



OPTIMADE: an API for exchanging materials data

Casper W. Andersen^{1,†}, Rickard Armiento^{2,†}, Evgeny Blokhin^{3,4,†}, Gareth J. Conduit^{5,†}, Shyam Dwaraknath^{6,†}, Matthew L. Evans^{5,7,†}, Ádám Fekete^{8,7,†}, Abhijith Gopakumar^{9,†}, Saulius Gražulis^{10,11,†}, Andrius Merkys^{10,†}, Fawzi Mohamed^{12,†}, Corey Oses^{13,14,†}, Giovanni Pizzi^{1,†}, Gian-Marco Rignanese^{7,†}, Markus Scheidgen^{12,15,†}, Leopold Talirz^{1,16,†}, Cormac Toher^{13,14,†}, Donald Winston^{6,†}, Rossella Aversa^{17,18}, Kamal Choudhary¹⁹, Pauline Colinet^{13,14}, Stefano Curtarolo^{13,14}, Davide Di Stefano²⁰, Claudia Draxl¹⁵, Suleyman Er²¹, Marco Esters^{13,14}, Marco Fornari^{22,13}, Matteo Giantomassi⁷, Marco Govoni²³, Geoffroy Hautier⁷, Vinay Hegde⁹, Matthew K. Horton⁶, Patrick Huck⁶, Georg Huhs¹⁵, Jens Hummelshøj²⁴, Ankit Kariyaa²⁵, Boris Kozinsky^{26,27}, Snehal Kumbhar¹, Mohan Liu⁹, Nicola Marzari¹, Andrew J. Morris²⁸, Arash Mostofi²⁹, Kristin A. Persson^{6,30}, Guido Petretto⁷, Thomas Purcell¹², Francesco Ricci⁷, Frisco Rose^{13,14}, Matthias Scheffler¹², Daniel Speckhard^{12,15}, Martin Uhrin¹, Antanas Vaitkus¹⁰, Pierre Villars⁴, David Waroquiers⁷, Chris Wolverton⁹, Michael Wu⁶, and Xiaoyu Yang³¹

¹Theory and Simulation of Materials (THEOS), and National Centre for Computational Design and Discovery of Novel Materials (MARVEL), École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

²Materials Design and Informatics unit, Department of Physics, Chemistry and Biology, Linköping University, Sweden

³Tilde Materials Informatics, Straßmannstraße 25, 10249, Berlin, Germany

⁴Materials Platform for Data Science, Sepapaja 6, 15551, Tallinn, Estonia

⁵Theory of Condensed Matter Group, Cavendish Laboratory, 19 J.J. Thomson Avenue, Cambridge, CB3 0HE, United Kingdom

⁶Lawrence Berkeley Labs, Berkeley, CA, USA

⁷UCLouvain, Institut de la Matière Condensée et des Nanosciences (IMCN), Chemin des Étoiles 8, Louvain-la-Neuve 1348, Belgium

⁸Department of Physics, Kings College London, London, United Kingdom

⁹Department of Materials Science and Engineering, Northwestern University, Evanston, IL 60208, USA

¹⁰Institute of Biotechnology, Life Science Center, Vilnius University, Saulėtekio av. 7, LT-10257 Vilnius, Lithuania

¹¹Institute of Computer Science, Faculty of Mathematics and Informatics, Vilnius University, Naugarduko g. 24, LT-03225 Vilnius, Lithuania

¹²Fritz-Haber-Institut der Max-Planck-Gesellschaft, Faradayweg 4-6, 14195, Berlin, Germany

¹³Center for Autonomous Materials Design, Duke University, Durham, NC 27708, USA

¹⁴Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA

¹⁵Humboldt-Universität zu Berlin, Institut für Physik and IRIS Adlershof, 12489 Berlin, Germany

¹⁶Laboratory of Molecular Simulation (LSMO), École Polytechnique Fédérale de Lausanne, 1951 Sion, Switzerland

¹⁷National Research Council-Istituto Officina dei Materiali (CNR-IOM), 34136, Trieste, Italy

¹⁸Karlsruhe Institute of Technology (KIT), Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany

¹⁹Materials Measurement Laboratory, National Institute of Standards and Technology, Gaithersburg, MD, 20899, USA

²⁰Ansys, 300 Rustat House, 62 Clifton Rd, Cambridge CB1 7EG, United Kingdom

²¹Dutch Institute for Fundamental Energy Research (DIFFER), De Zaaie 20, 5612 AJ, Eindhoven, The Netherlands

²²Department of Physics and Science of Advanced Materials Program, Central Michigan University, Mount Pleasant, Michigan 48859, USA

²³Argonne National Laboratory, Lemont, IL 60439, USA

²⁴Toyota Research Institute (TRI), Los Altos, California 94022, USA

²⁵Department of Computer Science, University of Copenhagen, Universitetsparken 1, 2100 Copenhagen, Denmark

²⁶John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA

²⁷Robert Bosch LLC, Research and Technology Center North America, 255 Main St, Cambridge, Massachusetts 02142, USA

²⁸School of Metallurgy and Materials, University of Birmingham, Edgbaston, Birmingham, B15 2TT, United Kingdom

²⁹Departments of Materials and Physics, and the Thomas Young Centre, Imperial College London, Exhibition Road, London SW7 2AZ, United Kingdom

³⁰Department of Material Science and Engineering, Hearst Mining Memorial Building, UC Berkeley, Berkeley CA 94720, USA.

³¹Computer Network Information Center, Chinese Academy of Sciences, Beijing, China

†These authors contributed equally to this work, and can be contacted at dev@optimade.org

ABSTRACT

The Open Databases Integration for Materials Design (OPTIMADE) consortium has designed a universal application programming interface (API) to make materials databases accessible and interoperable. We outline the first stable release of the specification, v1.0, which is already supported by many leading databases and several software packages. We illustrate the advantages of the OPTIMADE API through worked examples on each of the public materials databases that support the full API specification.

Introduction

Data has become a crucial resource in many fields of our everyday life, and materials science is no exception. Experimental data has long been meticulously curated in handbooks and databases, with the first edition of Landolt-Börnstein¹ being published in 1883. Nowadays, various commercial and non-commercial experimental databases, such as the Inorganic Crystal Structure Database (ICSD)², are widely used throughout the field.

High-throughput electronic structure calculations, themselves enabled by algorithmic improvements and growing computational resources, have significantly increased the availability of useful data from computational simulations of materials. Since the pioneering work of Ceder *et al.*³, a large number of high-throughput first-principles studies have been reported in the literature (for a review, see Ref. 4), with results typically collated in databases. This explosion in the amount of available data has kickstarted a new paradigm of data-driven materials science⁵, creating opportunities for concurrent materials design boosted by databases that can be queried online by humans and machines via an application programming interface (API)^{6–9}.

It is extremely beneficial to be able to access information from multiple materials databases, as they often differ in fidelity and focus across material classes and properties. However, retrieving data from multiple databases is difficult as each database has its own specific API. Moreover, as the APIs of individual databases inevitably evolve, existing clients must also evolve and are required to translate the responses from the new API to the internal representation of the client, which can require significant effort.

Motivated by these considerations, the providers of several of materials databases joined forces to design a common API specification to enable seamless access and interoperability across materials databases. The effort started at the workshop “Open Databases Integration for Materials Design”, held at the Lorentz Center in Leiden, Netherlands in October 2016, and continued at followup workshops held at CECAM in Lausanne, Switzerland in June 2018, June 2019, and June 2020. The result is the OPTIMADE specification (v1.0)¹⁰; OPTIMADE defines a RESTful API that is queried with URLs, with responses adhering to the JSON:API specification¹¹. Specification development adheres to Semantic Versioning¹² to avoid surprises and enable backwards-compatibility where possible, whilst not impeding further development. By learning from the development of existing APIs, the OPTIMADE API has been designed so that it can be implemented across a broad range of database back-ends and sizes.

In this paper, we first review the query format of existing databases to motivate the construction of the OPTIMADE API specification. We then illustrate the use of the API with a set of worked examples. Databases that already fully support the OPTIMADE API are enumerated in Table 1. We further highlight libraries that could accelerate uptake and assist materials data curators to support the OPTIMADE API format. Finally, we discuss future prospects and ongoing development of the OPTIMADE API.

How to query a database

Leading materials databases are a veritable treasure trove of information, but they only become useful once a human, or machine, can access them. In this section we review the current range of APIs used by various databases to enable access to an example compound, SiO₂, which serves to highlight the esoteric range of APIs that a user must navigate in order to make use of multiple materials databases. We then demonstrate the universal nature of the OPTIMADE API that permits seamless access to all materials databases that support it.

Current generation of APIs

We first compare and contrast the APIs that must be used to request records on an exemplar system, SiO₂, from three different databases: AFLOW, Materials Project, and the Crystallographic Open Database (COD). All three queried databases support requests using a representational state transfer through a web service (RESTful), at the following URLs[‡]:

AFLOW [http://aflow.org/API/aflux/?species\(Si,O\),nspecies\(2\)](http://aflow.org/API/aflux/?species(Si,O),nspecies(2))

Materials Project <https://www.materialsproject.org/rest/v2/materials/SiO2/vasp/structure>

COD <https://www.crystallography.net/cod/result.php?formula=O2%20Si>

The three APIs vary syntactically (in format), taxonomically (having different names for terms), and semantically (in the conflicting definitions of chemical formula as an intensive or extensive property). AFLOW returns all structures with both Si and O present, whereas both Materials Project and COD deliver any structure with a formula unit of SiO₂. The wide range of query formats that will deliver non-overlapping structures significantly complicates access to all available data for SiO₂, without even considering the differing representations of the structures returned.

The inconsistent format of the query is further complicated by the difficulty of accessing other structures with the SiO₂ formula. Focusing on just AFLOW, two possible queries that users more familiar with the other APIs might attempt are

[http://aflow.org/API/aflux/?compound\(SiO2\)](http://aflow.org/API/aflux/?compound(SiO2))
which returns no response;

[http://aflow.org/API/aflux/?compound\(O2Si1\)](http://aflow.org/API/aflux/?compound(O2Si1))
now lists the elements in alphabetical order as required by AFLOW, and includes the “1” after element symbols, so that “SiO₂” becomes “O₂Si₁”. This returns entries where the unit cell is SiO₂, but does not return Si₂O₄ or simulation cells containing more formula units;

while the exemplar [http://aflow.org/API/aflux/?species\(Si,O\),nspecies\(2\)](http://aflow.org/API/aflux/?species(Si,O),nspecies(2)) returns all entries with at least one Si and an O, so returns the systems of interest, but also many more besides.

The distinction between the requests, which are different for each database, requires the user to become expert in many different APIs. This again emphasises the need for a single well-designed and standardized API to access all materials databases, which is provided by the OPTIMADE API.

OPTIMADE API

The OPTIMADE API provides a holistic tool to access the information in any compatible materials database. To retrieve information about materials from a particular database, the user submits a request/query through a URL. Each database provider will publish a base URL that serves the OPTIMADE API, for example <https://example.com/optimade/>. The same query, using the OPTIMADE API, opens access to all databases. Both human-readable and machine-readable versions of the OPTIMADE API specification are available online at GitHub¹³, with releases archived at Zenodo¹⁰. It is also registered as a standard on FAIRsharing.org¹⁴.

The philosophy of the OPTIMADE query is to enable the structural formula to be specified in a straightforward and intuitive manner. The three queries from the previous section can each be performed on standardised, versioned endpoints (`/v1/structures`) shown in green that permit a common filter format with well-defined terms (`?filter=chemical_formula_reduced="O2Si"`) shown in purple: `<optimade_implementation_url>/v1/structures?filter=chemical_formula_reduced="O2Si"`.

The list of databases confirmed and tested in this paper to support the OPTIMADE API is shown in Table 1. They are all publicly accessible, providing users with open access to large repositories of computational and experimental data. The OPTIMADE consortium provides an open, federated list of implementations at <https://providers.optimade.org>; new implementations can register themselves via a simple pull request on GitHub. Further databases are known to have partial implementations of the OPTIMADE API, including JARVIS^{15,16} and MatCloud^{17,18}. Some software frameworks, such as AiiDA^{19–21}, also enable users to access their personal data through an OPTIMADE API.

[‡]The Materials Project requires the user to supply an [API key](#) specified as a parameter for a GET request (e.g., `?API_KEY=YOUR_API_KEY`) or as a POST variable (e.g., `{'API_KEY': 'YOUR_API_KEY'}`).

Worked example

To illustrate the effective use of the OPTIMADE API we now provide a worked example. We explore materials containing Group 14 elements (the carbon family), starting with a general search before drilling down to specific materials. The Group 14 elements are of particular interest as their atomic orbitals regularly hybridise, enabling a variety of bonding with differing geometries. The hybridised orbitals enable these elements to form the backbone of a wide range of compounds, both inorganic and organic, that underpin plastics, drugs, and semiconductors. Group 14 therefore forms both a diverse and important family of compounds that heavily populates databases, so are an ideal case study to demonstrate the OPTIMADE API.

Common features of the response

Whilst our previous exploration of the Group 14 compounds considered only SiO₂, the flexibility of the OPTIMADE API allows us to start with a search over all materials in Group 14, comprising carbon (C), silicon (Si), germanium (Ge), tin (Sn), and lead (Pb). We start with a simple API call that searches for all materials that contain at least one element in Group 14:

```
/v1/structures?filter=elements HAS ANY "C", "Si", "Ge", "Sn", "Pb"
```

This string can be appended to the base URL of any of the available implementations, to gather results in a standardised form. The base URL can be found on the providers dashboard at <https://www.optimade.org/providers-dashboard> (see also Table 1).

As an example, this query is run through the Theoretical Crystallography Open Database (TCOD)²² with the following URL:

```
https://www.crystallography.net/tcod/optimade/v1/structures?filter=elements HAS ANY "C", "Si", "Ge", "Sn", "Pb"
```

The JSON response is summarized below, where some lines have been omitted for brevity and the full response is given in a supplementary file (response1.json).

The first tranche of the JSON response comprises the “data” field that contains a list of entries returned for the query; a truncated version of this field is shown above, displaying a few salient properties of just one of the ten entries from the full response. The response for a particular material comprises three sections:

```
1 {
2   "data": [
3     {
4       "attributes": {
5
6         "dimension_types": [1, 1, 1],
7         "elements": ["O", "Sn", "Ta"],
8
9
10        "lattice_vectors": [
11          [4.02510400822, 0, 0],
12          [0, 4.02510400822, 0],
13          [0, 0, 4.02510400822]
14        ],
15
16
17        "_tcod_a": 4.02510400822,
18        "_tcod_alpha": 90,
19        "_tcod_b": 4.02510400822,
20        "_tcod_beta": 90,
21        "_tcod_c": 4.02510400822,
```

attributes The physical properties of the material comprises both mandatory information such as `elements` and `lattice_vectors`, as well as optional, additional database-specific information prefixed with the database name (e.g., `_tcod_`, here used to provide lattice parameters). This ensures that all databases return the most important and common information in a standardized format, as well as allowing them to include additional database-specific data. Importantly, the OPTIMADE specification provides a standardized way for database implementations to be self-documenting, via introspective `/info` endpoints. We see in the `elements` section that here we have returned a material comprising the element of interest, Sn, as well as O and Ta.

```
40   "id": "10000003",
41   "links": {
42     "self": "http://www.crystallography.net/tcod/10000003.cif"
43   },
```

id and links The unique ID for the entry within the database, and a self-link to the database-specific representation/rendering of the entry (in this case, a link to a Crystallographic Information File^{23,24}).

```
44     "relationships": {
45       "references": {
46         "data": {
47           "id": "10000003",
48           "type": "references"
49         }
50       }
51     }
```

relationships The relationships section links the user to other entries in the database and beyond, here the bibliographic references.

The additional nine materials not shown here all comprised of compounds containing either C, Si, Ge, Sn, or Pb, supplemented by a variety of other elements. The foot of the response contains information about the underlying database, comprised of three sections:

```
684     "links": {
685       "base_url":
686       ↪ "http://www.crystallography.net/tcod/optimade/v1.0.0/",
687       "first": {
688         "href": "http://www.crystallography.net/tcod/optimade/v1.0.0_
689         ↪ /structures?page_limit=10&filter=elements HAS ANY
690         ↪ \"C\", \"Si\", \"Ge\", \"Sn\", \"Pb\""
691       },
692       "next": {
693         "href": "http://www.crystallography.net/tcod/optimade/v1.0.0_
694         ↪ /structures?page_limit=10&page_offset=10&filter=elements
695         ↪ HAS ANY \"C\", \"Si\", \"Ge\", \"Sn\", \"Pb\""
696       }
697     },
698     }
```

links The response returned the first ten (i.e., the default page limit) entries that matched the query, however, more materials are available within the database. This section contains pagination links to the current and next page of results for this query, as well as relevant external links.

```
693     "meta": {
694       "api_version": "1.0.0",
695       "data_available": 2906,
696       "data_returned": 2631,
697       "implementation": {
698         "maintainer": {
699           "email": "cod-bugs@ibt.lt"
700         },
701         "name": "Theoretical Crystallography Open Database",
702         "source_url": "svn://www.crystallography.net/cod/tags/tcod_
703         ↪ /v0.2.1/cgi-bin/optimade.pl@261229",
704         "version": "v0.2.1"
705       },
706       "last_id": "10000031",
707       "more_data_available": true,
708       "provider": {
709         "description": "Theoretical Crystallography Open Database",
710         "name": "Theoretical Crystallography Open Database",
711         "prefix": "tcod"
712       },
713       "query": {
714         "representation": "/structures?filter=elements HAS ANY
715         ↪ \"C\", \"Si\", \"Ge\", \"Sn\", \"Pb\""
716       },
717       "time_stamp": "2021-02-24T16:36:05Z"
718     }
719   }
```

meta This section provides metadata associated with the request, such as number of results, the details of the database provider, the implementation and the representation and timestamp of the submitted query.

A benefit of the OPTIMADE API is that the structure of the response is common to all materials databases. The responses differ only in the optional and database-specific information prefixed with the database name (here `_tcod_`). Table 1 lists several materials databases that have active OPTIMADE API implementations, and the large number of results (N_1) that they return for this particular filter.

Exploring the database

Requesting the example filter above from the TCOD database returns 2,631 materials entries, but the same filter could return millions (see Table 1). For some requests the volume of the materials data could become unmanageable so the specification allows for the use of several pagination methods laid out by JSON:API¹¹. These approaches all provide a link to the “next” page of data:

```
/v1/structures?page_limit=10&page_offset=10&filter=elements HAS ANY "C", "Si", "Ge", "Sn", "Pb"
```

with parameter `page_offset=10` to allow the user to select which page to enter, and the parameter `page_limit=10` to control the number of materials returned per individual request.

The most useful way to explore an OPTIMADE database is to apply a filter; the specification mandates that several relevant properties must be queryable. For example, we can perform a more focused search for materials comprising at least one element in Group 14, and a maximum of two elements (a binary material), with the filter

```
/v1/structures?filter=elements HAS ANY "C", "Si", "Ge", "Sn", "Pb" AND nelements=2
```

This query returns 296 materials from the TCOD database, with the response summarized below, where some lines have been omitted for brevity and the full response is given in a supplementary file (`response3.json`). The number of matching entries (N_2) across all implementations for this filter are shown in Table 1.

```
1  {
2    "data" : [
3      {
4        "attributes" : {

34         "elements" : [
35           "Ge",
36           "O"
37         ],

62       },

78     },

791   ],

825 }
```

We can now see that the first structure, and indeed all structures, returned comprise at least one element in Group 14, here Ge, and a maximum of one other element (a binary material), here O. Additional filters can be applied to further refine the materials returned, or to construct more complex queries. For example, ternary structures that contain at least one of the elements C, Si, Ge, or Sn, but do not contain Pb (e.g. for applications where Pb toxicity would be a concern), can be retrieved using the filter

```
/v1/structures?filter=elements HAS ANY "C", "Si", "Ge", "Sn" AND NOT elements HAS "Pb" AND elements LENGTH 3
```

The number of entries matching this filter are denoted as N_3 in Table 1.

These simple examples demonstrate how useful chemical queries are expressible with the OPTIMADE API, allowing users to refine their queries and to suit their specific application. Further functionality of the OPTIMADE API can be found in the specification¹⁰.

Related libraries

The wider usage of the OPTIMADE API is a key goal for the consortium; to this end, several open source libraries have been developed to help users of the OPTIMADE API (either implementation developers, or client end-users), of which three are described below. The first two libraries both offer tools that aid the implementation of the API for materials database developers, with the first also containing tools to construct and validate queries, while the third library is intended for end users of OPTIMADE-compliant APIs.

optimade-python-tools

[optimade-python-tools](#)²⁵ is a Python 3.6+ package that contains a complete set of tools for implementing an OPTIMADE-compliant API, as well as several utilities that can be used by client code. The package is listed on the Python Package Index (PyPI) as [optimade](#) and can be found on GitHub at [Materials-Consortia/optimade-python-tools](#). Current (v0.13) functionality of the package includes:

- [pydantic](#)²⁶ data models for all objects defined in the OPTIMADE specification that can be used in server or client code;
- an extensible reference server implementation leveraging the modern Python tools [pydantic](#)²⁶ and [FastAPI](#)²⁷. This reference server forms the basis of the OPTIMADE implementations for the Materials Project²⁸, NOMAD²⁹, Materials Cloud⁹ and [odbx](#)^{30,31} providers. These tools are also used to generate a machine-readable OpenAPI version of the OPTIMADE specification;
- a Python parser for the OPTIMADE filter language written with the Lark parsing library³²;
- a validator for OPTIMADE implementations, available standalone or as a GitHub action at [Materials-Consortia/optimade-validator-action](#). This validator is run against the federated list of OPTIMADE providers every day, with a live dashboard indicating compliance with the specification at <https://www.optimade.org/providers-dashboard/>;
- filter transformers from the abstract syntax tree to queries to popular database back-ends. Currently, the supported back-ends are MongoDB (via [pymongo](#))^{33,34} and [Elasticsearch](#)³⁵;
- adapter classes to interface with other popular libraries, namely [pymatgen](#)³⁶, [ASE](#)³⁷, [AiiDA](#)^{19–21} and [JARVIS](#)¹⁵ as well as converting OPTIMADE structures to the [CIF](#)^{23,24}, XYZ, and more domain standardised file formats.

OPTIMADE::Filter

[OPTIMADE::Filter](#)³⁸ is a Perl library for the syntactical analysis of the OPTIMADE filter language. Apart from the construction of abstract syntax trees, the library can translate simple filter strings to SQL queries.

pymatgen optimade module

[pymatgen](#)³⁶ is a Python library for materials science. A user-friendly `OptimadeRester` client has been added to a new OPTIMADE module within `pymatgen` to provide a way to query OPTIMADE structure resources in a way familiar to existing users of `pymatgen` and the Materials Project API.

Summary

The latest OPTIMADE API specification v1.0¹⁰ offers holistic access to many leading crystal structure databases, namely: AFLOW, COD, TCOD, Materials Cloud, Materials Project, NOMAD, [odbx](#), Open Materials Database (*omdb*), and OQMD. Open client implementations are also available^{39,40} that enable aggregated searches over many databases as well as user friendly graphical widgets that can create an OPTIMADE filter to empower the user with even easier access to data. OPTIMADE provides researchers easy access to over 10,000,000 results for different materials, providing benchmarking opportunities whilst offering a huge opportunity for screening and machine learning studies. The ability of the OPTIMADE API to search databases, expose links between databases, and deliver multiple results makes it well-positioned to significantly enhance the impact of pre-existing data silos. This should empower researchers to scan through new and unexpected material families, and train models from all available data that can understand deep correlations.

The OPTIMADE API is flexible and will be extended to more use cases going forward. The development and adoption of the OPTIMADE API relies on the involvement of a large number of scientists so contributions from the community are strongly encouraged⁴¹, and questions on development, registration of a provider, or usage can be directed to the [user forum](#)⁴². Proposed developments include the addition of more materials properties, the inclusion of molecular dynamics simulations and of experimental results, and extensions beyond electronic-structure calculations. The future development of APIs, including OPTIMADE, should herald an era of effective use of big, open data in materials science.

Data availability

The OPTIMADE specification is developed openly on GitHub¹³, with releases archived on Zenodo¹⁰. Version 1.0.0 was released on 1 July 2020 and is licensed under Creative Commons Attribution 4.0 International (CC-BY 4.0) available at <https://github.com/Materials-Consortia/OPTIMADE/blob/v1.0.0/LICENSE>.

Code availability

All associated code is hosted under the Materials-Consortia organisation on GitHub at <https://github.com/Materials-Consortia>.

References

1. Springer Materials. <http://materials.springer.com/>.
2. Zagorac, D., Müller, H., Ruehl, S., Zagorac, J. & Rehme, S. Recent developments in the inorganic crystal structure database: theoretical crystal structure data and related features. *J. Appl. Crystallogr.* **52**, 918–925, [10.1107/S160057671900997X](https://doi.org/10.1107/S160057671900997X) (2019).
3. Ceder, G. *et al.* Identification of cathode materials for lithium batteries guided by first-principles calculations. *Nature* **392**, 694–696, [10.1038/33647](https://doi.org/10.1038/33647) (1998).
4. Curtarolo, S. *et al.* The high-throughput highway to computational materials design. *Nat. Mater.* **12**, 173, [10.1038/nmat3568](https://doi.org/10.1038/nmat3568) (2013).
5. Himanen, L., Geurts, A., Foster, A. S. & Rinke, P. Data-Driven Materials Science: Status, Challenges, and Perspectives. *Adv. Sci.* **6**, 1900808, [10.1002/adv.201900808](https://doi.org/10.1002/adv.201900808) (2019).
6. Taylor, R. H. *et al.* A RESTful API for exchanging materials data in the AFLOWLIB.org consortium. *Comput. Mater. Sci.* **93**, 178–192, [10.1016/j.commatsci.2014.05.014](https://doi.org/10.1016/j.commatsci.2014.05.014) (2014).
7. Rose, F. *et al.* AFLUX: The LUX materials search API for the AFLOW data repositories. *Comput. Mater. Sci.* **137**, 362–370, [10.1016/j.commatsci.2017.04.036](https://doi.org/10.1016/j.commatsci.2017.04.036) (2017).
8. Ong, S. P. *et al.* The Materials Application Programming Interface (API): A simple, flexible and efficient API for materials data based on REpresentational State Transfer (REST) principles. *Comput. Mater. Sci.* **97**, 209–215, [10.1016/j.commatsci.2014.10.037](https://doi.org/10.1016/j.commatsci.2014.10.037) (2015).
9. Talirz, L. *et al.* Materials Cloud, a platform for open computational science. *Sci. Data* **7**, 299, [10.1038/s41597-020-00637-5](https://doi.org/10.1038/s41597-020-00637-5) (2020).
10. Andersen, C. W. *et al.* The OPTIMADE Specification. *Zenodo* [10.5281/zenodo.4195050](https://doi.org/10.5281/zenodo.4195050) (2020).
11. JSON:API v1.0. <https://jsonapi.org/format/1.0/>.
12. Preston-Werner, T. Semantic versioning 2.0.0. <https://semver.org/spec/v2.0.0.html> (2013).
13. <https://github.com/Materials-Consortia/OPTIMADE/blob/master/optimade.rst>.
14. FAIRsharing.org: OPTIMADE; Open Databases Integration for Materials Design. <https://doi.org/10.25504/FAIRsharing.xvfqAC>.
15. <https://jarvis.nist.gov/optimade/>.
16. Choudhary, K. *et al.* The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design. *npj Comput. Mater.* **6**, 1–13, [10.1038/s41524-020-00440-1](https://doi.org/10.1038/s41524-020-00440-1) (2020).
17. <http://matcloud.cnic.cn/>.
18. Yang, X. *et al.* Matcloud: A high-throughput computational infrastructure for integrated management of materials simulation, data and resources. *Comput. Mater. Sci.* **146**, 319–333 (2018).
19. Pizzi, G., Cepellotti, A., Sabatini, R., Marzari, N. & Kozinsky, B. AiiDA: automated interactive infrastructure and database for computational science. *Comput. Mater. Sci.* **111**, 218–230, [10.1016/j.commatsci.2015.09.013](https://doi.org/10.1016/j.commatsci.2015.09.013) (2016).
20. Huber, S. P. *et al.* AiiDA 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance. *Sci. Data* **7**, 300, [10.1038/s41597-020-00638-4](https://doi.org/10.1038/s41597-020-00638-4) (2020).
21. Uhrin, M., Huber, S. P., Yu, J., Marzari, N. & Pizzi, G. Workflows in AiiDA: Engineering a high-throughput, event-based engine for robust and modular computational workflows. *Comput. Mater. Sci.* **187**, 110086, [10.1016/j.commatsci.2020.110086](https://doi.org/10.1016/j.commatsci.2020.110086) (2021).
22. Merkys, A. *et al.* A posteriori metadata from automated provenance tracking: integration of AiiDA and TCOD. *J. Cheminform.* **9**, 56, [10.1186/s13321-017-0242-y](https://doi.org/10.1186/s13321-017-0242-y) (2017).
23. <https://www.iucr.org/resources/cif>.
24. Bernstein, H. J. *et al.* Specification of the crystallographic information file format, version 2.0. *J. Appl. Crystallogr.* **49**, 277–284, [10.1107/S1600576715021871](https://doi.org/10.1107/S1600576715021871) (2016).

25. <https://github.com/Materials-Consortia/optimade-python-tools>.
26. <https://github.com/samuelcolvin/pydantic>.
27. <https://github.com/tiangolo/fastapi>.
28. Jain, A. *et al.* Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *APL Mater.* **1**, 011002, [10.1063/1.4812323](https://doi.org/10.1063/1.4812323) (2013).
29. Draxl, C. & Scheffler, M. NOMAD: The FAIR concept for big data-driven materials science. *MRS Bull.* **43**, 676–682, [10.1557/mrs.2018.208](https://doi.org/10.1557/mrs.2018.208) (2018).
30. <https://odbx.science>.
31. <https://github.com/ml-evs/odbx.science>.
32. <https://github.com/lark-parser/lark>.
33. <https://github.com/mongodb/mongo>.
34. <https://github.com/mongodb/mongo-python-driver>.
35. <https://github.com/elastic/elasticsearch>.
36. Ong, S. P. *et al.* Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis. *Comput. Mater. Sci.* **68**, 314–319, [10.1016/j.commatsci.2012.10.028](https://doi.org/10.1016/j.commatsci.2012.10.028) (2013).
37. Bahn, S. R. & Jacobsen, K. W. An object-oriented scripting interface to a legacy electronic structure code. *Comput. Sci. Eng.* **4**, 56–66, [10.1109/5992.998641](https://doi.org/10.1109/5992.998641) (2002).
38. <https://github.com/Materials-Consortia/OPTIMADE-Filter>.
39. <https://optimade.science/>.
40. <https://materialscloud.org/optimadeclient>.
41. OPTIMADE mailing list. <mailto:dev@optimade.org>.
42. OPTIMADE user forum. <https://matsci.org/optimade>.
43. Curtarolo, S. *et al.* AFLOWLIB.ORG: A distributed materials properties repository from high-throughput *ab initio* calculations. *Comput. Mater. Sci.* **58**, 227–235, [10.1016/j.commatsci.2012.02.002](https://doi.org/10.1016/j.commatsci.2012.02.002) (2012).
44. Toher, C. *et al.* The AFLOW fleet for materials discovery. In Andreoni, W. & Yip, S. (eds.) *Handbook of Materials Modeling*, 1–28, [10.1007/978-3-319-42913-7_63-1](https://doi.org/10.1007/978-3-319-42913-7_63-1) (Springer International Publishing, Cham, Switzerland, 2018).
45. Gražulis, S. *et al.* Crystallography Open Database - an open-access collection of crystal structures. *J. Appl. Crystallogr.* **42**, 726–729, [10.1107/S0021889809016690](https://doi.org/10.1107/S0021889809016690) (2009).
46. Gražulis, S. *et al.* Crystallography Open Database (COD): an open-access collection of crystal structures and platform for world-wide collaboration. *Nucleic Acids Res.* **40**, D420–D427, [10.1093/nar/gkr900](https://doi.org/10.1093/nar/gkr900) (2012).
47. Jain, A. *et al.* A high-throughput infrastructure for density functional theory calculations. *Comput. Mater. Sci.* **50**, 2295–2310, [10.1016/j.commatsci.2011.02.023](https://doi.org/10.1016/j.commatsci.2011.02.023) (2011).
48. Mathew, K. *et al.* Atomate: A high-level interface to generate, execute, and analyze computational materials science workflows. *Comput. Mater. Sci.* **139**, 140–152, [10.1016/j.commatsci.2017.07.030](https://doi.org/10.1016/j.commatsci.2017.07.030) (2017).
49. Ghiringhelli, L. M. *et al.* Towards efficient data exchange and sharing for big-data driven materials science: metadata and data formats. *npj Comput. Mater.* **3**, [10.1038/s41524-017-0048-5](https://doi.org/10.1038/s41524-017-0048-5) (2017).
50. Evans, M. L. & Morris, A. J. matador: a Python library for analysing, curating and performing high-throughput density-functional theory calculations. *J. Open Source Softw.* **5**, 2563, [10.21105/joss.02563](https://doi.org/10.21105/joss.02563) (2020).
51. Armiento, R. Database-Driven High-Throughput Calculations and Machine Learning Models for Materials Design. In Schütt, K. T. *et al.* (eds.) *Machine Learning Meets Quantum Physics*, Lecture Notes in Physics, 377–395, [10.1007/978-3-030-40245-7_17](https://doi.org/10.1007/978-3-030-40245-7_17) (Springer International Publishing, Cham, 2020).
52. <http://openmaterialsdb.se>.
53. Saal, J. E., Kirklin, S., Aykol, M., Meredig, B. & Wolverton, C. Materials design and discovery with high-throughput density functional theory: The Open Quantum Materials Database (OQMD). *JOM* **65**, 1501–1509, [10.1007/s11837-013-0755-4](https://doi.org/10.1007/s11837-013-0755-4) (2013).

Acknowledgements

The authors acknowledge support CECAM in Lausanne (Switzerland) and Lorentz Center in Leiden (Netherlands) for hosting OPTIMADE workshops. CWA, SK, NM, GPi, MU, LT acknowledge financial support by the MARVEL National Centre of Competence in Research funded by the Swiss National Science Foundation (grant agreement ID 51NF40-182892), the European Centre of Excellence MaX “Materials design at the Exascale” (grant no. 824143) and by the the “MARKETPLACE” H2020 project (grant agreement ID 760173). RAr acknowledges financial support from the Swedish e-Science Research Centre (SeRC) and the Swedish Research Council (VR) project no. 2016-04810 and 2020-05402. RAv has received funding from the European Union’s Horizon 2020 Research and Innovation Program under grant agreement No. 654360 NFFA-EUROPE. GJC would like to acknowledge financial support from the Royal Society. SD, DWi, MWu, MH, PH, KP acknowledge financial support by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, Materials Sciences and Engineering Division under Contract No. DE-AC02-05-CH11231 (Materials Project program KC23MP). CD, MLE, GH, GMR, MSchef, MSchei, and DWa acknowledge support from the European Union’s Horizon 2020 research and innovation program under the European Union’s Grant agreement No. 951786 (NOMAD CoE). MLE acknowledges the EPSRC CDT in Computational Methods for Materials Science for funding under grant number EP/L015552/1. AG, ML, VH, and CW acknowledges support by Toyota Research Institute, and the award 70NANB14H012 from U.S. Department of Commerce, National Institute of Standards and Technology as part of the Center for Hierarchical Materials Design (CHiMaD). MGo acknowledges support by MICCoM, as part of the Computational Materials Sciences Program funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division. SG, AMe, AV were supported by funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No 689868 (SOLSA). AJM acknowledges networking support via the EPSRC Collaborative Computational Projects, CCP9 (EP/M022595/1) and CCP-NC (EP/T026642/1) and the EPSRC High-End Computing Consortium, UKCP (EP/P022561/1). GMR acknowledges support from the F.R.S.-FNRS. MSchei acknowledges support from the European Union’s Horizon 2020 research and innovation program under the European Union’s Grant agreement No. 676580 (NoMaD) and financial support from the Max Planck research network on big-data-driven materials science (BiGmax). CO, CT, PC, MEs, FRo, SC acknowledge support by DOD-ONR (N00014-17-1-2090, N00014-17-1-2876) and by the National Science Foundation under DMREF Grant No. DMR-1921909.

Author contributions statement

All authors participated in discussions that lead to the OPTIMADE specification version 1.0, in particular, by taking part to the OPTIMADE workshops that took place at the Lorentz Center in Leiden (Netherlands) in 2016, and at the CECAM in Lausanne (Switzerland) in 2018, 2019, and 2020. They also revised and approved the final version of the paper.

CWA, RAr, EB, GJC, SD, MLE, AF, AG, SG, AMe, FM, CO, GPi, GMR, MSchei, LT, CT, and DWi provided the key ideas for the design of the API, took part to the discussions about the issues opened on the Materials Consortia GitHub repositories, writing its specifications, and prepared the paper.

CWA and MLE, with help from AF, SD and MSchei, write and maintain the `optimade-python-tools` library, which is based on initial work by DWi and MWu, with contributions from RAr, KC, AG, AMe, FM, LT and TP. This library forms the basis of the Materials Cloud, Materials Project, NOMAD and odbx providers as well as the client hosted on Materials Cloud.

CD, NM, GH, GMR, and MSchef started the OPTIMADE workshop series, GMR and NM suggesting to dedicate these to the creation of a common API for databases of materials properties. RAr, GJC, CD, SG, NM, GH, GPi, GMR, MSchef, MSchei, and CT were on the organizing committee. CWA implemented the OPTIMADE server for the AiiDA framework and the Materials Cloud databases with contributions from GPi, LT, SK.

RAr has contributed very actively to the OPTIMADE specification GitHub repository. He also implemented the OPTIMADE server in the high-throughput toolkit (*httk*) open source framework used for the *omdb* database.

SD implemented the OPTIMADE server for the Materials Project provider, with assistance from MKH and PH. MKH wrote the OPTIMADE client in *pymatgen*.

MLE implemented the OPTIMADE server for the odbx provider, with assistance from AJM.

AG and ML implemented OPTIMADE server for OQMD provider with assistance from VH SG, AMe, and LT contributed to the implementation of the OPTIMADE::Filter Perl library for the syntactical analysis of the OPTIMADE filter language.

MSchei implemented the OPTIMADE server for the NOMAD provider with contributions from FM.

CT, CO, PC, ME, FRo, and SC implemented the OPTIMADE server for the AFLOW provider.

AV contributed to the long-term maintenance of the COD, to make it suitable for presentation via the OPTIMADE API.

Competing interests

GJC is a shareholder and Director of Intellegens Ltd. GH, GPe, GMR and DW are shareholders and Directors of Matgenix SRL. The other authors declare no competing commercial interests.

Figures & Tables

Provider	N₁	N₂	N₃
AFLOW ^{43,44}	700,192	62,293	382,554
Crystallography Open Database (COD) ^{45,46}	416,314	3,896	32,420
Theoretical Crystallography Open Database (TCOD) ²²	2,631	296	660
Materials Cloud ^{9,19,20}	886,518	801,382	103,075
Materials Project ^{28,36,47,48}	27,309	3,545	10,501
Novel Materials Discovery Laboratory (NOMAD) ^{29,49}	3,359,594	532,123	1,611,302
Open Database of Xtals (odbx) ^{31,50}	55	54	0
Open Materials Database (omdb) ^{51,52}	19,317	396	3,303
Open Quantum Materials Database (OQMD) ⁵³	153,113	11,011	70,252

Table 1. Materials databases with active OPTIMADE API implementations and the number of entries they return for the filters presented in this paper. The columns N_i provide hyperlinks for running the filters against each provider's implementation. The OPTIMADE website provides an up-to-date record of the implementation status of the databases at <https://www.optimade.org/providers-dashboard/>. AFLOW, Materials Project, odbx, omdb, and OQMD comprise computational materials data generated using database specific workflows^{28,44,50,51,53}. For the purposes of this table, Materials Cloud results were aggregated across all provided sub-databases. COD and TCOD comprise experimental and theoretical crystal structure data extracted from the literature. Materials Cloud comprises materials data from computational workflows; sub-databases group data by research project and can be contributed by users⁹. NOMAD aggregates computational data from multiple sources including from several of the repositories listed here.