

MIT Open Access Articles

Python Materials Genomics (pymatgen): A robust, open-source python library for materials analysis

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ong, Shyue Ping, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L. Chevrier, Kristin A. Persson, and Gerbrand Ceder. "Python Materials Genomics (pymatgen): A Robust, Open-Source Python Library for Materials Analysis." *Computational Materials Science* 68 (February 2013): 314–319.

As Published: <http://dx.doi.org/10.1016/j.commatsci.2012.10.028>

Publisher: Elsevier

Persistent URL: <http://hdl.handle.net/1721.1/101936>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License



Python Materials Genomics (pymatgen) : A Robust, Open-Source Python Library for Materials Analysis.

Shyue Ping Ong

Department of Materials Science and Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

William Davidson Richard

Department of Materials Science and Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

Anubhav Jain

Lawrence Berkeley National Lab, 1 Cyclotron Rd Berkeley, CA 94720

Geoffroy Hautier

Université catholique de Louvain, Place de l'Université 1, 1348 Louvain-La-Neuve Belgium

Michael Kocher

Lawrence Berkeley National Lab, 1 Cyclotron Rd Berkeley, CA 94720

Shreyas Cholia

Lawrence Berkeley National Lab, 1 Cyclotron Rd Berkeley, CA 94720

Dan Gunter

Lawrence Berkeley National Lab, 1 Cyclotron Rd Berkeley, CA 94720

Vincent L. Chevrier

3M, Electronics Markets Materials Division, St. Paul, Minnesota 55144, USA

Email addresses: shyue@mit.edu (Shyue Ping Ong), wrichard@mit.edu (William Davidson Richard), ajain@lbl.gov (Anubhav Jain), geoffroy.hautier@uclouvain.be (Geoffroy Hautier), mpkocher@lbl.gov (Michael Kocher), scholia@lbl.gov (Shreyas Cholia), dkgunter@lbl.gov (Dan Gunter), vincentchevrier@gmail.com (Vincent L. Chevrier), kapersson@lbl.gov (Kristin A. Persson), gceder@mit.edu (Gerbrand Ceder)
URL: <http://ceder.mit.edu> (Gerbrand Ceder)

Kristin A. Persson

Lawrence Berkeley National Lab, 1 Cyclotron Rd Berkeley, CA 94720

Gerbrand Ceder

*Department of Materials Science and Engineering, Massachusetts Institute of
Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA*

Abstract

We present the Python Materials Genomics (pymatgen) library, a robust, open-source Python library for materials analysis. A key enabler in high-throughput computational materials science efforts is a robust set of software tools to perform initial setup for the calculations (e.g., generation of structures and necessary input files) and post-calculation analysis to derive useful material properties from raw calculated data. The pymatgen library aims to meet these needs by (1) defining core Python objects for materials data representation, (2) providing a well-tested set of structure and thermodynamic analyses relevant to many applications, and (3) establishing an open platform for researchers to collaboratively develop sophisticated analyses of materials data obtained both from first principles calculations and experiments. The pymatgen library also provides convenient tools to obtain useful materials data via the Materials Project’s REpresentational State Transfer (REST) Application Programming Interface (API). As an example, using pymatgen’s interface to the Materials Project’s REST API and *phasediagram* package, we demonstrate how the phase and electrochemical stability of a recently synthesized material, Li_4SnS_4 , can be analyzed using a minimum of computing resources. We find that Li_4SnS_4 is a stable phase in the Li-Sn-S phase diagram (consistent with the fact that it can be synthesized), but the narrow range of lithium chemical potentials for which it is predicted to be stable would suggest that it is not intrinsically stable against typical electrodes used in lithium-ion batteries.

Keywords: materials, project, design, thermodynamics, high-throughput

1. Introduction

First principles calculations have the potential to greatly accelerate the design and optimization of new materials. In the past decade, electronic structure calculation codes[1, 2, 3, 4] have reached a level of maturity such that it is now possible to reliably automate and scale first principles calculations across any number of compounds, subject only to the limits of available computing resources. Indeed, there are currently several parallel initiatives that employ high-throughput first principles calculations in materials design. For example, the Materials Project (<http://www.materialsproject.org>) aims to calculate the properties of all known inorganic materials and make this data publicly available to the materials community to accelerate innovation in materials research. The Materials Project is based on the high-throughput framework developed by Jain et al.[5] and subsequently extended by collaborators at the Lawrence Berkeley Laboratory and National Energy Research Scientific Computing Center (NERSC). This framework has been used to screen over 80,000 inorganic compounds for a variety of applications, including Li-ion and Na-ion batteries.[6, 7, 8, 9, 10] Similarly, Curtarolo et al.[11] have developed the AFLOW (Automatic Flow) software framework for high-throughput calculation of crystal structure properties of alloys, intermetallics and inorganic compound and applied it to the investigation of the effect of structure on the stability of binary alloys[12] and superconductors,[13] and the search for topological insulators.[14] Yet another example of high-throughput materials design can be found in the CatApp developed by Hummelshoj et al.,[15] which provides a web application to access activation energies of elementary surface reactions and is part of a larger database of surface reaction data being developed under the Quantum Materials Informatics Project (<http://www.qmip.org>). On the molecular front, the Clean Energy Project[16] uses high-throughput computational chemistry to look for the best organic molecules for various applications, including organic semiconductors[17] and polymers for the membranes used in fuel cells for electricity generation. In this paper, we describe the Python Materials Genomics (pymatgen) library, a robust, open-source Python library for materials analysis. A key enabler in high-throughput computational materials science efforts is a robust set of software tools to perform initial setup for the calculations (e.g., generation of structures and necessary input files) and post-calculation analysis to derive useful material properties from raw calculated data. The aims of pymatgen are as follows:

1. Define core Python objects for materials data representation.
2. Provide a well-tested set of structure and thermodynamic analysis tools relevant to many applications.
3. Establish an open platform for researchers to collaboratively develop sophisticated analyses of materials data obtained both from first principles calculations and experiments.

The pymatgen library is currently used in the Materials Project for structure generation, manipulation and thermodynamic analysis. As such, it has been robustly tested over the large database of compounds in the Materials Project database. However, it should be noted that while the pymatgen library supports the Materials Project, it is designed to be a standalone library, and most of its analysis tools are flexible enough to be used by any materials researcher with other electronic structure codes and sources of data. The latest stable version of pymatgen (version 2.0.0 as of this paper) can be obtained via the Python Package Index at <http://pypi.python.org/pypi/pymatgen>, while the “bleeding edge” developmental version can be obtained from the official GitHub repo at <http://github.com/materialsproject/pymatgen>.

2. Overview of pymatgen

The pymatgen library is written in the Python programming language, and leverages the large number of available standard and scientific programming libraries, including the widely used numpy and scipy libraries.[18] It is compatible with Python version 2.7.x, but a transition to Python 3 is planned when the necessary libraries become available. It is primarily based on the object-oriented programming paradigm to facilitate code reuse and ensure modularity in design. In terms of development, we adopt a test-driven approach, and pymatgen includes unit tests for all non-trivial classes and methods. We also place an emphasis on clear and concise documentation, which is available at <http://materialsproject.github.com/pymatgen/>. Figure 1 provides an overview of the pymatgen library. A typical workflow would involve a user converting data (structure, calculations, etc.) from various sources (first principles calculations, crystallographic and molecule input files, Materials Project, etc.) into Python objects using pymatgen’s io packages, which are then used to perform further structure manipulation or analyses. The pymatgen library is structured in modular Python packages. The main packages are as follows:

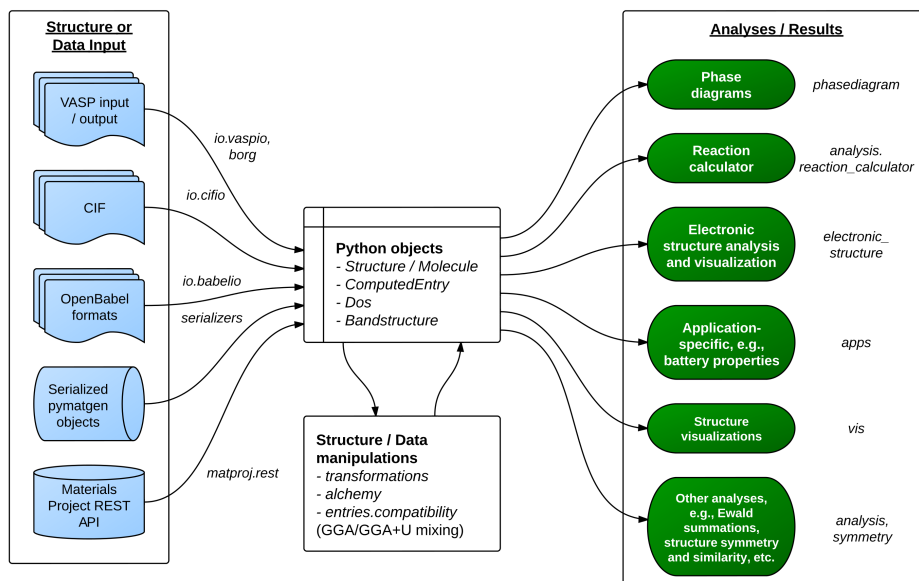


Figure 1: Overview of the pymatgen library. Text in italics represent names of Python packages, modules or classes.

1. The *core* package, as its name implies, provides the core definitions of various objects used by the rest of the library. Core objects include representations of elements in the periodic table (Element class in the *core.periodic_table* module), infinite periodic lattices (Lattice in the *core.lattice* module), periodic and non-periodic sites (Site and PeriodicSite classes in the *core.structure* module), molecules and structures (Molecule and Structure classes in the *core.structure* module) and compositions (Composition class in the *core.structure* module). The core objects encapsulates information relevant to many materials applications. For example, the Element class includes useful properties such as electronegativity, atomic numbers and atomic masses, etc.
2. The *electronic_structure* package defines objects representing various electronic structure analyses, including density of states (*electronic_structure.dos* module) and bandstructures (*electronic_structure.bandstructure* module). Plotting capabilities for these analyses are also provided using the matplotlib library (see Figure 2).
3. The *entries* package defines the basic ComputedEntry object (in the *computed_entries* module) for performing analyses. The ComputedEntry object is essentially a container for materials information. At the most basic level, a ComputedEntry comprise a composition and an energy, which are necessary for phase diagram generation (using the *phasediagram* package) and calculating reaction energies (using the *analysis.reaction_calculator* package). However, a ComputedEntry is designed to be flexible enough to encompass any data of interest for a material, such as its structure, spacegroup, etc. The ComputedEntry object is also designed to be agnostic to the source of the information, e.g., the energy can be obtained from VASP,[1] ABINIT[3, 4] or any other electronic structure calculation. A similar ExpEntry object (in the *exp_entries* module) is also available as a container for experimental thermochemical data to be used in analyses.
4. The *io* (input/output) package provides facilities to read and write common structure and molecule file formats as well as input and output files for various electronic structure codes. Support for the commonly used Crystallographic Information File (CIF) format is provided using the PyCifRW library,[19] and support for a large number of molecular file formats is supported via an adaptor to the OpenBabel library.[20] Among the *io* modules for electronic structure codes, the *vaspio* module is currently the most mature and supports most Vienna Ab initio

Simulation Package (VASP)[1] input and output files. VASP input parameters based on those used in the Materials Project as well as the originating MIT high-throughput project[5] are provided in the *vaspio_set* module. Limited support is currently available for Gaussian[2] input files as well, though we expect this to improve considerably in future. In addition, pymatgen also provides an adaptor (the *aseio* module) to provide conversion between pymatgen’s Structure object and the Atoms object used by the Atomic Simulation Environment (ASE).[21] A trivial use of the *io* package is for the conversion between various file formats (e.g, converting CIF files to VASP POSCAR files). A more powerful use is converting flat files into Python objects (such as Structure or ComputedEntry), which can then be used for further structure manipulation or analysis. The pymatgen library is highly extensible in terms of electronic structure code support, and parsers for ABINIT and other first principles codes are currently under development.

5. The *serializers* package implements customized modules for the serialization of pymatgen objects. Serialization allows users to save pymatgen objects easily for subsequent reuse. In pymatgen, most non-trivial objects implement a *to_dict* property, which is a Python dictionary representation that can be serialized in the lightweight JavaScript Object Notation (JSON) format, and a *from_dict* static method that regenerates that object from a JSON representation. The JSON representation can be easily stored on a user’s hard disk or inserted into a database such as the MongoDB used by the Materials Project.

In addition to the above packages, several packages have been implemented to aid structure manipulation and transformation and to perform thermodynamic analyses. These packages are outlined in the following sections.

3. Compound generation and structure transformations

Pymatgen provides a powerful framework for performing compound generation and structure transformations via the *transformations* package. A transformation is essentially a well-defined algorithm for generating new compounds and structures from existing structures. For example, a common approach to developing new materials from existing materials involve the substitution of existing species in the structure for others. Users can, for instance, use the data-mined substituted rules developed by Hautier et al.[22] to ob-




Figure 2: Bandstructure of Fe₂O₃, plotted using data from the Materials Project and pymatgen's *electronic_structure* package. Up spins are in blue while down spins are in red.

tain new materials. Such a manipulation can be performed using the `SubstitutionTransformation` class in the `transformations.standard_transformations` module. Other supported transformations include the partial or complete removal of a species in a structure, ordering of disordered structures, and generation of supercells and primitive cells. In addition, `pymatgen` also provides the facility to perform high-throughput compound generation and electronic structure run generation via the `alchemy` package. Using the `alchemy` package, a developer can define a sequence of transformations to be applied to a set of structures to generate a corresponding set of new structures. The set of structures can be conveniently provided as a directory of CIF files, VASP POSCAR files, etc. These structures can then be output to the necessary input formats for electronic structure calculations. Furthermore, the `alchemy` package provides a means to store the history of all transformations applied on a structure, allowing one to trace back the origins of a new structure. The `alchemy` package is currently used in the CrystalToolkit of the Materials Project (http://www.materialsproject.org/apps/crystal_toolkit/) to perform structure manipulations with unlimited undo and redo capabilities (see Figure 3).

4. Analysis tools

The `pymatgen` library provides many tools for high-throughput, automated assimilation of data from electronic structure calculations, and for subsequent analysis of the assimilated data.

4.1. Data assimilation and processing

The `borg` package can automatically traverse a directory tree to search for calculations and assimilate calculation data, utilizing multiple processors where available using Python's multiprocessing package. A predefined algorithm for converting VASP runs into a list of `ComputedEntry` objects has been implemented. `ComputedEntry` objects, which are essentially containers for calculated data, serve as the basic unit for subsequent analysis. Sometimes, some post-processing of the list of `ComputedEntry` objects is necessary before they can be reliably used in analyses. In the `pymatgen` library, the `entries.compatibility` module implements the scheme for mixing energies calculated using different functionals, in particular, those calculated using the generalized gradient approximation (GGA) and the +U extension to it (GGA+U)[23, 24, 25] as outlined by Jain et al.[26] While standard GGA

Home Apps Support About References Profile for shyyap@gmail.com - Logout

CrystalToolkit

CrystalToolkit is a structure editor that generates new structures from existing structures in the Materials Project or from an uploaded POSCAR or CIF file. Batch processing of structures is supported.

You may also use the CrystalToolkit from the Materials Explorer by performing a search, selecting materials using the checkboxes and clicking the Edit button.

Tool Selection Select one or more tools. Transformations are applied in the order shown.

- NonFile Conversion
- Generate supercell
- Substitutions
- Species Removal
- Fractional Species Removal
- Order Theoretical Structure

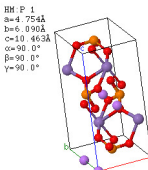
Substitutions Enter substitutions as colon-separated species pairs, with pairs separated by commas. e.g. Li:Na, O:S substitutes Na for Li and S for O. Species names must correspond to the species names in the structure table below. (i.e., if Li is listed, you must enter Li and not just Li).

Output format VASP files include POSCAR, INCAR, and KPOINTS files used in the Materials Project's structural relaxation. (VASP Files)

Apply changes Undo last change Preview output Download output Start Over

Li₂Mn₂P₂O₁₀ P6mm (62)

HH P 1
a= 4.754 Å
b= 0.904 Å
c=10.463 Å
α= 90.0°
β= 90.0°
γ= 90.0°



Space Filling: off 20% 100% polyhedra
Spin: off around x axis around y axis
Background color:

Lattice Vector Coordinates			
1	4.754	0	0
2	0	0.904	0
3	0	0	10.463

Atom Fractional coordinates			
Li	0	0	0
Li	0.5	0	0.5
Li	0	0.5	0
Li	0.5	0.5	0.5
O	0.789	0.547	0.334
O	0.211	0.453	0.666
O	0.711	0.547	0.834
O	0.789	0.953	0.334
O	0.289	0.047	0.166
O	0.744	0.25	0.096
O	0.289	0.453	0.166
O	0.211	0.047	0.866
O	0.289	0.75	0.904
O	0.756	0.25	0.966

CrystalToolkit version 1.1, Database version 0.5.3. Powered by pymatgen version 1.9.0.

questions? contact feedback [at] materialsproject.org
copyright © 2010-2011 Materials Project, a collaboration between MIT and LBNL. Terms of Use

(a) CrystalToolkit

Home Apps Support About References Login or Register

PhaseDiagramApp

The Phase Diagram App uses density functional theory calculations from our database to generate 3K compositional and grand potential phase diagrams of two to four components. For details of our methodology and usage tips, please refer to the Phase Diagram App manual.

HTK-compatible browser recommended. (e.g., Chrome, Safari, Firefox (28.0 and above) and IE8. Non-interactive plots only for IE8.)

Usage Examples >

Elements (e.g., Fe P): Open element (const. μ):

Generate Phase Diagram

Fe-Li-O-P Phase Diagram open to O

Show Unstable Label points Mark Comp. Show Data Table Interpreting PDs
μ_O = -6.54 eV - 6.42 eV

15 Stable Compounds

Name	Form. Enthalpy [eV/atom]	HT
FePO ₃	-0.595	0.000
Fe ₂ O ₃	0.000	0.000
Fe ₃ P ₂ O ₇	-0.744	0.000
Fe ₂ P ₂ O ₇	-0.737	-0.000
Li ₂ O	0.000	0.000
Li ₂ P ₂ O ₇	-1.054	0.000
Li ₃ P ₂ O ₇	-1.038	0.000
Li ₂ FeO ₄	-0.149	0.000
Li ₂ Fe ₂ O ₇	-0.170	0.000
LiFe ₂ O ₄	-0.102	0.000
LiFe ₂ P ₂ O ₇	-0.278	0.000
LiFe ₂ P ₂ O ₈	-0.893	0.000

24 Unstable Compounds

Name	Form. Enthalpy [eV/atom]	HT
Fe	1.211	1.211
Fe ₂ O ₃	0.101	0.101
Fe ₃ P	1.946	1.961
Fe ₂ P ₂ O ₇	-0.396	0.048
Fe ₂ P	0.010	0.010
Fe ₃ P	1.297	1.757
FeP	1.500	2.244
Fe ₂ P	1.866	2.503
Fe ₂ P ₂	2.850	2.798
Fe ₂ P ₂ O ₇	-0.693	0.051
Li	1.984	1.984
LiFe ₂ P ₂ O ₇	-0.160	0.351

PDApp version 0.3. Database version 0.5.3. Powered by pymatgen version 1.9.0.
To cite the PDApp, please refer to these works.

questions? contact feedback [at] materialsproject.org
copyright © 2010-2011 Materials Project, a collaboration between MIT and LBNL. Terms of Use

(b) PhaseDiagramApp

Figure 3: The CrystalToolkit and PhaseDiagramApp in the Materials Project, utilizing pymatgen's *alchemy* and *phasediagram* packages respectively.

is reasonably accurate for calculating energy differences between delocalized states, it generally fails when the degree of electronic localization varies greatly between the products and reactants, such as in a redox reaction.[27] For the latter, the addition of a Hubbard U parameter generally improves the accuracy of calculated reaction energies considerably. The “mixing” scheme adjusts the GGA+U energies using known experimental binary formation enthalpies in a way that makes them compatible with GGA energies. In addition, it also adjusts the energy of well-known gaseous elements such as O₂, N₂, etc. to correct for well-known tendency of GGA to overbind such molecules.[28] Jain et al. demonstrated that this “mixing” scheme provides reasonably accurate results for formation enthalpies and phase diagrams.[26] With some modifications, this module could be used to combine energies obtained with any set of different functionals. It should be noted that the set of pseudopotentials and Hubbard U parameters used by the Materials Project are different from those originally used by Jain et al.; the pseudopotentials used by the Materials Project generally include more electrons in the valence shell and the Hubbard U parameters have been fitted using the approach of Wang et al.[28] for this set of pseudopotentials. Thus, the necessary “mixing” scheme corrections have been refitted for the Materials Project parameter set. Two Compatibility classes, `MaterialsProjectCompatibility` and `MITCompatibility`, are provided, and it is recommended that users use the appropriate class to process their runs prior to other analyses. The Materials Project parameters and corrections are provided in the Supplementary Information.

4.2. Calculating reactions

The *analysis.reaction_calculator* module provides classes for the analysis of reactions, including reaction balancing and calculation of reaction energies. A user can calculate reactions energies from computed data (using `ComputedEntry` objects) or experimental data (using `ExpEntry` objects). These features are currently used in the `ReactionCalculator` of the Materials Project to provide calculated reaction energies and comparison of those energies with experimental reaction energies where available.

4.3. Phase diagrams

The *phasediagram* package provides facilities to generate and plot phase diagrams. The methodology and algorithms are based on those developed by Ong et al.[29, 30] Both “standard” compositional and grand canonical phase

diagrams (representing phase equilibria in systems open to one or more components) are supported. Phase diagrams representing the thermodynamic phase equilibria of multicomponent systems reveal fundamental material aspects regarding the processing and reactions of materials. For example, two key considerations in designing a new material are its stability and potential synthesis routes. By comparing a new material’s energy relative to competing phases in the phase diagram, a user can assess the new material’s phase stability and the predicted phase equilibria at a particular composition. The *phasediagram* package is currently used in the Phase Diagram App of the Materials Project (see Figure 3(b)) to generate phase diagrams from calculated materials data. Currently, only 0K phase diagrams are available in the Materials Project as only energies are available at this point.

5. Integration with the Materials Project REST API

One of the key impediments to materials design is the availability of materials information. The Materials Project aims to meet this need by providing open, public access to a large database of calculated data on known materials. Currently, there are several user-friendly “apps” available on the Materials Project that use this data. In order to reach a broader materials community, we have created an application programming interface (API) based on a subset of the principles of REpresentational State Transfer (REST).[31] Like the APIs of many well-known sites, the Materials Project API uses some shared knowledge about the form and semantics of URIs, as well as pre-determined media types, to reduce the number of round-trips needed to discover and use the interface. By convention, this deviation from the full set of requirements of a REST API are indicated by the term “RESTful”. The Materials Project RESTful API allows users to directly access Materials Project data via the Hypertext Transfer Protocol (HTTP), and provides a powerful way for users to programmatically query for materials information instead of relying on browser-based interfaces. We hope that this capability to query the database directly spurs the creation of applications for other materials properties and applications. Under RESTful design, each object is represented as a unique resource and can be queried in a uniform manner. A RESTful HTTP service exposes a consistent set of semantics that uses HTTP methods (GET, POST, PUT, DELETE, etc.) in conjunction with unique Uniform Resource Identifiers (URIs) to access the underlying resources. This allows for the creation of an API using a combination of HTTP

methods and URIs. For the purposes of the Materials Project, this means that each object (such as a material) can be represented by a unique URI (e.g., [http://www.materialsproject.org/rest/\[unique-id\]](http://www.materialsproject.org/rest/[unique-id])) and an HTTP verb can be used to act on that object. In most cases, this action returns structured data that represents the object or the result of an operation against the object. The JavaScript Object Notation (JSON) is used as the data interchange format to represent this structured data. In conjunction with this beta release, the pymatgen library has released the *matproj* package to provide a convenient set of tools for users to obtain pymatgen objects and data via the Materials Project RESTful API. For instance, users can obtain formation energies, calculated VASP energies and calculation parameters for all database entries with given formula. Users can also query the Materials Project’s database of experimental thermochemistry data. In addition, users can obtain pymatgen Structure and ComputedEntry objects, which can serve as the starting point for further structure manipulation or thermodynamic analysis.

6. Application example - Phase stability of a new material

To illustrate the power of the pymatgen library, we will present a practical example of how it can be used to determine the phase stability of a new material. One of the main obstacles to performing phase stability analyses on new materials is that the phase stability of a particular material depends on its energy relative to that of competing phases. Without relying on an external database of pre-computed materials data, such an effort requires the materials researcher to identify all competing phases of interest and invest a fairly large amount of computational resources to calculating their energies. For example, to determine the phase stability of a new $A_xB_yC_z$ phase, the materials researcher needs to calculate the energies of all known A, B, C, A_xB_y , A_xC_y , B_xC_y and $A_xB_yC_z$ phases. Using pymatgen’s analysis tools and interface with the Materials Project, this can be accomplished with a minimal of computational resources. We will use the example of Li_4SnS_4 , a material of interest as a lithium superionic conductor, which has been just recently synthesized by Kaib et al.[32] Li_4SnS_4 is isostructural with Li_4GeS_4 , a well-known compound for which data is available in the Materials Project. The structure of Li_4SnS_4 itself is currently not reported in the Inorganic Crystal Database (ICSD)[33] or the Materials Project. Using the Materials Project’s CrystalToolkit, we perform a Sn for Ge substitution on Li_4GeS_4 ,

download the necessary input files based on the parameters used in the Materials Project and perform first principles calculations using VASP to obtain the ground state structure and energy for Li_4SnS_4 . We then combine the data from the Li_4SnS_4 VASP calculation (assimilated using pymatgen’s *borg* package) with the data available on other Li-Sn-S phases in the Materials Project (using the *matproj.rest* module) to generate the Li-Sn-S phase diagram shown in Figure 4 (see Supplementary Information for the code snippet used to perform this analysis). We may observe that Li_4SnS_4 is indeed predicted to be stable in the Li-Sn-S phase diagram. A slightly more advanced




Figure 4: The Li-Sn-S phase diagram generated using a single Li_4SnS_4 calculation and data from the Materials Project using pymatgen’s *borg*, *matproj* and *phasediagram* packages.

example is given in Figure 5, where we have plotted the Li-Sn-S phase diagram in the Li and S chemical potential space. From Figure 5, we may observe that Li_4SnS_4 is stable only over a narrow range of μ_{Li} , indicating

that when used as an electrolyte in a lithium-ion battery, this material will likely react via lithium absorption at the anode and lithium desorption at the cathode to form a solid electrolyte interphase. In this example, we have




Figure 5: The chemical potential stability map for the Li-Sn-S system generated using a single Li_4SnS_4 calculation and data from the Materials Project using pymatgen’s *borg*, *matproj* and *phasediagram* packages. μ_{Li}^0 and μ_S^0 refers to the calculated bulk energies of lithium and sulfur respectively. This map shows the ranges of μ_{Li} and μ_S for which each phase is stable. Note that chemical potential values to the right of the Li_2S line (shaded in grey) are forbidden ranges.

determined the phase stability of the newly-synthesized Li_4SnS_4 compound by performing only a single first principles ground state calculation and using pymatgen’s analysis tools and interface to the Materials Project’s REST API. This determination would otherwise have required the user to obtain all relevant crystal structure, generate the necessary input files, and perform first principles calculations on more than 30 structures (based on the number of Li-Sn-S phases in the Materials Project), incurring significantly greater computational expense. There is also significant time savings for the user, in that he/she needs only to obtain initial candidate structures and perform calculations and analysis for the particular phase he/she is interested

in (Li_4SnS_4), while querying the Materials Project for pre-computed data for the other phases.

7. Conclusion

The Python Materials Genomics (pymatgen) library is a robust, open-source python library for materials data analysis. By defining core Python objects for materials data representation, providing a well-tested set of structure and thermodynamic analysis tools relevant to many materials applications, and establishing an open platform for researchers to collaboratively develop sophisticated analyses of materials data, the pymatgen library is a key enabler of the Materials Project, powering several of the Project’s web applications. The pymatgen library also provides convenient tools to obtain useful materials data via the Materials Project’s REST API. Using pymatgen’s interface to the Materials Project’s REST API and *phasediagram* package, we demonstrate how the phase and electrochemical stability of a recently synthesized material, Li_4SnS_4 , can be analyzed using a minimum of computing resources. We hope that the open source dissemination of this library and the free access to the Materials Project data will spur the creation of more property apps that use computed data, and ultimately to a broader impact of computational modeling on the materials community.

8. Acknowledgments

This work was supported in part by the Department of Energy’s Basic Energy Sciences program under Grant No. DE-FG02-96ER45571. We also thank the National Energy Research Scientific Computing Center (NERSC) for providing invaluable computing resources and IT support for this project. A. Jain acknowledges funding from the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under US Department of Energy Contract No. DE-AC02-05CH1123.

Appendix A. Supplementary information

Appendix A.1. VASP calculation and compatibility parameters used in the Materials Project.

The figures and tables below summarizes the current VASP calculation and compatibility parameters used in the Materials Project, as of the time of


```

from pymatgen.apps.borg.hive import VaspToComputedEntryDrone
from pymatgen.apps.borg.queen import BorgQueen
from pymatgen.phasediagram.pdmaker import PhaseDiagram
from pymatgen.phasediagram.plotter import PDPlotter
# Assimilate user VASP calculation (assumed to be located in
# current directory) into ComputedEntry object
drone = VaspToComputedEntryDrone()
queen = BorgQueen(drone, rootpath=".")
entries = queen.get_data()
# entries now contain a single ComputedEntry for Li4SnS4
# Obtain all existing Li-Sn-S phases using the
# Materials Project RESTful API
adaptor = MPRester("USER_API_KEY")
mp_entries = adaptor.get_entries_in_chemsys(["Li", "Sn", "S"])
# Combine entry from calculated run with Materials Project entries
entries.extend(mp_entries)
# Filter entries using the MaterialsProjectCompatibility
compat = MaterialsProjectCompatibility()
entries = compat.process_entries(entries)
# Generate and plot Li-Sn-S phase diagram
pd = PhaseDiagram(entries)
plotter = PDPlotter(pd)
plotter.show() # Produces Figure 4 of this paper

```

- [1] G. Kresse and J. Furthmuller, *Physical Review B* **54**, 11169 (1996).
- [2] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, J. A. Montgomery Jr., T. Vreven, K. N. Kudin, J. C. Burant, J. M. Millam, S. S. Iyengar, J. Tomasi, V. Barone, B. Mennucci, M. Cossi, G. Scalmani, N. Rega, G. A. Petersson, H. Nakatsuji, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, M. Klene, X. Li, J. E. Knox, H. P. Hratchian, J. B. Cross, V. Bakken, C. Adamo, J. Jaramillo, R. Gomperts, R. E. Stratmann, O. Yazyev, A. J. Austin, R. Cammi, C. Pomelli, J. W. Ochterski, P. Y. Ayala, K. Morokuma, G. A. Voth, P. Salvador, J. J. Dannenberg, V. G. Zakrzewski, S. Dapprich, A. D. Daniels, M. C. Strain, O. Farkas, D. K. Malick, A. D. Rabuck, K. Raghavachari, J. B. Foresman, J. V. Ortiz, Q. Cui, A. G. Baboul, S. Clifford,

- J. Cioslowski, B. B. Stefanov, G. Liu, A. Liashenko, P. Piskorz, I. Komaromi, R. L. Martin, D. J. Fox, T. Keith, M. A. Al-Laham, C. Y. Peng, A. Nanayakkara, M. Challacombe, P. M. W. Gill, B. Johnson, W. Chen, M. W. Wong, C. Gonzalez, and J. A. Pople, "Gaussian 03, Revision C.02," .
- [3] X. Gonze, B. Amadon, P.-M. Anglade, J.-M. Beuken, F. Bottin, P. Boulanger, F. Bruneval, D. Caliste, R. Caracas, M. Côté, T. Deutsch, L. Genovese, P. Ghosez, M. Giantomassi, S. Goedecker, D. R. Hamann, P. Hermet, F. Jollet, G. Jomard, S. Leroux, M. Mancini, S. Mazevet, M. J. T. Oliveira, G. Onida, Y. Pouillon, T. Rangel, G.-M. Rignanese, D. Sangalli, R. Shaltaf, M. Torrent, M. J. Verstraete, G. Zerah, and J. W. Zwanziger, *Computer Physics Communications* **180**, 2582 (2009).
- [4] X. Gonze, G. M. Rignanese, M. J. Verstraete, J. M. Beuken, Y. Pouillon, R. Caracas, F. Jollet, M. Torrent, G. Zerah, M. Mikami, P. Ghosez, M. Veithen, J. Y. Raty, V. Olevano, F. Bruneval, L. Reining, R. W. Godby, G. Onida, D. R. Hamann, and D. C. Allan, *Zeitschrift für Kristallographie* **220**, 558 (2005).
- [5] A. Jain, G. Hautier, C. J. Moore, S. P. Ong, C. C. Fischer, T. Mueller, K. A. Persson, and G. Ceder, *Computational Materials Science* **50**, 2295 (2011).
- [6] G. Hautier, A. Jain, S. P. Ong, B. Kang, C. Moore, R. Doe, and G. Ceder, *Chemistry of Materials* **23**, 3495 (2011).
- [7] G. Hautier, A. Jain, H. Chen, C. Moore, S. P. Ong, and G. Ceder, *Journal of Materials Chemistry* **21**, 17147 (2011).
- [8] G. Ceder, G. Hautier, A. Jain, and S. P. Ong, *MRS Bulletin* **36**, 185 (2011).
- [9] S. P. Ong, V. L. Chevrier, G. Hautier, A. Jain, C. Moore, S. Kim, X. Ma, and G. Ceder, *Energy & Environmental Science* **4**, 3680 (2011).
- [10] Y. Mo, S. P. Ong, and G. Ceder, *Chemistry of Materials* **24**, 15 (2012).
- [11] S. Curtarolo, W. Setyawan, G. L. Hart, M. Jahnatek, R. V. Chepulskii, R. H. Taylor, S. Wang, J. Xue, K. Yang, O. Levy, M. J. Mehl, H. T.

- Stokes, D. O. Demchenko, and D. Morgan, *Computational Materials Science* **58**, 218 (2012).
- [12] S. Curtarolo, D. Morgan, and G. Ceder, *Computer Coupling of Phase Diagrams and Thermochemistry* **29**, 163 (2005).
- [13] A. Kolmogorov, M. Calandra, and S. Curtarolo, *Physical Review B* **78**, 1 (2008).
- [14] K. Yang, W. Setyawan, S. Wang, M. Buongiorno Nardelli, and S. Curtarolo, *Nature Materials* **11**, 614 (2012).
- [15] J. S. Hummelshøj, F. Abild-Pedersen, F. Studt, T. Bligaard, and J. K. Nørskov, *Angewandte Chemie (International ed. in English)* **51**, 272 (2012).
- [16] J. Hachmann, R. Olivares-Amaya, S. Atahan-Evrenk, C. Amador-Bedolla, R. S. Sanchez-Carrera, A. Gold-Parker, L. Vogt, A. M. Brockway, and A. Aspuru-Guzik, *The Journal of Physical Chemistry Letters* **2**, 2241 (2011).
- [17] A. N. Sokolov, S. Atahan-Evrenk, R. Mondal, H. B. Akkerman, R. S. Sánchez-Carrera, S. Granados-Focil, J. Schrier, S. C. B. Mannsfeld, A. P. Zoombelt, Z. Bao, and A. Aspuru-Guzik, *Nature communications* **2**, 437 (2011).
- [18] T. E. Oliphant, *Computing in Science & Engineering* **9**, 10 (2007).
- [19] J. R. Hester, *Journal of Applied Crystallography* **39**, 621 (2006).
- [20] N. M. O'Boyle, M. Banck, C. a. James, C. Morley, T. Vandermeersch, and G. R. Hutchison, *Journal of cheminformatics* **3**, 33 (2011).
- [21] S. R. Bahn and K. W. Jacobsen, *Comput. Sci. Eng.* **4**, 56 (2002).
- [22] G. Hautier, C. Fischer, V. Ehrlacher, A. Jain, and G. Ceder, *Inorganic chemistry* , 656 (2010).
- [23] V. I. Anisimov, J. Zaanen, and O. K. Andersen, *Physical Review B* **44**, 943 (1991).

- [24] V. I. Anisimov, F. Aryasetiawan, and A. I. Lichtenstein, *J. Phys.: Condens. Matter* **9**, 767 (1997).
- [25] A. I. Liechtenstein, V. I. Anisimov, and J. Zaanen, *Physical Review B* **52**, R5467 (1995).
- [26] A. Jain, G. Hautier, S. Ong, C. Moore, C. Fischer, K. Persson, and G. Ceder, *Physical Review B* **84**, 045115 (2011).
- [27] F. Zhou, M. Cococcioni, C. A. Marianetti, D. Morgan, and G. Ceder, *Physical Review B* **70**, 235121 (2004).
- [28] L. Wang, T. Maxisch, and G. Ceder, *Physical Review B* **73**, 195107 (2006).
- [29] S. P. Ong, L. Wang, B. Kang, and G. Ceder, *Chemistry of Materials* **20**, 1798 (2008).
- [30] S. P. Ong, A. Jain, G. Hautier, B. Kang, and G. Ceder, *Electrochemistry Communications* **12**, 427 (2010).
- [31] R. T. Fielding and R. N. Taylor, *ACM Transactions on Internet Technology* **2**, 115 (2002).
- [32] T. Kaib, S. Haddadpour, M. Kapitein, P. Bron, C. S. Der, H. Eckert, B. Roling, and S. Dehnen, *Chem. Mater.* **8**, 2 (2012).
- [33] G. Bergerhoff, R. Hundt, R. Sievers, and I. D. Brown, *Journal of chemical information and computer sciences* **23**, 66 (1983).