



HAL
open science

Materials Databases: The Need for Open, Interoperable Databases with Standardized Data and Rich Metadata

François-Xavier Coudert

► **To cite this version:**

François-Xavier Coudert. Materials Databases: The Need for Open, Interoperable Databases with Standardized Data and Rich Metadata. *Advanced Theory and Simulations*, 2019, 2 (11), pp.1900131. <10.1002/adts.201900131>. <hal-02492782>

HAL Id: hal-02492782

<https://hal.science/hal-02492782v1>

Submitted on 5 Aug 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Materials databases: the need for open, interoperable databases with standardized data and rich metadata

François-Xavier Coudert*

Chimie ParisTech, PSL University, CNRS, Institut de Recherche de Chimie Paris, 75005 Paris, France

(Dated: July 8, 2019)

Driven by the recent rapid increase in the number of materials databases published (open and commercial), I discuss here some perspectives on the growing need for standardized, interoperable, open databases. The field of computational materials discovery is quickly expanding, and recent advances in data mining, high throughput screening, and machine learning highlight the potential of open databases.

One of the recent important trends in materials science is the emergence of several large-scale online databases of materials, trying to bring together experimental data with computational techniques in order to understand the behavior of materials families and design novel materials through data-mining? Maybe the most visible effort in this area is the Materials Project,[1] a database of computed information on known and predicted properties, part of the US-funded Materials Genome Initiative[2] launched in 2011. However, this trend is more general, and an increasing amount of research in the field focuses on the generation of these databases, their extension with additional data, and the use of these databases for analysis, screening, and prediction. This was clearly exemplified at recent materials meetings and materials modeling workshop, such as the MOFSIM 2019[3] meeting whose excellent discussion provoked this short comment.

I. EXISTING DATABASES

The need for aggregation of curated data in the physical and chemical sciences was recognized very early, and possibly the best-known database in our field is the CRC Handbook of Chemistry and Physics[4] (also known as the “rubber book”), which has been published since 1914. When it comes to materials, databases ranging in size and scope have emerged since the beginning of computing, and have advanced at the same pace as both computer hardware and networking capabilities. The Cambridge Structural Database (CSD),[5] launched in 1965, is one of the first numerical scientific databases, operating as a repository for validated experimental crystal structures of organic and organometallic compounds. Additional structural databases have emerged focusing on other categories of materials, including the Inorganic Crystal Structure Database (ICSD),[6] the Protein Data Bank (PDB),[7] the American Mineralogist Crystal Structure Database,[8] and the Crystallography Open Database (COD).[9]

In addition to these structural databases, databases of materials properties have also been compiled, either

by standardization institutes, learned societies, or commercial entities. We can cite examples of the NIST databases for materials and fluids properties, the glass property database SciGlass, Pearson’s Crystal Data, or more specific examples like the Polymer Gas Separation Membrane Database from the Membrane Society of Australasia. Moreover, databases of hypothetical structures — predicted by theory or computations — have also appeared over time. In the field of porous materials, for example, databases of computed zeolitic structures have been published more than 15 years ago.[10, 11] More recently, this effort has intensified — probably in response to the increase in capacity of computations, as well as the ease of hosting large datasets online. Databases of various scales, containing hypothetical (enumerated) metal–organic frameworks, have been published.[12] Other groups have worked on refining experimental databases to include additional data (such as atomic charges), making them suitable for computational applications and screening.[13, 14] All these databases, however, are hosted independently as archive files, with heterogeneous file formats, on individual research groups’ websites.

We note, however, that there have been some recent initiatives in order to integrate data from different sources into larger, coherent databases. This is particularly the case of computational data, whose volume increases with high-performance computing capabilities. The goals differ for the various initiatives, but in general, they aim at providing large-scale platforms for open science and data sharing, as well as improve discoverability and searchability of existing data. A first example is the Materials Project,[1] the aim of which is to “remove guesswork from materials design in a variety of applications” by computing properties of all known materials (and many that are not yet synthesized, too) through electronic structure analyses — a project funded as part of the bigger Materials Genome Initiative[2]. It aggregates structural data from other existing databases, as well as physical and chemical properties (band structures, elastic constants, piezoelectric tensors, electrode properties) computed as part of the project itself. Other platforms have been built to sharing results and resources in computational materials sciences, such as the ioChem-BD digital repository[15] and the more recent Materials

* fx.coudert@chimieparistech.psl.eu

Cloud[16].

II. THE CURRENT STATE OF AFFAIRS

As shown above, the number of existing databases is increasing at a fast pace. Yet, the datasets themselves are often hosted using makeshift solutions, in different places. They are typically too large to be hosted as supporting information by the publisher of the associated research paper, and they might be expanded and refined over time, with the publication of updates not related to a particular peer-reviewed paper. In the most common case, they are hosted on the web server of the research group, on institutional repositories offered by some universities, or on free data hosting solutions such as Figshare,[17] GitHub,[18] etc. This leads to the available data being dispersed between several platforms, and in some cases, raises the question of long-term availability of the data (e.g., when it is hosted on a group web server or a commercial data hosting solution). Moreover, there exists no universal way to access these data, unlike what exists for open archives where protocols such as OAI-PMH have been developed for interoperability and discoverability of data sources.

This fragmentation of the landscape of materials databases is accompanied by a large heterogeneity in the formats used: while for crystalline material structures, the CIF (Crystallographic Information File) format is predominant, the data made available as part of the CIF file is not homogeneous between different groups. The use of symmetry operators, for example, is not always consistent, with some databases being stored with symmetry systematically lowered to $P1$. In addition to this heterogeneity in data format, there is also a general lack of availability of metadata, meaning that most of the databases do not contain information about how the data was generated, gathered, curated, and possibly updated. Yet, this metadata can be crucial in exploiting databases, in order to identify identical or related data items, to understand how databases evolve over time, and to allow further investigation of specific data items. Metadata enables researchers to answer simple queries such as: Where does this structure come from? Where was it first reported, and under what conditions was it synthesized? How was this computational property calculated? What are the conditions of reuse of this data?

Here, before highlighting some requirements for truly open materials databases, we want to introduce the FAIR data principles. The FAIR principles are a set of guidelines in order to make data findable, accessible, interoperable and reusable (Figure 1). They have been formalized in 2016 by a consortium of scientists and organizations,[19] and were formally endorsed by the G20 Leaders at their 2016 summit in Hangzhou (杭州), China, in order to “promote open science and facilitate appropriate access to publicly funded research results”. [20] Requirements for FAIR data are:




FIG. 1. FAIR data principles:[19] make data findable, accessible, interoperable and reusable.

- **Findable:** data are indexed in a searchable resource, has a stable unique identifier, and is described with rich metadata.
- **Accessible:** data are retrievable using a standardized communications protocol, that is open, free, and universally implementable.
- **Interoperable:** data are represented using an open, well-defined format; data and metadata are interlinked.
- **Reusable:** data contains relevant metadata about its origin, a clear and accessible data usage license, and meet the community standards of its domain

III. REQUIREMENTS FOR OPEN DATABASES

Based on these formal requirements, and during the discussions at MOFSIM 2019 and other workshops, there appears to be a need for more open and interoperable materials databases. We outline below some of the requirements, drawing both on the FAIR principles and shared experiences and discussions.

Open databases. A majority of the available data on materials is produced as part of academic research, and largely funded by public monies. Moreover, for data associated with public research, there is a general consensus that it should be accessible to all. Recent years have seen the addition of “data availability” requirements in several journals, and it is now a required part for any funding application. This creates a need for open databases, where content can be hosted regardless of its origin (contrary to institution-wide repositories) and where it is accessible to all. Moreover, the database should provide clear information about its users’ rights when it comes to reusing the data, mining it, and republishing derived products.

Because open databases do not charge their users for subscription or access, they allow the dissemination of knowledge to categories of users that would otherwise find it difficult to obtain access: researchers in developing countries, nongovernmental organizations, independent researchers, journalists, even enthusiastic citizens.

Operating such open databases of course requires funding, and it is important to note that a number of national and supranational initiatives have been launched in that direction (as discussed in the introduction).

Interoperable databases. It is relatively clear that, given the vastly different needs of scientists working in different areas of materials science, there can be no “one size fits all” database, i.e., no single centralized database that fulfills the needs of every different community. So, how can a good balance be reached in developing specific topical databases while retaining some uniformity, in order to avoid a complete fragmentation of the field? It turns out, this problem is one that has been worked on for many years in a related area, that of document (or papers) archives. While there are many different open archives on the internet, they have been developed in a way that allows interoperability between them. Specifically, the Open Archives Initiative (OAI) has standardized a Protocol for Metadata Harvesting (OAI-PMH) through which each archive exposes its metadata, in a common format, allowing for cross-database search and discoverability.[21]

In order to achieve this goal, several design choices are needed. One is the use of a well-documented, standardized Application Programming Interface (API). Through the use of that API, the data does not have to be retrieved with a database-specific client or web portal,[22] but can be written in any programming language without inside knowledge of how the database operates internally. This means, in turn, that code that is developed for one specific database will work seamlessly with all others.

Another is the inclusion of data in standard, publicly-documented file formats. Given that most current databases are currently structural databases, a part of this problem has already been addressed in the several past decades: crystal structures are uniformly reported in CIF format (although the details available are not always consistent), macromolecular structures are consistently in PDB format, etc. However, there is currently no unified format for storing the properties of these materials. This is made difficult by the fact that properties are rather diverse in their mathematical nature: some are dimensionless but others have units; some are integers or half-integers, others vary continuously; some are scalars, others are matrices or higher-order tensors. Moreover, they sometimes need to be accompanied by additional information: expected uncertainty, reference orientation, etc.

With rich metadata and interlinked datasets. Metadata can be defined, in its simplest form, as “data that provides information about other data”. Like in any database, in a materials database metadata is crucial in assessing the data present, answering questions such as: How was this data gathered, by whom, when? In which conditions was a given property measured? For compu-

tational information, what was the theoretical method used, what is the level of description of the system? This is particularly important in databases of computational properties, where there can be a clear influence — and sometimes even a systematic bias — of the computational method chosen on the physical and chemical data calculated. If metadata is present in the databases, it opens the door to large-scale systematic explorations of various theoretical methods, and their comparison with experimental results obtained with different techniques, too.

Moreover, metadata can also provide much-needed links between several different interoperable databases. If a unique identifier is given for each dataset, and databases are interlinked through their metadata, it provides a much simpler exploration for users. It makes it easy to determine, e.g., if two properties from two datasets are independent or come from the same original calculation. It also allows greater discoverability, making it possible to find other properties in other databases, related to any given entry.

Curation that preserves the scientific record. The requirements listed above do not stop a fixed point of time, but instead must be considered throughout the databases’ timeline. For example, metadata can record the time of measurement of a given data, but also its time of inclusion in the database, and its further history. Indeed, with any database of significant size, it is expected that curation of the data is an important topic, and the dataset will be modified to remove errors, updated to reflect new measurements, and sometimes data will be removed for a variety of legitimate reasons. However, for the sake of research reproducibility and conserving the scientific record, it is important that such modifications be recorded in the metadata — just like corrections and retractions are publicly announced and archived for scientific articles. To my knowledge, this is not currently the case in existing databases, although the Materials Project is publishing “release logs”[23] which are kept on a separate page, but not recorded in the database as metadata.

Long-term availability. Finally, this discussion cannot be concluded without addressing the issue of long-term availability of the deposited data, meaning that it is necessary, over time, to build institutional support with long-term commitments. This also requires planning for what happens if and when the hosting institutions decide to “pull the plug” on the project. There, having an open database with an API for direct access to bulk data is a benefit, because it means other interested parties can duplicate the database and take over hosting. Other lessons can be learned from open archives, and initiatives in that field such as CLOCKSS, a long-term preservation project for articles and books with highly-redundant mirroring.[24]

-
- [1] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, *APL Mater.* **1**, 011002 (2013).
- [2] A. White, *MRS Bull.* **37**, 715 (2012).
- [3] MOFSIM2019 meeting, April 10–12 2019, Ghent (Belgium).
- [4] J. R. Rumble, *CRC Handbook of Chemistry and Physics, 99th Edition* (CRC Press, 2018).
- [5] C. R. Groom, I. J. Bruno, M. P. Lightfoot, and S. C. Ward, *Acta Cryst. B* **72**, 171 (2016).
- [6] M. Hellenbrandt, *Cryst. Rev.* **10**, 17 (2014).
- [7] H. M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov, and P. E. Bourne, *Nucleic Acids Research* **28**, 235 (2000).
- [8] R. T. Downs and M. Hall-Wallace, *American Mineralogist* **88**, 247 (2003).
- [9] S. Gražulis, A. Daškevič, A. Merkys, D. Chateigner, L. Lutterotti, M. Quirós, N. R. Serebryanaya, P. Moeck, R. T. Downs, and A. Le Bail, *Nucleic Acids Research* **40**, D420 (2012).
- [10] Y. Li, J. Yu, D. Liu, W. Yan, R. Xu, and Y. Xu, *Chem. Mater.* **15**, 2780 (2003).
- [11] D. J. Earl and M. W. Deem, *Ind. Eng. Chem. Res.* **45**, 5449 (2006).
- [12] C. E. Wilmer, M. Leaf, C. Y. Lee, O. K. Farha, B. G. Hauser, J. T. Hupp, and R. Q. Snurr, *Nature Chem* **4**, 83 (2012).
- [13] Y. G. Chung, J. Camp, M. Haranczyk, B. J. Sikora, W. Bury, V. Krungleviciute, T. Yildirim, O. K. Farha, D. S. Sholl, and R. Q. Snurr, *Chem. Mater.* **26**, 6185 (2014).
- [14] D. Nazarian, J. S. Camp, Y. G. Chung, R. Q. Snurr, and D. S. Sholl, *Chem. Mater.* **29**, 2521 (2016).
- [15] Available online at <https://www.iochem-bd.org/>.
- [16] Available online at <https://www.materialscloud.org/>.
- [17] Available online at <https://www.github.com/>.
- [18] Available online at <https://www.figshare.com/>.
- [19] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, J. Bouwman, A. J. Brookes, T. Clark, M. Crosas, I. Dillo, O. Dumon, S. Edmunds, C. T. Evelo, R. Finkers, A. Gonzalez-Beltran, A. J. Gray, P. Groth, C. Goble, J. S. Grethe, J. Heringa, P. A. 't Hoen, R. Hooft, T. Kuhn, R. Kok, J. Kok, S. J. Lusher, M. E. Martone, A. Mons, A. L. Packer, B. Persson, P. Rocca-Serra, M. Roos, R. van Schaik, S.-A. Sansone, E. Schultes, T. Sengstag, T. Slater, G. Strawn, M. A. Swertz, M. Thompson, J. van der Lei, E. van Mulligen, J. Velterop, A. Waagmeester, P. Wittenburg, K. Wolstencroft, J. Zhao, and B. Mons, *Sci. Data* **3**, e1002295 (2016).
- [20] “G20 Leaders’ Communique Hangzhou Summit”, available online at http://europa.eu/rapid/press-release_STATEMENT-16-2967_en.htm.
- [21] Available online at <http://www.openarchives.org/OAI/openarchivesprotocol.html>.
- [22] Although of course such clients can exist, providing a user-friendly way to query the database! The existence of a public API makes it possible for advanced users to develop their own portals, bringing added value to the database.
- [23] <https://discuss.materialsproject.org/t/materials-project-database-release-log/1609>.
- [24] Available online at <https://clockss.org>.