrace federated storage of data, embracing the goals of FAIR data^[67] to make all data produced more discoverable. The use of established open standards such as InChI, InChI key, and SMILES link data produced in individual instances to the global data commons with minimal ambiguity.

One of the primary challenges in computational chemistry is to develop a software infrastructure capable of executing codes reproducibly such that others can look at every aspect of what you did and build upon it. Even quite subtle details can hamper this, such as tolerances used in convergence criterion or software library versions. The system described in this paper uses Docker, along with conversion to other software containers, to package a code along with all of its dependencies. The driver scripts within these containers use Python to drive the execution of the code from JSON input and to convert the output. This is then layered within an execution framework, database, and user management/file management system. Together, they can be deployed from tagged container versions in order to reproducibly create the same environment.

All of this infrastructure is not without cost, and so, when considering the publication of data and its more permanent dissemination, it is important to consider simpler approaches. This led to the development of integration with Binder and the use of static repositories with exported data. The open specification of the formats and the open-source Python modules mean that anyone can access and process the data with a rather minimal Jupyter deployment. Even without any of the software, the data specification coupled with input specification offers starting points to manually run calculations with the Docker containers locally. Standalone web widgets enable viewing of processed molecular orbitals without any processing component and could be extended further in the future to do more within the web browser's JavaScript environment.

9 | CONCLUSIONS AND OUTLOOK

The development of computational chemistry program packages, with its heavy reliance on legacy code, has so far rarely included the infrastructure to develop a data standard suitable for sharing data with the wider community, input specification, workflow transparency, and reproducibility, as well as analysis and visualization of the outputs. Such an infrastructure should generally be a shared resource, and one that is open, so that it can live beyond a single project, company, or research group. It is necessary for the developers of the different parts of the solution to communicate more effectively and enhance the pathways for getting data in and out of each step. Ultimately, the entire workflow needs to use appropriate data standards and feature validation/peer review to ensure that each code agrees on the form, units, and conventions used.

In this paper, we have demonstrated the utility and potential of a web-based end-to-end software platform to enable automation, transparency, and data-centric research in quantum chemistry along these very lines. The software platform, data, and metadata described have been developed, released, and disseminated under permissive OSI-approved open-source licenses. The embrace of open development practices allows this work to be extended in future projects and use cases. The codes feature a modular design to optimize the ease with which its elements can be reused. We have engaged the community—and will continue to do so—to help develop standards, workflows, and working solutions that showcase what can be achieved by combining modern web browser and data server technology. This is at the core of the projects described and underpins the necessary improvements to better disseminate research. The use of web standards and technologies makes the research more accessible, moving beyond static images to interactive web visualization, web platforms, and snapshots of data that can be interactively queried, built on community efforts such as Binder.

As a community, it is important to embrace open source, open data, open standards, and open access to reproducible research.

ACKNOWLEDGMENTS

This material is based on work supported by the U.S. Department of Energy, Office of Science, Chicago Operations Office, under Award Number DE-SC0017193. This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Marcus D. Hanwell was affiliated with Kitware for the development of the software platform described (Contract No. DE-SC0017193) and made minor corrections to the manuscript after changing affiliation to Brookhaven National Laboratory (Contract No. DE-SC0012704). The development of ChemML is supported by the National Science Foundation (NSF) CAREER program (grant No. OAC-1751161) and the New York State Center of Excellence in Materials Informatics (grants No. CMI-1140384 and CMI-1148092). Mojtaba Haghaghatlari gratefully acknowledges support by Phase-I and Phase-II Software Fellowships (grant No. ACI-1547580-479590) of the NSF Molecular Sciences Software Institute (grant No. ACI-1547580) at Virginia Tech.^[68,69] This research used resources of the National Energy Research Scientific Computing Center (NERSC), a US Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

AUTHOR CONTRIBUTIONS

Marcus Hanwell: Conceptualization; data curation; funding acquisition; project administration; software; validation; visualization; writing-original draft; writing-review and editing. **Chris Harris:** Data curation; software; visualization; writing-original draft; writing-review and editing. **Alessandro Genova:** Data curation; software; visualization; writing-original draft; writing-review and editing. **Mojtaba Haghaghatlari:** Data curation; software; visualization; writing-original draft; writing-review and editing. **Muammar El Khatib:** Data curation; software; visualization; writing-original draft; writing-review and editing. **Patrick Avery:** Data curation; software; visualization; writing-original draft; writing-review and editing. **Johannes Hachmann:** Conceptualization; data curation; software; supervision; visualization; writing-original draft; writing-review and editing. **Wibe Albert de Jong:** Conceptualization; software; supervision; visualization; writing-original draft; writing-review and editing.

PEER-REVIEW DETAILS

Open Peer-Reviewer Details for this article are available here:

- Peer-Reviewer Report #1 DOI: 10.22541/au.159906593.30810770
- Peer-Reviewer Report #2 DOI: 10.22541/au.159906594.42729025
- Author Response DOI: 10.22541/au.159906595.50733741

ORCID

Marcus D. Hanwell  <https://orcid.org/0000-0002-5851-5272>

Chris Harris  <https://orcid.org/0000-0002-1113-3728>

Alessandro Genova  <https://orcid.org/0000-0001-8965-7495>

Mojtaba Haghaghatlari  <https://orcid.org/0000-0002-3779-2246>

Muammar El Khatib  <https://orcid.org/0000-0002-0977-8439>

Patrick Avery  <https://orcid.org/0000-0003-2254-1345>

Johannes Hachmann  <https://orcid.org/0000-0003-4501-4118>

Wibe Albert de Jong  <https://orcid.org/0000-0002-7114-8315>

REFERENCES

- [1] Licenses & Standards | Open Source Initiative, <https://opensource.org/licenses> (accessed: April 2020).
- [2] The 3-Clause BSD License | Open Source Initiative, <https://opensource.org/licenses/BSD-3-Clause> (accessed: April 2020).
- [3] Welcome to Python.org, <https://www.python.org/> (accessed: April 2020).
- [4] Open Chemistry, <https://github.com/OpenChemistry> (accessed: April 2020).
- [5] JSON, <https://www.json.org/> (accessed: April 2020).
- [6] JSON-LD—JSON for Linking Data, <https://json-ld.org/> (accessed: April 2020).
- [7] HDF Group—HDF5, <https://support.hdfgroup.org/HDF5/> (accessed: April 2020).
- [8] MessagePack: It's like JSON, but fast and small, <https://msgpack.org/> (accessed: April 2020).
- [9] BSON (Binary JSON) Serialization, <http://bsonspec.org/> (accessed: April 2020).
- [10] PostgreSQL: Documentation: 9.4: JSON Types, <https://www.postgresql.org/docs/9.4/datatype-json.html> (accessed: April 2020).
- [11] The most popular database for modern apps, <https://www.mongodb.com/> (accessed: April 2020).
- [12] PostgreSQL: The world's most advanced open source database, <https://www.postgresql.org/> (accessed: April 2020).
- [13] XYZ file format—Wikipedia, https://en.wikipedia.org/wiki/XYZ_file_format (accessed: April 2020).
- [14] W. Phadungsukanan, M. Kraft, J. A. Townsend, P. Murray-Rust, *Aust. J. Chem.* **2012**, *4*, 15. <https://doi.org/10.1186/1758-2946-4-15>.
- [15] P. Murray-Rust, H. S. Rzepa, *Aust. J. Chem.* **2011**, *3*, 44. <https://doi.org/10.1186/1758-2946-3-44>.
- [16] P. Murray-Rust, J. A. Townsend, S. E. Adams, W. Phadungsukanan, J. Thomas, *Aust. J. Chem.* **2011**, *3*, 43. <https://doi.org/10.1186/1758-2946-3-43>.
- [17] W. A. de Jong, A. M. Walker, M. D. Hanwell, *J. Cheminform.* **2013**, *5*, 25. <https://doi.org/10.1186/1758-2946-5-25>.
- [18] M. D. Hanwell, W. A. de Jong, C. J. Harris, *J. Cheminform.* **2017**, *9*, 55. <https://doi.org/10.1186/s13321-017-0241-z>.
- [19] N. M. OBoyle, M. Banck, C. A. James, C. Morley, T. Vandermeersch, G. R. Hutchison, *J. Cheminform.* **2011**, *3*, 3. <https://doi.org/10.1186/1758-2946-3-33>.
- [20] RDKit: Open-Source Cheminformatics Software, <https://www.rdkit.org/> (accessed: April 2020).
- [21] N. M. OBoyle, A. L. Tenderholt, K. M. Langner, *J. Comput. Chem.* **2008**, *29*, 839. <https://doi.org/10.1002/jcc.20823>.
- [22] L. A. Hjorth, M. J. Jørgen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargas, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbørg, J. Kubal, K. Kaasbørg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, K. W. Jacobsen, *J. Phys. Condens. Matter* **2017**, *29*, 273002. <https://doi.org/10.1088/1361-648x/aa680e>.
- [23] S. van der Walt, S. C. Colbert, G. Varoquaux, *Comput. Sci. Eng.* **2011**, *13*, 22. <https://doi.org/10.1088/1361-648x/aa680e>.
- [24] Eigen, <https://eigen.tuxfamily.org/> (accessed: April 2020).
- [25] OpenChemistry/chemicaljson, <https://github.com/OpenChemistry/chemicaljson> (accessed: April 2020).
- [26] MolSSI/QCSchema, <https://github.com/MolSSI/QCSchema> (accessed: April 2020).
- [27] M. D. Hanwell, D. E. Curtis, D. C. Lonie, T. Vandermeersch, E. Zurek, G. R. Hutchison, *Aust. J. Chem.* **2012**, *4*, 17. <https://doi.org/10.1186/1758-2946-4-17>.
- [28] R. M. Parrish, L. A. Burns, D. G. A. Smith, A. C. Simmonett, A. E. DePrince, E. G. Hohenstein, U. Bozkaya, A. Y. Sokolov, R. Di Remigio, R. M. Richard, J. F. Gonthier, A. M. James, H. R. McAlexander, A. Kumar, M. Saitow, X. Wang, B. P. Pritchard, P. Verma, H. F. Schaefer, K. Patkowski, R. A. King, E. F. Valeev, F. A. Evangelista, J. M. Turney, T. D. Crawford, C. D. Sherrill, *J. Chem. Theory Comput.* **2017**, *13*, 3185. <https://doi.org/10.1021/acs.jctc.7b00174>.
- [29] S. R. Heller, A. McNaught, I. Pletnev, S. Stein, D. Tchekhovskoi, *Aust. J. Chem.* **2015**, *7*, 23. <https://doi.org/10.1186/s13321-015-0068-4>.
- [30] OpenSMILES Home Page, <http://opensmiles.org/> (accessed: April 2020).
- [31] S. Kim, J. Chen, T. Cheng, A. Gindulyte, J. He, S. He, Q. Li, B. A. Shoemaker, P. A. Thiessen, B. Yu, L. Zaslavsky, J. Zhang, E. E. Bolton, *Nucleic Acids Res.* **2018**, *47*, D1102. <https://doi.org/10.1093/nar/gky1033>.
- [32] A. J. Williams, *Collaborative Computational Technologies for Biomedical Research*, John Wiley & Sons Inc., Hoboken, New Jersey **2011**, p. 363. <https://doi.org/10.1002/9781118026038.ch22>.
- [33] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, C. Willing, Jupyter Development Team, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas* (Eds: F. Loizides, B. Scmidt), IOS Press, Amsterdam, Netherlands **2016**, p. 87. <https://doi.org/10.3233/978-1-61499-649-1-87>.
- [34] Project Jupyter, <https://www.jupyter.org/> (accessed: April 2020).
- [35] Girder: a data management platform—Girder 3.0.5 documentation, <https://girder.readthedocs.io/> (accessed: April 2020).
- [36] JupyterLab Documentation—JupyterLab 2.2.5 documentation, <https://jupyterlab.readthedocs.io/> (accessed: April 2020).
- [37] Zenodo—Research. Shared, <https://zenodo.org/> (accessed: April 2020).
- [38] P. Jupyter, M. Bussonnier, J. Forde, J. Freeman, B. Granger, T. Head, C. Holdgraf, K. Kelley, G. Nalvarte, A. Osherooff, M. Pacer, Y. Panda, F. Perez, B. Ragan-Kelley, C. Willing, in *Proc. the 17th Python in Science Conference, SciPy* **2018**, <https://doi.org/10.25080/majora-4af1f417-011>.
- [39] The Binder Project, <https://mybinder.org/> (accessed: April 2020).
- [40] Figshare, <https://knowledge.figshare.com/> (accessed: April 2020).
- [41] OpenChemistry/mongochemserver, <https://github.com/OpenChemistry/mongochemserver> (accessed: April 2020).
- [42] OpenChemistry/mongochemdeploy, <https://github.com/OpenChemistry/mongochemdeploy> (accessed: April 2020).
- [43] D. Merkel, *Linux J.* **2014**, *239*. <https://www.linuxjournal.com/content/docker-lightweight-linux-containers-consistent-development-and-deployment> (accessed: April 2020).
- [44] G. M. Kurtzer, V. Sochat, M. W. Bauer, *PLoS One* **2017**, *12*, e0177459. <https://doi.org/10.1371/journal.pone.0177459>.
- [45] L. Gerhardt, W. Bhimji, S. Canon, M. Fasel, D. Jacobsen, M. Mustafa, J. Porter, V. Tsulaia, *J. Phys.: Conf. Ser.* **2017**, *898*, 082021. <https://doi.org/10.1088/1742-6596/898/8/082021>.
- [46] M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, H. J. J. V. Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus, W. A. de Jonga, *Comput. Phys. Commun.* **2010**, *181*, 1477. <https://doi.org/10.1016/j.cpc.2010.04.018>.
- [47] M. Haghightlari, G. Vishwakarma, D. Altarawy, R. Subramanian, B. U. Kota, A. Sonpal, S. Setlur, J. Hachmann, *WIREs Comput. Mol. Sci.* **2020**, *10*, e1458. <https://doi.org/10.1002/wcms.1458>, <https://doi.org/10.1002/wcms.1458>.

[48] J. Hachmann, M. A. F. Afzal, M. Haghighatlari, Y. Pal, *Mol. Simul.* **2018**, *44*, 921. <https://doi.org/10.1080/08927022.2018.1471692>.

[49] J. S. Smith, O. Isayev, A. E. Roitberg, *Chem. Sci.* **2017**, *8*, 3192. <https://doi.org/10.1039/C6SC05720A>.

[50] PyTorch: An Imperative Style, High-Performance Deep Learning Library, <https://papers.nips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library>

[51] X. Gao, F. Ramezanghorbani, O. Isayev, J. S. Smith, A. E. Roitberg, *J. Chem. Inf. Model.* **2020**, *60*, 3408. <https://doi.org/10.1021/acs.jcim.0c00451>.

[52] React—A JavaScript library for building user interfaces, <https://reactjs.org/> (accessed: April 2020).

[53] AngularJS—Superheroic JavaScript MVW Framework, <https://angularjs.org/> (accessed: April 2020).

[54] Vue.js, <https://vuejs.org/> (accessed: April 2020).

[55] Stencil, <https://stenciljs.com/> (accessed: April 2020).

[56] N. Rego, D. Koes, *Bioinformatics* **2015**, *31*, 1322. <https://doi.org/10.1093/bioinformatics/btu829>.

[57] Vtk.js, <https://kitware.github.io/vtk-js/index.html> (accessed: April 2020).

[58] Avogadro and Quantum Chemistry Viz in JupyterLab—Kitware Blog, <https://blog.kitware.com/open-chemistry-avogadro-electronic-structure-in-jupyterlab/>

[59] OpenChemistry/jupyter-examples, <https://github.com/OpenChemistry/jupyter-examples> (accessed: April 2020).

[60] NCI/CADD Chemical Identifier Resolver, <https://cactus.nci.nih.gov/chemical/structure> (accessed: April 2020).

[61] M. Solà, F. DeProft, F. M. Bickelhaupt, *ChemistryOpen* **2019**, *8*, 814. <https://doi.org/10.1002/open.201900197>.

[62] T. Sperger, I. A. Sanhueza, F. Schoenebeck, *Acc. Chem. Res.* **2016**, *49*, 1311. <https://doi.org/10.1021/acs.accounts.6b00068>.

[63] R. A. Mata, M. A. Suhm, *Angew. Chem. Int. Ed.* **2017**, *56*, 11011. <https://doi.org/10.1002/anie.201611308>.

[64] NIST Chemistry WebBook, <https://webbook.nist.gov/chemistry/> (accessed: April 2020).

[65] spacetelescope/notebooks, <https://github.com/spacetelescope/notebooks> (accessed: April 2020).

[66] D. Smith, D. Altarawy, L. Burns, M. Welborn, L. N. Naden, L. Ward, S. Ellis, T. Crawford, *WIREs Comput. Mol. Sci.* **2020**, *e1491*. <https://doi.org/10.1002/wcms.1491>.

[67] GO FAIR, <https://www.go-fair.org/> (accessed: April 2020).

[68] A. Krylov, T. L. Windus, T. Barnes, E. Marin-Rimoldi, J. A. Nash, B. Pritchard, D. G. A. Smith, D. Altarawy, P. Saxe, C. Clementi, T. D. Crawford, R. J. Harrison, S. Jha, V. S. Pande, T. Head-Gordon, *J. Chem. Phys.* **2018**, *149*, 180901. <https://doi.org/10.1063/1.5052551>.

[69] N. Wilkins-Diehr, T. D. Crawford, *Comput. Sci. Eng.* **2018**, *20*, 26. <https://doi.org/10.1109/MCSE.2018.05329813>.

How to cite this article: Hanwell MD, Harris C, Genova A, et al. Open Chemistry, JupyterLab, REST, and quantum chemistry. *Int J Quantum Chem.* 2020;e26472. <https://doi.org/10.1002/qua.26472>