

## RESEARCH ARTICLE

# Sim2Ls: FAIR simulation workflows and data

Martin Hunt<sup>1</sup>, Steven Clark<sup>2</sup>, Daniel Mejia<sup>3</sup>, Saaketh Desai<sup>4</sup>, Alejandro Strachan<sup>3,5\*</sup>

**1** Halicioğlu Data Science Institute, University of California at San Diego, La Jolla, California, United States of America, **2** San Diego Supercomputer Center, University of California at San Diego, La Jolla, California, United States of America, **3** Network for Computational Nanotechnology, Purdue University, West Lafayette, Indiana, United States of America, **4** Center for Integrated Nanotechnologies, Sandia National Laboratories, Albuquerque, New Mexico, United States of America, **5** School of Materials Engineering, Purdue University, West Lafayette, Indiana, United States of America

\* [strachan@purdue.edu](mailto:strachan@purdue.edu)

## Abstract

Just like the scientific data they generate, simulation workflows for research should be findable, accessible, interoperable, and reusable (FAIR). However, while significant progress has been made towards FAIR data, the majority of science and engineering workflows used in research remain poorly documented and often unavailable, involving *ad hoc* scripts and manual steps, hindering reproducibility and stifling progress. We introduce *Sim2Ls* (pronounced simtools) and the *Sim2L* Python library that allow developers to create and share end-to-end computational workflows with well-defined and verified inputs and outputs. The *Sim2L* library makes *Sim2Ls*, their requirements, and their services discoverable, verifies inputs and outputs, and automatically stores results in a globally-accessible simulation cache and results database. This simulation ecosystem is available in nanoHUB, an open platform that also provides publication services for *Sim2Ls*, a computational environment for developers and users, and the hardware to execute runs and store results at no cost. We exemplify the use of *Sim2Ls* using two applications and discuss best practices towards FAIR simulation workflows and associated data.

## OPEN ACCESS

**Citation:** Hunt M, Clark S, Mejia D, Desai S, Strachan A (2022) *Sim2Ls*: FAIR simulation workflows and data. PLoS ONE 17(3): e0264492. <https://doi.org/10.1371/journal.pone.0264492>

**Editor:** Parag A. Deshpande, Indian Institute of Technology Kharagpur, INDIA

**Received:** December 13, 2021

**Accepted:** February 10, 2022

**Published:** March 10, 2022

**Copyright:** This is an open access article, free of all copyright, and may be freely reproduced, distributed, transmitted, modified, built upon, or otherwise used by anyone for any lawful purpose. The work is made available under the [Creative Commons CC0](https://creativecommons.org/licenses/by/4.0/) public domain dedication.

**Data Availability Statement:** All data generated from the use of *Sim2Ls* is automatically cached by nanoHUB and indexed in the ResultsDB that can be queried by all nanoHUB users at <https://nanohub.org/developer/api/endpoint/dbexplorer>. nanoHUB accounts are free and can be opened at: <https://nanohub.org/register/>. The *Sim2L* library is available for online simulation in the open platform nanoHUB <https://nanohub.org>, and for download at <https://github.com/hubzero/simtool>. Documentation is available at <https://simtool.readthedocs.io/en/stable/>.

## 1. Introduction

Scientific progress is based on the ability of researchers to independently reproduce published results and verify inferences [1, 2]. These results are nearly universally obtained via complex, multi-step, workflows involving experiments and/or simulations with multiple inputs, data collection, and analysis. It is often the case that, even when the authors carefully document their procedures, reproducing published results requires a significant investment of time even for experts. This is true both in experimental and computational work, it slows down progress and results in wasted resources. A related issue hindering innovation is the fact that the majority of the data generated during research is not made available to the community and the fraction that is used in publications, generally skewed, is often not findable or queryable. This is particularly problematic with the increasingly important role machine learning is playing in physical sciences and engineering [3, 4]. Guidelines to making data findable, accessible, interoperable, and reusable (FAIR) have been put forward [5] and a variety of concrete efforts to

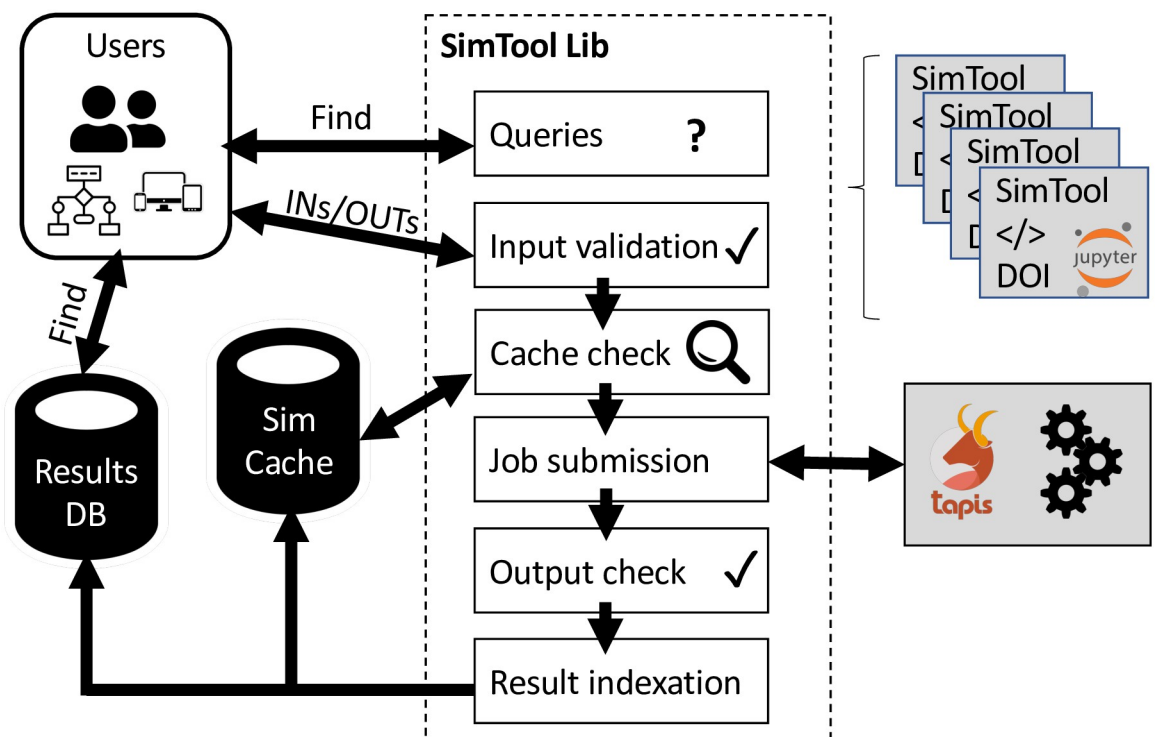
**Funding:** This work was partially supported by the Network for Computational Nanotechnology, a project of the US National Science Foundation, through a grant awarded to MH, SC, DM, SD, and AS (EEC-1227110). This work was also partially supported by Sandia National Laboratories, a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration through a contract awarded to SD (DE-NA0003525). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

**Competing interests:** The authors have declared that no competing interests exist.

tackle these issues have been launched in recent years. Examples in the physical sciences range from open and queryable repositories of materials properties, both computational and experimental [6–10], to publications devoted to scientific data [11] as well as infrastructure to publish and share models [12–14].

FAIR principles apply not just to scientific data but also to research workflows used to generate them, this is particularly true for computational workflows where documentation, automatization, and reproducibility are easier than in experiments [15]. Growing interest in making workflows available are reflected by the increasing popularity of Git repositories [16] and Jupyter [17]. Notable examples of reproducible workflows include *ab initio* calculations performed in the Materials Project [18], openKIM property calculations [19], osteoarthritis image processing [20]. In addition, several publishers require either data availability statements or all data and codes to be made available [21]; some have also developed lists of suggested repositories, see, for example Ref. [22]. Despite these laudable efforts, the majority of research workflows used in published research are described in incomplete terms and using technical English as opposed to using specialized tools. Furthermore they often involve ad hoc analysis scripts and manual steps that conspire against automation and reproducibility. This is in part due to the lack of general tools for the development and publication of computational tools with well defined, verifiable, and discoverable inputs and outputs and the automatic storage of results.

To address these gaps we introduce *Sim2Ls*, a library to create and share end-to-end computational workflows with verified inputs and outputs, see Fig 1 for a schematic representation of the ecosystem. These workflows have verified inputs and outputs, could launch large-scale simulations in high-performance computing resources, employ a simulation cache to re-use previous runs, and index results in a database to enable querying. The *Sim2L* library is



**Fig 1. Simtool ecosystem diagram.**

<https://doi.org/10.1371/journal.pone.0264492.g001>

available via the US National Science Foundation's nanoHUB [13] which also provides services for workflow publication, free and open online simulations of published *Sim2Ls*, automatic caching of simulation runs, and indexing of the outputs in a queryable database. This ecosystem makes *Sim2Ls*, their services, requirements, and the results they produce FAIR. The *Sim2L* library is available at <https://github.com/hubzero/simtool> and its documentation at <https://simtool.readthedocs.io/en/stable/>. At the time of writing there are currently 8 published *Sim2Ls* in nanoHUB some of which have served nearly 100 users who performed thousands of simulations, see for example Ref. [23]. The remainder of this paper discusses the elements of a *Sim2L* and the *Sim2L* library (Section II), provides examples of their use (Section III), followed by general *Sim2L* design guidelines for developers (Section IV) and conclusions.

## II. *Sim2Ls* and the *Sim2L* library

### A. Elements of a *Sim2L*

*Sim2Ls* are developed and stored in Jupyter notebooks. As depicted in Fig 2, the main components of a *Sim2L* are: i) declaration of input and output variables using YAML [24], ii) notebook parameterization cells that use PaperMill [24], iii) the computational workflow connecting inputs to outputs, including all pre and post-processing and computations (this step can involve accessing external data resources and launching parallel simulation to external high-performance computing systems), and iv) population of all the output fields. Each element of a *Sim2L* is described in detail in the following paragraphs and subsection IIB describes the *Sim2L* library through which users interact with *Sim2Ls*.

**Description.** The *Sim2L* notebook should contain a cell tagged as DESCRIPTION. The plain text content of the cell should provide an overview of the *Sim2L* requirements (inputs)

#### Inputs and outputs declaration using YAML

```
In [ ]:
%yaml INPUTS

material:
  type: Text
  value: Ni
  maxlen: 2
  description: Element to be simulated
  options: ['Sc', 'Ti', 'V', 'Cr', 'Mn', 'Fe', 'Co', 'Ni', 'Cu', 'Zn',
           'Y', 'Zr', 'Nb', 'Mo', 'Tc', 'Ru', 'Rh', 'Pd', 'Ag', 'Cd',
           'Hf', 'Ta', 'W', 'Re', 'Os', 'Ir', 'Pt', 'Au', 'Hg']

crystal_structure:
  type: Text
  value: fcc
```

```
In [ ]:
%yaml OUTPUTS

final_snapshot:
  type: Image
  description: Snapshot of the final structure (check for coexistence between ...)

melting_temperature:
  type: Number
  description: Melting temperature predicted by the simulation
  units: K

melting_temperature_ci:
  type: Number
  description: 95% confidence interval in melting temperature prediction (from
```

#### Parameterization

```
In [ ]:
parameters x

# PARAMETERS (for testing)
# These get replaced when SimTool is run
material = 'Ni'
crystal_structure = 'fcc'
lattice_parameter = 3.5203
mass = 58.6934
Tsolid = 800
```

#### Simulation(s) setup & execution

```
In [ ]:
write_string = '#!/bin/sh
lammpsInput=$1
kinMode=$2
logname=$3

./etc/environ.sh

if [ -n "${ANACONDA_CHOICE}" ]; then
umuse -e ${ANACONDA_CHOICE}
fi

use -e -r openkim-2.1.3
use -e -r lammps-07Aug19
```

#### Postprocessing

```
In [ ]:
logfile = log(logname)
time, temperature, potential_energy, volume = logfile.get("Step", "Temp", "PotEng", "Vo
time = np.array(time)
temperature = np.array(temperature)

#establish trend to print converged or not converged
idx = np.where(time >= 0.0*(inputs['time'] + 10000))
y = temperature[idx[0]]
x = time[idx[0]]/1000

def f(x, A, B):
    return Ax + B

popt, pcov = curve_fit(f, x, y)
slope = popt[0]
intercept = popt[1]
```

#### Output save

```
In [ ]:
db = DB(OUTPUTS)

In [ ]:
db.save('final_snapshot', file='final.jpg')

In [ ]:
db.save('melting_temperature', melting_temperature)
db.save('melting_temperature_ci', melting_temperature_ci)
```

Fig 2. Elements of a *Sim2L*.

<https://doi.org/10.1371/journal.pone.0264492.g002>

and services (outputs) provided, this information is reported when using the *Sim2L* library to query for available *Sim2Ls*.

**Inputs.** One of the fundamental aspects of a *Sim2L* is that all independent input variables (those that users will be allowed to control) need to be declared and enumerated as a fixed list. Importantly, developers can specify acceptable ranges for numerical variables. All inputs and their values are checked before execution and only simulations with all valid inputs are accepted. *Sim2L* developers should decide which parameters will be adjustable by users and which ones will be hard-coded. The hard-coded parameters and the ranges associated with the various adjustable inputs should be designed to result in meaningful simulations. Importantly, by selecting the list of adjustable parameters and their ranges, developers can focus their *Sim2Ls* on specific tasks and minimize the chance of erroneous runs due to unphysical or otherwise inappropriate input parameters. This is an important feature of *Sim2Ls* as most research codes do not perform such checks. In addition, while most scientific software has broad applicability, *Sim2Ls* enable developers to design workflows for specific tasks, and the explicit declaration of input and outputs enables queries to *Sim2L* results. As will be discussed in Section IV, this is important to make the workflows and their data findable and reusable.

*Sim2Ls* accept ten types of input variables: `Boolean`, `Integer`, `Number`, `Array`, `Text`, `Choice`, `List`, `Dictionary`, `Image`, and `Element`. All input types have shared characteristics: type, description, and value. The `Integer` and `Number` types additionally accept minimum and maximum values and the `Number` and `Array` types can have units. Unit conversion between user-provided input data and simulation input data is performed automatically using the `Pint` [25] library. An `Image` refers to a file using one of several popular formats including PPM, PNG, JPEG, GIF, TIFF, and BMP. The `Element` type allows specification of several chemical element properties using only the periodic table identifier, this is powered by the `mendelev` [26] library. The `Array`, `Text`, `List`, `Dictionary`, and `Image` types may be provided as files or Python variables of the proper type. All *Sim2L* inputs must be enumerated in a single notebook cell using YAML. The tool *Introduction to SimTools* includes all possible input types and exemplifies their use [27].

**Parameterization.** The *Sim2L* notebook must contain a cell tagged as `parameters`. The cell should contain an assignment statement for each input. The example given in Fig 2 sets specific values to the input variables, this is useful for developers during testing. The function `getValidatedInputs` from the *Sim2L* library should be used in the parameterization cell to set default values; this is exemplified in Ref. [27]. When using the *Sim2L* library to execute a simulation the parameter values will be replaced by those provided by the user.

**Workflow.** Following the `parameters` cell, *Sim2Ls* should include the workflow required to generate the outputs (described below) from the inputs. This workflow can include multiple simulations, including parallel runs in HPC resources. Within nanoHUB the `submit` command [28] enables users to launch simulations in various venues outside the execution host that powers the notebook. Importantly, this workflow should contain all the pre- and post- processing steps required to turn inputs into outputs. While these steps are often considered unimportant and poorly described in many publications, they can significantly affect results [29].

**Outputs.** Another key aspect of a *Sim2L* is that all outputs of interest must be enumerated as a fixed list. It should be noted that there is a difference between a *Sim2L* output and the simulation results. A scientific application may produce many more results than what is reported by a *Sim2L* as outputs. Like inputs, outputs are not optional, if an output is declared it must be saved during the simulation or the *Sim2L* library will return an error. Output types are the same as the ten input types described above. All *Sim2L* outputs must be enumerated in a single notebook cell using YAML. Developers might be tempted to include important outputs in files

with ad hoc formatting. This is discouraged as it precludes the results from being discoverable and queryable and hinders the re-use of simulations.

**Files.** The *Sim2L* notebook may contain an optional cell tagged `FILES`. The cell contains a list of auxiliary files required by the *Sim2L* notebook. Examples would be additional Python files containing utility methods to support the simulation. In some cases it might be useful to provide files containing static data.

### B. Interacting with *Sim2Ls*: *Sim2L* library

Users and developers interact with *Sim2Ls* using the *Sim2L* library, see Fig 3. This library enables users to find deployed *Sim2Ls*, their requirements (inputs), and services (outputs); it also provides a mechanism to executing them.

**Exploring *Sim2Ls* and setting up runs.** The `findSimTools` command enables users to discover available *Sim2Ls* and descriptions. This command can be combined with the `getSimToolInputs` and `getSimToolOutputs` to find *Sim2Ls* that provide the services of interest.

The *Sim2L* library also facilitates simulation by providing an object used to declare all required inputs. This object is passed back to the *Sim2L* library for parameterization and execution. Upon completion of the simulation a second object gives access to all declared simulation outputs. After the successful execution of a *Sim2L*, the resulting notebook (including all inputs and outputs) is automatically stored in nanoHUB's simulation cache.

When a new *Sim2L* run is requested, the *Sim2L* library checks the cache before execution. If a perfect match is found, the *Sim2L* library pulls the result from the cache. This not only saves compute cycles (with the consequent energy savings and reduction in carbon footprint) but also provides users with results nearly instantaneously. The simulation cache is particularly

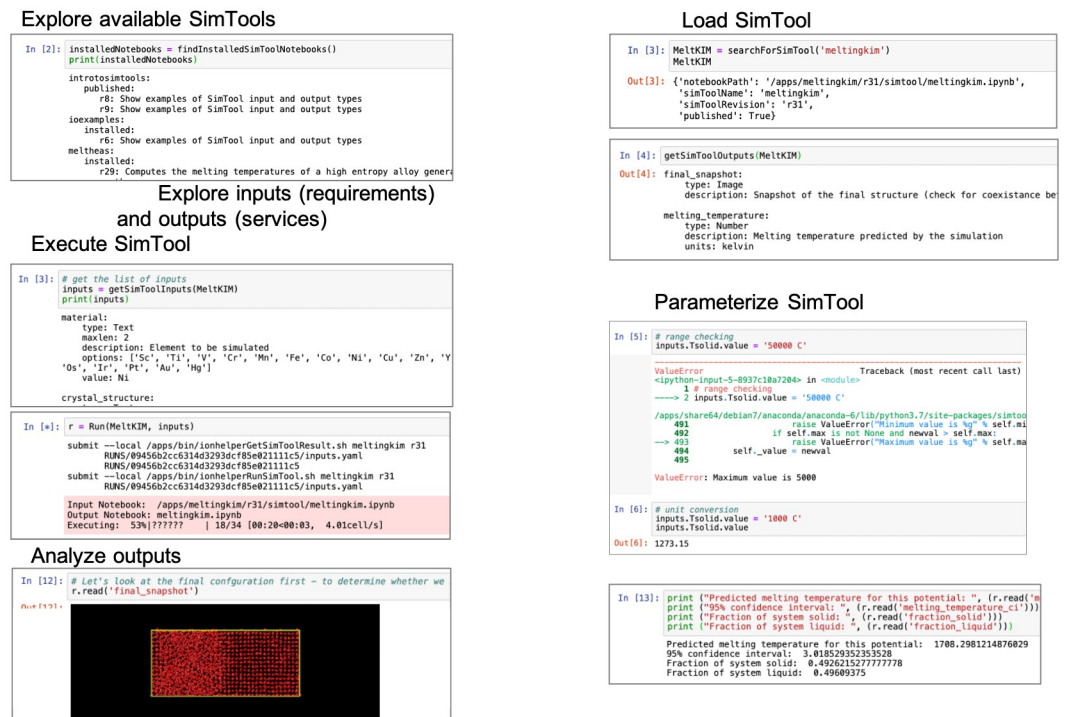


Fig 3. Discovering, parameterizing, executing, and analyzing *Sim2Ls*.

<https://doi.org/10.1371/journal.pone.0264492.g003>

useful for computationally intensive tools and for classroom use when many users perform identical simulations.

**Execution.** The papermill [24] library is used to execute the code contained in the *Sim2L* notebook. The constrained nature of *Sim2Ls* means that only the *Sim2L* notebook, self declared additional files, and optional user provided input files need be provided to run a simulation. This well-defined structure lends itself to being able to run simulations in a variety of venues. By default, simulations are executed within the HUBzero tool session environment [28]. Another option is to build Docker or Singularity containers that mimic the HUBzero environment. Such containers can then be distributed to other locations and executed. This strategy is used to execute *Sim2Ls* utilizing MPI or other parallel computational methods. The use of off-site execution utilizes the submit command and requires only minimal additional specification including maximum wall time and number of cores to be provided. The use of containers allows our team to deploy simulation execution to various resources without modifying the *Sim2L* itself, eliminating the need for developer customization.

The following lists the *Sim2L* library functions available to interact with *Sim2Ls*.

- `findSimTools`—find all installed and published *Sim2Ls*. In addition to name and revision a brief description is returned for each *Sim2L*.
- `searchForSimTool`—search for a particular *Sim2L*. The search may include a specific revision requirement.
- `getSimToolInputs`—get definition of each input for given *Sim2L*. Definition includes name and type for each variable plus type dependent information such as units, minimum, maximum, description, default value.
- `getSimToolOutputs`—get definition of each output for given *Sim2L*. Definition includes name and type for each variable plus type dependent information such as units and description.
- `Run`—method to run specific *Sim2L* with provided input values. In addition more information may be provided if the simulation is to be executed remotely. There is also an option to control data exchange with the results cache.

### C. Publishing a *Sim2L*, simulation caching, and results database

Once tested by developers the process of tool publication makes them available to nanoHUB users. Every published nanoHUB tool is assigned a digital object identifier (DOI) which is updated as new versions are released. The tool publication process enables developers to specify authorship, acknowledgments, provide appropriate references, and describe the tool. Optional supporting material can be included with the tool. Once published, nanoHUB tools, including *Sim2Ls*, are indexed by Google Scholar and Web of Science. Published *Sim2Ls* can be invoked by users from any Jupyter notebook running in nanoHUB which enables them to be invoked in high-throughput or machine learning workflows called Apps, see Section III.

As mentioned above, every successful *Sim2L* run performed on behalf of users is stored in nanoHUB's simulation cache and the *Sim2L* outputs indexed and stored in the results database (resultsDB). Thus, when a user requests a simulation previously performed it is retrieved from the cache. This results in faster response time for the user and saves computational resources. Finally, the resultsDB can be queried using an API. Thus, every *Sim2L* performed in nanoHUB is automatically stored and the results are queryable.

### III. Sim2L examples

#### A. Melting temperature calculations using molecular dynamics

The *Melting point simulation using OpenKIM Sim2L* [23] in nanoHUB calculates the melting temperature of metals using molecular dynamics simulations. Users specify the element of interest, the model to describe atomic interactions, any additional simulation parameters, and the *Sim2L* calculates the melting temperature of the material of choice using the two-phase coexistence method [30]. In this approach one seeks to achieve the coexistence between a liquid and a crystal phase, by definition the temperature at which this occurs is the melting temperature of the system. The tool creates a simulation cell, assigns initial temperature values to two halves of the simulation box and, after a short equilibration, performs a molecular dynamics simulation under constant pressure and enthalpy. The choice of ensemble results in the system temperature naturally evolving towards the melting temperature and if coexistence is observed once the system reaches steady-state, the system temperature corresponds to the melting temperature. If the entire cell ends up as a solid, the initial temperatures were too low and should be adjusted upward. Conversely, if the entire system melts, the initial temperatures were too high. The *Sim2L* sets up, executes, and analyzes the simulation results. The simulation reports the fraction of solid and liquid phases, the time evolution of the instantaneous temperature and the overall system temperature with a confidence interval. In addition, the *Sim2L* analyzes the data to report whether a meaningful melting temperature can be extracted from the simulation. Below is a description of the key inputs and outputs, focusing on the use of the *Sim2L* library to standardize this melting point calculation protocol. The *Sim2L* is available for online simulation in nanoHUB [23].

**1. Inputs.** *Material.* Users input the element for which they wish to calculate the melting temperature. This input can be one of 29 metals, listed explicitly in the ‘options’ keyword of the *Sim2L* input. This explicit listing allows users to quickly inspect this *Sim2L* input and determine the list of allowed elements. The complete list of elements can be found in Ref. [23].

*Mass.* The *Sim2L* requires the atomic mass of the material as an input, this is of type ‘Element’. This allows users to either specify a numeric mass value or simply specify the symbol of the element, which the *Sim2L* library uses to automatically obtain the mass.

*Crystal structure and lattice parameter.* The crystal structure can be specified to be face centered cubic (FCC), body centered cubic (BCC) or hexagonal close packed (HCP). The ‘options’ feature for this input prevents users from selecting any other crystal. The *Sim2L* expects the lattice parameter to be a number between 2 and 10 Å. However, by leveraging the Pint unit conversion tool [25], the *Sim2L* library allows users to specify the lattice parameter in any units. The *Sim2L* library automatically handles unit conversion and checks whether the converted value belongs to the range of the *Sim2L* input. Thus, a user input of ‘0.5 nm’ is automatically converted to 5 Å, but a user input of ‘5 nm’ will result in an error as the input is internally converted to 50 Å, beyond the range allowed by the *Sim2L*.

*Solid and liquid temperatures.* The *Sim2L* inputs also include the initial temperatures to assign to the solid and the liquid regions of the simulation. Users can enter temperatures in any units as the *Sim2L* library automatically converts them to Kelvin.

*Run time.* Users can also specify the time for which the coexistence calculation is carried out. The run time determines whether the simulation is converged or not. Short run times can result in non-steady state conditions and unreliable calculations. The run-time is also internally converted to femtoseconds, with a default of 50000 fs or 50 ps.

*Interatomic model.* Every molecular dynamics simulation requires an interatomic model to describe the interactions between atoms. The *meltingkim Sim2L* obtains the user-specified interatomic model from the OpenKIM repository [31].

**2. Workflow and outputs.** The *Sim2L* takes in all the user inputs, creates an input file for the parallel molecular dynamics code LAMMPS [32], executes the simulations, and post-processes the results to determine the melting temperature if the simulation was successful. We describe the workflow in some detail to exemplify the various steps and decisions required, even for a relatively simple and standard calculation. The *Sim2L* documents all these steps facilitating reproducibility and accelerating progress as researchers can re-use parts or the entirety of the workflow.

The *Sim2L* first creates a system with the requested crystal structure and lattice parameter and initializes the solid and liquid with atomic velocities matching the specified temperatures. The user specified OpenKIM interatomic model is then downloaded from OpenKIM using their API. The KIM model name is included in the LAMMPS input file such that OpenKIM can interface with LAMMPS and modify any LAMMPS setting (units, atom style etc.) to run the simulation.

Once the system is initialized, the simulation cell parameters and atomic positions are relaxed via energy minimization. The system is then equilibrated under atmospheric pressure for 10 ps, using two independent thermostats applied to the solid and liquid regions to keep the regions at the user-specified temperatures. Following the thermalization, the system is evolved via molecular dynamics under constant pressure and enthalpy (no thermostats), for the run time specified by the user. This phase of the simulation results either in the coexistence of solid and liquid phases (success) or a single phase; the latter indicates that initial temperatures need to be modified and a new simulation must be performed.

The raw output from LAMMPS is also post-processed by the *Sim2L* to provide users with information about the simulation. Not just the systems temperature but also if both solid and liquid are coexisting in equilibrium. The final atomic configuration from the simulation is analyzed to establish whether solid-liquid coexistence exists at the end of the simulation or the system evolved into a single phase. This is done by analyzing the local environment of each atom using the polyhedral template matching algorithm [33] as implemented in OVITO [34]. Each atom is classified into one of many crystal structures based on its neighborhood, with any atom having an unknown neighborhood identified as liquid.

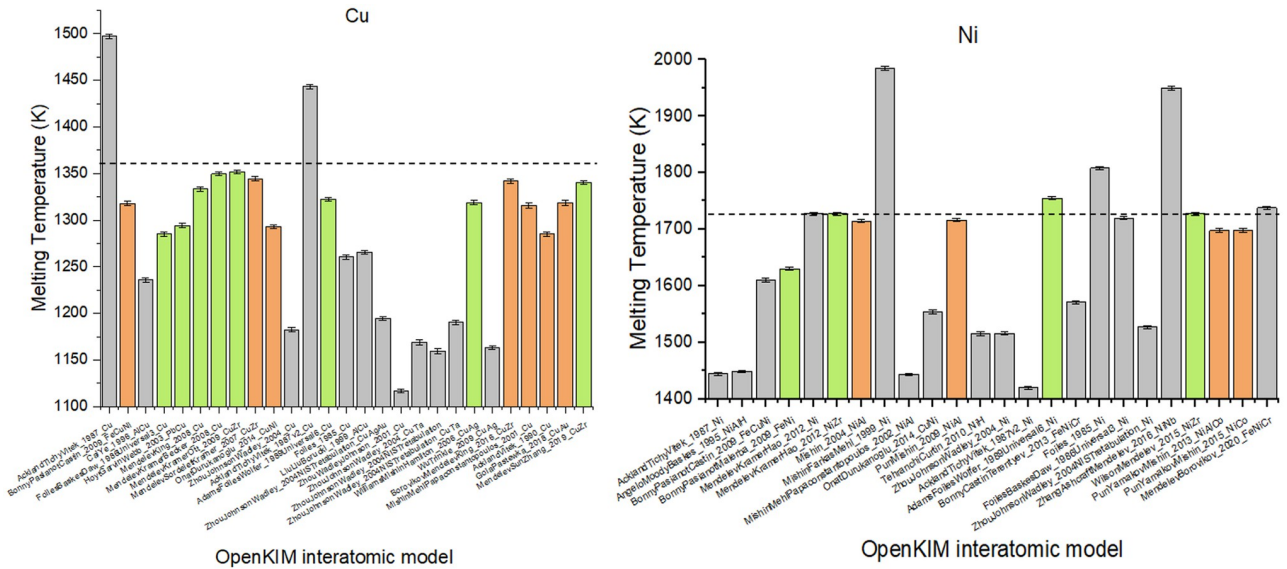
Based on this analysis a boolean output variable, 'coexistence', is determined. If 35% to 65% of the atoms are identified to belong to the initial crystal structure and if 35% to 65% of the atoms are identified as liquid, the system is deemed to have achieved coexistence and the output variable is set to TRUE. The *Sim2L* also outputs a snapshot of the final atomic configuration, for the users to visually inspect coexistence, see bottom panel in Fig 3.

The second test to establish a successful melting temperature the *Sim2L* performs is to check if the system is in equilibrium. To do this, it computes the the slope of the instantaneous temperature vs. time over the 20 ps of the simulation. If the absolute value of the slope is less or equal to 10 K/ps equilibrium is declared and a second boolean variable, 'steady state', is set to TRUE. Lastly, the temperature obtained from the last 20 ps of the simulation is reported as an output and fluctuations of the instantaneous temperature are used to determine the 95% confidence interval on the melting temperature.

The *Sim2L* then saves the melting temperature, the confidence interval, the 'coexistence' and 'steady state' flags, and the fraction of atoms belonging to each crystal structure. This is performed using the save() command from the *Sim2L* library that allows these results to be stored in the Results Database, for easy access later.

**3. Invoking the *Sim2L* and example results.** The tool [23] also contains a Jupyter notebook to invoke the *Sim2L* and demonstrate its use. This driver notebook exemplifies the use of `getSimToolInputs()` and `getSimToolOutputs()` functions to understand the *Sim2L* inputs and outputs, following which the user specifies some or all of the inputs. For





**Fig 4. Melting temperature for copper and nickel, using all OpenKIM interatomic models (bars) compared to experimental results (dashed lines).** Green bars represent calculations which achieved coexistence and steady state, orange bars are calculations which achieved coexistence but not steady state, indicating that longer run times can successfully determine the melting temperature. Gray bars are calculations that did not result in coexistence.

<https://doi.org/10.1371/journal.pone.0264492.g004>

unspecified inputs, the *Sim2L* uses default values, which are also displayed when the `getSimToolInputs()` function is called. The `Run()` function invokes the *Sim2L* by passing it all user inputs, which the *Sim2L* then uses to launch the the LAMMPS simulation and save the outputs to the Results Database. The `getResultSummary()` function can then be used to get a dataframe with the results from the simulation.

The workflow notebook additionally showcases an example of using the *Sim2L* to calculate the melting temperature of a list of elements in an automated manner. We define functions to query repositories such as Pymatgen [35] for elemental properties to be passed as inputs. We also query OpenKIM to find interatomic models appropriate for the element. This example demonstrates how using the *Sim2L* as a fundamental compute unit can help users develop complex workflows and script multiple runs, while utilizing *Sim2L* library capabilities such as unit conversion, input validation, and result caching.

As an illustration of this capability, Fig 4 shows the predicted melting temperatures for copper and nickel using all the interatomic models available for that metal on the OpenKIM repository. Each bar shows the melting temperature predicted for a particular model, with the error bar indicating the uncertainty in the calculation.

### B. P-N junction diode properties using semiconductor device simulations

The *P-N junction Sim2L* [36] uses the device simulator PADRE (Pisces And Device REplacement software) [37] to explore basic concepts of P-N junctions. A P-N junction consists of a P-doped and an N-doped semiconductor arranged in series and has the electrical characteristics of a diode. Despite their simplicity, their operation involves several fundamental concepts of solid state physics and these devices provide useful pedagogical examples and are at the heart of many electronic devices. The *Sim2L* models these devices solving a coupled set of partial differential equations describing its electrostatics (Poisson equation), drift-diffusion (carrier continuity equation), and energy balance (carrier temperature equation). Users of the *Sim2L* can

change doping concentrations for each section of the device, modify materials, the operating temperature, and tune additional properties and explore the resulting I-V characteristics, electronic band structures and hole/electron recombination. The *Sim2L* verifies the input parameters, creates the required input files required for PADRE, and stores energy bands, carrier densities, net charge distribution, voltage-current (IV) characteristic, and other properties as output variables.

**1. Inputs.** *Dimensions.* The *Sim2L* inputs include dimensions of each section in the device, P/N doped and intrinsic regions. PADRE expects these values in microns, however, *Sim2L* users can use any unit that represents distance, the *Sim2L* library process the values and transform the values accordingly.

*Mesh refinement.* The *Sim2L* also expects values for the meshing required by the regions mentioned before, all values are expected to be positive and dimensionless.

*Doping concentration.* Doping levels are required for the P/N type regions, and values are expected in  $cm^{-3}$  units, these values can be expressed on any scientific notation supported by YAML.

*Material.* The material properties used for the simulation depend on the parameter passed to the *Sim2L*, the material input is a string, and supports selected semiconductors ('Si', 'Ge', 'GaAs', and 'InP').

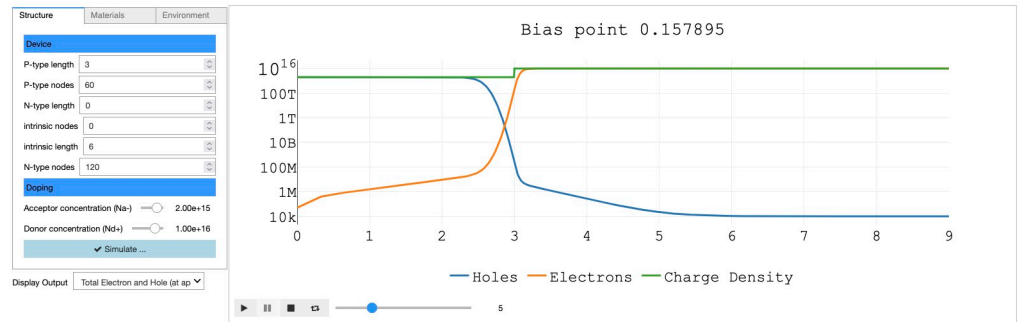
*Additional inputs.* The *Sim2L* also expects inputs for temperature, carriers lifetime, applied voltage, intrinsic region impurity, and environmental options. All the parameters, units, ranges and restrictions are defined on YAML on the cell tagged as 'INPUTS'.

**2. Workflow and outputs.** The *Sim2L* translates inputs as an ASCII text file required by PADRE to run the charge transport analysis. PADRE can calculate DC, AC small signal, and transient solutions. The input file generated defines the structure, material models, numerical methods, and solutions. Meshes for each region of the device are defined based on the length and doping level of each region. The transport model includes Shockley-Read-Hall generation/recombination process, concentration-dependent mobility model, field-dependent mobility, and impact ionization process. The *Sim2L* first calculates the solution for the equilibrium state, and then the solutions are calculated for bias applied to the anode. The bias is increased on the specified range, and the step size defined by the inputs. PADRE's outputs are saved as text files, files are post-processed and the results saved as the *Sim2L* outputs. *Sim2L* outputs provide users with the characteristics and quantities representative of the device. The most relevant outputs of the *Sim2L* are described next.

*Energy bands.* The *Sim2L* calculates electron and hole energies, conduction band ( $E_c$ ), valence band ( $E_v$ ), intrinsic Fermi energy ( $E_i$ ), and Fermi levels along the dimension of the device. Together these outputs represent the band diagram that describes the operation of the device under the desired conditions. The *Sim2L* not only calculates energies at equilibrium but also under different bias potentials, see Fig 5 This can be used to visualize evolution of the band diagram as voltage is increased.

*Device characteristics and related outputs.* The *Sim2L* calculates and outputs current-voltage characteristics as well as capacitance. In addition, doping densities, electric fields, charge densities, potentials and recombination rates as function of position are tool outputs, see Fig 5.

**3. Running the *Sim2L* via an App.** *Sim2Ls* can be invoked from Python scripts, including Jupyter notebooks, or from graphical user interfaces. The P-N junction tool includes an easy-to-use GUI implemented in a Jupyter notebook [36]. This App enables users to set inputs and visualize the device band structure, recombination rates, as well other *Sim2L* outputs. The workflow within the App calculates the electric field and potentials using the depletion approximation. This approximation assumes that the depletion region around the junction has well-defined edges and transitions between regions are abrupt. The workflow only includes



**Fig 5. P-N junction App invoked a *Sim2L* that performs the device simulations.** The App enables users to easily setup the simulation and visualize results. The example shows the charge density of the diode for an applied voltage of 0.157895 eV, the slider allows user to visualize different applied voltages.

<https://doi.org/10.1371/journal.pone.0264492.g005>

approximations for junctions without intrinsic regions, in equilibrium, and only use the ideal silicon intrinsic doping. The approximation is displayed together with the simulation results for educational purposes.

#### IV. Discussion and outlook

This section discusses important aspects of the simulation ecosystem for developers to consider when designing *Sim2Ls*. While nanoHUB makes *Sim2Ls* and their data automatically accessible (via DOIs, standard licenses and APIs), these additional considerations are important to facilitate findability, interoperability, and reuse and the *Sim2Ls* themselves and the data they produce.

##### Descriptions and metadata

*Sim2L* abstracts are required as part of the publication process and the *Sim2L* itself has a [description] field that can be queried when searching for *Sim2Ls*. Detailed descriptions help users find the appropriate tools. In addition, concise and accurate descriptions of inputs (requirements) and outputs (services) help with findability.

##### Narrow focus vs. general *Sim2Ls*

We believe narrowly defined *Sim2Ls*, i.e. workflows designed to accomplish specific tasks, contribute to the usability of the tool and the findability and reuse of the results produced. The success of large repositories of ab initio materials data is due, at least in part, to the specific nature of the quantities included [9].

Many physics-based simulation codes have a very broad applicability, and *Sim2Ls* can be used to establish workflows for specific tasks. For example, molecular dynamics simulations can be used to explore mechanical properties, chemical reactions, shock physics, thermal transport, in materials ranging from metals to bio-inspired composites. *Sim2Ls* can be used by researchers in all those fields to document and share specific workflows targeting specific properties.

##### Input and outputs

The choice of inputs and outputs and their descriptions is critical to make *Sim2Ls* and their data FAIR. While files are allowed as input and outputs, their use should be very limited since

they can defeat the purpose of queryable inputs, outputs, and results. For example, a *Sim2L* could take the input file of a physics-simulator as the only input and produce a single output that contains a tar file of all results. This is strongly discouraged. Inclusion of results files from the simulator as a *Sim2L* output *in addition* to outputs that focus on the quantities of interest may be useful to enable users to perform a detailed exploration of their runs and even identify problems with certain simulations. Another acceptable use of output files are well defined file types like PDB for molecular structures or CIF files for crystals.

### The results database (ResultsDB)

All cache simulations in nanoHUB are indexed and stored in the ResultsDB and can be explored via an easy-to-use API [38]; these elements will be described in a subsequent publication. The ability to query and re-use community-generated *Sim2L* results highlights the importance of carefully defining inputs and outputs quantities and types and designing complete end-to-end workflows that generate all relevant quantities of interest.

### Web services

*Sim2Ls* can be launched from within the nanoHUB simulation environment, either a terminal or a Jupyter notebook. In addition, *Sim2Ls* can be queried and launched from nanoHUB web services as will be described subsequently.

## V. Conclusion

In summary, *Sim2Ls* are a key component of the nanoHUB ecosystem to deliver simulations and their data. Queryable descriptions, requirements, and services (including metadata) and the use of standard technologies make both the workflows and data FAIR. The declaration of inputs and outputs, including metadata, together with the simulation cache and ResultsDB means that all data generated can be explored, analyzed, and repurposed. *Sim2Ls* are available in the open platform nanoHUB both for developers and users. nanoHUB provides a complete scientific software development environment and compute power free of charge and online to lower the barrier of access to advanced simulations and to level the playing field in computational science.

## Acknowledgments

Stimulating discussions with Michael Zentner and Gerhard Klimeck are gratefully acknowledged.

## Author Contributions

**Conceptualization:** Alejandro Strachan.

**Funding acquisition:** Alejandro Strachan.

**Software:** Martin Hunt, Steven Clark, Daniel Mejia, Saaketh Desai.

**Supervision:** Alejandro Strachan.

**Visualization:** Daniel Mejia, Saaketh Desai.

**Writing – original draft:** Steven Clark, Daniel Mejia, Saaketh Desai, Alejandro Strachan.

**Writing – review & editing:** Steven Clark, Daniel Mejia, Saaketh Desai, Alejandro Strachan.

## References

1. Baker Monya. Reproducibility crisis. *Nature*, 533(26):353–66, 2016.
2. Goodman Steven N, Fanelli Daniele, and Ioannidis John PA. What does research reproducibility mean? *Science translational medicine*, 8(341):341ps12–341ps12, 2016. <https://doi.org/10.1126/scitranslmed.aaf5027> PMID: 27252173
3. Butler Keith T, Davies Daniel W, Cartwright Hugh, Isayev Olexandr, and Walsh Aron. Machine learning for molecular and materials science. *Nature*, 559(7715):547–555, 2018. <https://doi.org/10.1038/s41586-018-0337-2> PMID: 30046072
4. Himanen Lauri, Geurts Amber, Foster Adam Stuart, and Rinke Patrick. Data-driven materials science: status, challenges, and perspectives. *Advanced Science*, 6(21):1900808, 2019. <https://doi.org/10.1002/advs.201900808> PMID: 31728276
5. Wilkinson Mark D, Dumontier Michel, Aalbersberg IJsbrand Jan, Appleton Gabrielle, Axton Myles, Baak Arie, et al. The fair guiding principles for scientific data management and stewardship. *Scientific data*, 3(1):1–9, 2016. <https://doi.org/10.1038/sdata.2016.18> PMID: 26978244
6. Saal James E, Kirklın Scott, Aykol Muratahan, Meredig Bryce, and Wolverton Christopher. Materials design and discovery with high-throughput density functional theory: the open quantum materials database (oqmd). *Jom*, 65(11):1501–1509, 2013. <https://doi.org/10.1007/s11837-013-0755-4>
7. Curtarolo Stefano, Setyawan Wahyu, Hart Gus LW, Jahnatek Michal, Chepulskii Roman V, Taylor Richard H, et al. Aflow: an automatic framework for high-throughput materials discovery. *Computational Materials Science*, 58:218–226, 2012. <https://doi.org/10.1016/j.commatsci.2012.02.005>
8. Blaiszik Ben, Ward Logan, Schwarting Marcus, Gaff Jonathon, Chard Ryan, Pike Daniel, et al. A data ecosystem to support machine learning in materials science. *MRS Communications*, 9(4):1125–1133, 2019. <https://doi.org/10.1557/mrc.2019.118>
9. Jain Anubhav, Persson Kristin A, and Ceder Gerbrand. Research update: The materials genome initiative: Data sharing and the impact of collaborative ab initio databases. *APL Materials*, 4(5):053102, 2016. <https://doi.org/10.1063/1.4944683>
10. O'Mara Jordan, Meredig Bryce, and Michel Kyle. Materials data infrastructure: a case study of the citri-nation platform to examine data import, storage, and access. *Jom*, 68(8):2031–2034, 2016. <https://doi.org/10.1007/s11837-016-1984-0>
11. Nature news. URL <https://www.nature.com/sdata/>.
12. OpenKIM. Open Knowledgebase of Interatomic Models <https://openkim.org/>, 2018. URL <https://openkim.org/>.
13. Strachan Alejandro, Klimeck Gerhard, and Lundstrom Mark. Cyber-enabled simulations in nanoscale science and engineering. *Computing in Science & Engineering*, 12(2):12–17, 2010. <https://doi.org/10.1109/MCSE.2010.38>
14. Pizzi Giovanni, Cepellotti Andrea, Sabatini Riccardo, Marzari Nicola, and Kozinsky Boris. Aiida: automated interactive infrastructure and database for computational science. *Computational Materials Science*, 111:218–230, 2016. <https://doi.org/10.1016/j.commatsci.2015.09.013>
15. Lamprecht Anna-Lena, Garcia Leyla, Kuzak Mateusz, Martinez Carlos, Arcila Ricardo, Del Pico Eva Martin, et al. Towards fair principles for research software. *Data Science*, 3(1):37–59, 2020. <https://doi.org/10.3233/DS-190026>
16. Spinellis Diomidis. Git. *IEEE software*, 29(3):100–101, 2012. <https://doi.org/10.1109/MS.2012.61>
17. Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E Granger, Matthias Bussonnier, Jonathan Frederic, et al. *Jupyter Notebooks—a publishing format for reproducible computational workflows.*, volume 2016. 2016.
18. Jain Anubhav, Hautier Geoffroy, Moore Charles J, Ong Shyue Ping, Fischer Christopher C, Mueller Tim, et al. A high-throughput infrastructure for density functional theory calculations. *Computational Materials Science*, 50(8):2295–2310, 2011. <https://doi.org/10.1016/j.commatsci.2011.02.023>
19. Karls Daniel S, Bierbaum Matthew, Alemi Alexander A, Elliott Ryan S, Sethna James P, and Tadmor Elad B. The openkim processing pipeline: A cloud-based automatic material property computation engine. *The Journal of Chemical Physics*, 153(6):064104, 2020. <https://doi.org/10.1063/5.0014267>
20. Bonaretti Serena, Gold Garry E, and Beaupre Gary S. pykneer: An image analysis workflow for open and reproducible research on femoral knee cartilage. *Plos one*, 15(1):e0226501, 2020. <https://doi.org/10.1371/journal.pone.0226501> PMID: 31978052
21. Science journals: editorial policies. URL <https://www.sciencemag.org/authors/science-journals-editorial-policies>.
22. Scientific Data. Scientific data recommended repositories, Mar 2019. URL [https://figshare.com/articles/dataset/Scientific\\_Data\\_recommended\\_repositories\\_June\\_2015/1434640/16](https://figshare.com/articles/dataset/Scientific_Data_recommended_repositories_June_2015/1434640/16).

23. Martin Hunt, Alejandro Strachan, and Saaketh Desai. Melting point simulation using openkim, Mar 2019. URL <https://nanohub.org/resources/meltingkim>.
24. Papermill Developers. Parameterize, execute, and analyze notebooks, a. URL <https://papermill.readthedocs.io>.
25. Pint Developers. Pint: Operate and manipulate physical quantities in python, b. URL <https://pint.readthedocs.io>.
26. Łukasz Mentel. mendeleeev—a python resource for properties of chemical elements, ions and isotopes. URL <https://github.com/lmmentel/mendeleeev>.
27. Saaketh Desai, Stephen Clark, and Alejandro Strachan. Introduction to simtools, April 2020. URL <https://nanohub.org/tools/introtosimtools>.
28. McLennan Michael, Clark Steven, Deelman Ewa, Rynge Mats, Vahi Karan, McKenna Frank, et al. Bringing scientific workflow to the masses via pegasus and hubzero. *parameters*, 13:14, 2013.
29. Alzate-Vargas Lorena, Fortunato Michael E, Haley Benjamin, Li Chunyu, Colina Coray M, and Strachan Alejandro. Uncertainties in the predictions of thermo-physical properties of thermoplastic polymers via molecular dynamics. *Modelling and Simulation in Materials Science and Engineering*, 26(6):065007, 2018. <https://doi.org/10.1088/1361-651X/aace68>
30. Morris James R, Wang CZ, Ho KM, and Chan CT. Melting line of aluminum from simulations of coexisting phases. *Physical Review B*, 49(5):3109, 1994. <https://doi.org/10.1103/PhysRevB.49.3109> PMID: 10011167
31. Tadmor Ellad B, Elliott Ryan S, Sethna James P, Miller Ronald E, and Becker Chandler A. The potential of atomistic simulations and the knowledgebase of interatomic models. *Jom*, 63(7):17, 2011. <https://doi.org/10.1007/s11837-011-0102-6>
32. Plimpton Steve. Fast parallel algorithms for short-range molecular dynamics. *Journal of computational physics*, 117(1):1–19, 1995. <https://doi.org/10.1006/jcph.1995.1039>
33. Larsen Peter Mahler, Schmidt Søren, and Schiøtz Jakob. Robust structural identification via polyhedral template matching. *Modelling and Simulation in Materials Science and Engineering*, 24(5):055007, 2016. <https://doi.org/10.1088/0965-0393/24/5/055007>
34. Stukowski Alexander. Visualization and analysis of atomistic simulation data with ovito—the open visualization tool. *Modelling and Simulation in Materials Science and Engineering*, 18(1):015012, 2009. <https://doi.org/10.1088/0965-0393/18/1/015012>
35. Ong Shyue Ping, Richards William Davidson, Jain Anubhav, Hautier Geoffroy, Kocher Michael, Cholia Shreyas, et al. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013. <https://doi.org/10.1016/j.commatsci.2012.10.028>
36. Daniel Mejia. pntoy using simtool infrastructure, Feb 2021a. URL <https://nanohub.org/resources/st4pnjunction>.
37. M. R. Pinto, C. S. Rafferty, R. K. Smith, and J. Bude. Ulsi technology development by predictive simulations. In *Proceedings of IEEE International Electron Devices Meeting*, pages 701–704, 1993. <https://doi.org/10.1109/IEDM.1993.347216>
38. Daniel Mejia. Database Results Explorer API. <https://nanohub.org/developer/api/endpoint/dbexplorer>, 2021b. [Online; accessed 20-August-2021].