

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

Common workflows for computing material properties using different quantum engines

### Permalink

<https://escholarship.org/uc/item/4w70n173>

### Authors

Huber, Sebastiaan P

Bosoni, Emanuele

Bercx, Marnik

et al.

### Publication Date

2021-05-11

Peer reviewed

# Common workflows for computing material properties using different quantum engines

Sebastiaan P. Huber,<sup>1,\*</sup> Emanuele Bosoni,<sup>2</sup> Marnik Bercx,<sup>1</sup> Jens Bröder,<sup>3,4</sup> Augustin Degomme,<sup>5</sup> Vladimir Dikan,<sup>2</sup> Kristjan Eimre,<sup>6</sup> Espen Flage-Larsen,<sup>7,8</sup> Alberto Garcia,<sup>2</sup> Luigi Genovese,<sup>5</sup> Dominik Gresch,<sup>9</sup> Conrad Johnston,<sup>10</sup> Guido Petretto,<sup>11</sup> Samuel Poncé,<sup>1</sup> Gian-Marco Rignanese,<sup>11</sup> Christopher J. Sewell,<sup>1</sup> Berend Smit,<sup>12</sup> Vasily Tseplyaev,<sup>3,4</sup> Martin Uhrin,<sup>1</sup> Daniel Wortmann,<sup>3</sup> Aliaksandr V. Yakutovich,<sup>12,1</sup> Austin Zadoks,<sup>1</sup> Pezhman Zarabadi-Poor,<sup>13,14</sup> Bonan Zhu,<sup>15,14</sup> Nicola Marzari,<sup>1</sup> and Giovanni Pizzi<sup>1,†</sup>

<sup>1</sup>*Theory 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*

<sup>2</sup>*Institut de Ciència de Materials de Barcelona, ICMAB-CSIC, Campus UAB, 08193 Bellaterra, Spain*

<sup>3</sup>*Peter Grünberg Institut and Institute for Advanced Simulation, Forschungszentrum Jülich, D-52425 Jülich, Germany*

<sup>4</sup>*Department of Physics, RWTH Aachen University, D-52056, Aachen, Germany*

<sup>5</sup>*Univ. Grenoble-Alpes, CEA, IRIG-MEM-L\_Sim, 38000 Grenoble, France*

<sup>6</sup>*nanotech@surfaces laboratory, Swiss Federal Laboratories for Materials Science and Technology (Empa), CH-8600 Dübendorf, Switzerland*

<sup>7</sup>*SINTEF Industry, Materials Physics, Oslo, Norway*

<sup>8</sup>*University of Oslo, Department of Physics, Norway*

<sup>9</sup>*Microsoft Station Q, University of California, Santa Barbara, California, 93106-6105, USA*

<sup>10</sup>*Atomistic Simulation Centre, School of Mathematics and Physics, Queen's University Belfast, United Kingdom*

<sup>11</sup>*UCLouvain, Institut de la Matière Condensée et des Nanosciences (IMCN), Chemin des Étoiles 8, Louvain-la-Neuve 1348, Belgium*

<sup>12</sup>*Laboratory of Molecular Simulation (LSMO), Institut des sciences et ingénierie chimiques (ISIC), École Polytechnique Fédérale de Lausanne (EPFL) Valais, CH-1951, Sion, Switzerland*

<sup>13</sup>*Department of Chemistry, Claverton Down, University of Bath, BA2 7AY, Bath, United Kingdom*

<sup>14</sup>*The Faraday Institution, Didcot OX11 0RA, United Kingdom*

<sup>15</sup>*Department of Chemistry, University College London, 20 Gordon St, Bloomsbury, London WC1H 0AJ, United Kingdom*

(Dated: May 12, 2021)

The prediction of material properties through electronic-structure simulations based on density-functional theory has become routinely common, thanks, in part, to the steady increase in the number and robustness of available simulation packages. This plurality of codes and methods aiming to solve similar problems is both a boon and a burden. While providing great opportunities for cross-verification, these packages adopt different methods, algorithms, and paradigms, making it challenging to choose, master, and efficiently use any one for a given task. Leveraging recent advances in managing reproducible scientific workflows, we demonstrate how developing common interfaces for workflows that automatically compute material properties can tackle the challenge mentioned above, greatly simplifying interoperability and cross-verification. We introduce design rules for reproducible and reusable code-agnostic workflow interfaces to compute well-defined material properties, which we implement for eleven different quantum engines and use to compute three different material properties. Each implementation encodes carefully selected simulation parameters and workflow logic, making the implementer's expertise of the quantum engine directly available to non-experts. Full provenance and reproducibility of the workflows is guaranteed through the use of the AiiDA infrastructure. All workflows are made available as open-source and come pre-installed with the Quantum Mobile virtual machine, making their use straightforward.

## I. INTRODUCTION

The use of density-functional theory (DFT) to compute the properties of systems at the atomic level has become widespread [1, 2], as both the number of quantum engines that implement it and the available computational power continue to increase. However, despite its large-scale deployment both in academia and in industry,

the application of DFT is still not a trivial operation. Accurate predictions require expert knowledge of not just DFT itself but also of the specific code used to perform the calculations (throughout this work we will use the terms quantum engine and code interchangeably). Although the diversity of available simulation packages improves the accuracy and reliability of results by virtue of cross-verification[3], different codes use diverse computational methods and interfaces, making it difficult even for experts to master more than just a few of them. This may result in software being used not for its applicability to a particular problem, but merely due to circumstantial reasons. Furthermore, the fact that the correct usage of

\* mail@sphuber.net

† giovanni.pizzi@epfl.ch

DFT-based codes requires expert knowledge directly limits its application and potential for scientific discovery.

Although DFT is used to compute many material properties of varying complexity, a large percentage of all performed calculations are defined by relatively simple recipes. Therefore, in addition to implementing new functionalities and improving the accuracy of existing ones, the effort of domain and code experts should be focused on providing robust workflows with common interfaces that can be used by experts and non-experts alike. If these are designed properly such that they are reusable, they can be employed as modular blocks in building more complex workflows, e.g. in a multi-scale approach. On top of reusability, in order to guarantee that results can be validated, it is crucial that these common workflows are reproducible.

The Atomistic Simulation Environment (ASE) [4] initiated an effort to provide a single interface for various quantum engines, which was later extended with the Atomistic Simulation Recipes (ASR) [5]. In this article, we address the additional challenges that one faces when trying to develop a common workflow interface, focusing particularly on the requirements of reusability and reproducibility, and we provide a solution based on AiiDA, an informatics infrastructure and workflow management system [6]. As a proof-of-concept, we define a common workflow interface specifically for the optimization of solid-state structures and molecular geometries, together with its implementation in eleven quantum codes: ABINIT[7–9], BigDFT[10], CASTEP[11], CP2K[12, 13], FLEUR[14], Gaussian[15], NWChem[16], ORCA[17, 18], QUANTUM ESPRESSO[19, 20], SIESTA[21, 22] and VASP[23, 24]. This particular common workflow interface, referred to as the “common relax workflow” throughout this work, allows a user to optimize a structure using any of these codes without having to define code specific parameters. The computed results are returned in a single unified format with identical units making the results directly comparable and reusable regardless of the underlying quantum engine used.

Each implementation of the common relax workflow interface provides at least three protocols (‘fast’, ‘moderate’ and ‘precise’) that allow a user to specify the desired computational accuracy in an intuitive and general way. The mapping between these levels of protocols and code specific parameters are up to the respective code experts to define. Through these protocols, expert knowledge of appropriate numerical parameters is thus encoded directly into the workflows, reducing the risk of incorrect or unreliable results and opening up the use of the quantum engines also to non-experts. Despite the ease-of-use of the workflows, the workflow interface design (which will be discussed later) maintains full flexibility and allows users to override any of the sensible defaults provided by the protocols. Furthermore, since AiiDA tracks the full provenance graph of executed workflows, storing all parameters used in workflow steps, the appropriateness of the inputs and the correctness of the results can also be

checked *a posteriori*.

To demonstrate the concept of modularity and potential for cross-verification, we use the common relax workflow to compute the equation of state (EOS) and the dissociation curve (DC), which are commonly computed properties for bulk compounds and diatomic molecules, respectively. Each of these properties is computed by a single workflow that exclusively leverages the common relax workflow as a modular building block, allowing any of the quantum engines to be used without specifying any code-specific parameters. The EOS and the DC are computed for a few compounds with different geometric, electronic and magnetic properties. As we will show later, the results computed by the various quantum engines show good agreement. We stress here that the focus of this paper is not on the validation of the results, but rather on the demonstration of the feasibility of a common workflow interface, directly enabling the reusability of complex workflows and the cross-verification of their results. We hope this will motivate readers to generalize these concepts and apply them to a broader and more complex range of problems.

The implementations of the common relax workflow interface of all quantum engines described in this paper are made available as free open-source software at <https://github.com/aidata/aiaa-common-workflows> under the MIT license. In addition, all workflows, as well as the seven quantum engines with a free open-source software license, come pre-installed in the Quantum Mobile[25] virtual machine (and quantum engines with a more restrictive license can be manually installed on any computational resource and configured to be used with AiiDA). This makes it straightforward to fully reproduce all the results presented in this paper (see the Supplementary Information for complete instructions).

## II. RESULTS AND DISCUSSION

### A. Reusability and reproducibility

Workflows, by definition, consist of multiple steps or multiple subprocesses that are executed in series, in parallel, or in a combination thereof, to obtain the final result. Ideally, workflows can themselves be used as modular blocks, becoming steps of higher-level workflows. To keep this process practical and tractable, workflows should be designed to be as modular and reusable as possible. Additionally, as workflows become ever more complex, so does their reproducibility. In this paper, we focus on two particular concepts that address these requirements: optional transparency and scoped provenance.

### 1. Optional transparency

In software design, the term *transparency* is often used to mean that a consumer of an interface should not be bothered with the inner details of the implementation (the details are invisible, or transparent). In terms of computational workflows, this can be taken to mean that a useful generic turn-key solution should have a simple interface, requiring as few inputs as possible from the user. Apart from physical inputs (e.g., the initial crystal structure in a relaxation workflow) and flags to determine which type of simulation to run (e.g., relax only atomic positions or also the periodic cell), any other input that is only needed as a numerical parameter by the underlying implementation should be automatically determined by the workflow.

However, this transparency of the interface comes at a cost. Complex workflows often consist of multiple subprocesses, each requiring their own inputs. Oftentimes at least some of these inputs cannot be automatically determined by the main workflow, as they are circumstantial and will be dependent on *how* and *where* the workflow is run. An example is when one of the subprocesses is executed on a high-performance computing (HPC) cluster and therefore requires specific environmental settings, such as the required resources and parallelization flags. A transparent interface is closed to these inputs being set (as shown schematically in Fig. 1a) and, as such, the workflow will be tied to a very specific environment for execution. Therefore, it will not be portable and consequently not reusable. But even if the inputs of the workflow could be automatically determined, an expert user may still want to override them. Transparent interfaces precluding this level of control diminish the reusability of workflows.

The solution to the aforementioned problem is to make the interface for all workflows fully *opaque* and expose all inputs of their subprocesses. That is to say, the workflow should make it possible to define each and every input that any of its subprocesses takes, as shown in Fig. 1b. By doing so, a user has access to all the inputs of the subprocesses, whether they could have been automatically determined by the workflow or not. Certainly, there are situations where the workflow can consciously decide not to expose certain inputs, as it is part of its task to determine them based on other inputs or intermediate results.

We are now confronted with two conflicting requirements, where a workflow interface must be both *transparent* for simplicity, yet at the same time fully *opaque* for reusability. The solution is to create an interface that is *optionally transparent*, i.e., it is opaque when needed but can still be used in a transparent manner whenever possible. Exposing the inputs of subprocesses is the first crucial step towards obtaining this goal, but it is not the only one. In addition, the workflow needs to specify sensible defaults such that the interface remains simple to operate with just a minimal set of inputs. An even better

solution is offered by what we refer to as *input generators*. An input generator for a workflow is a function that, based on a minimal set of essential inputs, generates the full set of inputs required by the workflow and all of its subprocesses. The advantage of this approach is that an expert user has the ability to inspect the full set of inputs that have been generated and even modify them before actually executing the workflow. This is the approach that we will take in the rest of this work.

### 2. Scoped provenance

It is commonly accepted that science is facing a reproducibility crisis in that many studies can often not be reproduced [26]. In recent years, guidelines have been developed to address this problem, such as the FAIR principles [27] that aim to make data, among other things, more reusable. For workflows to become FAIR as well, it is critical that they store the *provenance* of the data that they produce at each execution [28]. Concretely, this means that a workflow should store not only its own inputs and outputs but also those of all the subprocesses that it invokes. Recording the provenance of data that is produced at each step of a workflow is crucial to enable the reproducibility and intelligibility of the final result. However, the full provenance is not always required. Therefore, for complex workflows that produce large provenance graphs, it becomes important to be able to investigate the provenance within different granularity levels, i.e. different scopes. We refer to the possibility of inspecting the provenance at different levels as *scoped provenance*, which we illustrate in Fig. 2.

The next section explains in detail how we put the concepts of optional transparency and scoped provenance into practice.

### 3. Common workflow design

To ensure that the common workflows satisfy the requirements of optional transparency and scoped provenance, we have chosen to implement them using AiiDA [6], a scalable computational infrastructure for automated reproducible workflows and data provenance. The workflows are implemented as AiiDA *work chains* [29], whose data provenance and that of all their subprocesses is automatically stored by AiiDA in a relational database. AiiDA provides an application programming interface (API) to query the provenance graph at various levels of granularity, satisfying the scoped provenance requirement. The optional transparency criterion is made possible by the design of AiiDA's workflow language specification [29]. All processes in AiiDA are implemented in Python and, most importantly, the process specification (see Listing 1) is defined programmatically, allowing inspection of inputs and outputs before executing the workflow. In addition, it allows workflows to easily reuse




Figure 1. **Difference between a transparent and opaque workflow interface.** a) A schematic depiction of a process ( $P_A$ ) that consists of three subprocesses ( $S_A$ ,  $S_B$  and  $S_C$ ), that each require two inputs ( $I_1$  and  $I_2$ ). In this abstract example, the top-level process takes no inputs which is just for clarity; normally the top-level process takes at least one input based on which the inputs for the subprocesses are determined. Note that, although for simplicity the same symbol is used for these inputs, they do not necessarily represent identical inputs across the subprocesses, even though in practice the names could actually overlap. Only two inputs per process are arbitrarily chosen here for illustrative purposes. The interface of  $P_A$  does not expose the inputs of its subprocesses, but instead will decide them internally. This means that a user of  $P_A$  cannot customise the inputs of any of subprocesses. b) A schematic depiction of the same process  $P_A$  as in a), but in this case exposing the inputs of its subprocesses. Since the names of the inputs can potentially overlap, inputs are exposed in namespaces to prevent name clashes. A user of  $P_A$  can now directly set the inputs through the top-level interface. If any of the inputs of the subprocesses *should not* be defined by the user (due to being part of the workflow’s task to define it) the workflow can decide to not expose that particular input.

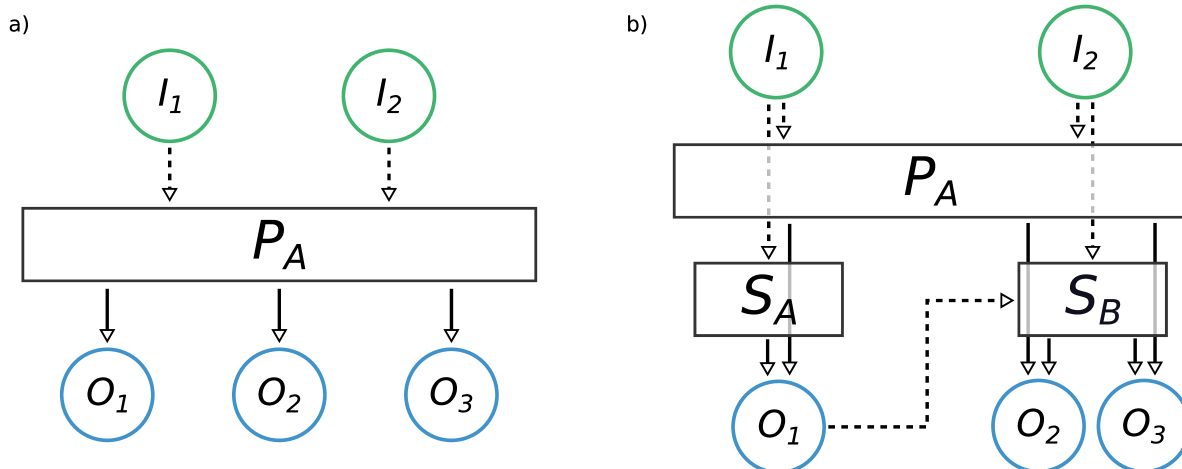


Figure 2. **Scope provenance.** a) Schematic provenance of a workflow ( $P_A$ ) that takes two inputs ( $I_1$  and  $I_2$ ) and produces three outputs ( $O_1$ ,  $O_2$ ,  $O_3$ ). b) A more detailed view of the complete provenance of  $P_A$ , which actually runs two subprocesses ( $S_A$  and  $S_B$ ). Input  $I_1$  is passed by  $P_A$  to  $S_A$ , which results in  $O_1$ . This intermediate output  $O_1$  is passed by  $P_A$  to  $S_B$  as an input, in addition to  $I_2$ , which results in the outputs  $O_2$  and  $O_3$ . The latter are returned by  $P_A$  as the final outputs together with the  $O_1$  intermediate result.

subworkflows as modular blocks, without making their interface inaccessible, by *exposing* the inputs and outputs [29, 30].

```

1 class ProcessA(WorkChain):
2
3     @classmethod
4     def define(cls, spec):
5         spec.input('I_1')
```

```

6         spec.expose_inputs(SubProcessA, namespace='S_A
7         spec.expose_inputs(SubProcessB, namespace='S_B
8
9     class SubProcessA(WorkChain):
10
11         @classmethod
12         def define(cls, spec):
13             spec.input('I_1')
14             spec.input('I_2')
15
16     class SubProcessB(WorkChain):
17
18         @classmethod
19         def define(cls, spec):
20             spec.input('I_1')
21             spec.input('I_2')

```

Listing 1. The definition of a process `ProcessA` implemented as a subclass of an AiiDA `WorkChain`. The process runs two subprocesses (`SubProcessA` and `SubProcessB`). The process declares an input `I_1` in its specification; in addition, the inputs of subprocesses are not redefined, but `ProcessA` simply exposes them in its own specification. The inputs of the subprocesses are exposed in separate namespaces so that inputs with same name do not shadow each other and remain all accessible (see Listing 3 for an example of how these are passed).

Launching a process in AiiDA is performed by passing the process class as an argument to the `submit` function and passing the inputs as keyword arguments as shown in Listing 2.

```

1 from aida.engine import submit
2 inputs = {
3     'I_1': 1,
4     'I_2': 2,
5 }
6 submit(SubProcessA, **inputs)

```

Listing 2. Example of how `SubProcessA` is launched. The `**` marker is Python syntactic sugar to unwrap the `inputs` dictionary into keyword arguments to the `submit` function. Note that the values of the inputs are simple integers just for clarity of the example.

Listing 3 shows an example of how the top-level `ProcessA` can be launched, defining its own inputs as well as those of its subprocesses.

```

1 from aida.engine import submit
2 inputs = {
3     'I_1': 1,
4     'S_A': {
5         'I_1': 1,
6         'I_2': 2,
7     },
8     'S_B': {
9         'I_1': 1,
10        'I_2': 2,
11    }
12 }
13 submit(ProcessA, **inputs)

```

Listing 3. Example of how `ProcessA` is launched. The inputs of the subprocesses can be passed in dictionaries that are nested in the main inputs dictionary, where the keys correspond to the namespace in which the inputs are exposed in the process specification (see Listing 1).

The concept of exposing inputs of subprocesses ensures that the inputs of any subprocess can be controlled from the top-level workflow, regardless of the level of nesting.

This directly satisfies the requirement of providing an opaque interface for expert users that need maximal control. However, the interface quickly risks becoming complex, as multiply layered workflows will require deeply nested input dictionaries. The workflow needs to optionally provide a transparent version of the interface to enable also non-expert users to easily use the workflow.

To solve this issue for the common workflows, we implement an input generator for each workflow. Input generators are not a native AiiDA concept but are a design pattern that emerged from the needs of defining and developing common workflows. Each common workflow defines a class method `get_input_generator` that returns an instance of an object that acts as the input generator. The input generator in turn defines the class method `get_builder`, which implements the common input interface and returns an instance of a ‘builder’. A builder is simply a container that wraps the generated inputs with additional information (such as the workflow class it pertains to), together with additional convenience functionality such as automatic input validation.

Since processes in AiiDA are implemented and executed directly in Python, their functionality can be easily extended. In addition, by being implemented in the same Python class as the workflow for which the inputs are generated, it is straightforward to keep the two aligned during workflow development. The `get_builder` method of the input generator takes a minimal amount of required arguments and returns a complete set of inputs for the corresponding workflow. Listing 4 shows how the input generator simplifies the usage of `ProcessA` for users, as now they only need to define a single input.

```

1 from aida.engine import submit
2 builder = ProcessA.get_input_generator().get_builder(
3     I_1=1)
4 submit(builder)

```

Listing 4. Example of how the launching of `ProcessA` is simplified by generating the inputs through the input generator. The `get_input_generator` class method returns an instance whose `get_builder` method can be called, to obtain a fully defined builder based on just a single input `I_1`. The builder, containing all the required inputs, can then be passed directly to the `submit` function. Since the builder also contains the process class for which it is defined, the process class itself no longer has to be explicitly passed to the `submit` function. Here, we are assuming that the inputs of the subprocess can be automatically determined by some algorithm implemented in the input generator; the number of minimally required inputs in this example is just one for simplicity.

An alternative approach to the problem could have been to make all subprocess inputs optional and let the workflow generate them at runtime (see Fig. 1a). Although this is a valid approach, in our experience it turns out to be much less flexible in particular for experienced users, as internal parameters cannot be set from the outside. Indeed, the input generator not only makes complex workflows accessible to non-experts, but it also gives maximal flexibility to advanced users. By generating in-

puts before execution, they can still be modified according to the user’s needs before they are passed to the workflow for execution, giving direct access to all parameters and achieving the goals of optional transparency.

## B. The common relax workflow

As a proof of concept of the principles explained in Sec. II A, we present a common interface to a workflow that performs a geometry optimization of both molecular and extended systems, which is implemented for eleven quantum engines. Structural relaxation towards the most energetically favorable configuration is a common task in materials science, and all selected quantum engines can perform it. Nevertheless, the quantum engines use a wide variety of algorithms to optimize forces on the atoms and stress on the cell. This, therefore, presents an ideal yet challenging test scenario to develop a workflow with a common interface.

### 1. Design strategy

The design of the interface is guided by the idea to employ *optional transparency* to create a workflow interface that is suitable both for expert and non-expert users. This interface must be simple and general (code-agnostic) but at the same time retain full flexibility in changing code-specific parameters. Our adopted solution consists in the creation of code-specific workflows (implemented as AiiDA work chains named `<Code>CommonRelaxWorkChain`, where `<Code>` indicates the name of the underlying quantum engine), whose interface design is not restricted. A common interface is achieved by ensuring that each work chain provides also an input generator whose interface is identical for every `<Code>`, as shown in Fig. 3.

Listing 5 shows an actual code example of how the input generator can be obtained from a work chain implementation. The `get_builder` method of the input generator will transform the inputs, that respect the common interface, into the inputs that are expected by the corresponding code-specific work chain implementation. The inputs are returned in the form of a “builder” which can be directly submitted to AiiDA to start running the workflow.

```

1 from aida.engine import submit
2 generator = <Code>CommonRelaxWorkChain.
   get_input_generator()
3 builder = generator.get_builder(structure=..., protocol
   =..., ...)
4 submit(builder)

```

Listing 5. Submission of the relax common workflow employing the quantum engine `<Code>`. The arguments accepted by `get_builder` are identical for every `<Code>`, establishing a common interface.

We note that not only the names but also the (Python) types of the inputs and outputs are standardized to en-




Figure 3. **Schematic diagram of the common relax workflow interface.** Any implementation consists of two parts: the `<Code>CommonRelaxWorkChain` and a `<Code>CommonRelaxInputGenerator`.

The `<Code>CommonRelaxWorkChain` is an AiiDA `WorkChain` that implements the logic necessary to perform the structure optimization and has an input interface that is code-specific. However, the outputs that it returns respect the schema of the common interface, where  $S_r$  is the relaxed structure,  $F$  are the forces on each atom,  $T$  is the stress on the cell,  $E_t$  is the total energy and  $M_t$  is the total magnetization of the system. Each `<Code>CommonRelaxWorkChain` provides its own `<Code>CommonRelaxInputGenerator` which, unlike the workflow, implements the common interface for the inputs (note that not all common inputs are shown for clarity). Here **structure** is the structure that is to be optimized, the **protocol** is a string that defines how the inputs are determined, and the **engine** is a dictionary that specifies what code(s) to use. The `<Code>CommonRelaxInputGenerator` translates the common inputs into the code-specific inputs that the corresponding `<Code>CommonRelaxWorkChain` expects (indicated with ?). Since the creation of the code-specific inputs and the launching of the workflow are two separate actions, the generated inputs can be adapted at will.

sure that the interface is truly generic. These are described in the next sections.

### 2. Common inputs

The second step of the design is the identification of the minimal set of arguments for the input generator, reflecting the inputs of the most generic relaxation process. We identified three fundamental inputs that we therefore implemented as mandatory arguments of the `get_builder` method.

- **structure.** The structure to relax. (type: an AiiDA `StructureData` instance, the common data format to specify crystal structures and molecules in AiiDA [31]).
- **protocol.** In the context of this work, this means a single string summarizing the computational accuracy of the underlying DFT calculation and relaxation al-

gorithm. Three protocol names are defined and implemented for each code: ‘fast’, ‘moderate’ and ‘precise’. The details of how each implementation translates a protocol string into a choice of parameters is code dependent, or more specifically it depends on the implementation choices of the corresponding AiDA plugin. For this work we have tried to follow these definitions: a possibly unconverged (but still meaningful) run that executes rapidly for testing (‘fast’); a safe choice for prototyping and preliminary studies (‘moderate’); and a set of converged parameters that might result in an expensive simulation but provides converged results (‘precise’). The reason for not mandating the details of the protocols in the common-workflow specifications is due to the variety of basis sets, input potentials and algorithms, requiring the specification of diverse and heterogeneous parameters in different codes. For the eleven implementations presented in this work, we report in the supplementary material the detailed parameter choices and the translation done for each respective code. We note here that the choice of the exchange-correlation functional could, in the future, become an additional optional input. In this work, we decided to use the Perdew–Burke–Ernzerhof (PBE)[32] functional as the default choice. (type: a Python string).

- **relax\_type.** The type of relaxation to perform, ranging from the relaxation of only atomic coordinates to the full cell relaxation for extended systems. The complete list of supported options is: ‘none’, ‘positions’, ‘volume’, ‘shape’, ‘cell’, ‘positions\_cell’, ‘positions\_volume’, ‘positions\_shape’. Each name indicates the physical quantities allowed to relax. For instance, ‘positions\_shape’ corresponds to a relaxation where both the shape of the cell and the atomic coordinates are relaxed, but not the volume; in other words, this option indicates a geometric optimization at constant volume. On the other hand, the ‘shape’ option designates a situation when the shape of the cell is relaxed and the atomic coordinates are re-scaled following the variation of the cell, not following a force minimization process. The term “cell” is short-hand for the combination of ‘shape’ and ‘volume’. The option ‘none’ indicates the possibility to calculate the total energy of the system without optimizing the structure. Not all the described options are supported by each code involved in this work; only the options ‘none’ and ‘positions’ are shared by all the eleven codes. The supported options might be extended in the future. (type: a Python string).

In addition to these mandatory arguments, the computational resources must be passed to the work chain in order to make the interface transferable between different computational environments. For this task, a specific argument of `get_builder` has been designed called **engines**.

- **engines.** It specifies the codes and the corresponding computational resources for each step of the relaxation process. Typically one single executable is sufficient to perform the relaxation. However, there are cases in which two or more codes in the same simulation package are required to achieve the final goal, as for example in the case of FLEUR. (type: a Python dictionary).

Other inputs have been recognized as common optional features that also a non-expert user might want to have control over:

- **threshold\_forces.** A real positive number indicating the target threshold for the forces in eV/Å. If not specified, the protocol specification will select an appropriate value. (type: Python float).
- **threshold\_stress.** A real positive number indicating the target threshold for the stress in eV/Å<sup>3</sup>. If not specified, the protocol specification will select an appropriate value. (type: Python float).
- **electronic\_type.** An optional string to signal whether to perform the simulation for a metallic or an insulating system. It accepts only the ‘insulator’ and ‘metal’ values. This input is relevant only for calculations on extended systems. In case no such option is specified, the calculation is assumed to be metallic which is the safest assumption. (type: Python string).
- **spin\_type.** An optional string to specify the spin degree of freedom for the calculation. It accepts the values ‘none’ or ‘collinear’. These will be extended in the future to include, for instance, non-collinear magnetism and spin-orbit coupling. The default is to run the calculation without spin polarization. (type: Python string).
- **magnetization\_per\_site.** An input devoted to the initial magnetization specifications. It accepts a list where each entry refers to an atomic site in the structure. The quantity is passed as the spin polarization in units of electrons, meaning the difference between spin up and spin down electrons for the site. This also corresponds to the magnetization of the site in Bohr magnetons ( $\mu_B$ ). The default for this input is the Python value `None` and, in case of calculations with spin, the `None` value signals that the implementation should automatically decide an appropriate default initial magnetization. The implementation of such choice is code-dependent and described in the supplementary material of this manuscript. (type: `None` or a Python list of floats).
- **reference\_workchain.** A previously performed `<Code>CommonRelaxWorkChain`. When this input is present, the interface returns a set of inputs which ensure that results of the new `<Code>CommonRelaxWorkChain` can be directly compared to the `reference_workchain`. This is necessary to create, for instance, meaningful equations of state.



Its use will be clarified in the Sec. II C 1. (type: a previously completed `<Code>CommonRelaxWorkChain`).

The arguments of the input generator described above fully satisfy the needs for the creation of a “ready-to-submit” `<Code>CommonRelaxWorkChain`, constructing all its necessary inputs (see Listing 5). These inputs are code-specific and, as discussed earlier, can be modified before submission by an expert user who is familiar with the internals of the `<Code>CommonRelaxWorkChain`.

### 3. Inspection of valid inputs

The arguments of the `get_builder` method represent high-level parameters that describe how the geometry optimization should be performed or how the system is to be treated. Each argument has a fixed number of accepted values, but not every code implementation may necessarily support all of them, as some values might correspond to features not supported by the code. In order to be able to inspect which options are supported by a workflow implementation, the input generator offers a number of methods. An example is shown in Listing 6.

```
1 input_gen = <Code>CommonRelaxWorkChain.  
   get_input_generator()  
2 input_gen.get_relax_types()
```

Listing 6. Call to the inspection method that returns information on the available relaxation types for the `<Code>` implementation of the common relax workflow.

The `get_relax_types` method returns the supported values for `relax_type` for the corresponding workflow implementation. Inspection methods are implemented for all codes and all the arguments of `get_builder` except the threshold values, the `structure`, the `reference_workchain`, and the `magnetization_per_site`. Associated to the `engines` argument, there are the methods `get_engine_types` and `get_engine_type_schema`, which return the steps required by the relaxation and information on the code type necessary for each step of the relaxation, respectively.

The described inspection methods allow to introspect, in a fully machine-readable and automatic way, what the valid options for a particular common workflow implementation are. This is particularly relevant to facilitate future development of a graphical user interface (GUI) for the submission of the common relax workflow. The GUI will be able to create the necessary input fields, with a list of accepted values, by programmatically introspecting the input generator interface.

### 4. Common outputs

To allow direct comparison and cross-verification of the results, the outputs of `<Code>CommonRelaxWorkChain` are standardized for all implementations and are defined as follows:

- **forces.** The final forces on all atoms in eV/Å. (type: an AiiDA `ArrayData` of shape  $N \times 3$ , where  $N$  is the number of atoms in the structure).
- **relaxed\_structure.** The structure obtained after the relaxation. It is not returned if the `relax_type` is ‘none’. (type: AiiDA’s `StructureData`).
- **total\_energy.** The total energy in eV associated to the relaxed structure (or initial structure in case no relaxation is performed). The total energy is not necessarily defined in a code-independent way (e.g., it does not have a common zero). We require, however, that the partial derivative of the returned energy with respect to the change of the coordinate  $i$  of atom  $j$  is always the  $i$ -th coordinate of the force on the atom  $j$ . We also stress that in general, even for calculations performed with the same code, there is no guarantee to have comparable energies in different runs if the inputs are generated with the input generator (because, for instance, the selected k-points depend on the input structure volume). However, in combination with the input argument `reference_workchain` mentioned earlier and discussed in Sec. II C 1, energies from different relaxation runs become comparable, and their energy difference is well defined. (type: AiiDA `Float`).
- **stress.** The final stress tensor in eV/Å<sup>3</sup>. Returned only when a variable-cell relaxation is performed. (type: AiiDA `Float`).
- **total\_magnetization.** The total magnetization in  $\mu_B$  (Bohr-magneton) units. Returned only for magnetic calculations. (type: AiiDA `Float`).

### 5. A simple test case: ammonia

As a first test case of the various implementations of the common relax workflow, we present the optimization of a simple molecular structure: ammonia. The thermodynamically stable polymorph of ammonia has a trigonal pyramidal shape, which makes the structure polar. However, ammonia also exists in a metastable planar form [33]. The optimized structure and its associated total energy have been calculated with the common relax workflow implementation for all eleven quantum engines discussed in this paper for both phases of ammonia, using the ‘precise’ protocol for the input generation.

The analysis of the energy difference between the planar and pyramidal configurations of ammonia is presented in Fig. 4. As mentioned in the introduction, comparing results among codes is not the focus of this paper. However, it is worth mentioning that the small discrepancies between codes in Fig. 4 are not surprising, considering that the treatment of polar molecules with codes designed for extended systems is not a trivial task. In particular, some codes always need to use periodic boundary conditions, introducing non-physical interactions among replicas in the calculation. Even for large




Figure 4. **Energy difference between the planar and pyramidal phases of ammonia.** Energy difference ( $\Delta E$ ) between the planar and pyramidal phases of ammonia, calculated with the eleven quantum engines reported in the horizontal axis (QE stands for QUANTUM ESPRESSO). A relaxation of the structure has been performed independently by every code before computing the energies (except for FLEUR\*, because for FLEUR the relaxation failed due to overlapping muffin-tin spheres, a method-specific issue requiring error handling. The energy difference for FLEUR was calculated through the common workflow without relaxation, using the relaxed output structures of QUANTUM ESPRESSO instead).

enough simulation cells, the long-range electrostatic potential due to periodic images of the polar molecule affects the energy of the system. Strategies such as the use of improved Poisson solvers [34, 35] and more sophisticated dipole corrections [36, 37] can be introduced in order to circumvent this problem. Since in this paper we are focusing only on showing the concept and feasibility of common workflows, no dipole correction is considered, and the simulation box is set to a  $(15 \text{ \AA})^3$  cube, without performing a proper convergence study on the cell size. However, extensions of this work can add optional flags to the input generator of the workflow to activate appropriate dipole corrections if needed and implemented by the underlying quantum engine. The data presented here also illustrates the potential of the common interface for the cross-verification of results, especially considering the variety of basis sets and algorithms of the eleven quantum engines. The present work offers the possibility to compare results from quantum-chemistry-oriented and electronic-structure codes (both pseudopotentials- and all-electrons-based) with minimum effort.

## 6. Provenance

Since all workflows are implemented using AiiDA, the full provenance is automatically stored when the workflow is executed, as discussed in Sec. II A. Fig. 5 shows




Figure 5. **Schematic provenance graph for a relaxation workflows.** Schematic provenance graph for a relaxation workflow powered by two different quantum engines (top: SIESTA; bottom: QUANTUM ESPRESSO). The node of the `<Code>CommonRelaxWorkChain` is highlighted with the label “RelaxWorkChain” and with a red edge. All the AiiDA work chains called during the relaxation are represented by orange rectangles. The dark red rectangles are calculations, meaning calls to an external executable that performs a calculation, for instance a call to the `pw.x` of QUANTUM ESPRESSO. All the ellipses represents data nodes, meaning nodes in the database that contain data, like, for instance, the initial structure, the total energy and so on. In blue is the data node representing the code utilized for the calculations. In pink are represented calls to Python functions that modify some data in order to create others.

a schematic provenance graph for a relaxation workflow powered by two different quantum engines. Note that only a subselection of the total number of inputs and outputs are shown for clarity, but all subprocesses are displayed and the connections between the nodes give an idea of the internal complexity of the workflows. Notably, the figure shows how the different workflow implementations can follow considerably different logical paths while ultimately returning the same quantities according to the same common interface.

The action of taking the arguments of the common

interface and transforming them in code-specific inputs (operated by the input generators) is not tracked. This is not crucial since only the code specific inputs fully determine the calculation results. Also, we do not consider protocols as immutable objects, but rather as flexible input suggestions that an expert user might want to change.

### C. Code-agnostic workflows

Optimizing the geometry of a solid-state structure or molecule is a common core building block of materials-science workflows. Creating a common workflow interface for this particular task allows higher level workflows that reuse it to become code agnostic. This makes it possible to run the workflow with any quantum engine that implements the common relax workflow interface, without having to explicitly specify any input that is specific to the quantum engine. We discuss two examples of such workflows in the following sections.

#### 1. Equation of State

An example of a workflow that uses structure optimization as a building block is the EOS workflow. The equation of state of a solid-state system is obtained by computing the total energy of the system at various volumes. We present here the implementation of an EOS workflow that uses the common relax workflow to perform the optimization of the system at each volume and compute its total energy. This workflow serves as an example to explain how the unified interface of the common relax workflow can be used to create code-agnostic workflows. It has been named `EquationOfStateWorkflow` and its schematic representation is shown in Fig. 6.

The `EquationOfStateWorkflow` takes a structure as input ( $S_0$  in Fig. 6) and scales the volume a number of times ( $N_i$ ), with the scaled structures centered around the volume of the input structure. The workflow calls the common relax workflow for each scaled structure to compute its total energy. The common relax workflow interface is entirely accessible at the level of the inputs of the `EquationOfStateWorkflow`. This means that one can specify arguments accepted by the input generator (which are code-agnostic and are labeled `generator_inputs` in Fig. 6), but also, optionally, some code-specific `overrides` for the inputs produced by the generator. Therefore, on the one hand, by virtue of the common interface being code-agnostic, the `EquationOfStateWorkflow` is also independent of the quantum engine that is used for the underlying calculations. On the other hand, the possibility to specify explicit `overrides` should fulfill the needs of expert users and fully satisfy the optional transparency requirement for reusable workflows.

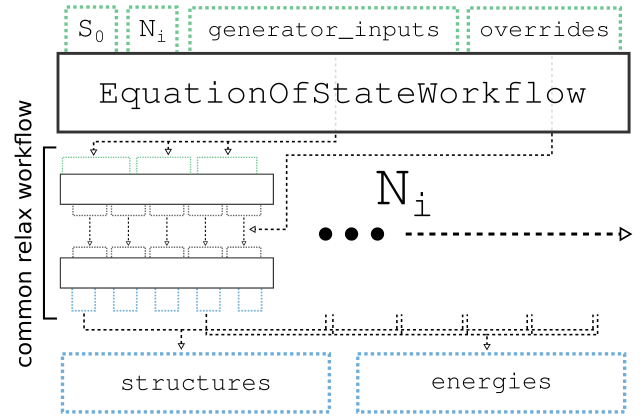


Figure 6. **Schematic diagram of the code-agnostic EOS workflow.** Schematic diagram of the implementation of the code-agnostic `EquationOfStateWorkflow`. The `EquationOfStateWorkflow` takes a number of arguments:  $S_0$  is the structure of the system at equilibrium volume and  $N_i$  are the number of volumes for which to compute the total energy. The `generator_inputs` will be passed directly to the inputs generator of the chosen common relax workflow implementation, which is called  $N_i$  times, once for each system volume. Note that the inset marked as “common relax workflow” corresponds directly to the schematic of the common relax workflow in Fig. 3. This highlights that the `EquationOfStateWorkflow` directly reuses the common relax workflow as its main building block. Which implementation of the common relax workflow is to be used is communicated to the `EquationOfStateWorkflow` by a single input, which is not shown for clarity. The `overrides` port allows an expert user to override certain inputs that are automatically determined by the generator, thus making the `EquationOfStateWorkflow` optionally transparent.

The total energies and optimized structure, as produced by the common relax workflow runs, are collected and returned by the `EquationOfStateWorkflow` as its outputs. Like the `<Code>CommonRelaxWorkChain` itself, the code-agnostic `EquationOfStateWorkflow` is implemented as an AiiDA work chain. This provides fully automated provenance tracking of all tasks performed inside the workflow, ensuring full reproducibility of the computed results.

It should be noted that the actual logic of the `EquationOfStateWorkflow` is slightly more complicated than depicted in Fig. 6. The common relax workflows are not all launched in parallel, but a single workflow is first performed for one of the scaled volumes. This first workflow is subsequently used as an additional input for the `reference_workchain` argument to the input generator for the common relax workflows for the remaining volumes. The input generator can use this reference to the first workflow to ensure that, if needed, parameters are kept constant between images in order for the energy differences to be meaningful. An example is the number of k-points used to sample the Brillouin zone (that is typically chosen by the input generators so as to get as

close as possible to a target density, and thus is volume-dependent if a `reference_workchain` is not specified).

The `EquationOfStateWorkflow` has been used to compute the EOS for a number of solid-state systems with varying electronic and magnetic properties: silicon (Si), aluminium (Al), germanium telluride (GeTe) and body-centered cubic (BCC) iron (Fe) both in a ferromagnetic and anti-ferromagnetic configuration. The results are shown in Fig. 7 and Fig. 8.

Fig. 7 reports the EOS results for the Si, Al and GeTe crystals. The curves for Si and Al have been obtained with all quantum engines, except ORCA and Gaussian, which are mainly specialized for non-periodic systems and the Gaussian AiiDA plugin does not yet support PBC. At each volume, the atomic positions are optimized while keeping the volume and cell shape fixed. The GeTe compound crystallizes at normal conditions in a trigonal phase (space group  $R\bar{3}m$ ) [38]. For this material, a correct calculation of the EOS requires the cell shape to be optimized at fixed volume in order to minimize the non-hydrostatic contributions of the stress tensor. This has been achieved in the common interface setting the `relax_type` to `positions_shape` (see Sec. II B 2), which is only supported by five out of eleven quantum engines. This is the reason why only five curves are shown in the right panel of Fig. 7. All calculations are carried out without spin-polarization and with the `precise` protocol.

The common interface also allows calculations on magnetic systems. Fig. 8 shows the EOS of BCC Fe, for both a ferromagnetic (left panel) and anti-ferromagnetic (right panel) ordering of atomic spin moments. At each volume, the atomic positions are optimized while keeping the volume and cell shape fixed. The central panel in Fig. 8 shows the total magnetization of the relaxed structure at each volume in the ferromagnetic case. The total magnetization in the anti-ferromagnetic case is zero at every volume and therefore not reported in the picture. The initial structure passed to the workflow is the same for the ferromagnetic and anti-ferromagnetic configurations and it is close to the equilibrium volume of the ferromagnetic case. This explains why the volume with minimum energy for the anti-ferromagnetic case is not placed in the middle of the analyzed volumes range. It is noteworthy that the BCC structure is the most thermodynamically stable configuration only in the ferromagnetic arrangement. The results show good overall agreement among codes. However, the scope of this section is only to demonstrate the variety of systems and physical quantities that can be analyzed with the code-agnostic EOS workflow.

## 2. Dissociation curve

In a similar fashion to the EOS workflow, a code-agnostic workflow for the calculation of the dissociation curve of a diatomic molecule has been implemented. In

this case, no relaxation is performed at all by the common relax workflow (accomplished by setting the `relax_type` equal to `'none'`) and it simply computes the energy of the system at various atomic distances. The same approach of the EOS workflow is used regarding the `reference_workchain` argument, meaning that the calculation at the first distance is used as a reference for the creation of inputs for the calculation at all the other distances.

Results are presented in Fig. 9 for the  $H_2$  dissociation curve obtained with the code-agnostic workflow. An initial anti-ferromagnetic configuration has been chosen as a starting point for each energy calculation. The results show good agreement among codes. DFT is not the most appropriate method for the calculation of dissociation curves in diatomic molecules, since these systems expose well known problems of DFT, like the delocalization error (self-interaction error) and static correlation [39]. The present test case wants to demonstrate the possibility to create code-agnostic workflows that support both electronic-structure codes and quantum-chemistry-oriented codes. In the future, the common relax workflow could be extended to allow calculations powered by different methods in addition to DFT, elevating the present work to a useful tool for comparing different levels of theory in the study of crystals and molecules.

## D. Conclusions

We have described how it is possible for domain experts to provide robust and reusable workflows that automatically compute materials properties, in order to exploit the ever-increasing computational power and popularity of DFT-based quantum engines, with the goal of accelerating materials discovery and characterization. For the workflows to be reusable, it is critical that they have *optionally transparent* interfaces and that the full provenance of executed workflows is automatically stored. We have demonstrated a concrete implementation of these two requirements using the workflow management system of the AiiDA informatics infrastructure [6]. We defined a common interface for a workflow that optimizes the geometry of a solid-state system or molecule, that was subsequently implemented for eleven popular quantum engines, with very diverse basis-set choices and algorithms. Using this common relax workflow, we have shown how higher-level workflows can reuse it to compute relevant material properties, such as the EOS and DC, while keeping a fully code-agnostic interface. Our results show how optionally transparent common workflow interfaces directly enable the cross-verification of results produced by different quantum engines. In addition, they empower a broader audience to use these methods in a robust way, encoding the experts' knowledge into reproducible code and hopefully stimulating new collaborations and more accurate materials science simulations.

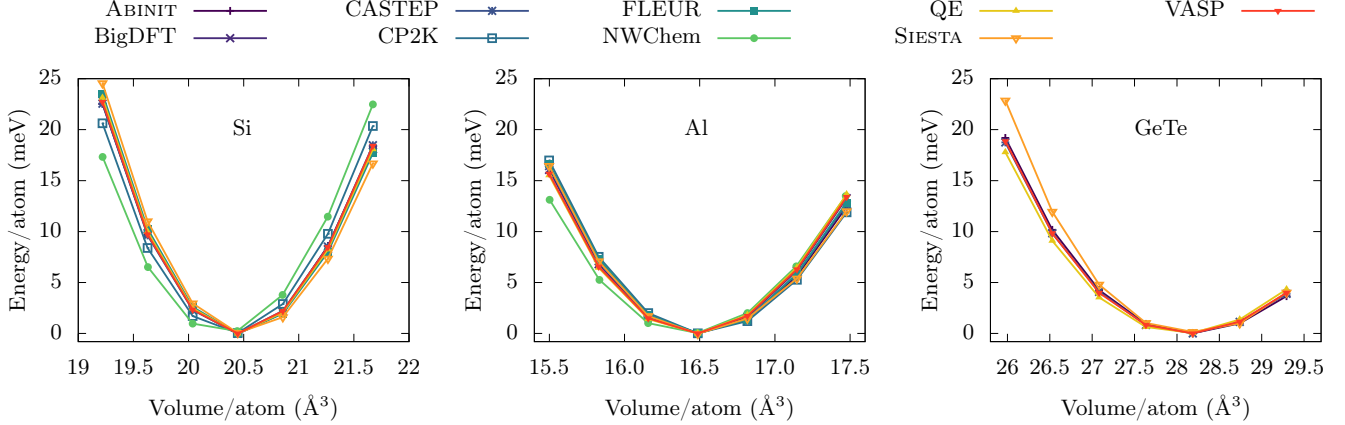


Figure 7. **EOS for Si, Al and GeTe.** Results obtained with the code-agnostic `EquationOfStateWorkflow`. For each code, the energy is shifted to set the minimum energy to zero. The EOS has been computed with all codes discussed in this work, except ORCA and Gaussian, which are mainly specialized for non-periodic systems. In addition, for GeTe, results are missing for BigDFT, CP2K, FLEUR and NWChem (see Table II in the Supplementary Information for more details). The label QE stands for QUANTUM ESPRESSO.




Figure 8. **EOS and total magnetization for BCC Fe.** Results obtained with the code-agnostic `EquationOfStateWorkflow`. The left panel reports the EOS obtained with a ferromagnetic initialization of the atomic moments. The corresponding total magnetization is reported for each volume in the central panel. The right panel reports the EOS obtained with an anti-ferromagnetic initialization of the atomic moments. The label QE stands for QUANTUM ESPRESSO. Results are missing for BigDFT, CP2K, Gaussian, NWChem and ORCA (see Table II in the Supplementary Information for more details).

### III. METHODS

#### A. Quantum Mobile implementation

The common interface and all corresponding code-agnostic workflows described in this paper allow anyone to run calculations to perform the same task with different codes, without knowing the details of each implementation. This is true assuming that the user can access a working executable of each code. The executable can be installed on the same machine as AiiDA

and the common workflows, or more typically in a remote computer (HPC cluster or supercomputer), since AiiDA allows automatic connection to external machines. Compiling and installing eleven different quantum engines can be a burden even for experienced users, and even more for non-experts. As one of the goals of this work is to facilitate the access to quantum codes to a broader audience, we also make available all codes related to this project in Quantum Mobile[25] version 21.05.1, which can be [downloaded here](#). Quantum Mobile is an open-source virtual machine based on Ubuntu, that comes with a large number of codes, tools and de-




Figure 9. **Dissociation curve of the  $\text{H}_2$  molecule.** Results obtained with a code-agnostic workflow. For all codes and all volumes, the magnetization is initialized to  $-1 \mu_B$  for one atom and to  $1 \mu_B$  for the other atom. The label QE stands for QUANTUM ESPRESSO. The VASP curve has been shifted by  $-24.89$  eV. Results are missing for BigDFT, FLEUR and NWChem (see Table II in the Supplementary Information for more details).

dependencies that are commonly needed to run materials-science atomistic simulations. In particular, it contains a pre-configured AiiDA installation together with the plugins interfacing AiiDA to all eleven quantum engines described here. In addition, since version 21.05.1 Quantum Mobile also includes the common-workflow interfaces and implementations discussed here, released as the [aiida-common-workflows package v0.1 on PyPI](#). Crucially, Quantum Mobile also includes the executables for the following open-source quantum engines: ABINIT, BigDFT, CP2K, FLEUR, NWChem, QUANTUM ESPRESSO and SIESTA (as well as a few more). Although CASTEP and ORCA provide free academic licenses, it requires users to have their own license which prevents pre-installation in Quantum Mobile. The remaining three codes discussed here (CASTEP, Gaussian and VASP) are commercial, therefore they cannot be redistributed freely without infringing their licenses. Nevertheless, a complete set of instructions is provided in the supplementary material, to guide users who already have access to these codes (on any computer of their choice) to configure them with AiiDA. In this way, the common workflows (all instead available open-source in the Quantum Mobile) can be run seamlessly also for commercial codes. Thanks to this setup, common workflows with these codes can be run with almost no preliminary step required. Only few codes require minimal adjustment that are described in the subsections “*running in the Quantum Mobile virtual machine*” of the supplementary material. A detailed list of instructions on how to run the test cases presented in this manuscript in the Quantum Mobile is reported in the Supplementary Material.

## CODE AVAILABILITY

The source code of the common workflows is released under the MIT open-source license and is made available on GitHub ([github.com/aiida-team/aiida-common-workflows](https://github.com/aiida-team/aiida-common-workflows)). It is also distributed as an installable package through the Python Package Index ([pypi.org/project/aiida-common-workflows](https://pypi.org/project/aiida-common-workflows)).

## DATA AVAILABILITY

The data and the scripts used to create all the images in this work are available on the Materials Cloud Archive [40]. Note that the data includes the entire AiiDA provenance graph of each workflow execution, as well as the curated data that is extracted from that database in order to produce the images.

## ACKNOWLEDGMENTS

This work is supported by the MARVEL National Centre of Competence in Research (NCCR) funded by the Swiss National Science Foundation (grant agreement ID 51NF40-182892) and by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 824143 (European MaX Centre of Excellence “Materials design at the Exascale”) and grant agreement No. 814487 (INTERSECT project). The authors thank M. Giantomassi and J.-M. Beuken for their contributions in adding support for PseudoDojo tables to the [aiida-pseudo plugin](#). The authors also thank X. Gonze, M. Giantomassi, M. Robert, C. Pickard, P. Hasnip, J. Hutter, M. Iannuzzi, D. Wortmann, S. Blügel, J. Hess, F. Neese and P. Delugas for providing useful feedback on the various quantum engine implementations. SP acknowledges support from the European Unions Horizon 2020 Research and Innovation Programme, under the Marie Skłodowska-Curie Grant Agreement SELPH2D No. 839217. EFL acknowledges the support of the Norwegian Research Council (project number 262339) and computational resources provided by Sigma2. PZP thanks the Faraday Institution CAT-MAT project (EP/S003053/1, FIRG016) for financial support. KE acknowledges the Swiss National Science Foundation (grant number 200020-182015). GPI and KE acknowledge the swissuniversities “Materials Cloud” (project number 201-003). Work at ICMAB is supported by the Severo Ochoa Centers of Excellence Program (MICINN CEX2019-000917-S), by PGC2018-096955-B-C44 (MCIU/AEI/FEDER, UE), and by Gen-Cat 2017SGR1506. BZ thanks the Faraday Institution FutureCat project (EP/S003053/1, FIRG017) for financial support. JB and VT acknowledge support by the Joint Lab Virtual Materials Design (JLVMD) of the Forschungszentrum Jülich.

## AUTHOR CONTRIBUTIONS

GPI and NM conceived the idea of a common interface among quantum engines and supervised the project. SPH, EB and GPI coordinated the project, implemented the interface design for the common relax workflow and the code-agnostic workflow implementations that use it as a modular block. SPH, MU and DG conceived the design of reusable modular workflow interfaces through programmatic process specification and implemented it in AiiDA. CS created the Quantum Mobile virtual machine to include the AiiDA plugins containing the implementation of the common relax workflow interface for all quantum engines, as well as an installation of the quantum engines for those that are open-source and can be freely redistributed. SP, ZP and GPe developed the ABINIT implementation of the common relax workflow which relies on the `aiida-abinit` plugin developed and maintained by SP, AZ, and GPe, and the work was supervised by GMR. AD developed the BigDFT implementation of the common relax workflow which relies on the `BigDFTCommonRelaxWorkChain` and `BigDFTBaseWorkChain` of the `aiida-bigdft` plugin package, supervised by LG. BZ developed the CASTEP implementation of the common relax workflow and its protocols which relies on the `CastepCommonRelaxWorkChain` of the package `aiida-castep` developed by the same author. AVY developed the CP2K implementation of the common relax workflow and its protocols which rely on the `Cp2kBaseWorkChain` of the `aiida-cp2k` package, developed by AVY and others, and the work was supervised by BS. JB developed the FLEUR implementation of the common relax workflow and its protocols which rely on the `FleurBaseCommonRelaxWorkChain`, `FleurCommonRelaxWorkChain`, `FleurScfWorkChain` and `FleurBaseWorkChain` of the `aiida-fleur` package[41], developed by JB, VT, DW and others, under the supervision of DW. KE developed the Gaussian implementation of the common relax workflow which relies on the `GaussianBaseWorkChain` of the package `aiida-gaussian`, developed by KE and PZP. CJ developed the NWChem implementation of the common relax workflow and its protocols which relies on the `NwchemBaseWorkChain` of the package `aiida-nwchem`, developed by CJ and others. PZP developed the ORCA implementation of the common relax workflow which relies on the `OrcaBaseWorkChain` of the `aiida-orca` package developed by the same author. SPH developed the QUANTUM ESPRESSO implementation of the common relax workflow with the support of MB, which relies on the `PwCommonRelaxWorkChain` and `PwBaseWorkChain` of the `aiida-quantumespresso` package, developed by SPH, MB, GPI and others. EB developed the SIESTA implementation of the common relax workflow which relies on the `SiestaBaseWorkChain` of the package `aiida-siesta`[22], developed by AG, VD and EB. EFL developed the VASP implementation of the common re-

lax workflow which relies on the `CommonRelaxWorkChain`, `VerifyWorkChain` and `VaspWorkChain` of the AiiDA plugin `aiida-vasp`, developed by EFL and others.

## COMPETING INTERESTS

The authors declare no competing interests.

# Supplementary Information: Common workflows for computing material properties using different quantum engines

(Dated: May 12, 2021)

## IV. COMMON DEFINITIONS

The following terms are common to a number of quantum engines and are used throughout this Supplementary Information.

- The "k-points distance" is a quantity used by almost every implementation to set the k-points mesh for the calculation of integrals in reciprocal space. It indicates the distance between adjacent k-points in the selected mesh. In other words, the number of k-points along each reciprocal axis  $i$  is the closest integer to  $|b_i|/\text{"k-points distance"}$  where  $b_i$  indicated the reciprocal lattice vector along  $i$ .
- The "forces threshold" is the target threshold for the forces during the relaxation process. A relaxation is considered converged when forces on all atoms are below "forces threshold". This is the same quantity referred to as `threshold_forces` in the main text of the manuscript.
- The "stress threshold" is the target threshold for the stress. A relaxation with variable cell is considered converged when the maximum stress component is smaller than "stress threshold". It is the same quantity referred to as `threshold_stress` in the main text of the manuscript.
- The term "PW cutoff" indicates the plane wave cutoff energy, a quantity commonly used to set the size of the basis set in plane-wave DFT calculations.
- The  $\Delta$  factor quantifies the difference between two EOS calculations[42]. It has become the standard for the comparison between DFT methods and codes[3].
- The "BFGS algorithm" refers to the Broyden–Fletcher–Goldfarb–Shanno algorithm for solving nonlinear optimization problems. The "L-BFGS algorithn" is its Limited-memory approximation[43].
- The "FIRE algorithm" refers to the Fast Inertial Relaxation Engine[44].

## V. OVERVIEW OF SUPPORTED FEATURES AND COMPLETED WORKFLOWS

Tab. V provides an overview of which arguments and what values for those arguments are supported by the implementations of the common relax workflow interface for the various quantum engines.

Code	relax_type	electronic_type	spin_type
ABINIT	all	insulator, metal	'none', 'collinear'
BigDFT	none, positions	insulator, metal	'none', 'collinear'
CASTEP	all	(ignored)	'none', 'collinear'
CP2K	none, positions, positions_cell	insulator, metal	'none', 'collinear'
FLEUR	none, positions	(ignored)	'none', 'collinear'
Gaussian	none, positions	(ignored)	'none', 'collinear'
NWChem	none, positions, positions_cell, cell	insulator, metal	'none'
ORCA	none, positions	(ignored)	'none', 'collinear'
QUANTUM ESPRESSO	all	insulator, metal	'none', 'collinear'
SIESTA	none, positions, positions_shape, cell	(ignored)	'none', 'collinear'
VASP	all	(ignored)	'none', 'collinear'

Table S1. Supported features by the various common relax workflow implementations. In addition to the tabulated features, all codes support the three protocols 'fast', 'moderate' and 'precise'.

Tab. V shows which workflows, as described in the main text, have been successfully completed for each quantum engine, with an explanation for those that are missing.



Code	EOS			Fe-f	Fe-af	DC H <sub>2</sub>	OPT	
	Si	Al	GeTe				NH <sub>3</sub> -pl	NH <sub>3</sub> -py
ABINIT	✓	✓	✓	✓	✓	✓	✓	✓
BigDFT	✓	✓	● <sup>a</sup>	✗ <sup>c</sup>	✗ <sup>c</sup>	✗ <sup>c</sup>	✓	✓
CASTEP	✓	✓	✓	✓	✓	✓	✓	✓
CP2K	✓	✓	● <sup>a</sup>	✗ <sup>d</sup>	✗ <sup>d</sup>	✓	✓	✓
FLEUR	✓	✓	● <sup>a</sup>	✓	✓	✗ <sup>e</sup>	✓	✓
Gaussian	● <sup>f</sup>	● <sup>f</sup>	● <sup>f</sup>	● <sup>f</sup>	● <sup>f</sup>	✓	✓	✓
NWChem	✓	✓	● <sup>a</sup>	● <sup>b</sup>	● <sup>b</sup>	● <sup>b</sup>	✓	✓
ORCA	● <sup>g</sup>	● <sup>g</sup>	● <sup>g</sup>	● <sup>g</sup>	● <sup>g</sup>	✓	✓	✓
QUANTUM ESPRESSO	✓	✓	✓	✓	✓	✓	✓	✓
SIESTA	✓	✓	✓	✓	✓	✓	✓	✓
VASP	✓	✓	✓	✓	✓	✓	✓	✓

Table S2. Overview of which workflow, presented in the main article, could be run by each of the eleven quantum engines. A checkmark (✓) means that the workflow was successfully completed, a circle (●) means the quantum engine does not support the required functionality or structure and a crossmark (✗) means the AiiDA plugin for the quantum engine does not yet support a required feature or it could not successfully run the workflow. For the latter two cases, a detailed explanation will be given below in the caption. Fe-f and Fe-af stand for ferromagnetic and anti-ferromagnetic iron, respectively. (a) The code does not support geometry optimizations at constant volume of the cell. (b) The code does not support running metallic systems which require smeared occupations. (c) BigDFT and its AiiDA plugin both support metallic systems, however, this requires specific mixing inputs which are still experimental. Although they give good results for the Al test case, the results are not trustworthy for the Fe examples and are therefore not included in the paper. The solution to address the issue is known: a two-steps approach is necessary, restarting the first calculation with a lower electronic temperature. This two-steps approach has been successfully used in earlier versions of PyBigDFT to perform the  $\Delta$  test, however, the AiiDA plugin and implementation of the common workflow interface for BigDFT do not yet support it. (d) The results for CP2K for the iron structure seemed to converge to an incorrect ground state for as of yet unknown reasons and therefore the results have been omitted. (e) In certain rare situations, FLEUR is unable to retain the specified initial magnetization when generating the start density. This sometimes leads to calculations wrongly converging to non-magnetic metastable solutions. This can be fixed by explicitly breaking the magnetic symmetry by means of a small electric field or enforcing initial occupation values, but neither option is currently supported by the implementation of the common workflow interface for FLEUR. (f) The code is not designed for extended systems although there is some support for it, but the AiiDA plugin does not implement it. (g) The code does not support extended systems.

## VI. CODE-SPECIFIC DESIGN CHOICES

We collect here the design choices of every common relax workflow implementation presented in the main text. These choices include details on the relaxation algorithm, an explanation of the supported features, and the specifications of every protocol implemented.

The textual descriptions in the next sections are meant to provide a quick reference of the most important numerical choices for each code, but cannot be fully complete without becoming cumbersome to read. This, however, is not an issue: the exact inputs of each run are captured in full detail by AiiDA in the provenance graph [40]. In addition the implementation of all plugins and workflows is open source and can be inspected to verify the exact choices and rules used to determine the input parameters.

### A. ABINIT

#### 1. Protocols

List of protocols supported and their description:

- **fast.** This protocol has low precision at minimal computational cost for testing purposes. k-points distance is  $0.25 \text{ \AA}^{-1}$ . Tolerance on the potential residual is  $1 \cdot 10^{-7} \text{ Ha}$ . No additional memory is allowed for basis set enlargement. Forces threshold is  $5 \cdot 10^{-5} \text{ Ha/Bohr}$ . Stress threshold is  $5 \cdot 10^{-3} \text{ Ha/Bohr}^3$ . Fermi-Dirac smearing with broadening  $0.008 \text{ Ha}$  and 2 times the number of atoms additional bands is set in the case of calculations on metals.
- **moderate.** This is the default protocol with normal precision and moderate computational cost. k-points distance is  $0.20 \text{ \AA}^{-1}$ . Tolerance on the potential residual is  $1 \cdot 10^{-9} \text{ Ha}$ . Forces threshold is  $5 \cdot 10^{-5} \text{ Ha/Bohr}$ . Stress threshold

is  $5 \cdot 10^{-3}$  Ha/Bohr<sup>3</sup>. Fermi-Dirac smearing with broadening 0.008 Ha and 2 times the number of atoms additional bands is set in the case of calculations on metals.

- **precise.** This protocol should yield fully converged results and is recommended for production calculations that require more precision than provided by the moderate protocol. k-points distance is  $0.1 \text{ \AA}^{-1}$ . Tolerance on the potential residual is  $1 \cdot 10^{-10}$  Ha. Forces threshold is  $5 \cdot 10^{-5}$  Ha/Bohr. Stress threshold is  $5 \cdot 10^{-3}$  Ha/Bohr<sup>3</sup>. Fermi-Dirac smearing with broadening 0.005 Ha and 2 times the number of atoms additional bands is set in the case of calculations on metals.

The PW cutoff is determined automatically thanks to the `djrepo` table provided by the PseudoDojo initiative [45]. An additional 15% memory is booked for the plane wave expansion basis in case of an ‘`positions_cell`’ and ‘`positions_volume`’ relaxation while only 5% is used when performing ‘`positions_shape`’ for numerical stability. In all other relaxation types, no additional memory is booked. In all cases, automatic multilevel parallelization of the calculations is performed directly by the ABINIT software [7, 8]. The workflow relies on Projector Augmented-Wave Method (PAW) [46] pseudopotentials from PseudoDojo [45] in the “standard” configuration with PBE exchange-correlation.

## 2. Supported calculation modes

The `relax_type` argument supports all the possible values, meaning that atomic positions and the cell size and shape can be optimized, as well as any combination of those. All protocols use a L-BFGS algorithm to minimize the forces on the atoms and stress on the cell.

The `electronic_type` argument supports the values ‘`metal`’ and ‘`insulator`’. The ‘`insulator`’ calculations are performed with fixed occupations, whereas, in ‘`metal`’ calculations, Fermi-Dirac smearing is employed.

The `spin_type` argument supports the values ‘`none`’ and ‘`collinear`’. In case of ‘`collinear`’ calculations, unless `magnetization_per_site` is explicitly defined, the maximum theoretical magnetic moment for each element is specified in the z-direction.

The `reference_workchain` argument guarantees that the used k-points mesh is identical to the one of the reference workchain.

## 3. Running in the Quantum Mobile virtual machine

All that is needed to run the ABINIT common workflows is to install the appropriate PseudoDojo [45] family, which is done as follows:

```
aida-pseudo install pseudo-dojo -v 1.0 -f jthxml
```

## B. BigDFT

### 1. Protocols

BigDFT input generator switches from cubic scaling computation to linear scaling computation if the number of atoms in the system is more than 200. This value is meant to be adapted more finely in future releases to provide a good heuristic.

For cubic computations, inputs are generated by the input generator.

- **fast.** This protocol is meant for testing and should provide fast and less accurate results. `hgrids` is set to 0.45 bohrs and k-points are generated using a real space equivalent length of 20 bohrs.
- **moderate.** This protocol the main one for providing reasonable accuracy while staying fast enough. Main difference with fast is that `hgrids` is set to 0.3, while default k-points distance is set to 40 bohrs (real space). Convergence criteria are also stricter.
- **precise.** This protocol aims at providing the most accurate results without taking costs into account. `hgrids` is set to 0.15. k-points are computed by PyBigDFT according to the number of atoms for this precision level.

For linear computations, BigDFT provides "linear\_fast", "linear\_moderate", and "linear\_accurate" input sets, which are mapped to 'fast', 'moderate', and 'precise' protocols respectively. They use 0.3 hgrids settings with more relaxed or strict convergence criterion.

In order to have more accurate results, it has been decided to override, when possible, the default pseudopotentials used in BigDFT with the soft and norm-conserving HGH pseudopotentials (including nonlinear core-correction) generated by S.Saha [47]. These pseudopotentials are not yet available in the UPF format and, for the moment, are included directly in PyBigDFT (in PSP8 format). In the near future, these pseudopotentials will be converted in UPF format and made available to AiiDA users in the standard way (UpfData class and corresponding pseudopotential families). BigDFT has to be compiled with PSPIO in order to support UPF pseudopotentials. The version of BigDFT in the Quantum Mobile already includes this support.

## 2. Supported calculation modes

As of now, the BigDFT implementation only supports the relaxation of the atomic coordinates. Therefore it handles the `threshold_forces` argument, but not the `threshold_stress` one. The default applied algorithm for relaxation in the BigDFT plugin is FIRE. Cell relaxation is not supported internally in BigDFT, but we foresee the creation of a future AiiDA workflow that drives externally the cell relaxation.

IMPORTANT NOTE: BigDFT does not yet handle non-orthorhombic cells, a transformation of the input cell is attempted if an invalid cell is provided in input. The computation is then performed on the resulting orthorhombic super-cell. The Si test-case used for in this paper has to be transformed this way before computation. The next release of BigDFT will overcome these limitations.

The `spin_type` argument supports the values 'none' and 'collinear'. If the 'collinear' option is provided, the computation is launched with spin-polarized enabled. By default, the plugin will initialize spin through a round-robin scheme over the atoms to match the computed polarity moment. This can be overridden by using the `magnetization_per_site` argument, to provide the values for each atom directly.

The `electronic_type` argument supports both 'metal' and 'insulator'. If metal is specified, specific mixing inputs are added to account for the nature of the system. These inputs are still experimental and they give good results for Al but not for the Fe test case, that is therefore not included in the paper. The solution to address the issue is known: a two-steps approach is necessary, restarting the first calculation with a lower electronic temperature. This two-steps approach has been successfully used in earlier versions of PyBigDFT to perform the  $\Delta$  test, however it is not yet included in the implementation of the present project. A future internal workflow will automatically perform the two steps allowing BigDFT to return correct results for Fe through the common interface.

## 3. Running in the Quantum Mobile virtual machine

BigDFT is provided in latest Quantum mobile releases. There is no specific configuration to run it, once set up with aiiida.

## C. CASTEP

### 1. Protocols

List of protocols supported and their description:

- **fast**. This protocol is intended for coarse calculations and quick tests. Many parameters are downgraded from the `moderate` settings. The main difference is that a different on-the-fly generated (OTFG) pseudopotential library QC5 is used, which is designed to have converged results using a PW cutoff  $\approx 300$  eV for most elements at the cost of transferability. The `medium` basis set precision settings is applied, with a reduced electronic energy tolerance  $1 \cdot 10^{-5}$  eV per atom, and a increased k-point distance of  $0.25 \text{ \AA}^{-1}$  (equivalent to  $0.03979 \text{ } 2\pi \text{ \AA}^{-1}$  in CASTEP's convention).
- **moderate**. This protocol is intended for general use and expected to give physically sensible results. The k-point distance is set to  $0.15 \text{ \AA}^{-1}$ . Note that CASTEP uses a convention without the explicit  $2\pi$  factor when converting between the real and reciprocal space distances, so this setting is equivalent to a `kpoints_mp_spacing` of  $0.02387 \text{ } 2\pi \text{ \AA}^{-1}$ . The `fine` basis precision setting is used. With this option, CASTEP internally chooses the PW cutoff based on the convergence data of the core-corrected ultrasoft pseudopotentials to be used, which are on-the-fly generated

during the calculation using the built-in C19 library of generation settings. If the calculation only involves elements for which the pseudopotentials are very soft, such as silicon, the default low PW cutoff can potentially cause problems during variable cell geometry optimisation. To avoid this issue, a minimum PW cutoff of 326 eV is imposed in the protocol. The electronic energy convergence tolerance is set to  $10^{-6}$  eV per atom with the default tolerance window of three steps. The geometry optimisation uses a forces threshold of 0.05 eV/Å in conjunction with a energy tolerance of  $2 \cdot 10^{-5}$  eV per atom, a stress threshold of 0.1 GPa (if applicable), and a atomic position change tolerance of 0.001 Å, with a convergence window of two ionic steps.

- **precise.** This protocol is intended for high accuracy calculations and obtaining converged results. Most parameters are based on the `moderate` settings, with the `basis_precision` setting increased to `precise`. The electronic energy tolerance is reduced to  $1 \cdot 10^{-8}$  eV per atom. The geometry optimisation thresholds are reduced to 0.03 eV/Å for forces, 0.05 GPa for stress and  $10^{-5}$  eV per atom for the total energy. The `grid_scale` setting that controlling the FFT grid density is set to 2, and the `fine_grid_scale` is increased to 3. This is intended for minimising any FFT aliasing errors. The k-points spacing is further reduced to  $0.1 \text{ \AA}^{-1}$  (equivalent to  $0.01591 \text{ } 2\pi \text{ \AA}^{-1}$  in CASTEP's convention) for a improved sampling of the reciprocal space.

## 2. Supported calculation modes

The `relax_type` argument supports all the agreed values. All protocol defaults to the L-BFGS algorithm with line search for performing geometry optimisation. The only exception is when the `relax_type` is `'positions_shape'`. In this case the two-point steepest descent optimiser (TPSD) is selected as it gives superior performance as of CASTEP version 19.1.1.

The `electronic_type` argument supports both `'metal'` and `'insulator'`, although the value is internally ignored, and density mixing solver is used in both cases with a Gaussian smearing of 0.2 eV applied to the occupations of the electronic bands.

The `spin_type` argument supports values `'none'` and `'collinear'`. The initial per-site magnetization can be passed to bias the electronic solver into specific spin arrangements. If no per-site magnetization value is passed, CASTEP defaults to have zero magnetisation per site and prints a warning message in the output file. Since it is unlikely that the spin symmetry can be broken spontaneously, it is recommended that the user always pass the per-site initial magnetisation explicitly. While CASTEP itself supports non-collinear and spin-orbit calculations, they are not yet enabled through the common interface presented in this work. This is because the common interface `magnetization_per_site` input does not allow yet non-collinear initialization of the spins and, moreover, specially generated norm-conserving potentials are needed for spin-orbit calculations.

The `reference_workchain` argument makes sure that the k-points mesh that is used is the same as that used in the defined workchain for consistent energy comparisons.

## 3. Running in the Quantum Mobile virtual machine

The source code of CASTEP can be compiled using the following commands on Quantum Mobile. We assume that the `aiida-castep` plugin has been installed as one of the dependencies of `aiida-common-workflows`.

If internet connection is available, the following command can be used to install OpenBLAS in Quantum Mobile:

```
1 sudo apt-get update
2 sudo apt-get install libopenblas-dev
```

Listing 7. Commands to install the OpenBLAS library.

The availability of OpenBLAS may improve the execution speed of the code, but it is not essential. It is also possible to link with intel MKL library, but it is beyond the scope of this guide.

Assuming a source archive of CASTEP version 19.1.1 is placed in the working directory, paste the following commands into the terminal to compile a binary and setup the `Code` node for AiiDA:

```
1 cat > install-castep.sh << EOF
2 set -e
3 set -x
4 tar zxvf CASTEP-19.11.tar.gz
5 cd CASTEP-19.11
6 apt list --installed | grep libopenblas-dev > /dev/null
7 if [ \ $? -eq 0 ]; then
8 mathlib=openblas
9 else
10 mathlib=default
```

```

11 fi
12 set +e
13 make MATHLIB=\$mathlib FFT=fftw3 COMMS_ARCH=mpi SUBARCH=mpi TARGETCPU=portable -j `nproc` clean
14 set -e
15 make MATHLIB=\$mathlib FFT=fftw3 COMMS_ARCH=mpi SUBARCH=mpi TARGETCPU=portable -j `nproc`
16 make MATHLIB=\$mathlib FFT=fftw3 COMMS_ARCH=mpi SUBARCH=mpi TARGETCPU=portable -j `nproc` install
17 EOF
18 bash install-castep.sh
19
20 cat > castep-code.yaml << EOF
21 label: castep-19.1.1
22 description: CASTEP 19.1.1 Compiled locally
23 input_plugin: castep.castep
24 on_computer: True
25 computer: localhost
26 remote_abs_path: /home/max/CASTEP-19.11/bin/bin/linux_x86_64_gfortran7.0--mpi/castep.mpi
27 prepend_text: 'export OMP_NUM_THREADS=1'
28 append_text: ''
29 EOF
30 workon aiida
31 verdi code setup --config castep-code.yaml

```

Listing 8. Compiling a CASTEP executable on Quantum Mobile and settings up the Code node.

During the compilation, the user will be prompted to confirm the path to the libraries, and directory where the compiled binary to be installed. In both cases, press the Enter to use the default option.

Free-of-charge source code licenses for CASTEP can be obtained for academic use. This option is available to the researchers world-wide. More details can be found one the [official website](#).

## D. CP2K

### 1. Protocols

Every protocol employs Gaussian and plane waves[48] method to compute energy and forces. A multi-grid approach is used to represent the electron density and the product Gaussian functions. Several parameters are required to define the multi-grid. CUTOFF (Ry) parameter is the plane wave cutoff of the finest grid used to map the electron density. REL\_CUTOFF (Ry) determines how Gaussians are mapped into the multi-grid. NGRIDS parameter defines the total number of grids. EPS\_DEFAULT acts as a default value for many other parameters trying to achieve the precision in energy up to the value of EPS\_DEFAULT. All protocols have a default maximum force threshold of  $4.5 \cdot 10^{-4}$  Hartree per bohr and root mean square force threshold of  $3.0 \cdot 10^{-4}$  Hartree per bohr.

Here is the list of supported protocols and the values of the parameters:

- **fast**. This protocol is intended for testing purposes and uses very loose settings. The k-points distance is set to  $1.0 \text{ \AA}^{-1}$ . The multi-grid parameters are: 400 Ry CUTOFF, 30 Ry REL\_CUTOFF, and 4 NGRIDS. Target accuracy for the SCF convergence is set to  $1 \cdot 10^{-6}$  and EPS\_DEFAULT is set to  $1 \cdot 10^{-10}$ .
- **moderate**. This protocol is intended for general use and is expected to provide sensible results. The k-points distance is set to  $0.5 \text{ \AA}^{-1}$ . The multi-grid parameters are: 600 Ry CUTOFF, 40 Ry REL\_CUTOFF, and 4 NGRIDS. Target accuracy for the SCF convergence is set to  $1 \cdot 10^{-7}$  and EPS\_DEFAULT is set to  $1 \cdot 10^{-12}$ .
- **precise**. This protocol is intended for high-accuracy calculations. The k-points distance is set to  $0.1 \text{ \AA}^{-1}$ . The multi-grid parameters are: 1000 Ry CUTOFF, 50 Ry REL\_CUTOFF, and 3 NGRIDS. Target accuracy for the SCF convergence is set to  $1 \cdot 10^{-8}$  and EPS\_DEFAULT is set to  $1 \cdot 10^{-16}$ .

### 2. Supported relax types and calculation modes

The list of currently supported `relax_type` is: ‘none’, ‘positions’, ‘positions\_cell’. All protocols use the standard BFGS algorithm to optimize the atomic positions and the cell.

The `electronic_type` argument supports the values ‘metal’ and ‘insulator’. For ‘insulator’, the calculation is performed using the orbital transformation method with fixed occupations. The default total magnetic moment is equal to 0 for an even number and 1 for an odd number of electrons. For ‘metal’ a diagonalisation with Fermi-Dirac smearing is used at electronic temperature of 500 K and 20 molecular orbitals are added for each spin. In case smearing is employed, the total magnetization is flexible.

The `spin_type` argument supports the values ‘none’ and ‘collinear’. In the case of the ‘collinear’ option, the spin-polarized calculation are enabled. Additionally, the user can specify the `magnetization_per_site`, which would be ignored in the ‘none’ case. If the user provides the `magnetization_per_site`, the total multiplicity is derived from it.

### 3. Running in the Quantum Mobile virtual machine

No additional setup is required to run CP2K on Quantum Mobile: the code comes already preinstalled and the necessary data files are shipped together with the `aiida-common-workflows` package.

## E. FLEUR

The execution of the FLEUR implementation of the common relax workflows relies on the `FleurBaseRelaxWorkChain`, `FleurRelaxWorkChain`, `FleurScfWorkChain` and `FleurBaseWorkChain` of the `aiida-fleur` package [41].

### 1. Protocols

The FLEUR program sets most internal parameters during generation of the full input to appropriate values. The protocols for FLEUR contain two kinds of values, the first influences the behavior of the underlying workflows, while the second explicitly sets some internal parameters during input generation. All protocols have a default force threshold of  $1 \cdot 10^{-3}$  Hartree per bohr, which can be overwritten by user through the dedicated input of the relax common workflow. List of `protocols` supported and their description:

- **fast**. In this protocol parameters are set in a way to reduce the computational cost. The protocol should still yield reasonable results but should not be used for production calculations, merely for testing. The most important parameters are: k-point distance is set to  $0.4 \text{ \AA}^{-1}$ , self-consistency convergence threshold for charge density to  $2 \cdot 10^{-7}$  electrons/bohr<sup>3</sup>, k-max basis cut-off to  $3.2 \text{ bohr}^{-1}$ , maximal number of relaxation calls to three, maximal number of iterations per SCF workflow to 240.
- **moderate**. This protocol provides a reasonable trade-off between accuracy and the computational cost. The most important parameters are: k-point distance is set to  $0.2 \text{ \AA}^{-1}$ , self-consistency convergence threshold for charge density to  $2 \cdot 10^{-8}$  electron/bohr<sup>3</sup>, k-max basis cut-off to  $4.0 \text{ bohr}^{-1}$ , maximal number of relaxation calls to five, maximal number of iterations per SCF workflow to 240.
- **precise**. This protocol is based on the moderate protocol but various parameters have been changed to improve the precision at the expenses of an increased computational cost. This protocol should yield fully converged results and is recommended for production calculations that require higher precision than provided by the moderate one. The calculation parameters are: k-point distance is set to  $0.1 \text{ \AA}^{-1}$ , self-consistency convergence threshold for charge density to  $2 \cdot 10^{-9}$  electron/bohr<sup>3</sup>, k-max basis cut-off to  $5.0 \text{ bohr}^{-1}$ , maximal number of relaxation calls to ten, maximal number of iterations per SCF workflow to 360.

For all other parameters within the Full-potential Linearized-Augmented-Plane-Wave method (FLAPW) we rely on the default choices of FLEUR and the FLEUR input generator. Consistency of these is only enforced within workflows where total energies need to be compared through the `reference_workchain` argument. Regarding the choice of k-point mesh, the use of a single k-point along directions having no periodic conditions is enforced. For example, a film calculation will always have a single k-point along z-direction.

### 2. Supported calculation modes

The `relax_type` argument supports the ‘positions’ and ‘none’ values, meaning that atomic positions can be optimized. All protocols use the standard BFGS algorithm implemented in FLEUR to minimize forces acting on atoms. However, first several iterations will be run using the straight mixing and BFGS is used only after the largest force is small enough. This is done to suppress the common problem of muffin-tin-overlap because BFGS tends to

propose too large displacements during the first steps, moving atoms far away from the equilibrium positions and hence overlapping them.

The `electronic_type` argument supports the values ‘metal’ and ‘insulator’. Simulations for ‘insulator’ and ‘metal’ are treated in the current protocol in the same way. In general, for calculation of metals with FLEUR it might be necessary to use a denser k-point mesh or introduce a larger smearing around the Fermi energy. The k-point distance of  $0.15 \text{ \AA}^{-1}$  for the precise protocol is expected to be a reasonable compromise to treat both `electronic_types` in the same way.

The `spin_type` argument supports the values ‘none’ and ‘collinear’. In case of ‘collinear’ calculations, unless `magnetization_per_site` is explicitly defined, a non-zero starting magnetization is chosen where the value for each site is determined by the input generator of FLEUR. If `magnetization_per_site` is defined, starting magnetic moments for given sites are set in the input generator, which might also require explicitly breaking the symmetry of the crystal structure.

The `reference_workchain` argument guarantees that the FLAPW parameters and species parameters, including basis set cutoffs, muffin-tin radii are consistent to allow high accuracy energy differences between these simulations.

### 3. Running in the Quantum Mobile virtual machine

Since everything needed to run the FLEUR common workflows is already installed on Quantum Mobile, no further steps are needed.

## F. Gaussian

The Gaussian workflows have been tested and all the results shown have been calculated using the Gaussian 09 Revision D.01 [15] but the presented implementation should work with most versions of Gaussian, as the input and output syntax of the demonstrated calculations has remained the same.

### 1. Protocols

Gaussian has many internal checks to automatically set most computational parameters at appropriate values. The protocol specification tries to take advantage of this by only varying the basis set size, integration grid and optimization tolerance. Additionally, symmetry is disabled (route parameter `NoSymm`) in all cases to make the protocols more generally applicable. List of supported `protocols` and their descriptions:

- **fast**. A fast protocol that is mainly used for testing, it uses a small `Def2SVP` [49] basis set and a loose (`opt=loose`) geometry optimization tolerance.
- **moderate**. A protocol with moderate accuracy, using the `Def2TZVP` [49] basis set, `ultrafine` integration grid and the default geometry optimization tolerance.
- **precise**. A protocol using the `Def2QZVP` [50] basis set, `superfine` integration grid and tight (`opt=tight`) optimization tolerance.

### 2. Supported relax types and calculation modes

The supported `relax_type` values are ‘none’, which just performs a force calculation without any structural optimization (Gaussian keyword `force`), and ‘positions’, which uses the standard Gaussian optimization algorithm to optimize the atomic positions.

The `electronic_type` input is ignored (‘metal’ and ‘insulator’ are treated the same way).

The `spin_type` ‘none’ and ‘collinear’ are supported and correspond to a restricted (RKS) and unrestricted Kohn-Sham (UKS) calculation, respectively. In case of ‘collinear’ calculations, if no `magnetization_per_site` is explicitly passed, the lowest allowed spin multiplicity is specified (1 in case of even and 2 in case of odd electrons). If `magnetization_per_site` is passed, the site-specific spin information is disregarded and only total spin guess (sum of `magnetization_per_site`) is processed to determine the input spin multiplicity (rounded to the closest allowed spin multiplicity). In case of a calculation of the open-shell singlet, HOMO and LUMO are mixed to break the spin-symmetry (Gaussian route parameter `guess=mix`).

The `reference_workchain` argument is ignored.

### 3. *Running in the Quantum Mobile virtual machine*

After the user sets up their Gaussian code in the standard AiiDA manner, the common workflows are available without any further setup.

## G. NWChem

### 1. *Protocols*

List of `protocols` supported and their description:

- **fast.** The minimum k-point distance is set to  $1.0 \text{ \AA}^{-1}$  and an SCF energy convergence tolerance of  $1.0e^{-5}$  Ha is set. The geometry convergence is set to `loose`, corresponding to a forces threshold on any atom of  $4.5e^{-3}$  Ha bohr $^{-1}$ , and a root mean square (RMS) gradient of  $3.0e^{-3}$  Ha bohr $^{-1}$
- **moderate.** The minimum k-point distance is set to  $0.2 \text{ \AA}^{-1}$  and an SCF energy convergence tolerance of  $1.0e^{-7}$  Ha is set. The geometry tolerance is set to `default`, corresponding to a forces threshold on any atom of  $4.5e^{-4}$  Ha bohr $^{-1}$ , and an RMS gradient of  $3.0e^{-4}$  Ha bohr $^{-1}$
- **precise.** The minimum k-point distance is set to  $0.1 \text{ \AA}^{-1}$  and an SCF energy convergence tolerance of  $1.0e^{-9}$  Ha is set. The geometry tolerance is set to `tight`, corresponding to a forces threshold on any atom of  $1.5e^{-5}$  Ha bohr $^{-1}$ , and an RMS gradient of  $1.0e^{-5}$  Ha bohr $^{-1}$

For all protocols, the PBE exchange-correlation functional is used with Hamann and Troullier-Martins norm-conserving pseudopotentials. In the case of molecular calculations run using the common interface, the plane wave code is also used rather than the main DFT module in NWChem which employs localised basis sets. This for consistency between protocols. The main differences are that the gamma-point only code is used, and that a larger cutoff of 140 Ha is set for all protocols.

### 2. *Supported relax types and calculation modes*

The `relax_type` supported are: `'none'`, `'positions'`, `'positions_cell'`, and `'cell'`. The relaxation of the atomic coordinates and cell vectors is performed using the `DRIVER` module, which implements a quasi-newton optimization with line searches and approximate energy Hessian updates.

If `electronic_type` is set as `'insulator'`, the wavefunction is optimized using Grassmann L-BFGS algorithm. If smearing is required, by setting `electronic_type` to `'metal'`, a band-by-band optimiser is used with Fermi-Dirac smearing.

Spin polarized calculations are not supported yet through the common interface, they will be enabled in the near future.

### 3. *Running in the Quantum Mobile virtual machine*

Other than adding the NWChem code in the standard AiiDA manner, the common workflows are available without any further setup.

## H. ORCA

### 1. *Protocols*

The following protocols are defined to be used to a wide range of systems considering the required accuracy/cost by the user. These protocols provide higher accuracy via increasing the size of the basis set and tightening the convergence



criteria for self-consistent field and geometry convergence. We used recent versions of Ahlrichs basis sets[49, 50] in the presented protocols. It is noteworthy that the latter criteria can be easily altered via invoking proper keywords.

- **fast**. It uses Def2-SVP[49] basis along with scf and geometry optimization convergence criteria of Strong and LOOSOPT, respectively, to provide quick setup for testing purposes.
- **moderate**. We used triple zeta Def2-TZVP[49] basis set and decreased the electronic and geometry optimization to Tight and NORMALOPT which provides a fair level of accuracy.
- **precise**. Increased basis set size to quadruple zeta (Def2-QZVP[50]) along with tighter scf and geometry tolerance of VeryTight and TIGHTOPT, respectively.

## 2. Supported relax types and calculation modes

The ORCA implementation supports the ‘none’ and ‘positions’ values for the `relax_type`. The former one is intended for single point calculation and the latter one for the geometry optimization.

‘metal’ and ‘insulator’ have the same effect as `electronic_type` input; both will be ignored.

Restricted and unrestricted Kohn-Sham calculations can be requested by setting the `spin_type` to ‘none’ and ‘collinear’, respectively. In the latter case if `magnetization_per_site` is NOT provided the spin multiplicity of 1 and 2 will be set for even and odd number of electrons in the system, respectively, if `magnetization_per_site` is given, the site-specific spin information is ignored and only total spin guess (sum of `magnetization_per_site`) is processed to determine the input spin multiplicity (rounded to the closest allowed spin multiplicity).

This implementation does not use the `reference_workchain`. Also, as ORCA internally sets thresholds for forces based on the selected geometry optimization convergence criteria, we do not explicitly define them.

## 3. Running in the Quantum Mobile virtual machine

ORCA is a free-ware for academic users for academic usage and can be obtained via registering on the official website (<https://orcaforum.kofo.mpg.de>). All other uses can request a license at [www.faccts.de](http://www.faccts.de). ORCA comes as two versions with static and shared libraries and requires an MPI engine that can be either OpenMPI or MPICH. The shared version is recommended to be used within Quantum Mobile as it requires less disk space. Afterwards, ORCA can be setup following the instructions provided in detail in AiiDA documentations and be used for running ORCA implementation of the present work. It should be noted that in case of using shared version of ORCA, `LD_LIBRARY_PATH` environmental variable should be set to the location of ORCA in Quantum Mobile.

# I. QUANTUM ESPRESSO

## 1. Protocols

All of the QUANTUM ESPRESSO[19, 20] protocols use the Standard Solid State Pseudopotentials (SSSP)[51] v1.1. Below is a list the supported `protocols` and their description:

- **fast**. The fast protocol is designed to yield reasonable results at minimal computational cost and should only be used for testing and demonstration purposes. It uses the efficiency configuration of the SSSP and the following precision settings: k-points distance is  $0.5 \text{ \AA}^{-1}$ , self-consistency convergence threshold is  $0.4 \cdot 10^{-9}$  Ry per atom, energy threshold for the ionic convergence is  $1 \cdot 10^{-4}$  Ry per atom and forces threshold is  $1 \cdot 10^{-3}$  Ry/bohr. The convergence on the stress is left to the default of the code which is 0.05 GPa.
- **moderate**. The moderate protocol is the default protocol used for production calculations. It uses the efficiency configuration of the SSSP and the following precision settings: k-points distance of  $0.15 \text{ \AA}^{-1}$ , self-consistency convergence threshold of  $0.2 \cdot 10^{-9}$  Ry per atom, energy threshold for the ionic convergence of  $1 \cdot 10^{-5}$  Ry per atom and forces threshold of  $1 \cdot 10^{-4}$  Ry/bohr. The convergence on the stress is left to the default of the code which is 0.05 GPa.
- **precise**. The precise protocol should yield fully converged results and is recommended for production calculations that require more precision than provided by the moderate protocol. It uses the precision configuration of the SSSP and the following precision settings: k-points distance is  $0.1 \text{ \AA}^{-1}$ , self-consistency convergence threshold is

$0.1 \cdot 10^{-9}$  Ry per atom, energy threshold for the ionic convergence is  $0.5 \cdot 10^{-5}$  Ry per atom and forces threshold is  $0.5 \cdot 10^{-4}$  Ry/bohr. The convergence on the stress is left to the default of the code which is 0.05 GPa.

The protocol settings listed above are the result of a rigorous study that will be detailed in a forthcoming publication.

## 2. Supported calculation modes

The `relax_type` argument supports all the agreed values, meaning that atomic positions and the cell size and shape can be optimized, as well as any combination of those. All protocols will use the standard BFGS algorithm to minimize the forces on the atoms and stress on the cell.

The `electronic_type` argument supports the values ‘metal’ and ‘insulator’. For ‘insulator’, the calculation is performed with fixed occupations, whereas for ‘metal’ a Marzari-Vanderbilt smearing [52] is employed with a broadening of 0.01 Ry.

The `spin_type` argument supports the values ‘none’ and ‘collinear’. In case of ‘collinear’ calculations, unless `magnetization_per_site` is explicitly defined, a non-zero starting magnetization is chosen where the value for each site is element dependent.

The `reference_workchain` argument guarantees that the k-points mesh that is used is identical to that used in the defined workchain.

## 3. Running in the Quantum Mobile virtual machine

All that is needed to run the QUANTUM ESPRESSO common workflows is to install the SSSP families, which is done as follows:

```
1 aida-pseudo install sssp -p efficiency
2 aida-pseudo install sssp -p precision
```

Listing 9. Install the required SSSP pseudopotential families for the QUANTUM ESPRESSO common workflows.

## J. SIESTA

The siesta implementation of the common relax workflow can be used only with a version of the SIESTA code that support the use of pseudopotentials in "psml" format. The [wiki page](#) of the SIESTA code can be checked for up-to-date information on the compatible versions.

### 1. Protocols

All the SIESTA protocols use the same set of pseudopotentials: the "standard" scalar relativistic set of PseudoDojo[45], which are norm-conserving pseudopotentials. We warn users that the optimal basis set for a SIESTA calculation depends on the chemical environment, therefore an optimization procedure should be carried on before the study of any new system. The statically chosen basis for each chemical species included in the protocols must be considered a good starting point, but not assumed to be optimal in any possible environment. In fact, they have been tested only on crystal elements through the  $\Delta$  test. More about the tests of this protocols will be detailed in a forthcoming publication.

The list of supported protocols and their description follows.

- **fast.** Protocol running SIESTA with standard inputs for testing. This includes an automatically generated basis “DZ” (Double-Zeta, two radial functions per angular momentum channel) with `energy-shift` of 200 meV, a `mesh-cutoff` of 50 Ry and a k-points distance of  $0.2 \text{ \AA}^{-1}$ . Tolerance for the density matrix is set to  $10^{-3}$ . The forces and stress thresholds are 0.04 eV/Ang and 1 GPa respectively.
- **moderate.** A protocol combining simple basis selection, moderate computational resources and satisfactory results in the comparison with all-electron references. It globally sets basis “DZP” (Double-Zeta-Polarized, two radial functions per angular momentum channel plus an orbital obtained from the polarization of the highest occupied orbital) with `energy-shift` of 50 meV, a `mesh-cutoff` of 200 Ry and a k-points distance of  $0.1 \text{ \AA}^{-1}$ . However custom basis (manually selected basis orbitals radii) are employed for the Ca, Sr and Ba elements, since the DZP

choice resulted in too-large radius for the “s” orbitals. Tolerance for the density matrix is set to  $10^{-4}$ . The forces and stress thresholds are 0.04 eV/Å and 1 GPa respectively.

- **precise.** Protocol with stringent settings and optimized basis for crystal elements. It globally sets a `mesh-cutoff` of 500 Ry and a k-points distance of  $0.1 \text{ \AA}^{-1}$ . Basis sizes and orbital radii are customized for each chemical species through an optimization procedure based on the basis enthalpy minimization for crystal elements. Tests based on the  $\Delta$  test confirm the effectiveness of the choice. Tolerance for the density matrix is set to  $10^{-5}$ . The forces and stress thresholds are 0.005 eV/Å, 0.7 GPa respectively.

## 2. Supported relax types and calculation modes

The `relax_type` supported are: ‘none’, ‘positions’, ‘positions\_shape’ and ‘cell’. The relaxation of atomic coordinates is performed with the conjugate-gradient algorithm until reaching the forces threshold value. In the case of variable-cell relaxation, the minimization targets zero hydrostatic pressure through the conjugate-gradient algorithm.

The `electronic_type` ‘metal’ and ‘insulator’ are treated in the same way. SIESTA calculations on metallic systems usually require a denser k-points grid compared to insulators, however, at least for the `precise` and `moderate` protocols, the selected k-points distance is expected to generate a sufficiently dense mesh to treat both `electronic_types` in the same way.

The `spin_type` ‘none’ and ‘collinear’ are supported. In case of ‘collinear’ calculations, if no `magnetization_per_site` is explicitly passed, a ferromagnetic arrangement is imposed as initial magnetization, with maximum atomic moment on each atom.

The `reference_workchain` argument generates a calculation with the same k-points mesh and same real space mesh of the `reference_workchain`. The real space mesh is set using the SIESTA keyword `mesh-sizes`.

## 3. Running in the Quantum Mobile virtual machine

In the Quantum Mobile, the creation of a pseudopotential family with name `nc-sr-04_pbe_standard_psml` is necessary in order to run SIESTA calculations through the common interface. This is achieved simply typing in the command line:

```
1 verdi data psml uploadfamily /usr/local/share/siesta/psml-files-qm/nc-sr-04_pbe_standard/ nc-sr-04
   _pbe_standard_psml "pseudos from PseudoDojo"
```

Listing 10. Command to set up the pseudopotential family required by the SIESTA implementation of the relax common workflow.

## K. VASP

All results in this work have been produced with VASP 5.4.4 and AiiDA-VASP 2.1.0. The potentials used are based on the standard PBE potentials supplied with the same VASP version. The precise protocol was used to generate all results. The Projector Augmented-Wave method (PAW) was used[24].

### 1. Protocols

List of protocols supported and their differences:

- **fast.** Low precision, minimal computational cost. For testing purposes. k-points distance is  $0.25 \text{ \AA}^{-1}$ . Maximum forces threshold is  $1.0 \times 10^{-1} \text{ eV/\AA}$ . `PREC` is `Single`. `EDIFF` is  $1.0 \times 10^{-4}$ . The protocol imposes at least six electronic steps per self-consistent cycle. A conjugate gradient algorithm for relaxing the positions is implemented.
- **moderate.** Standard precision, moderate computational cost. k-points distance is  $0.15 \text{ \AA}^{-1}$ . Maximum forces threshold is  $1.0 \times 10^{-2} \text{ eV/\AA}$ . `PREC` is `Normal`. `EDIFF` is  $1.0 \times 10^{-5}$ . The protocol imposes at least four electronic steps per self-consistent cycle. A quasi-Newton algorithm for relaxing the positions.
- **precise.** Elevated precision, high computational cost. k-points distance is  $0.10 \text{ \AA}^{-1}$ . Maximum forces threshold is  $1.0 \times 10^{-3} \text{ eV/\AA}$ . `EDIFF` is  $1.0 \times 10^{-6}$ . `PREC` is `Accurate`. A quasi-Newton algorithm for relaxing the positions.

For all the workflows, the following was used: Gaussian smearing with width of 0.2 eV for the integration and a static PW cutoff of 550 eV. Otherwise default VASP settings has been used, as defined for the version used. Also note that we do not use an explicit stress cutoff, but rely strictly on the maximum force cutoff.

We would like to stress the fact that the statically chosen PW cutoff and k-point grids are purely an artifact of needing higher cutoffs for the calculation of the H<sub>2</sub> dissociation curve and at the same time focusing on simplicity and transferability of the same protocol definition across the demonstrators in this work. The purpose of this work is not to present state-of-the-art accuracy or precision. And in order to make the main content to the point, we did not want to utilize dedicated convergence workchains to enable a more tailored and production ready result. However, we hope the readers will, after reading, appreciate that adding such a workchain would be straightforward. For similar reasons, as can be seen, we use the same Gaussian smearing and a rather wide smearing width for all demonstrators.

## 2. Supported calculation modes

The `relax_type` argument supports all defined values in this work, while the `spin_type` argument supports the values ‘none’ and ‘collinear’. In case of ‘collinear’ calculations, `magnetization_per_site` can be supplied to indicate initial magnetic moment per site in units of Bohr magneton. In the case this is not provided, the default in VASP is used, which is one Bohr magneton per atom. The keyword `threshold_stress` is ignored for the reason explained in the previous paragraph. Finally, the `reference_workchain` argument guarantees that the k-points mesh that is used is identical to that used in the reference workchain. The `electronic_type` is not used for VASP and we used Gaussian smearing for all calculations as described in the previous section.

## 3. Running in the Quantum Mobile virtual machine

Due to licensing VASP is not provided in the Quantum Mobile. However, it is straightforward to compile VASP and add it to Quantum Mobile for licence holders. The AiiDA-VASP plugin is however included in Quantum Mobile as part of this work.

## VII. INSTRUCTIONS TO RUN THE TEST CASES IN THE QUANTUM MOBILE.

We present here the full set of instructions that are needed to reproduce the results presented in the manuscript, using the Quantum Mobile. N.B.: since the results reported in the paper are obtained with very stringent choices of parameters (‘precise’ protocol), the commands listed below might start computationally-demanding simulations. In addition, the quantum-engine simulations are run with two processors by default. For Quantum Mobile v21.05.1 the included compiled binaries for ABINIT, CP2K, FLEUR and SIESTA can fail for certain runs if run with more than one processor and instead have to be run in serial. To specify the number of processors that should be used, the command line interface (CLI) provides the option “-n”. The option accepts a value for each engine of the calculation, meaning two integers must be provided for FLEUR (“-n 1 2” for instance) and one for any other code (“-n 4”). Moreover calculations on molecules (H<sub>2</sub> and ammonia) performed with codes designed for extended systems might result in calculations that require a lot of memory.

The following list of instructions shows how the results of this paper can be reproduced. The label <code> must be substituted with the name of one of the available quantum engines.

1. [Download Quantum Mobile v21.05.1](#) and follow the instructions for its installation. The installation just requires the Virtual Box software.
2. Start the virtual machine, open a terminal and type "workon aiiida" to activate the virtual environment where all the python packages are installed.
3. Make sure to perform the preliminary steps required by the quantum engine of your interest, they are listed in the previous section under *Running in the Quantum Mobile virtual machine* subsection for each code.
4. For the ammonia test case, run the command:

```
aiida-common-workflows launch relax -S NH3-pyramidal -p precise -- <code>
```

wait the end of the execution and then run:

```
aiida-common-workflows launch relax -S NH3-planar -p precise -- <code>
```

Both commands will return an output link called `total_energy`. Run:

```
verdi node attributes <pk>
```

to explore the value of the total energy for both calculation, `<pk>` is the node pk of `total_energy`. The difference in energy between the planar and pyramidal energy is the value reported in Figure 4 of the main text.

5. The calculation of the EOS for elements  $\text{El} = [\text{Si}, \text{Al}]$  is started with the command:

```
aiida-common-workflows launch eos -S El -p precise -- <code>
```

This command launches a relaxation of the atomic coordinates for a spin-less system with ‘metal’ setting. All these are defaults for the command line interface implemented in the package. Once the calculation is over, results are obtained running:

```
aiida-common-workflows launch plot-eos -- <pk>
```

where `<pk>` is the pk of the EOS workflow, reported at run time. The command directly plots the energy versus volume data. In case a print of the values is needed, the option `"-t -p 4 5"` can be added. The `"-p 4 5"` indicates that the volumes and energies should be printed with 4 and 5 decimal figures respectively.

6. The calculation of the EOS for GeTe is started with the command:

```
aiida-common-workflows launch eos -S GeTe -p precise -r positions_shape -- <code>
```

Once the calculation is finished, the results can be obtained with the same command explained in the previous point.

7. The calculation of the EOS for Fe in the ferromagnetic arrangement is started with:

```
aiida-common-workflows launch eos -S Fe -p precise -s collinear -- <code>
```

The initial magnetization used is different for each `<code>` since the default is used. The command for obtaining the results is always the same except that a further integer can be accepted by the `"-p"` option to indicate the decimal figures for the total magnetization (`"-p 4 5 4"`) for instance. The calculation for the anti-ferromagnetic case is done with:

```
aiida-common-workflows launch eos -S Fe -p precise -s collinear --magnetization_per_site -4 4 -- <code>
```

The only difference respect to the ferromagnetic case is the explicit choice of an initial magnetization.

8. The  $\text{H}_2$  dissociation curve is obtained with:

```
aiida-common-workflows launch dissociation-curve -S H2 -p precise -s collinear --magnetization-per-site -1 1 -- <code>
```

and the results are obtained with:

```
aiida-common-workflows launch plot-dissociation-curve -- <pk>
```

The same options of the `plot-eos` command can be used to custom the results analysis.

For simplicity, we illustrated how to run the test cases using the command line interface of the package. Another possibility is to create a submission script for each case. Detailed documentation on the topic can be found in the [aiida-common-workflow online documentation](#).

## BIBLIOGRAPHY

- 
- [1] K. Burke, *Perspective on density functional theory*, *The Journal of Chemical Physics* **136**, 150901 (2012).  
 [2] R. Jones, *Density functional theory: Its origins, rise to prominence, and future*, *Reviews of Modern Physics* **87**, 897 (2015).  
 [3] K. Lejaeghere, G. Bihlmayer, T. Bjorkman *et al.*, *Reproducibility in density functional theory calculations of solids*, *Science* **351**, aad3000 (2016).

- [4] A. H. Larsen, J. J. Mortensen, J. Blomqvist *et al.*, *The atomic simulation environment—a Python library for working with atoms*, *Journal of Physics: Condensed Matter* **29**, 273002 (2017).
- [5] M. Gjerding, T. Skovhus, A. Rasmussen *et al.*, *Atomic Simulation Recipes – a Python framework and library for automated workflows* (2021).
- [6] S. P. Huber, S. Zoupanos, M. Uhrin *et al.*, *AiiDA 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance*, *Scientific Data* **7** (2020).
- [7] X. Gonze, F. Jollet, F. A. Araujo *et al.*, *Recent developments in the ABINIT software package*, *Computer Physics Communications* **205**, 106 (2016).
- [8] X. Gonze, B. Amadon, G. Antonius *et al.*, *The Abinitproject: Impact, environment and recent developments*, *Computer Physics Communications* **248**, 107042 (2020).
- [9] A. H. Romero, D. C. Allan, B. Amadon *et al.*, *ABINIT: Overview and focus on selected capabilities*, *The Journal of Chemical Physics* **152**, 124102 (2020).
- [10] L. E. Ratcliff, W. Dawson, G. Fiscaro *et al.*, *Flexibilities of wavelets as a computational basis set for large-scale electronic structure calculations*, *The Journal of Chemical Physics* **152**, 194110 (2020).
- [11] S. J. Clark, M. D. Segall, C. J. Pickard *et al.*, *First principles methods using CASTEP*, *Zeitschrift für Kristallographie - Crystalline Materials* **220** (2005).
- [12] J. Hutter, M. Iannuzzi, F. Schiffrmann and J. VandeVondele, *cp2k: atomistic simulations of condensed matter systems*, *Wiley Interdisciplinary Reviews: Computational Molecular Science* **4**, 15 (2013).
- [13] T. D. Kühne, M. Iannuzzi, M. D. Ben *et al.*, *CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations*, *The Journal of Chemical Physics* **152**, 194103 (2020).
- [14] <https://www.flapw.de>, *The Jülich FLAPW code family*.
- [15] M. J. Frisch, G. W. Trucks, H. B. Schlegel *et al.*, *Gaussian 09 Revision D.01* (2016).
- [16] E. Aprà, E. J. Bylaska, W. A. de Jong *et al.*, *NWChem: Past, present, and future*, *The Journal of Chemical Physics* **152**, 184102 (2020).
- [17] F. Neese, *The ORCA program system*, *WIREs Computational Molecular Science* **2**, 73 (2011).
- [18] F. Neese, *Software update: the ORCA program system, version 4.0*, *WIREs Computational Molecular Science* **8** (2017).
- [19] P. Giannozzi, S. Baroni, N. Bonini *et al.*, *QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials*, *Journal of Physics: Condensed Matter* **21**, 395502 (2009).
- [20] P. Giannozzi, O. Andreussi, T. Brumme *et al.*, *Advanced capabilities for materials modelling with Quantum ESPRESSO*, *Journal of Physics: Condensed Matter* **29**, 465901 (2017).
- [21] J. M. Soler, E. Artacho, J. D. Gale *et al.*, *The SIESTA method for ab initio order-N materials simulation*, *Journal of Physics: Condensed Matter* **14**, 2745 (2002).
- [22] A. García, N. Papior, A. Akhtar *et al.*, *Siesta: Recent developments and applications*, *The Journal of Chemical Physics* **152**, 204108 (2020).
- [23] G. Kresse and J. Furthmüller, *Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set*, *Physical Review B* **54**, 11169 (1996).
- [24] G. Kresse and D. Joubert, *From ultrasoft pseudopotentials to the projector augmented-wave method*, *Physical Review B* **59**, 1758 (1999).
- [25] L. Talirz, S. Kumbhar, E. Passaro *et al.*, *Materials Cloud, a platform for open computational science*, *Scientific Data* **7** (2020).
- [26] C. Stoddart, *Is there a reproducibility crisis in science?*, *Nature* (2016).
- [27] M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg *et al.*, *The FAIR Guiding Principles for scientific data management and stewardship*, *Scientific Data* **3** (2016).
- [28] C. Goble, S. Cohen-Boulakia, S. Soiland-Reyes *et al.*, *FAIR Computational Workflows*, *Data Intelligence* **2**, 108 (2020).
- [29] M. Uhrin, S. P. Huber, J. Yu *et al.*, *Workflows in AiiDA: Engineering a high-throughput, event-based engine for robust and modular computational workflows*, *Computational Materials Science* **187**, 110086 (2021).
- [30] D. Gresch, Q. Wu, G. W. Winkler *et al.*, *Automated construction of symmetrized Wannier-like tight-binding models from ab initio calculations*, *Physical Review Materials* **2** (2018).
- [31] G. Pizzi, A. Cepellotti, R. Sabatini *et al.*, *AiiDA: automated interactive infrastructure and database for computational science*, *Computational Materials Science* **111**, 218 (2016).
- [32] J. P. Perdew, K. Burke and M. Ernzerhof, *Generalized Gradient Approximation Made Simple*, *Physical Review Letters* **77**, 3865 (1996).
- [33] D. C. Ghosh, J. Jana and R. Biswas, *Quantum chemical study of the umbrella inversion of the ammonia molecule*, *International Journal of Quantum Chemistry* **80**, 1 (2000).
- [34] A. Cerioni, L. Genovese, A. Mirone and V. A. Sole, *Efficient and accurate solver of the three-dimensional screened and unscreened Poisson's equation with generic boundary conditions*, *The Journal of Chemical Physics* **137**, 134108 (2012).
- [35] A. Castro, A. Rubio and M. J. Stott, *Solution of Poisson's equation for finite systems using plane-wave methods*, *Canadian Journal of Physics* **81**, 1151 (2003).
- [36] L. Bengtsson, *Dipole correction for surface supercell calculations*, *Physical Review B* **59**, 12301 (1999).
- [37] G. Makov and M. C. Payne, *Periodic boundary conditions in ab initio calculations*, *Physical Review B* **51**, 4014 (1995).
- [38] J. Goldak, C. S. Barrett, D. Innes and W. Youdelis, *Structure of Alpha GeTe*, *The Journal of Chemical Physics* **44**, 3323 (1966).
- [39] A. J. Cohen, P. Mori-Sanchez and W. Yang, *Insights into Current Limitations of Density Functional Theory*, *Science* **321**, 792 (2008).

- [40] S. P. Huber, E. Bosoni, M. Bercx *et al.*, *Common workflows for computing material properties using different quantum engines*, *Materials Cloud Archive* (2021), doi:10.24435/materialscloud:nz-01.
- [41] J. Bröder, D. Wortmann and S. Blügel, *Using the AiiDA-FLEUR package for all-electron abinitio electronic structure data generation and processing in materials science*, in *In Extreme Data Workshop 2018 Proceedings, IAS Series*, vol. 40, pp. p 43–48, Forschungszentrum Jülich, Jülich (2019).
- [42] K. Lejaeghere, V. V. Speybroeck, G. V. Oost and S. Cottenier, *Error Estimates for Solid-State Density-Functional Theory Predictions: An Overview by Means of the Ground-State Elemental Crystals*, *Critical Reviews in Solid State and Materials Sciences* **39**, 1 (2013).
- [43] D. C. Liu and J. Nocedal, *On the limited memory BFGS method for large scale optimization*, *Mathematical Programming* **45**, 503 (1989).
- [44] E. Bitzek, P. Koskinen, F. Gähler *et al.*, *Structural Relaxation Made Simple*, *Physical Review Letters* **97** (2006).
- [45] M. van Setten, M. Giantomassi, E. Bousquet *et al.*, *The PseudoDojo: Training and grading a 85 element optimized norm-conserving pseudopotential table*, *Computer Physics Communications* **226**, 39 (2018).
- [46] M. Torrent, F. Jollet, F. Bottin *et al.*, *Implementation of the projector augmented-wave method in the ABINIT code: Application to the study of iron under pressure*, *Computational Materials Science* **42**, 337 (2008).
- [47] S. Saha, *Soft and accurate norm conserving pseudopotentials and their application for structure prediction*, Ph.D. thesis, University of Basel (2017).
- [48] J. VandeVondele, M. Krack, F. Mohamed *et al.*, *Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach*, *Computer Physics Communications* **167**, 103 (2005).
- [49] F. Weigend and R. Ahlrichs, *Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy*, *Physical Chemistry Chemical Physics* **7**, 3297 (2005).
- [50] F. Weigend, F. Furche and R. Ahlrichs, *Gaussian basis sets of quadruple zeta valence quality for atoms H–Kr*, *The Journal of Chemical Physics* **119**, 12753 (2003).
- [51] G. Prandini, A. Marrazzo, I. E. Castelli *et al.*, *Precision and efficiency in solid-state pseudopotential calculations*, *npj Computational Materials* **4** (2018).
- [52] N. Marzari, D. Vanderbilt, A. D. Vita and M. C. Payne, *Thermal Contraction and Disorder of the Al(110) Surface*, *Physical Review Letters* **82**, 3296 (1999).