

Shared Metadata for Data-Centric Materials Science

Luca M. Ghiringhelli^{1,2}, Carsten Baldauf³, Tristan Berau⁴, Sandor Brockhauser¹, Christian Carbogno², Javad Chamanara⁵, Stefano Cozzini⁶, Stefano Curtarolo⁷, Claudia Draxl^{1,2}, Shyam Dwaraknath⁸, Ádám Fekete¹, James Kermode⁹, Christoph T. Koch¹, Markus Kühbach¹, Alvin Noe Ladines¹, Patrick Lambrix¹⁰, Maja-Olivia Lenz-Himmer², Sergey Levchenko², Micael Oliveira¹¹, Adam Michalchuk¹², Ron Miller¹³, Berk Onat⁹, Pasquale Pavone¹, Giovanni Pizzi¹⁴, Benjamin Regler², Gian-Marco Rignanes¹⁵, Jörg Schaarschmidt¹⁶, Markus Scheidgen¹, Astrid Schneidewind¹⁷, Tatyana Sheveleva⁵, Chuanxun Su¹⁸, Denis Usvyat¹⁹, Omar Valsson²⁰, Christof Wöll²¹, and Matthias Scheffler^{1,2}

¹Physics Department and IRIS Adlershof, Humboldt-Universität zu Berlin, Germany; ²The NOMAD Lab at the Fritz Haber Institute of the Max Planck Society and Humboldt University, Berlin, Germany; ³Fritz Haber Institute of the Max Planck Society, Berlin, Germany; ⁴Van 't Hoff Institute for Molecular Sciences and Informatics Institute, University of Amsterdam, Amsterdam 1098 XH, The Netherlands; ⁵TIB – Leibniz Information Centre for Science and Technology and University Library, Welfengarten 1B, 30167 Hanover, Germany; ⁶AREA Science Park, località Padriciano, 34149 Trieste, Italy; ⁷Center for Autonomous Materials Design and Department of Mechanical Engineering and Materials Science, Duke University, Durham NC 27708, USA; ⁸Lawrence Berkeley National Laboratory, Berkeley, CA, USA; ⁹Warwick Centre for Predictive Modelling, School of Engineering, University of Warwick, Coventry, CV4 7AL, United Kingdom; ¹⁰Department of Computer and Information Science, Linköping University, Sweden and The Swedish e-Science Research Centre, Linköping University, Sweden; ¹¹Max Planck Institute for the Structure and Dynamics of Matter, Hamburg, Germany; ¹²Federal Institute for Materials Research and Testing (BAM), 12489 Berlin, Germany; ¹³Department of Mechanical and Aerospace Engineering, Carleton University, Ottawa, ON K1S 5B6, Canada; ¹⁴Theory and Simulation of Materials (THEOS) and National Centre for Computational Design and Discovery of Novel Materials (MARVEL), École Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland; ¹⁵Institute of Condensed Matter and Nanosciences (IMCN), UCLouvain, Chemin des Étoiles 8, B-1348 Louvain-la-Neuve, Belgium; ¹⁶Institute of Nanotechnology, Karlsruhe Institute of Technology (KIT), 76344 Eggenstein-Leopoldshafen, Germany; ¹⁷Jülich Center for Neutron Science at MLZ, Forschungszentrum Jülich GmbH, Lichtenbergstraße 1, 85748 Garching, Germany; ¹⁸CAS Key Laboratory of Quantum Information, University of Science and Technology of China, Hefei 230026, People's Republic of China; ¹⁹Chemistry Department, Humboldt-Universität zu Berlin, Germany; ²⁰Department of Chemistry, University of North Texas, Denton, TX 76201, USA; ²¹Institute of Functional Interfaces, Karlsruhe Institute of Technology (KIT), 76344 Eggenstein-Leopoldshafen, Germany. (Dated: May 31, 2022)

The expansive production of data in materials science, their widespread sharing and repurposing requires educated support and stewardship. In order to ensure that this need helps rather than hinders scientific work, the implementation of the FAIR-data principles (*Findable, Accessible, Interoperable, and Reusable*) must not be too narrow. Besides, the wider materials-science community ought to agree on the strategies to tackle the challenges that are specific to its data, both from computations and experiments. In this paper, we present the result of the discussions held at the workshop on “Shared Metadata and Data Formats for Big-Data Driven Materials Science”. We start from an operative definition of metadata, and what features a FAIR-compliant metadata schema should have. We will mainly focus on computational materials-science data and propose a constructive approach for the *FAIRification* of the (meta)data related to ground-state and excited-states calculations, potential-energy sampling, and generalized workflows. Finally, challenges with the *FAIRification* of experimental (meta)data and materials-science ontologies are presented together with an outlook of how to meet them.

I. METADATA AND FAIR DATA PRINCIPLES

The amount of data that has been produced in materials science till today and its day-by-day increase are massive¹. The dawn of the data-centric era² requires that such data are not just stored, but also carefully annotated in order to find, access, and possibly reuse them. Terms of good practice to be adopted by the sci-

entific community for the management and stewardship of its data, the so-called FAIR-data principles, have been compiled by the FORCE11 group.³ Here, the acronym FAIR stands for *Findable, Accessible, Interoperable, and Reusable*,⁴ which applies not only to *data*, but also to *metadata*.

Before examining the meaning and importance of the four terms of the FAIR acronym, it is worth defining what metadata are with respect to data. To the pur-

pose, we start by introducing the concept of *data object*. A *data object* is the collective storage of information related to an elementary entry in a database. One can consider it as a row in a table, where the columns can be occupied by simple scalars, higher-order mathematical objects, strings of characters, or even full documents (or other media objects). In the materials-science context, a *data object* is the collection of attributes (the columns in the above-mentioned table) that represent a material or, even more fundamentally, a *snapshot* of the material captured by a single configuration of atoms, or it may be a set of measurements from well-defined *equivalent* samples. For instance, in computational materials science, the attributes of a *data object* could be both the inputs (e.g., the coordinates and chemical species of the atoms constituting the material, the description of the physical model used for calculating its properties), and the outputs (e.g., total energy, forces, electronic density of states, etc.) of a calculation. Logically and physically, inputs and outputs are at different levels, in the sense that the former determine the latter. Hence, one can consider the inputs as *metadata* describing the *data*, i.e., the outputs. In turn, the set of coordinates *A* that are metadata to some observed quantities, may be considered as data that depend on another set of coordinates *B*, and the forces acting on the atoms in that set *A*. So, the set of coordinates *B* and the acting forces are metadata to the set *A*, now regarded as data. Metadata can always be considered to be data as they could be objects of different, independent analyses than those performed on the calculated properties. In this respect, whether an attribute of a *data object* is data or metadata depends on the *context*. This simple example also depicts a *provenance* relationship between the data and their metadata.

The above discussion can be summarized in a more general definition of the term metadata: *Metadata are attributes that are necessary to locate, fully characterize, and ultimately reproduce other attributes that are identified as data.*⁵ *The metadata include a clear and unambiguous description of the data as well as their full provenance.*

The need for storing and characterizing data by means of metadata is determined by two main aspects, related to data usage. The first aspect is as old as science: *reproducibility*. In an experiment or computation, all the necessary information needed to reproduce the measured/calculated data (i.e., the metadata) should be recorded, stored, and retrievable. The second aspect becomes prominent with the demand for *reusability*⁴. Data can and should be also usable for purposes that were not anticipated at the time they were recorded. A useful way of looking at metadata is that they are attributes of *data objects* answering the “wh- questions”: who, what, when, where, why, and how. For example, “Who has produced the data?”, “What are the data expected to represent (in physical terms)?”, “When were they produced?”, “Where are they stored?”, “For what purpose were they produced?”, and “By means of which methods

were the data obtained?”. The latter two questions also refer to the concept of *provenance*, i.e., the logical sequence of operations that determine, ideally univocally, the data. Keeping track of the provenance requires the possibility to record the whole *workflow* that has lead to some calculated or measured properties (for more details, see Section VI).

From a practical point of view, the metadata are organized in a schema. We summarise what the FAIR principles imply in terms of a metadata schema as follows:

- *Findability* is achieved by assigning unique and Persistent Identifiers (PIDs) to data and metadata, describing data with rich metadata, and registering⁶ the (meta)data in searchable resources. Widely known examples of PIDs are digital object identifiers (DOIs) and (permanent) Uniform Resource Locators (URLs).
- *Accessibility* is enabled by application programming interfaces (APIs), which allow to query and retrieve single entries as well as entire archives.
- *Interoperability* implies the use of formal, accessible, shared, and broadly applicable languages for knowledge representation (these are known as formal ontologies and will be discussed in Section IX), use of vocabularies to annotate data and metadata, and inclusion of references.
- *Reusability* hints at the fact that data in materials science are often exploited only once for a focus-oriented research project, and many data are not even properly stored as they turned out to be irrelevant for the focus. In this sense, many data constitute a “waste” that can be recycled, provided that the data can be found and they are properly annotated. In this sense the “R” in FAIR could also mean *recyclable*. Another evocative interpretation could be *repurposable*, indicating that data may well be used, and valuable, for a different purpose than that for which they were initially created.

Establishing one or more metadata schemas that are FAIR-data-principles compliant, and that therefore enable the materials-science community to efficiently share the heterogeneously and decentrally produced data, needs to be a community effort. The workshop “Shared Metadata and Data Formats for Big-Data Driven Materials Science: A NOMAD-FAIR-DI Workshop” was organized and held in Berlin in July 2019 to ignite this effort. In the following sections, we describe the identified challenges and first plans, divided into different aspects that are crucial to be addressed in computational materials science. We close with an outlook on a metadata schema for experimental materials science and on the introduction of formal ontologies for materials-science databases.

In the next Section, we describe the identified challenges and first plans for FAIR metadata schemas for computational materials science, where we also summarize as an example the main ideas behind the metadata schema implemented in the Novel-Materials Discovery

(NOMAD) Laboratory for storing and managing millions of data objects produced by means of atomistic calculations (both *ab initio* and molecular mechanics), employing tens of different codes, which cover the overwhelming majority of what is actually used in terms of volume-of-data production in the community. We then follow with more detailed sections discussing the specific challenges related to *interoperability* and *reusability* for ground-state calculations (Section III), perturbative and excited-state calculations (Section IV), potential-energy sampling (molecular-dynamics and more, Section V), and workflows in general (Section VI) are addressed in detail in the following sections. Challenges related to the choice of file formats are discussed in Section VII. We close with an outlook of metadata schema(s) for experimental materials science and the introduction of formal ontologies for materials-science databases.

II. TOWARDS FAIR METADATA SCHEMAS FOR COMPUTATIONAL MATERIALS SCIENCE

The materials-science community has realized long ago that it is necessary to structure data by means of metadata schemas. In this Section, we describe the pioneering and recent examples of such schemas, and how a metadata schema becomes FAIR-data-principles compliant.

Perhaps the first systematic effort to build an exhaustive metadata schema for chemistry and materials science is the Chemical Markup Language (CML⁷⁻⁹), whose first public version was released in 1995. CML is a dictionary, encoded in XML for chemical metadata. XML is an extensible markup language, a general framework for encoding text documents in a format that is meant to be at the same time human and machine readable. CML is accessible (for reading, writing, and validation) via the Java library JUMBO (Java Universal Molecular/Markup Browser for Objects⁹). The general idea of CML is to represent with a common language all kinds of documents that contain chemical data, even though currently the language — as of the latest update in 2012¹⁰ — covers mainly the description of molecules (e.g., IUPAC name, atomic coordinates, bond distances) and of inputs/outputs of computational chemistry codes such as Gaussian03¹¹ and NWChem¹². Specifically, in the CML representation of computational chemistry calculations¹³, (ideally) all the information on a simulation that is contained in the input and output files is mapped onto a format that is in principle independent of the code itself. Such information is:

- Administrative data like the code version, libraries for the compilation, hardware, user submitting the job;
- Materials-specific (or materials-snapshot-specific) data like computed structure (e.g., atomic species, coordinates), the physical method (e.g., electronic exchange-correlation treatment, relativistic treatment), numerical settings (basis set, integration

grids, etc.);

- Computed quantities (energies, forces, sequence of atomic positions in case a structure relaxation or some dynamical propagation of the system is performed, etc ...).

The different types of information are hierarchically organized in *modules*, e.g., *environment* (for the code version, hardware, run date, etc.), *initialization* (for the exchange correlation treatment, spin, charge), *molgeom* (for the atomic coordinates and the localized basis set specification), *finalization* (for the energies, forces, etc.). The most recent release of the CML schema contains more than 500 metadata-schema items, i.e., unique entries in the metadata schema. Another long-standing activity is JCAMP-DX (Joint Committee on Atomic and Molecular Physical Data - Data Exchange)¹⁴, a standard file format for exchange of infrared spectra and related chemical and physical information was established in 1988 and then updated with IUPAC recommendations until 2004. It contains standard dictionaries for infrared spectroscopy, chemical structure, NMR¹⁵, and mass¹⁶ and ion-mobility spectrometry¹⁷. The European Theoretical Spectroscopy Facility (ETSF) File Format Specifications were proposed in 2007¹⁸⁻²⁰, in the context of the European Network of Excellence NANOQUANTA, in order to overcome widely known portability issues of input/output file formats across platforms. The initiative was inspired by the experience of the meteorology/climatology community, which developed the NetCDF (Network Common Data Form) library. The Electronic Structure Common Data Format (ESCDF) Specifications²¹ is the ongoing continuation of the ETSF project and is part of the CECAM Electronic Structure Library, a community-maintained collection of software libraries and data standards for electronic-structure calculations²².

The largest databases of computational materials-science data, AFLOW²³, Materials Cloud²⁴, Materials Project²⁵, the NOMAD Repository and Archive²⁶⁻²⁸, OQMD²⁹, and TCOD³⁰ offer application programming interfaces (APIs) that rely on dedicated metadata schemas. Similarly, AiiDa³¹⁻³³ and ASE³⁴, which are schedulers and workflow managers for computational materials-science calculations, adopt their own metadata schema. OpenKIM³⁵ is a library of interatomic models (force fields) and simulation codes that test the predictions of these models, complemented with the necessary first-principles and experimental reference data. Within OpenKIM, a metadata schema is defined for the annotation of the models and reference data. Some of the metadata in all these schemas are straightforward to map onto each other (e.g., those related to the structure of the studied system, i.e., atomic coordinates and species, and simulation-cell specification), others can be mapped with some care. The OPTIMADE (Open Databases Integration for Materials Design^{36,37}) consortium has recognized this potential and has recently released the first version of an API that allows users to access a common subset of metadata-schema items, independent of the schema

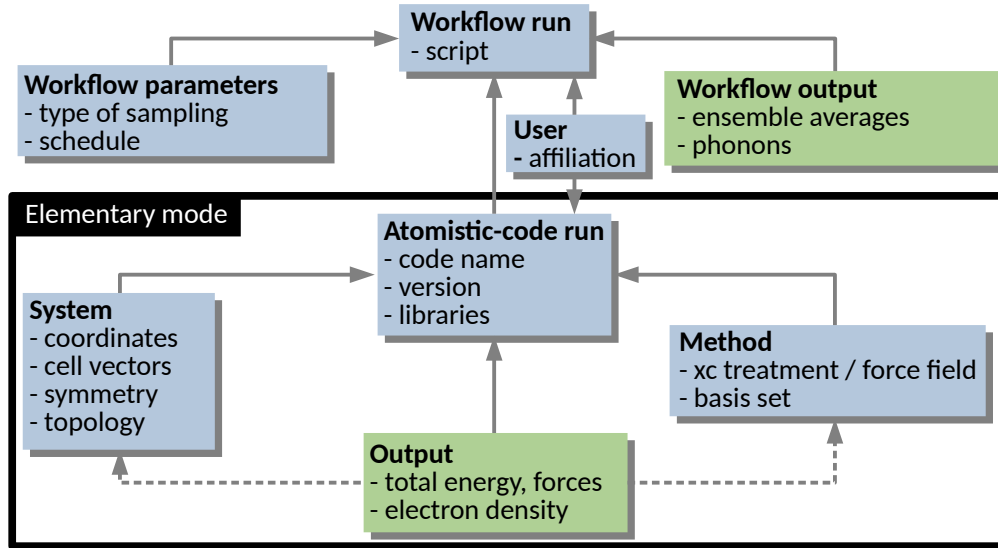


FIG. 1. Simplified schema of the *NOMAD Metainfo*. The rectangles symbolize the section-type metadata, for each section a few examples of therein contained quantity-type or (sub)sections metadata are listed. Sections are always written in bold font. The solid arrows stand for the *is contained in* relationship, while the dashed arrows are for the *has reference in* relationship.

adopted for any specific database/repository that is part of the consortium.

In order to clarify how a metadata schema can explicitly be FAIR-data-principles compliant, we describe as an example the main features of the *NOMAD Metainfo*, onto which the information contained in the input and output files of atomistic codes, both *ab initio* and force-field based, is mapped. The first released version of the *NOMAD Metainfo* is described in Ref. 21 and it has powered the NOMAD Archive since the latter went online in 2014, thus predating the formal introduction of the FAIR-data principles³.

Here, we give a simplified description, graphically aided by Fig. 1, which highlights the hierarchical/modular architecture of the metadata schema. The *elementary mode* in which an atomistic materials-science code is run (encompassed by the black rectangle) serves the computation of some observables (*Output*) for a given *System*, specified in terms of atomic species arranged by their coordinates in a box, and for a given physical model (*Method*), including specification of its numerical implementation. Sequences or collections of such runs are often defined via a *Workflow*. Examples of workflows are:

- Perturbative physical models (e.g., second-order Møller–Plesset, MP2, Green’s function based methods such as G_0W_0 , random-phase approximation, RPA) evaluated using self-consistent solutions provided by other models (e.g., density-functional theory, DFT, Hartree-Fock method, HF) applied on the same *System*;
- Sampling of some desired thermodynamic ensemble by means of, e.g., molecular dynamics;
- Global- and local-minima structure searches;
- Numerical evaluations of equations of state,

phonons, or elastic constants by evaluating energies, forces, and possibly other observables;

- Scans over the compositional space for a given class of materials (high-throughput screening).

The workflows can also be nested, e.g., a scan over materials typically contains local optimizations and might contain evaluation of phonons.

The *NOMAD Metainfo* organizes metadata into sections, which are represented in Fig. 1 by the labeled boxes. The sections are a *type* of metadata, which group other metadata, e.g., other sections or *quantity-type* metadata. The latter are metadata related to scalars, tensors, strings, which represent the physical quantities resulting from calculations or measurements. In a relational-database model, the sections would correspond to tables, where the *data objects* would be the rows, and the quantity-type metadata the columns. Each metadata-schema item has *attributes*³⁸, starting with its *name*, a string that must be globally unique, well-defined, intuitive, and as short as possible. Other attributes are the human-understandable *description*, which clarifies the meaning of the metadata, the *parent section*, i.e., the section the metadata belongs to, and the *type*, whether the metadata is, e.g., a section or a *quantity*. Another possible *type*, the *category* type, will be discussed below. For the quantity-type metadata, other important attributes are *physical units* and *shape*, i.e., the dimensions (scalar, vector of a certain length, a matrix with a certain number of rows and columns, etc.), and *allowed values*, for metadata that admit only a discrete and finite set of values.

All definitions in the *NOMAD Metainfo* have the following attributes:

- A globally unique qualified name;
- Human-readable/interpretable description and expected format (e.g., scalar, string of a given length, array of given size);
- Allowed values (intervals or explicit lists, or rules to check the validity);
- Provenance, which is realized in terms of a hierarchical and modular schema, where each *data object* is linked to all the metadata that concur to its definition. Related to provenance, an important aspect of *NOMAD Metainfo* is its *extensibility*. It stems from the recognition that reproducibility is an empirical concept, thus at any time, new, previously unknown or disregarded metadata may be recognized as necessary. The metadata schema must be ready to accommodate such extensions seamlessly.

The representation in Fig. 1 is very simplified for tutorial purposes. For instance, a workflow can be arbitrarily complex. In particular, it may contain a hierarchy of sub-workflows. In the currently released version of the *NOMAD Metainfo*, the elementary-code-run modality is fully supported, i.e., ideally all the information contained in a code run is mapped onto the metadata schema. However, the workflow modality is still under development. An important implication of the hierarchical schema is the mapping of any (complex) workflow onto the schema. That way, all the information obtained by its steps is stored. This is achieved by parsers, which have been written by the NOMAD team for each supported simulation code. One of the outcomes of the parsing is the assignment of a PID to each parsed *data object*, thus allowing for its localization, e.g., via a URL.

The *NOMAD Metainfo* is inspired by the CML, in particular in being hierarchical/modular. Each instance of each metadata-schema item is uniquely identified, so that it can be (and, in the NOMAD Archive, is) associated with a URL for its convenient accessibility. An instance of a metadata-schema item is the pairing of the item with a specific value from a *data object*, where the pairing operation is performed by a dedicated parser. The modularity and uniqueness together allow for a straightforward extensibility including customization, i.e., introduction of metadata-schema items that do not need to be shared among all users, but may be used by a smaller subset of users, without conflicts.

In Fig. 1, the solid arrows stand for the relationship *is contained in* between section-type metadata. Very few examples of quantity-type metadata are listed in each box/section. Such metadata are also in an *is-contained-in* relationship with the section they are listed in. The dashed arrows symbolize the relationship *has reference in*. In practice, in the example of an *Output* section, the quantity-type metadata contained in such a section are evaluated for a given system described in a *System* section and for a given physical model described in a *Method* section. So, the section *Output* contains a reference to the specific *System* and *Method* sections holding the nec-

essary input information. At the same time, the *Output* section *is contained in* a given *Atomistic-code run* section. These relationships among metadata already build a basic ontology, induced by the way computational data are produced in practice, by means of workflows and code runs. This aspect will be reexamined in Section IX.

We now come to the *category-type* metadata that allow for complementary, arbitrarily complex ontologies to be built by starting from the same metadata. They define a concept, such as “energy” or “energy component”, in order to specify that a given quantity-type metadata has a certain meaning, be it physical (such as “energy”) or computer-hardware related, or administrative. To the purpose, each section and quantity-type metadata is related to a category-type metadata, by means of an *is-a* kind of relationship. Each category-type metadata can be related to another category-type metadata by means of the same *is-a* relationship, thus building another ontology on the metadata, which can be connected with top-down ontologies such as EMMO³⁹ (see section IX for a short description of EMMO).

The current version of *NOMAD Metainfo* includes more than 400 metadata-schema items⁴⁰. To summarize, the *NOMAD Metainfo* addresses the FAIR-data principles in the following sense:

- *Findability* is enabled by unique names and a human-understandable description;
- *Accessibility* is enabled by the PID assigned to each metadata-schema item, which can be accessed via a ReSTful⁴¹ API, specifically developed for the *NOMAD Metainfo*. Essentially all NOMAD data are open access and users who wish to search and download data do not need to identify themselves. They only need to accept the CC BY license. Uploaders can decide for an embargo. These data are then shared with a selected group of colleagues. The embargo can last for up to 3 years.
- *Interoperability* is enabled by the extensibility of the schema and the category-type metadata, which can be linked to existing and future ontologies (see Section IX).
- *Reusability/Repurposability/Recyclability* is enabled by the modular/hierarchical structure that allows for accessing calculations at different abstraction scales, from the single observables in a code run to a whole complex workflow (see Section VI).

The usefulness and versatility of a metadata schema are demonstrated by the multiple access modalities it allows for. The *NOMAD Metainfo* schema is the basis of the whole NOMAD Laboratory infrastructure, which supports access to all the data in the NOMAD Archive, via the NOMAD API (also an implementation of the OP-TIMADE API³⁶ is supported). This API powers three different access modes of the Archive: the *Browser*⁴², which allows searches for single or groups of calculations, the *Encyclopedia*⁴³, which display the content of the Archive organized by *materials*, and the *Artificial-*

Intelligence (AI) Toolkit^{44,45}, which connects in Jupyter notebooks script-based queries and AI (machine learning, data mining) analyses of the filtered data. All the three services are accessible via a web browser running the dedicated GUI offered by NOMAD.

III. METADATA FOR GROUND-STATE ELECTRONIC-STRUCTURE CALCULATIONS

By ground-state calculations, we mean calculations of the electronic structure — e.g., single particle eigenvalues and wavefunctions, the electron density, the total energy and possibly its derivatives (forces, force constants) — for a fixed configuration of nuclei. This refers to a point located on the Born-Oppenheimer potential-energy surface, and is a necessary step in geometry optimization, molecular dynamics, the computation of vibrational (phonon) spectra or elastic constants, and more. Thus, ground-state calculations represent the most common task in computational materials science, and the involved approximations are relatively well established. For this reason, they are already extensively covered by the *NOMAD Metainfo*. Nevertheless, some challenges in defining metadata for such calculations still remain, as discussed below. In particular, density-functional theory (DFT) is the workhorse approach for the great majority of ground-state calculations in materials science. Highly accurate quantum-chemistry models are more computationally expensive than DFT and their use in applications is less widespread. However, they can provide accurate benchmark references for DFT, making high-quality quantum-chemical data essential also for DFT-based studies. Below we analyze the ground-state electronic structure calculations mainly in reference to DFT, but most of the stated principles are also valid for quantum-chemical calculations. A detailed discussion of the latter is deferred to Section IV C

A. Approximations to the DFT exchange-correlation functional

Approximations to the DFT exchange-correlation (xc) functionals are identified by a name or acronym (e.g., “PBE”), although sometimes this identification is not unique or complete. As metadata, we suggest to use the identifiers of the Libxc library^{46,47}, which is the largest bibliography of xc functionals. In order to be both human and computer friendly, the Libxc identifiers consist of a human-readable string that has a unique integer associated with it. Often, the above-noted identification needs some refinement, because xc functionals typically depend on a set of parameters and these may be modified for a given calculation. Obviously, there is a need to standardize the way in which such parameters are referenced. Just like it is possible to use the Libxc identifiers for the functionals themselves, one may also use the Libxc

naming scheme for their internal parameters. Obviously, code developers have to ensure that this information is contained in the respective input and/or output files. As Libxc provides version numbers of the xc functionals, it is important that this information is also available.

B. Basis sets

Complete and unambiguous specification of the basis set is crucial for judging the precision of a calculation. Ground-state calculations should include the full information about the basis sets used, including a DOI that a basis may be referred to. The use of repositories of basis sets, like the Basis Set Exchange repository⁴⁸, is therefore strongly recommended.

Basis sets can be coarsely divided into two classes, i.e., atom-position-dependent (atom-centered, bond-centered) and cell-dependent (such as plane waves) ones. Also a combination of both is possible, as, e.g., realized in augmented plane-wave or projector-augmented wave methods. For the atom-centered basis, the list of centers needs to be provided, and these may even contain positions where no actual atomic nucleus is located. The *NOMAD Metainfo* contains a rather complete set of metadata to describe atom-centered basis sets. A more complete description of cell-dependent basis sets can be found in the ESCDF, which is planned to be merged with the *NOMAD Metainfo*.

C. Energy reference

In order to enable *interoperability* and *reusability* of energies computed with different electronic-structure methods, it is necessary to define a “general energy zero”. An analysis of this problem and some clues on how to tackle it were already discussed by some of us in a previous work²¹. The following is a further attempt to advance and systematize ideas and solutions.

The problem of comparing energies is not restricted to computational materials science and chemistry. In fact, it also arises in experimental chemistry, as for instance, only enthalpy or entropy differences can be measured, but not absolute values. To solve this, chemists have defined a reference state for each element, called the *standard state*, which is defined as the element in its natural form at standard conditions, while the *heat of formation* is used to measure the change from the elements to the compound. In computational materials science and chemistry, we can adopt a similar approach. For each element we need to define a reference system as the zero of the energy scale. To do so, we introduce some definitions:

- A system is a defined set of one or more atoms, with a given geometry and, if periodic, a given unit cell. It can be an atom, a molecule, a periodic crystal, etc. If relevant, the charge, the spin-state or magnetic ordering needs to be specified.

- A reference system is a well-defined system to which other systems are compared to.
- A calculated energy is the energy obtained by a numerical simulation of a system with given input data and parameters, defining the Hamiltonian (i.e., DFT xc-functional approximation) or the many-electron model (e.g., Hartree-Fock, MP2, “coupled-cluster singles, doubles, and perturbative triples”, CCSD(T)), the basis set, and the numerical parameters.

Whether the reference system is an atom, an element in its natural form, some molecule or other system, does not matter, as long as it is well defined. Defining the system by atoms requires specifying how the orbitals are occupied, whether the atom is spherical, spin-polarized, etc. For each computational method and numerical settings, the energy per atom of the reference system must be calculated. The *standard energy* is then obtained by subtracting these values (multiplied by the number of constituents) from the calculated total energy. For example, to determine the energy of formation of a molecule like H_2O or a crystal like SiC , we calculate the difference in total energies $E(\text{H}_2\text{O}) - E(\text{H}_2) - \frac{1}{2}E(\text{O}_2)$, or $E(\text{SiC}) - E(\text{Si}) - E(\text{C})$, respectively. Here, H_2 and O_2 are isolated, neutral molecules while Si and C are free, neutral atoms. However, using the energy per atom of Si and C in their crystalline ground-state structure would be an option as well. We propose to tabulate the reference energies for the most common computational methods, so that they can be applied without further computations and preferably automatically by the codes themselves.

Finally, we need to define what is meant by a computational method. The Hamiltonian and DFT functional are clearly part of the definition as is the basis set and the potential shape (including pseudopotentials (PP) and effective core potentials). The specific implementation may also be relevant. Gaussian-based molecular-orbital codes may give the same energy for an identical setup (see Section IV C), while plane-wave DFT codes may not. One factor here is the choice of the PP. Irrespective of the used method, the computational settings determine the quality of a calculation. Most decisive here is the basis-set cut-off. For the plane-wave basis, convergence with respect to this parameter is straightforward. In any case, depending on the code, the method and details of the calculation, care needs to be taken to define all the adjustable parameters that significantly affect the energy when defining computational methods.

To tabulate standard energies, as suggested above, every computational method needs to be applied to all reference systems. This requires care in choosing the reference systems to ensure that an as-wide-as-possible range of codes and methods are actually suited for these calculations. It may be that some codes cannot constrain the occupancies of atoms, or keep them spherical, which would be a problem if spherical atoms were chosen as the reference. Clearly, periodic crystals such as

silicon are not suitable for molecular codes. It is possible, however, that some other codes could help bridging this gap. For example, FHI-aims⁴⁹ is not only capable of simulating crystalline system, but can also handle atoms and molecules and it can employ Gaussian-type orbitals (GTO) basis sets. Thus, FHI-aims is able to reproduce energy differences between atoms/molecules and crystals. In this way, it can support codes such as Gaussian or GAMESS.

IV. METADATA FOR EXTERNAL-PERTURBATION AND EXCITED-STATE ELECTRONIC-STRUCTURE CALCULATIONS

A direct link from the DFT ground state (GS) to excitations is provided by time-dependent DFT (TDDFT). Alternatively, charged and neutral electronic excitations are described by means of Green-function approaches from many-body perturbation theory (MBPT). This route is predominantly (but not exclusively) used for the solid state, while TDDFT and quantum-chemistry approaches are typically preferred for finite systems. For strongly correlated materials, in turn, dynamical mean-field theory (DMFT) is often the methodology of choice, potentially combined with DFT and Green-function methods. Lattice excitations, if not directly treated by DFT molecular dynamics, are often handled by density-functional perturbation theory (DFPT); for their interaction with light, also Green-function techniques are used. DFPT not only allows for the description of vibrational properties, but also for treating macroscopic electric fields, applied macroscopic strains, or combinations of these. The type of perturbation is intimately related to the physical properties of interest, e.g., harmonic and anharmonic phonons, effective charges, Raman tensors, dielectric constants, hyper-polarizabilities, and many others.

Characterizing the corresponding research data is a very complex and complicated task, for various reasons. First, such calculations rely on an underlying ground-state calculation, and thus carry along all uncertainties from it. Second, the methodology for excited states is scientifically and technically more involved by including many-body effects that govern diverse interactions. The methods thus rely on various, often not fully characterized approximations.

A. Diagrammatic techniques and TDDFT

The most common application of *GW* is to compute quasi-particle energies, i.e., energies that describe the removal or addition of a single electron. For this, the many-body electron-electron interaction is described by a two-particle operator, called the electronic self-energy. To compute this object, on the technical side we may need

an additional (auxiliary) basis set, not the same as the one used in the ground-state calculation, coming with additional parameters. Likewise, there are various ways for doing the analytical continuation of the Green’s function, as there are various ways for carrying out the required frequency integration, possibly employing a plasmon-pole model as an approximation. And there are also different ways of how to evaluate the screened Coulomb potential W . Most important is the flavor of GW , i.e., whether it is done in a single-shot manner, called G_0W_0 , or in a self-consistent way. If the latter, what kind of self-consistency (scf) is used — any type of partial scf, quasi-particle scf, or any other type which would remedy any starting-point dependence, i.e., the dependence of the results on the xc functional of the initial DFT (or Hartree-Fock or alike) used in the GS.

While GW is the method of choice for quasi-particle energies (and potentially also life times) within the realm of MBPT, we need to solve the Bethe-Salpeter equation (BSE) to tackle electron-hole interactions. This approach should typically be applied on top of a GW calculation, but often the quasi-particle states are approximated by DFT results adjusted by a scissors operator to widen the band gap in a similar way to the latter. In all cases, BSE carries along all subtleties from the underlying steps. In addition, it comes with its own issues, like the way of screening the Coulomb interaction (electron-hole this time), the representation of non-local operators, and alike.

DMFT, as a rather young and quickly developing field, naturally experiences a plentitude of “experimental” implementations, differing in many aspects, with one of the major obstacles being the quite vast amount of combinations of software. Some of the approaches are computationally light, allowing for the construction of model Hamiltonians based on DFT calculations; others are computationally too demanding and can be applied only to simple systems with a few orbitals; most of the methods rely on Green’s functions and self-energies. Diagrammatic extensions beyond standard DMFT methods employ various kinds of vertex functions. Other issues concern the definition of how to handle the Coulomb interactions, where the parameters can either be chosen empirically or can be calculated by first principles.

Specific issues of TDDFT concern, in a first place, the distinction between the linear-response regime and the time-propagation of the electronic states in presence of a time-dependent potential. For the former, the xc kernel plays the same role as the xc functional of the GS, raising (besides numerical precision) questions related to accuracy. For the latter, there are various ways and flavors for how to implement the time-evolution operator. Moreover, one can write this operator as a simple exponential or use more elaborate expressions, like the Magnus expansion or the enforced time reversal symmetry. Regarding the exponential, one can employ a Crank-Nicolson expansion, expand in a Taylor series or employ Houston states. Obviously, each of them comes with approxima-

tions and additionally, numerical issues.

In summary, all the variety captured by the different methods together with the related multitude of computational parameters, needs to be carefully reflected by the metadata schema. This is not only imperative for ensuring reproducible results but also for evaluating the accuracy of methods and commonly used approximations. Besides, further subtleties related to algorithms in the actual implementations in different codes requires the code developers to embark on this challenge.

B. Density-functional perturbation theory

Density-functional perturbation theory is used to obtain physical properties that are related to the (density-)response of the system to external perturbations, like the displacement potential according to lattice vibrations. Also in this case, the calculation relies on a preliminary GS run, inheriting all issues therefrom. After having chosen the type of perturbation, which requires method-dependent definitions and inputs, one needs to choose the order of perturbation: The linear response approach, that is implemented in many codes, allows for the determination of second-order derivatives of the total energy. Among these codes, some of them also allow for the calculation of third-order derivatives, like anharmonic vibrational effects. The variation of the Kohn-Sham orbitals can be obtained from the Sternheimer equation, where different methods are used for deriving its solution (iterative methods, direct linearization, integral formulation).

C. Quantum-chemistry methods

Quantum chemistry offers several methodological hierarchies for calculating quantities related to excited states, such as excitation energies, transition moments, ionization potentials, etc. As high-quality methods are computationally intensive, without additional approximations such methods can be applied to relatively small molecular systems only.

Among the standard quantum chemical approaches that can be routinely applied to study excited states of small to medium-sized molecules one can distinguish two large groups, i.e., single-reference and multi-reference methods. The single-reference coupled-cluster (CC) hierarchy for excited states can be formulated in terms of the so-called equation-of-motion approach or time-dependent linear response.

Generally, for well-behaving closed-shell molecules, the single-reference quantum-chemical methods can be used as a black box. The formalisms of the MP_n and CC models are uniquely defined and well documented. The GTO basis sets from the standard basis set families (Pople, Dunning, etc.) are also uniquely defined by the acronym. In practical implementations of these methods, of course

various thresholds are usually introduced for prescreening, convergence, etc., but the default values for these thresholds are routinely set very conservatively to guarantee a sub-microhartree precision of the final total energies. Problems might, however, arise due to the iterative character of most of the mentioned techniques, as convergence to a certain state (both in the ground-state and/or excited-state parts of the calculation) depends on starting guess, preconditioner, possible level shifts, type of convergence accelerator, etc. Unfortunately, the parameters that control the convergence are often not sufficiently well documented and might not be found in the output. Such problems mainly occur in open-shell cases (note that in the Delta methods at least one of the calculations has to involve an open-shell system). Sometimes a cross-check between several codes becomes essential to detect convergence faults.

When it comes to larger systems and approximate CC models are utilized, the importance of the involved tolerances and underlying protocols substantially increases. The approximations can include, for example, the density-fitting technique, local approximation, Laplace transform, and others. Important parameters here are the auxiliary basis set, the fitting metric, the type of fitting (local or non-local), and if local, how the fit-domains are determined, etc. The result of the calculations that use local correlation techniques are influenced by the choice of the virtual space and the corresponding truncation protocols and tolerances, the pair hierarchies and the corresponding approximations for the CC terms, etc. For Laplace-transform-based methods, the details of the numerical quadrature matter. Unfortunately, these subtleties are very specific and technical and even if given in the output, can hardly be properly understood and analyzed by non-specialists who are not involved in the development of the related methods. Therefore, the protocols behind the approximations are usually appropriately automatized, and the defaults are chosen such that for certain (benchmarking) sets of systems the deviations in the energy are substantially smaller than the expected error of the method itself (e.g., 0.01 eV for the excitation energy). However, for these methods, additional benchmarks and cross-checks between different programs and approaches would be very important.

Multi-reference methods come with quite a number of different flavors, where the most widely used ones are complete active-space self-consistent Field (CASSCF), complete active-space second-order perturbation theory (CASPT2), and multireference configuration interaction (MRCI). For difficult cases (e.g., strongly correlated systems), these methods might remain the only option to obtain qualitatively and quantitatively correct result. Unfortunately, compared to the single reference methods, they are computationally expensive and much less of a black box. First of all, for each calculation one has to specify the active space or active spaces. The results may depend dramatically on this choice. Furthermore, the underlying theory is not always uniquely defined by

the used acronym. For example, different formulations of CASPT2, MRCI, or other theories are not mutually equivalent depending on whether and how much internal contraction is used and additional approximations that neglect certain terms (e.g., many-electron density matrices) can be implicitly invoked. Besides, certain deficiencies of these methods, such as for example lack of size consistency in MRCI or intruder states in CASPT2, are often corrected by additional (sometimes empirical) schemes, which again are not always fully specified. All this makes the interpretation of deviations in results and cross-checks of these methods less conclusive.

To summarize, quantum-chemical methods offer an excellent toolbox for accurate *ab initio* calculations for molecules (especially so for small and medium sized ones). However, severe issues concerning reproducibility and replicability remain, in particular for extended and/or open-shell systems. This calls for a more detailed specification of the implemented techniques by the developers, for example, a better design of the outputs, and a thorough analysis and documentation of the employed methods and parameters by the users.

V. METADATA FOR POTENTIAL-ENERGY SAMPLING

Molecular dynamics (MD) simulations — i.e., following the time evolution of a system, employ either *ab initio* calculated forces and energies (aiMD) or molecular mechanics (MM) i.e., forces and energies are defined through empirical atomistic and coarse-grained potentials. The FAIR storing and sharing of their inputs and outputs comes with a number of specific challenges in comparison to single-point electronic-structure calculations.

Conceptually, aiMD and MM are similar, as a sequence of system configurations is evolved at discrete time steps. Positions, velocities, and forces at a given time step are used to evaluate positions and velocities, and hence forces in the new configuration, and so on. In practice, MM simulations are orders of magnitude faster than aiMD, enabling much longer time scales and/or much larger system sizes. Even though the trend towards massive parallelization will enable aiMD in the near future system to handle sizes comparable to today’s standards for MM simulations, the latter will probably always enable larger systems. However, with machine-learned potentials and active learning techniques for their training, aiMD and MM may grow together in the future.

In this Section, we focus on challenges more specific to MM simulations, having in mind large length scales, long time scales, and complex phase-space-exploration algorithms and workflows. They can be summarized as follows:

- (i) In many cases, the investigated systems feature thousands of atoms with complex short- and long-range order and disorder, e.g., describing microstructural evolution such as crack propagation.

This requires large, complex simulation cells with a range of chemical species to be correctly described and categorized.

- (ii) Force-fields exist in a wide variety of flavors that require proper classification. On top of that, they allow for granular fine-tuning of the interactions, even for individual atoms. Faithfully representing complex force fields thus requires to also capture the chemical-bonding topology that is often needed to define the actual interactions.
- (iii) The large length and long time scales presently come together with a multitude of simulation protocols, which use specific boundary conditions, thermostats, constraints, integrators, etc. The various approaches enable the computations of additional observables to be computed and as statistical averages or correlations. Representing these properties, implies the need to efficiently store and access large volumes of data, e.g., trajectories, including positions, and possibly also velocities and forces, for each atom at each time step.

For the purpose of illustration, we start by identifying some typical use cases, then describe what is currently implemented in the NOMAD infrastructure and what is missing. The examples we adopt fall into two classes: (i) high throughput systems that are individually *simple* ($\sim 1\,000 - 10\,000$ particles) where the value of sharing comes from the ability to run analysis across many variants of, e.g., chemical composition or force field; (ii) sporadic simulations of very large systems or very long time scales which cannot readily be repeated by other researchers and thus are individually valuable to share. Examples of the first class, could be MD simulations in the *NVT* ensemble for liquid butane or bulk silicon, using well-defined standard force fields (e.g., CHARMM or Stillinger-Weber). Quantities of interest are typically computed during MD simulations (e.g., liquid densities). For flexibility, full trajectory files should also be stored but some important observables might be worth precomputing (e.g., radial distribution functions). The second class could include multi-billion atom MD simulations of dislocation formation⁵⁰ or solidification^{51,52} or very long time-scale simulations of protein folding⁵³. For more complex use cases, the current infrastructure as discussed in Section II is not yet sufficient. The challenges to be addressed are the need for support for (i) complex, heterogeneous, possibly multi-resolution systems; (ii) custom force fields; (iii) advanced sampling; (iv) classes of sampling besides MD (e.g., Monte Carlo, global structure prediction/search); (iv) larger simulations (i.e., need for sparsification of the stored data with minimal loss of information)

Complex systems include heterogeneous systems, e.g., adsorbate and surfaces, interfaces, solute (macro)molecules in solvent fluids, and multi-resolution systems, i.e., systems that are described at different

granularity. The representation of complex systems requires a hierarchy of structural components, from atoms, through moieties, molecules, and larger (super)structures. Annotating such complexity will require human intervention as well as algorithm for automatically recognizing the structural elements (see, e.g., Ref. 54).

Annotation of *force fields* into publicly accessible databases has been pioneered by OpenKIM³⁵ in materials science and MoSDeF⁵⁵ for soft matter. However, many simulations are performed with customized force fields. The field is already being augmented and will likely be further supported by machine-learning (ML) force fields. So far, the great majority of ML force fields are used only in the publication where they are defined. The *reusability*-oriented annotation of force fields, including ML ones, require also establishing a criterion for comparing them. Comparisons can be carried out by means of standardized benchmark datasets, with a well-defined set of properties. Differences among predicted properties can establish a metric for the similarity of the force fields.

Advanced sampling techniques (e.g, metadynamics⁵⁶, umbrella sampling⁵⁷, replica exchange⁵⁸, transition-path⁵⁹ and forward-flux⁶⁰ samplings) are typically supported by MD plugins, such as PLUMED⁶¹ and OpenPathSampling⁶². The input and output of these plugins will serve as the basis for the metadata related to these sampling techniques. In this regard, it would also be interesting to connect materials-science databases, such as the NOMAD Repository and Archive²⁶ or Materials Cloud Archive²⁴ to the PLUMED-NEST⁶³, the public repository of the PLUMED consortium⁶⁴, for example by allowing for automatic uploading of PLUMED input files to the PLUMED-NEST when uploading to the data repositories.

For *long time-* and *large length-scale* simulations, several questions arise: How should we deal with these simulations, where the extensive amount of data produced by MD simulations becomes overwhelmingly large to systematically store and share? Can we afford to store and share all of it? If the storage is limited or data retrieval is unpractically slow, how can we identify the significant and crucial part of the simulation to store it in a reduced form? Keeping the whole data locally and sharing the metadata with only the important parts of the simulations would be a viable alternative, assuming the different servers have enough redundancy. Standard analysis techniques such as similarity analysis and monitoring dynamics can also be used to identify the changes in structure and dynamics to store only the significant frames or specific regions in MD simulations (e.g., some QM/MM models uses large MM buffer-atom regions that may not be stored entirely). Further, the cost/benefit of storing versus running a new simulation must be weighed. On the other hand, researchers may soon face increased requirements from funding agencies to store their data for a number of years, in which case the present endeavour offers a convenient implementation. We note ongoing al-

gorithmic developments on compression algorithms for trajectories, see, e.g., Ref. 65.

VI. METADATA FOR COMPUTATIONAL WORKFLOWS

A computational workflow represents the coordinated execution of *repeatable* (computational) steps while accounting for *dependencies* and *concurrency* of tasks. In other words, a workflow can be thought as a script, a wrapper code that manages the scheduling of other codes, by controlling what should run in parallel, what sequentially and/or iteratively. This definition can be extended to workflows in experimental materials science or hybrid computational-experimental investigations, but, consistently with the previous sections, we limit the discussion to computational aspects only.

Once shared, workflows become useful building blocks that can be combined or modified for developing new ones. Furthermore, FAIR data can be reused as part of workflows completely unrelated to the workflows with which they were generated. An obvious example is artificial-intelligence-based data analytics, which can entail complex workflows involving data originally created for different purposes. During the last decade, the interest in workflow development has grown considerably in the scientific community⁶⁶ and various multi-purpose engines for managing calculation workflows, have been developed, including AFLOW^{23,67,68}, AiiDA^{31,69}, ASE³⁴, and Fireworks⁷⁰. Using these infrastructures, a number of workflows have been used for scientific purposes, like convergence studies⁷¹, equations of state (e.g., AFLOW Automatic Gibbs Library⁷² and the AiiDA common workflows ACWF⁷³), phonons⁷⁴⁻⁷⁷, elastic properties (e.g., the elastic-properties library for Inorganic Crystalline Compounds of the Materials Project⁷⁸, AFLOW Automatic Elasticity Library, AEL⁷⁹, ElaStic⁸⁰), anharmonic properties (e.g., the Anharmonic Phonon Library, APL⁸¹, AFLOW Automatic Anharmonic Phonon Library, AAPL⁸²), high-throughput in the compositional space (e.g., AFLOW Partial Occupation, POCC⁸³), charge transport (e.g., organic semiconductors^{84,85}), of covalent organic frameworks (COFs) for gas storage applications⁸⁶, of spin-dynamics simulations⁸⁷, high-throughput automated extraction of tight-binding Hamiltonians via Wannier functions⁸⁸, and high-throughput on-surface chemistry⁸⁹

There are two types of metadata associated to workflows. Thinking of a workflow as a code to be run, the first type of metadata characterizes the code itself. The second type is the annotation of a run of a workflow, i.e., its inputs and outputs. This type of metadata has been already described in Section II, together with a schematic list of possible workflow classes. It is important to realize that the inputs and outputs of the elementary-mode runs of the atomistic codes that are invoked in a workflow run are complemented by the inputs and outputs

of the overarching workflows. A simple example: In an equation-of-state type of workflow, the energy and volume per unit cell of each single configuration that is part of the workflow is the output of the elementary run of the code, while the energy-vs-volume equation of state, e.g., fit to the Birch-Murnaghan model, is an output of the workflow.

VII. FILE FORMATS

On an abstract level, a metadata schema is independent from its representation in computer memory, on a hard drive, or on just a piece of paper. But on a practical level, all data and metadata need to be managed, i.e., stored, indexed, accessed, shared, deleted, archived, etc. File formats used in the community address different requirement and intended use cases. Some file formats privilege human readability (e.g., XML, JSON, YAML) but are not very storage efficient, others are binary and overall optimized for efficient searches, but require interpreters to be understood by a person (e.g., HDF5). There are a few use-cases and data properties in the domain of computational material science that are worth mentioning. First, such data are very heterogeneous and contain many simple properties (e.g., the name of a used code, or a list of considered atoms) that are mixed with properties in the form of large vectors, matrices, or tensors (e.g., the density of states or wave functions). The number of different properties requires hierarchical organization (e.g., with XML, JSON, YAML, or HDF5). It is desirable that many properties are easily human readable (e.g., to quickly verify the sanity of a piece of data), on the other hand large matrices should be stored as efficiently as possible for storing, retrieving and searching. Second, there are use cases where random (non-sequential) access of individual properties is desirable (e.g., return all band structures from a set of DFT calculations). Third, computational-material-science (meta)data need to be archived (efficient storage, prevention of corruption, backups, etc.) on one side, but they also need to be shared via APIs, e.g., for search queries. This requires to transform (meta)data from one representation in one file format (e.g., BagIt and HDF5) to another representation in a different format (e.g., JSON or XML).

These use-cases and data properties lead to the following conclusions. Even on a technical level, (meta)data need to be handled independent of the file format. Pieces of it have to be managed in different formats, and we need to be able to freely transform from one representation into another. Furthermore, if many different resources (files, databases, etc.) are used to store (meta)data from a logically conjoined dataset, references to these resources become an important piece of metadata itself. We propose to use an abstract interface (e.g., implemented as a Python library) based on an abstract schema. This interface allows to manage (meta)data independent of the actual representation used underneath. Various imple-

mentations of such an abstract interface can then realize storage in various file formats and access to databases.

VIII. METADATA SCHEMA FOR EXPERIMENTAL MATERIALS SCIENCE

In contrast to computational materials science, in experimental materials science the atomic structure and composition is only approximately known. Several techniques are used to collect data that may be more or less directly interpreted in terms of the atomic and/or electronic structure of the material. In cases where the structure of the material is already known, careful characterization of properties helps to establish valuable relationships between structure and properties which, in turn, may help to refine theoretical models of these links. The inherent uncertainty in every measurement process causes the precision with which data can be reproduced to be, in most cases, lower than in theoretical/computational materials science. These uncertainties are present even in a well-characterized experimental setup, i.e., when a comprehensive set of metadata is used. In many cases it is not even the focus of an experiment to produce the most perfectly characterized data, but to invest just enough effort to address the specific question that drives the experiment.

The information available about the material whose properties are to be measured is also much less complete than in the computational world, where often the position of every atom is known. However, while physical measurements may be limited in their precision, the accuracy with which a physically observable quantity is obtained is by definition of being physically observable much higher than in computational materials science, where the accuracy of the obtained physical quantity may depend strongly on the validity of approximations being made.

The uncertainty in retrieving structure-property relationships in computational materials science, which depends on the suitability of the applied theoretical model and its computational implementation, translates in the real of experiments to an uncertainty in the atomic structure of the object that is being characterized and generally also some uncertainty in the measurement process itself. The metadata necessary to reproduce a given experimental data set must thus include detailed information about the material and its history together with all the parameters which are required to describe the state of the instrument used for the characterization. In most cases, both classes of metadata, i.e., those describing the material and those describing the instrument are going to be incomplete. While, for example, the full history of temperature, air pressure, humidity, and other relevant environmental parameters are not commonly tracked for the complete lifetime of a material (counter-examples exist, e.g., in pharmaceutical research), also information about the state of the instrument is not generally as compre-

hensive as it should ideally be (e.g., parameters are not recorded, or are not properly controlled, such as hysteresis effects in devices involving magnetic fields, or many mechanical setups).

To overcome part of the uncertainty in the data, one needs to collect as many metadata about the material and its history, as possible, including those that one has no immediate use for at the moment, but might potentially need in the future. Since most of the research equipment being used for characterization tasks is commercial instrumentation, collecting this metadata in an (ideally) fully automated fashion requires the manufacturer’s support. In many cases the formats in which scientific data are provided by these instruments is proprietary. Even if all the data to describe the instrument’s condition of operation are stored, large parts of them may get lost when using the vendor’s software to export the data to other formats; mostly because the “standard format” does not foresee storing vendor- and instrument-specific metadata.

At large research infrastructures like synchrotrons and neutron scattering facilities, where a significant fraction of instruments is custom-built, and data are often shared with external partners, standards for file formats and metadata structures are being agreed upon the NeXus standard. NeXus⁹⁰ defines hierarchies and rules on how metadata should be described and allows compliant storage using HDF5. Experimental research communities can learn from these activities and provide *NeXus-format application definitions* which describe necessary metadata that should be stored in a dataset, along with definitions for some optional metadata. This common file format for scientific data is slowly beginning to spread to other communities. Having a standard file format for different types of scientific data seems to be an important step forward towards FAIR data management, since it severely reduces the threshold to share data across communities. Note that NeXus provides a glossary and connected ontology which helps in machine interpretability and so in *reusability*.

While standard file formats are of very high value in making data *findable* and *accessible*, due to common use of keywords to describe a given parameter, they also make them more *interoperable*, since the barrier for reading the data is lowered. However, making experimental data truly reproducible requires in many cases more metadata to be collected. Only if the uncertainty with which data can be reproduced is well understood, they may also be *reusable*. As discussed in the previous paragraph, part of these metadata must be provided by manufacturers of commercially available components of the experimental setup. Often this just requires more exhaustive data export functions and/or proper, i.e. versioned descriptions, for all of the instrument-state-describing metadata which are being collected during the experiment. Additionally, it may be necessary to equip home-built lab equipment with additional sensors and functionalities for logging their signals.

Even with added sensors and automated logging of

all accessible metadata, in many cases, it is also necessary to compile and complete the record of metadata describing the current and past states of the sample that is being characterized by manually adding information and/or combining data from different sources. Tools for doing this in a machine-readable fashion are Electronic Lab Notebooks (ELNs) and/or Laboratory Information Management Systems (LIMS). Many such systems are already available^{91–102}, including open-source solutions that combine features of both ELN and LIMS into one software. Server-client-solutions that do not require a specific client, but may be accessed through any web browser have the advantage that information may be accessed and edited from any electronic device capable of interacting with the server. Such ease of access, combined with the establishment of rules and practices of holistic metadata recording about sample conditions and experimental workflows will also help to increase the reproducibility and with that the *reusability* of experimental data. The easier the use of such a system is, and the more apparent it makes the benefits of the availability of FAIR experimental data, the faster it will be adopted by the scientific community.

IX. OUTLOOK ON ONTOLOGIES IN MATERIALS SCIENCE

In data science, an ontology is a *formal representation of the knowledge of a community about a domain of interest, for a purpose*. As ontologies are currently less common in basic materials science than in other fields of science, let us explain these terms:

- *Formal representation* means that: 1) the ontology is a *representation*, hence it may be a simplification or a model of the target domain, and 2) the attribute *formal* communicates that the ontological terms and relationships between them must have a deterministic and unambiguous meaning. Furthermore, *formal representation* implies that the mechanism to specify the ontology must have a degree of logical processing capability, e.g., inference and reasoning should be possible. Crucially, the attribute *formal* refers to the fact that an ontology should be machine readable.
- *Knowledge* is the accumulated facts, information, and skills of the experts of the domain of interest that are represented in the ontology.
- The *community* influences the ontology in two aspects; 1) it implies an overall agreement between a group of experts/users of the knowledge as represented in the ontology and 2) it indicates that the ontology is not meant to convince a whole population nor wants to be universal. However, if it fulfills the requirements of bigger communities, the ontology will be adopted by broader audiences and will find its way towards standardization.
- The *domain of interest* is the common ground for

the community, e.g., a scientific discipline, a sub-ordinate of discipline, or a market section. It is often used as a boundary to limit the scope of the ontology. It is a proper tool to detect overlapping concepts, modularizing ontologies, and identifying extension and integration points.

- The *purpose* conveys the goals of the ontology designers so that the ontology is applicable to a set of situations. In many ontology design efforts, the purpose is formulated by a collection of so-called *competency questions*. These questions and the answers provided to them identify the intent and viewpoint of the designers and set the potential applications of the ontology.

In practice, ontologies are often mapped onto, and visualized by means of, directed acyclic graphs, where an edge is one of a well-defined set of relationships (e.g., *is a*, *has property*) and each node is a *class*, i.e., a concept which is specific to the domain of interest. Each node-edge-node *triple* is interpreted as a subject-predicate-object expression. For instance, in an ontology for catalysis, one could find the triples: “catalytic material – has property – selectivity”, and “selectivity – refers to – reaction product”. Ontologies address the *interoperability* requirement of FAIR data. By means of a machine-readable formal structure, which can be connected to an existing or *ex novo* derived metadata schema of a database, ontologies allow queries over various databases, even from different fields.

The literature already contains several ontologies created for representing (aspects of) materials science. The most ambitious project is probably EMMO³⁹, which stands for both European Materials Modelling Ontology, developed within the European Materials Modelling Council (EMMC), and Elemental Multiperspective Material Ontology. EMMO is designed to provide a formal way to describe the fundamental concepts of physics, chemistry, and materials science, to provide an all-purposes common ground for describing materials, models, and data that can be adapted by all sub-domains of condensed-matter physics and chemistry. The development of EMMO includes also a handful of *domain ontologies* that assume EMMO as top-level ontology¹⁰³. These domain ontologies span subjects such as “atomistic and electronic modeling”, “crystallography”, “mechanical testing”, and more. So far, however, EMMO and its domain ontologies have not been connected to existing databases.

Other domain-specific ontologies, not related to EMMO, have been developed. For instance, the Materials Ontology¹⁰⁴ was developed for the exchange of data among databases for thermal properties, the MatOnto ontology¹⁰⁵ addresses oxygen ion conducting materials in the fuel cell domain, the NanoParticle Ontology¹⁰⁶ maps properties of nanoparticles with the purpose of designing new nanoparticles with given properties, while the eNanoMapper ontology¹⁰⁷ focuses on assessing risks related to the use of nanomaterials from the engineering

point of view.

An application-oriented ontology is Materials Design Ontology (MDO)¹⁰⁸, developed under the guidance of the schemas from OPTIMADE³⁶, and therefore aimed at dealing with data from the various materials data repositories (AFLOW, Materials Project, etc.) on a common ground. In practice, MDO connects calculated structures with the calculated properties and the physical model adopted to calculate structures and properties. Furthermore, the provenance for each calculation, is also represented in MDO. It has recently been extended using text mining on thousands of journal articles¹⁰⁹.

The hierarchical structure of *NOMAD Metainfo* already includes ontological aspect. More specifically, it represents atomistic calculations, as performed by all the parsed simulation codes. *NOMAD Metainfo* contains already five types of relations between the metadata: (a) is subclass of, (b) is part of, (c) has reference, (d) has dimension and (e) has category. The latter relation, *has category* is introduced to describe conceptually physical quantities (e.g., “energy”, “velocity”, etc.). Recently¹¹⁰, this basic *NOMAD Metainfo* ontology has been expanded to include a representation of operations among arrays (in an ontology, any mathematical concept needs to be represented in order to properly operate with the physical quantities in complex queries). This extension allowed for the introduction of the notion of “similarity” relationship that has been applied as a proof of concept to the calculated electronic density of states, as stored in the NOMAD Archive, in order to identify materials with similar electronic structures^{111,112}.

Achievements and challenges of ontologies for materials science were discussed at the first “Workshop on Ontologies for Materials-Databases Interoperability (OMDI2021)”, held in Linköping and virtually on October 2021. The workshop was organized by the OPTIMADE consortium³⁷ and funded by Psi-k¹¹³. The main outcomes of the workshop were: a) the strengthening of the idea that the development of useful ontologies need a community effort; b) they need to build from the data, i.e., their development needs to be driven by existing data and the aim of connecting data from different sources; c) tools for text mining need to be developed^{109,114}, in order to map into ontologies the enormous wealth buried in decades of scientific literature. Another important outcome of the workshop was the utterance of an insightful warning: “is the field proposing solutions (i.e, the existing ontologies) still in search of a problem?”. In other words, the community realizes that it needs specific questions to be addressed (the *competence questions*) in order to shape the ontologies and then propose demonstrative applications of such ontologies to answer the agreed upon questions.

X. CONCLUSIONS

Defining — as completely as possible — a pool of metadata for all the methods and computed quantities described above, is crucial for processing, storing, and providing FAIR materials-science data. A key challenge is the mapping into a metadata schema of the full set of input parameters, including those hidden into the specific codes, and all the available output. This practice will facilitate reproducibility, benchmarking, and peer-review processes.

In particular, we emphasize the importance of developing a hierarchical and modular metadata schema in order to represent the complexity of materials science data and allow for access, reproduction, and repurposing of data, from single-structure calculations to complex workflows. Furthermore, the modularity of the schema enables its extensibility, which is vital for the long-term maintenance of the metadata infrastructure.

As an example, we presented the current status of the NOMAD metadata schema, which was designed to comply with the FAIR principles. By means of existing parsers that map a growing set of atomistic-simulation code packages into the hierarchical, modular NOMAD metadata schema, the NOMAD infrastructure already provides the community with a FAIR storage of materials science data. The challenges of fully covering the ground-state electronic calculations, and extending the schema to excited states, dynamical simulations, and complex workflows were examined in detail. By means of a community effort, all aspects of the different subfields, and all the practical details of each specific implementation can be mapped on the NOMAD metadata schema. Finally, we discussed the challenges of the *FAIRification* of experimental materials-science metadata and the creation of ontologies for materials science. Ontologies will unlock the *interoperability* of the FAIR data by enabling the access and reuse of data across materials-science areas, but also outside materials science.

ACKNOWLEDGEMENTS

We would like to thank all the participants to the workshop “Shared Metadata and Data Formats for Big-Data Driven Materials Science: A NOMAD-FAIR-DI Workshop”, as listed at <https://th.fhi-berlin.mpg.de/meetings/META2019/>, who have contributed with questions and comments to ideas discussed in this paper.

The organizers of and participants to the OMDI2021 workshop (see <https://liu.se/en/research/omdi2021> for the full list of names) are acknowledged for insightful discussions that inspired some of the concepts discussed in Section IX.

This work received funding by the European Union’s Horizon 2020 research and innovation program under the grant agreement N° 951786 (NOMAD CoE) and by the German Research Foundation (DFG) through the NFDI

consortium FAIRmat, project 460197019.

AUTHOR CONTRIBUTIONS STATEMENT

The present paper is inspired by and based on the minutes of the work-groups discussions at the workshop “Shared Metadata and Data Formats for Big-Data Driven Materials Science: A NOMAD–FAIR-DI Workshop”. Here, we report the composition of the original work groups, which reflect into the main contributions to the paper’s sections. **Metadata, metadata schemas and ontologies** (Sections I, II and IX): Patrick Lambrix, Javad Chamanara, Carsten Baldauf, Tatyana Sheveleva, Benjamin Regler, Alvin Noe Ladines, Christoph T. Koch, Christof Wöll, Stefano Cozzini, Astrid Schneidewind, Maja-Olivia Lenz-Himmer; **Ground-state calcula-**

tions (Section III): Micael Oliveira, Sergey Levchenko; **Perturbative and excited-states calculations** (Section IV): Claudia Draxl, Pasquale Pavone, Denis Usvyat; **Potential-energy sampling** (Section V): James Kermode, Tristan Bereau, Christian Carbogno, Omar Valsion, Markus Kühbach, Chuanxun Su, Ron Miller, Berk Onat; **Workflows** (Section VI): Stefano Curtarolo, Shyam Dwaraknath, Adam Michalchuk, Giovanni Pizzi, Gian-Marco Rignanese, Jörg Schaarschmidt; **Data formats** (Section: VII): Ádám Fekete, Markus Scheidgen; **Metadata for experiments** (Section VIII): Christoph T. Koch, Sandor Brockhauser, Astrid Schneidewind. Luca M. Ghiringhelli and Matthias Scheffler coordinated the formation of the work groups, participated to the discussions in several work groups, and prepared the first draft of the paper. All authors contributed to the final version of the paper.

¹ J. M. Rickman, T. Lookman, and S. V. Kalinin, *Acta Mater.* **168**, 473 (2019).

² T. Hey, S. Tansley, and K. Tolle, *The Fourth Paradigm: Data-Intensive Scientific Discovery* (Microsoft Research, 2009). The concept of a fourth paradigm was probably first discussed by J. Gray at a workshop 22on Januar 11, 2007 before he went missing at the Pacific on January 28, 2007.

³ 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**, 1 (2016).

⁴ Other terms for the “R” in FAIR are “repurposable” and “recyclable”. The former term indicates that data may be used for a different purpose than that for which they were initially created. The latter term hints at the fact that data in materials science are often exploited only once for supporting the thesis of a single publication and then they are stored and forgotten. In this sense, they would constitute a “waste” that can be recycled, provided that they can be found and they are properly annotated.

⁵ This definition is reminiscent of the definition given by NIST: “Structured information that describes, explains, locates, or otherwise makes it easier to retrieve, use, or manage an information resource. Metadata is often called data about information or information about information” (<https://www.nist.gov/itl/tig/projects/attribute-metadata>). With our definition, we highlight the role of data “reproducibility”, which is crucial in science.

⁶ According to ISO/IEC 11179, a metadata registry (MDR) is a database of metadata that supports the functional-

ity of registration. Registration accomplishes three main goals: identification, provenance, and monitoring quality. Furthermore, an MDR manages the semantics of the metadata, i.e., the relationships (connections) among them.

⁷ <http://www.xml-cml.org>.

⁸ P. Murray-Rust, J. A. Townsend, S. E. Adams, W. Phadungsukanan, and J. Thomas, *J. Cheminformatics* **3**, 1 (2011).

⁹ P. Murray-Rust and H. S. Rzepa, *J. Cheminformatics* **3**, 44 (2011).

¹⁰ <http://www.xml-cml.org/schema/schema3>.

¹¹ M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, J. A. Montgomery, Jr., T. Vreven, K. N. Kudin, J. C. Burant, J. M. Millam, S. S. Iyengar, J. Tomasi, V. Barone, B. Mennucci, M. Cossi, G. Scalmani, N. Rega, G. A. Petersson, H. Nakatsuji, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, M. Klene, X. Li, J. E. Knox, H. P. Hratchian, J. B. Cross, V. Bakken, C. Adamo, J. Jaramillo, R. Gomperts, R. E. Stratmann, O. Yazyev, A. J. Austin, R. Cammi, C. Pomelli, J. W. Ochterski, P. Y. Ayala, K. Morokuma, G. A. Voth, P. Salvador, J. J. Dannenberg, V. G. Zakrzewski, S. Dapprich, A. D. Daniels, M. C. Strain, O. Farkas, D. K. Malick, A. D. Rabuck, K. Raghavachari, J. B. Foresman, J. V. Ortiz, Q. Cui, A. G. Baboul, S. Clifford, J. Cioslowski, B. B. Stefanov, G. Liu, A. Liashenko, P. Piskorz, I. Komaromi, R. L. Martin, D. J. Fox, T. Keith, M. A. Al-Laham, C. Y. Peng, A. Nanayakkara, M. Challacombe, P. M. W. Gill, B. Johnson, W. Chen, M. W. Wong, C. Gonzalez, and J. A. Pople, “Gaussian 03,” Gaussian, Inc., Wallingford, CT, 2004, <http://www.gaussian.com>.

¹² M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma, H. J. Van Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus, and W. A. de Jong, *Comput. Phys. Commun.* **181**, 1477 (2010), <http://www.nwchem-sw.org/>.

¹³ An example of a CML mapping from a NWChem calculation can be found at <http://www.xml-cml.org/>

- examples/schema3/compchem.
- ¹⁴ R. S. McDonald and P. A. Wilks Jr, Appl. Spectrosc. **42**, 151 (1988).
 - ¹⁵ A. N. Davies and P. Lampen, Appl. Spectrosc. **47**, 1093 (1993).
 - ¹⁶ P. Lampen, H. Hillig, A. N. Davies, and M. Linscheid, Appl. Spectrosc. **48**, 1545 (1994).
 - ¹⁷ J. I. Baumbach, A. N. Davies, P. Lampen, and H. Schmidt, Pure Appl. Chem. **73**, 1765 (2001).
 - ¹⁸ X. Gonze, C. Almbladh, A. Cucca, D. Caliste, C. Freysoldt, M. Marques, V. Olevano, Y. Pouillon, and M. Verstraete, Psi-k Newsletters & Highlights, 53 (2007).
 - ¹⁹ X. Gonze, C.-O. Almbladh, A. Cucca, D. Caliste, C. Freysoldt, M. Marques, V. Olevano, Y. Pouillon, and M. J. Verstraete, Comput. Mater. Sci. **43**, 1056 (2008).
 - ²⁰ D. Caliste, Y. Pouillon, M. J. Verstraete, V. Olevano, and X. Gonze, Comput. Phys. Commun. **179**, 748 (2008).
 - ²¹ L. M. Ghiringhelli, C. Carbogno, S. Levchenko, F. Mohamed, G. Huhs, M. Lüders, M. Oliveira, and M. Scheffler, Npj Comput. Mater. **3**, 1 (2017).
 - ²² M. J. T. Oliveira, N. Papior, Y. Pouillon, V. Blum, E. Artacho, D. Caliste, F. Corsetti, S. de Gironcoli, A. M. Elena, A. García, V. M. García-Suárez, L. Genovese, W. P. Huhn, G. Huhs, S. Kokott, E. Küçükbenli, A. H. Larsen, A. Lazzaro, I. V. Lebedeva, Y. Li, D. López-Durán, P. López-Tarifa, M. Lüders, M. A. Marques, J. Minar, S. Mohr, A. A. Mostofi, A. O’Cais, M. C. Payne, T. Ruh, D. G. A. Smith, J. M. Soler, D. A. Strubbe, N. Tancogne-Dejean, D. Tildesley, M. Torrent, and V. W. Yu, J. Chem. Phys. **153**, 024117 (2020).
 - ²³ S. Curtarolo, W. Setyawan, S. Wang, J. Xue, K. Yang, R. H. Taylor, L. J. Nelson, G. L. W. Hart, S. Sanvito, M. Buongiorno Nardelli, N. Mingo, and O. Levy, Comp. Mat. Sci. **58**, 227 (2012), <http://aflowlib.org>.
 - ²⁴ L. Talirz, S. Kumbhar, E. Passaro, A. V. Yakutovich, V. Granata, F. Gargiulo, M. Borelli, M. Uhrin, S. P. Huber, S. Zoupanos, C. S. Adorf, C. W. Andersen, O. Schütt, C. A. Pignedoli, D. Passerone, J. VandeVondele, T. C. Schulthess, B. Smit, G. Pizzi, and N. Marzari, Sci. Data **7**, 1 (2020).
 - ²⁵ 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), <https://materialsproject.org>.
 - ²⁶ C. Draxl and M. Scheffler, MRS Bull. **43**, 676 (2018), <https://nomad-coe.eu/>.
 - ²⁷ C. Draxl and M. Scheffler, JPhys Materials **2**, 036001 (2019).
 - ²⁸ C. Draxl and M. Scheffler, in *Handbook of Materials Modeling: Methods: Theory and Modeling*, edited by W. Andreoni and S. Yip (Springer, Cham, 2020) pp. 49–73.
 - ²⁹ S. Kirklin, J. E. Saal, B. Meredig, A. Thompson, J. W. Doak, M. Aykol, S. Rühl, and C. Wolverton, Npj Comput. Mater. **1**, 1 (2015), <http://oqmd.org>.
 - ³⁰ A. Merkys, N. Mounet, A. Cepellotti, N. Marzari, S. Gražulis, and G. Pizzi, J. Cheminformatics **9**, 56 (2017), <http://www.crystallography.net/tcod>.
 - ³¹ G. Pizzi, A. Cepellotti, R. Sabatini, N. Marzari, and B. Kozinsky, Comput. Mater. Sci. **111**, 218 (2016), <http://www.aiiida.net>.
 - ³² S. P. Huber, S. Zoupanos, M. Uhrin, L. Talirz, L. Kahle, R. Häuselmann, D. Gresch, T. Müller, A. V. Yakutovich, C. W. Andersen, *et al.*, Sci. Data **7**, 1 (2020).
 - ³³ M. Uhrin, S. P. Huber, J. Yu, N. Marzari, and G. Pizzi, Comp. Mat. Sci. **187**, 110086 (2021).
 - ³⁴ A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dułak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, J. Kolsbjerg, Esben Leonhard Kubal, S. Kaasbjerg, Kristen Lysgaard, J. B. Maronsson, T. Maxson, L. Olsen, Thomas Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, M. Wilhelmisen, Lasse Walter, Z. Zeng, and K. W. Jacobsen, J. Condens. Matter Phys. **29**, 273002 (2017), <https://wiki.fysik.dtu.dk/ase>.
 - ³⁵ E. B. Tadmor, R. S. Elliott, J. P. Sethna, R. E. Miller, and C. A. Becker, JOM **63**, 17 (2011), <https://openkim.org>.
 - ³⁶ C. W. Andersen, R. Armiento, E. Blokhin, G. J. Conduit, S. Dwaraknath, M. L. Evans, Á. Fekete, A. Gopakumar, S. Gražulis, A. Merkys, F. Mohamed, C. Oses, G. Pizzi, G.-M. Rignanese, M. Scheidgen, L. Talirz, C. Toher, D. Winston, R. Aversa, K. Choudhary, P. Colinet, S. Curtarolo, D. Di Stefano, C. Draxl, S. Er, M. Esters, M. Fornari, M. Giantomassi, M. Govoni, G. Hautier, V. Hegde, M. K. Horton, P. Huck, G. Huhs, J. Hummelshøj, A. Kariryaa, B. Kozinsky, S. Kumbhar, M. Liu, N. Marzari, A. J. Morris, A. A. Mostofi, K. A. Persson, G. Petretto, T. Purcell, F. Ricci, F. Rose, M. Scheffler, D. Speckhard, M. Uhrin, A. Vaitkus, P. Villars, D. Waroquiers, C. Wolverton, M. Wu, and X. Yang, Sci. Data **8**, 217 (2021).
 - ³⁷ <https://www.optimade.org>.
 - ³⁸ In its most simple realization, a metadata schema is a *key-value* dictionary, where the key is a name identifying a given metadata. In *NOMAD Metainfo*, similarly to CML, the key is a complex entity grouping the several attributes.
 - ³⁹ <https://emmc.info/emmo-info>.
 - ⁴⁰ These are the common metadata, i.e., those that are code-independent. Hundreds more metadata are code-specific, i.e., mapping pieces of information in the codes’ input/output that are specific to a given code and not transferable to other codes. It can be browsed at <https://nomad-lab.eu/prod/v1/gui/analyze/metainfo>.
 - ⁴¹ A ReSTful API supports the access via web services, through common protocols, such as HTTP.
 - ⁴² <https://nomad-lab.eu/prod/v1/gui/search/entries>.
 - ⁴³ <https://nomad-lab.eu/prod/rae/encyclopedia/>.
 - ⁴⁴ L. M. Ghiringhelli, Nat. Rev. Phys. **3**, 724 (2021).
 - ⁴⁵ <https://nomad-lab.eu/AIToolkit>.
 - ⁴⁶ M. A. L. Marques, M. J. T. Oliveira, and T. Burnus, Comput. Phys. Commun. **183**, 2272 (2012).
 - ⁴⁷ S. Lehtola, C. Steigemann, M. J. Oliveira, and M. A. Marques, SoftwareX **7**, 1 (2018).
 - ⁴⁸ <http://www.basissetexchange.org/>.
 - ⁴⁹ V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, Comput. Phys. Commun. **180**, 2175 (2009).
 - ⁵⁰ L. A. Zepeda-Ruiz, A. Stukowski, T. Oppelstrup, and V. V. Bulatov, Nature **550**, 492 (2017).
 - ⁵¹ E. Miyoshi, T. Takaki, M. Ohno, Y. Shibuta, S. Sakane, and T. Aoki, Modell. Simul. Mater. Sci. Eng. **27**, 054003 (2019).
 - ⁵² Y. Shibuta, S. Sakane, E. Miyoshi, S. Okita, T. Takaki, and M. Ohno, Nat. Commun. **8**, 10 (2017).

- ⁵³ S. Piana, J. L. Klepeis, and D. E. Shaw, *Curr. Opin. Struct. Biol.* **24**, 98 (2014).
- ⁵⁴ A. Leitherer, A. Ziletti, and L. M. Ghiringhelli, *Nat. Commun.* **12**, 1 (2021).
- ⁵⁵ <https://mosdef.org/>.
- ⁵⁶ A. Laio and M. Parrinello, *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12562 (2002).
- ⁵⁷ G. M. Torrie and J. P. Valleau, *J. Comput. Phys.* **23**, 187 (1977).
- ⁵⁸ E. Marinari and G. Parisi, *EPL* **19**, 451 (1992).
- ⁵⁹ C. Dellago, P. Bolhuis, and P. L. Geissler, *Adv. Chem. Phys.* **123** (2002).
- ⁶⁰ R. J. Allen, C. Valeriani, and P. Rein Ten Wolde, *J. Phys. Condens. Matter* **21**, 463102 (2009).
- ⁶¹ G. A. Tribello, M. Bonomi, D. Branduardi, C. Camilloni, and G. Bussi, *Comput. Phys. Commun.* **185**, 604 (2014).
- ⁶² L. Greff da Silveira, M. Jacobs, G. Prampolini, P. R. Liv-otto, and I. Cacelli, *J. Chem. Theory Comput.* **14**, 4884 (2018).
- ⁶³ <https://www.plumed-nest.org/>.
- ⁶⁴ The PLUMED consortium, *Nat. Methods* **16**, 670 (2019).
- ⁶⁵ M. Brehm and M. Thomas, *J. Chem. Inf. Model.* **58**, 2092 (2018).
- ⁶⁶ E. Deelman, D. Gannon, M. Shields, and I. Taylor, *Future Gener. Comput. Syst.* **25**, 528 (2009).
- ⁶⁷ S. Curtarolo, W. Setyawan, G. L. W. Hart, M. Jahnátek, R. V. Chepulskii, R. H. Taylor, S. Wang, J. Xue, K. Yang, O. Levy, M. J. Mehl, H. T. Stokes, D. O. Demchenko, and D. Morgan, *Comp. Mat. Sci.* **58**, 218 (2012).
- ⁶⁸ C. E. Calderon, J. J. Plata, C. Toher, C. Oses, O. Levy, M. Fornari, A. Natan, M. J. Mehl, G. L. W. Hart, M. Buongiorno Nardelli, and S. Curtarolo, *Comp. Mat. Sci.* **108 Part A**, 233 (2015).
- ⁶⁹ G. Pizzi, in *Handbook of Materials Modeling : Methods: Theory and Modeling*, edited by W. Andreoni and S. Yip (Springer International Publishing, Cham, 2018) pp. 1–24.
- ⁷⁰ A. Jain, S. P. Ong, W. Chen, B. Medasani, X. Qu, M. Kocher, M. Brafman, G. Petretto, G.-M. Rignanese, G. Hautier, D. Gunter, and K. A. Persson, *Concurr Comput.* **7**, 5037 (2015), <https://materialsproject.org>.
- ⁷¹ J. Bröder, D. Wortmann, and S. Blügel, in *Extreme Data: Demands, Technologies, and Services Workshop Proceedings*, Vol. 40, edited by M. Schultz, D. Pleiter, and P. Bauer (Forschungszentrum Jülich, 2018) pp. 43–47.
- ⁷² C. Toher, J. J. Plata, O. Levy, M. de Jong, M. D. Asta, M. Buongiorno Nardelli, and S. Curtarolo, *Phys. Rev. B* **90**, 174107 (2014).
- ⁷³ S. P. Huber, E. Bosoni, M. Bercx, J. Bröder, A. Degomme, V. Dikan, K. Eimre, E. Flage-Larsen, A. Garcia, L. Genovese, *et al.*, *npj Computational Materials* **7**, 1 (2021).
- ⁷⁴ A. Togo and I. Tanaka, *Scr. Mater.* **108**, 1 (2015).
- ⁷⁵ G. Petretto, X. Gonze, G. Hautier, and G.-M. Rignanese, *Comput. Mater. Sci.* **144**, 331 (2018).
- ⁷⁶ G. Petretto, S. Dwaraknath, H. P. C. Miranda, D. Winston, M. Giantomassi, M. J. van Setten, X. Gonze, K. A. Persson, G. Hautier, and G.-M. Rignanese, *Sci. Data* **5**, 180065 EP (2018).
- ⁷⁷ N. Mounet, M. Gibertini, P. Schwaller, D. Campi, A. Merkys, A. Marrazzo, T. Sohler, I. E. Castelli, A. Cappelletti, G. Pizzi, and N. Marzari, *Nat. Nanotechnol.* **13**, 246 (2018).
- ⁷⁸ M. de Jong, W. Chen, T. Angsten, A. Jain, R. Notestine, A. Gamst, M. Sluiter, C. K. Ande, S. van der Zwaag, J. J. Plata, C. Toher, S. Curtarolo, G. Ceder, K. A. Persson, and M. D. Asta, *Sci. Data* **2**, 150009 (2015).
- ⁷⁹ C. Toher, C. Oses, J. J. Plata, D. Hicks, F. Rose, O. Levy, M. de Jong, M. D. Asta, M. Fornari, M. Buongiorno Nardelli, and S. Curtarolo, *Phys. Rev. Mater.* **1**, 015401 (2017).
- ⁸⁰ R. Golesorkhtabar, P. Pavone, J. Spitaler, P. Puschnig, and C. Draxl, *Comput. Phys. Commun.* **184**, 1861 (2013).
- ⁸¹ C. Oses, C. Toher, and S. Curtarolo, *MRS Bull.* **43**, 670 (2018).
- ⁸² J. J. Plata, P. Nath, D. Usanmaz, J. Carrete, C. Toher, M. de Jong, M. D. Asta, M. Fornari, M. Buongiorno Nardelli, and S. Curtarolo, *npj Comput. Mater.* **3**, 45 (2017).
- ⁸³ K. Yang, C. Oses, and S. Curtarolo, *Chem. Mater.* **28**, 6484 (2016).
- ⁸⁴ F. Symalla, P. Friederich, A. Massé, V. Meded, R. Coehoorn, P. Bobbert, and W. Wenzel, *Phys. Rev. Lett.* **117**, 276803 (2016).
- ⁸⁵ P. Friederich, V. Gómez, C. Sprau, V. Meded, T. Strunk, M. Jenne, A. Magri, F. Symalla, A. Colsmann, M. Ruben, and W. Wenzel, *Adv. Mater.* **29**, 1703505 (2017).
- ⁸⁶ R. Mercado, R.-S. Fu, A. V. Yakutovich, L. Talirz, M. Haranczyk, and B. Smit, *Chem. Mater.* **30**, 5069 (2018).
- ⁸⁷ P. Rüßmann, F. Bertoldo, and S. Blügel, *Npj Comput. Mater.* **7**, 1 (2021).
- ⁸⁸ V. Valerio, P. Giovanni, M. Antimo, J. R. Yates, N. Marzari, and A. A. Mostofi, *Npj Comput. Mater.* **6** (2020).
- ⁸⁹ S. Mishra, G. Catarina, F. Wu, R. Ortiz, D. Jacob, K. Eimre, J. Ma, C. A. Pignedoli, X. Feng, P. Ruffieux, *et al.*, *Nature* **598**, 287 (2021).
- ⁹⁰ <https://www.nexusformat.org/>.
- ⁹¹ <https://github.com/ispyb/>.
- ⁹² S. Delagenière, P. Brechereau, L. Launer, A. W. Ashton, R. Leal, S. Veyrier, J. Gabadinho, E. J. Gordon, S. D. Jones, K. E. Levik, S. M. McSweeney, S. Monaco, M. Nanao, D. Spruce, O. Svensson, M. A. Walsh, and G. A. Leonard, *Bioinformatics* **27**, 3186 (2011).
- ⁹³ S. Malbet-Monaco, G. A. Leonard, E. P. Mitchell, and E. J. Gordon, *Acta Crystallogr. D* **69**, 1289 (2013).
- ⁹⁴ S. Fisher, K. Levik, M. Williams, A. Ashton, and K. McAuley, *J. Appl. Crystallogr.* **48**, 927 (2015).
- ⁹⁵ A. De Maria Antolinos, P. Pernot, M. E. Brennich, J. Kieffer, M. W. Bowler, S. Delageniere, S. Ohlsson, S. Malbet Monaco, A. Ashton, D. Franke, D. Svergun, S. McSweeney, E. Gordon, and A. Round, *Acta Crystallogr. D* **71**, 76 (2015).
- ⁹⁶ <https://github.com/elabftw/>.
- ⁹⁷ N. Carpi, A. Minges, and M. Piel, *J. Open Source Softw.* **2**, 146 (2017).
- ⁹⁸ <https://www.labguru.com/>.
- ⁹⁹ <https://htxlab.embl.fr/>.
- ¹⁰⁰ G. Bricogne, L. Carter, C. Flensburg, R. Fogh, P. Keller, W. Paciorek, A. Sharff, C. Vornrhein, I. Cornaciu, M. Jamaroz, J. A. Marquez, M. Lehmann, and D. Musil, *Acta Crystallogr. A*, **74**, A248 (2018).
- ¹⁰¹ https://github.com/ComPlat/chemotion_ELN.
- ¹⁰² P. Tremouilhac, A. Nguyen, Y.-C. Huang, S. Kotov, D. S. Lütjohann, F. Hübsch, N. Jung, and S. Bräse, *J. Cheminformatics* **9**, 1 (2017).
- ¹⁰³ <https://github.com/emmo-repo/EMMO>.
- ¹⁰⁴ T. Ashino, *Data Sci. J.* **9**, 54 (2010).

- ¹⁰⁵ K. Cheung, J. Drennan, and J. Hunter, in *AAAI Spring Symposium: Semantic Scientific Knowledge Integration* (2008) pp. 9–14.
- ¹⁰⁶ D. G. Thomas, R. V. Pappu, and N. A. Baker, *J. Biomed. Inform* **44**, 59 (2011).
- ¹⁰⁷ J. Hastings, N. Jeliaskova, G. Owen, G. Tsiliki, C. R. Munteanu, C. Steinbeck, and E. Willighagen, *J. Biomed. Semantics* **6**, 1 (2015).
- ¹⁰⁸ H. Li, R. Armiento, and P. Lambrix, in *International Semantic Web Conference* (Springer, 2020) pp. 212–227.
- ¹⁰⁹ H. Li, R. Armiento, and P. Lambrix, *Data Sci. J.* **18**, 1 (2019).
- ¹¹⁰ Maja-Olivia Himmer-Lenz, PhD thesis, Humboldt Universität zu Berlin (2022) <https://doi.org/10.18452/24340>.
- ¹¹¹ M. Kuban, S. Rigamonti, M. Scheidgen, and C. Draxl, arXiv preprint arXiv:2201.02187 (2022).
- ¹¹² M. Kuban, Š. Gabaj, W. Aggoune, C. Vona, S. Rigamonti, and C. Draxl, arXiv preprint arXiv:2204.04056 (2022).
- ¹¹³ <https://psi-k.net/>.
- ¹¹⁴ E. A. Olivetti, J. M. Cole, E. Kim, O. Kononova, G. Ceder, T. Y.-J. Han, and A. M. Hiszpanski, *Appl. Phys. Rev.* **7**, 041317 (2020).