

Cite this: *Digital Discovery*, 2026, 5, 2310

Accelerating discovery across scientific disciplines through reproducible workflows with AiiDALab

Aliaksandr V. Yakutovich,^{†ab} Daniel Hollas,^{†c} Edan Bainglass,^{†d} Jusong Yu,^{†d} Corsin Battaglia,^{†bef} Miki Bonacci,^{†d} Lucas Fernandez Vilanova,^{†g} Stephan Henne,^{†g} Anders Kaestner,^{†h} Michel Kenzelmann,^{†h} Graham Kimbell,^{†b} Jakob Lass,^{†h} Fabio Lopes,^{†a} Daniel G. Mazzone,^{†h} Andres Ortega-Guerrero,^{†a} Xing Wang,^{†*d} Nicola Marzari,^{†di} Carlo A. Pignedoli,^{†*a} and Giovanni Pizzi,^{†*di}

With ever-increasing computational capabilities, robust and automated research workflows have become essential for orchestrating large numbers of interdependent simulations. However, significant technical expertise is still required to configure execution environments, define calculation inputs, interpret outputs, and manage the complexity of parallel code execution on remote machines. To address these challenges, we developed AiiDALab, a Jupyter-based web platform powered by the AiiDA computational infrastructure that provides a framework for managing and automating computational workflows while ensuring reproducibility through full provenance tracking. Through a collection of open-source user-friendly applications, AiiDALab enables scientists to set up, execute, and analyze complex computational workflows without interacting directly with the underlying technical details, allowing them to focus on their research questions. In this paper, we discuss how AiiDALab has matured over the past few years, expanding beyond computational materials science and its AiiDA origins. We present recent developments toward integrating with electronic laboratory notebooks (ELNs) for FAIR-compliant data management, adoption in large-scale facilities for secure access to experimental data and analytical tools, and applications in educational settings. Together with community-driven efforts to simplify onboarding, improve access to computational resources, and support large-scale data workflows, these advancements position AiiDALab as a powerful platform for accelerating scientific discovery and fostering collaboration across disciplines.

Received 17th December 2025
Accepted 1st April 2026

DOI: 10.1039/d5dd00567a

rsc.li/digitaldiscovery

1 Introduction

Computational modeling has established itself as the “third pillar of science” alongside experiment and theory.¹ With growing computational capabilities, the scope and complexity of simulation techniques have also increased. This highlighted a pressing need for software to manage the data and processes of computational workflows. As a result, hundreds of workflow management systems (WFMSs) proliferated in the past decades^{2–7} (see also ref. 8 for an extensive list of WFMSs and ref. 9 for terminology to categorize them). Among these WFMSs is AiiDA,^{10–12} whose core tenets include: (i) an extensible, plugin-based architecture, with a large ecosystem of community-maintained plugins and workflows; (ii) automatic provenance tracking capturing and enabling querying of inputs, outputs, workflow logic, and execution metadata, without requiring manual user intervention; and (iii) high-throughput and fault-tolerant workflow execution, supporting large-scale asynchronous execution of workflows with checkpointing and error handling, making it suitable for managing thousands of concurrent calculations on remote HPC resources. However,

^aNanotech@Surfaces Laboratory, Empa-Swiss Federal Laboratories for Materials Science and Technology, 8600 Dübendorf, Switzerland. E-mail: carlo.pignedoli@empa.ch

^bMaterials for Energy Conversion Laboratory, Empa-Swiss Federal Laboratories for Materials Science and Technology, Dübendorf, CH-8600, Switzerland

^cCentre for Computational Chemistry, School of Chemistry, University of Bristol, BS8 1TS Bristol, UK

^dPSI Center for Scientific Computing, Theory and Data, Paul Scherrer Institute, 5232 Villigen PSI, Switzerland. E-mail: xing.wang@psi.ch; giovanni.pizzi@psi.ch

^eDepartment of Information Technology and Electrical Engineering, ETH Zurich, 8092 Zurich, Switzerland

^fInstitute of Materials, School of Engineering, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

^gLaboratory for Air Quality/Environmental Technology, Empa-Swiss Federal Laboratories for Materials Science and Technology, 8600 Dübendorf, Switzerland

^hPSI Center for Neutron and Muon Sciences, Paul Scherrer Institute, 5232 Villigen PSI, Switzerland

ⁱTheory and Simulation of Materials (THEOS), National Centre for Computational Design, Discovery of Novel Materials (MARVEL), École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

[†] These authors contributed equally to the work.



while WFMSs like AiiDA can, in principle, make scientific workflows more robust and reproducible, executing workflows in practice (configuring calculation parameters, selecting computational resources) remains a challenge, often requiring deep domain expertise.

Human-computer interaction (HCI) research emphasizes that effective user interfaces (UIs) abstract underlying system complexity, allowing users to accomplish tasks without needing to understand internal mechanisms.^{13,14} By reducing cognitive load and presenting appropriate conceptual models, UIs lower the expertise required to operate sophisticated software systems. To reduce barriers to the adoption of workflow management systems (WFMSs) within the scientific community, it is therefore essential to develop UIs that simplify and streamline the steps required to configure and execute workflows, shielding users from underlying technical complexity while preserving flexibility and power.

In the case of AiiDA, this led to the development of AiiDALab,^{15,16} a containerized,¹⁷ Jupyter-based web platform¹⁸ whose focus is making workflow execution user-friendly and accessible through graphical user interfaces (GUIs) running in the user's web browser. AiiDALab is powered by a pre-configured AiiDA installation, providing a solid foundation for developing user-oriented AiiDALab GUI extensions, or apps, each exposing convenient tools for interacting with the corresponding AiiDA workflows. AiiDALab apps were originally developed within and for computational materials science, contributing to research in the fields of surface chemistry,^{19–24} graphene nanostructures,^{25–33} alloys,^{34,35} surface spectroscopy,^{36,37} and spin physics.^{28,29,38} The AiiDALab platform can be hosted online (on institutional servers or in the cloud) *via* JupyterHub and Kubernetes, following a standard recipe for deploying Jupyter-based containerized applications.³⁹ The platform can also be deployed locally on a personal machine. Both deployment schemes are discussed in Section 6. The AiiDALab team maintains several official deployments of the platform.

While GUIs can significantly lower the access barrier to advanced workflows, one needs to be aware, however, of the non-negligible effort required to design and develop GUIs for simulation software. Indeed, a GUI exposing all supported features of a given code or its associated workflows can quickly overwhelm the user, resulting in a negative impact on usability. To avoid this, while developing AiiDALab, we collaborated closely with our users, primarily experimental scientists, and adapted our interfaces following their constructive feedback. For example, to simplify interaction with AiiDA, we developed widgets (GUI components) to facilitate and automate workflow setup, supported by research on computational protocols to balance accuracy and speed.^{40,41} Such features provide users with more intuitive means of setting up computational workflows, reducing the need for deep domain knowledge of the underlying codes. The unified interface that we now provide enables users to focus on the scientific process itself, providing simplified access to the execution of common tasks (such as simulation preparation, execution, and post-processing and visualization of results). We also ensure consistent data representation and handling throughout the AiiDALab platform *via*

shared GUI components available for all app developers in the *aiidalab-widget-base*⁴² library.

An example of AiiDALab app development driven by user feedback is the comprehensive redesign of the On-Surface Chemistry app,⁴³ an AiiDALab app providing *ab initio* tools for studying molecules on surfaces using the CP2K⁴⁴ code. The app now integrates tools for simulating scanning tunneling microscopy (STM) and atomic force microscopy (AFM), leveraging tools developed in part by the AiiDALab team, including the *ppafm* code,⁴⁵ a high-resolution AFM simulation package with sub-molecular resolution based on the probe-particle model. Another example is given by the significant effort devoted to the development of AiiDALab's flagship Quantum ESPRESSO app,⁴⁶ or QE app for short. The complexity and richness of the QE app, covering many of the underlying features of the Quantum ESPRESSO (QE) *ab initio* code, highlights the full potential of the AiiDALab ecosystem. The QE app also offers a flexible plugin system for extending the app's support of additional QE features to meet specific user requirements. Moreover, to facilitate onboarding for new users, the QE app includes a built-in tutorial system to guide users through entire workflows, from setting up a crystal structure for density functional theory (DFT) simulations to job submission, monitoring, and subsequent analysis and visualization of workflow results. The integrated guidance system significantly reduces the need to consult external manuals or user guides, thus streamlining the onboarding process, and is now recommended for all AiiDALab app developers. Lastly, the QE app extends the FAIR⁴⁷ principles to encompass not only the data and processes (already provided by AiiDA), but also the app's analytic tools. The app implements means to save and restore its current state. This means two things: (i) the user can save a snapshot of selected parameters from which one can initialize a new instance, and (ii) the state of any tool (*e.g.*, plotting widgets) can be saved and shared with other users, enabling a quick reload of any simulation and its analysis.

The successful application of AiiDALab in materials science has inspired researchers from other scientific disciplines to adopt the platform. Collaborations between the AiiDALab team and such researchers have highlighted further challenges to address in order to facilitate the adoption of the platform by the wider scientific community. For example, in Section 2, we discuss the application of AiiDALab by researchers in atmospheric science and in battery research. In Section 4, we detail the use of AiiDALab for data management and analysis in applications that do not yet leverage AiiDA, highlighting the platform's features beyond an AiiDA GUI. In addition, a growing need in research to document inventories, protocols, preparation parameters, and results has led to the widespread adoption of Electronic Laboratory Notebooks (ELNs) and, more broadly, Laboratory Information Management Systems (LIMSS). This trend, along with the increasing demand to integrate experimental and computational workflows and to share the resulting data, has motivated the AiiDALab team to explore integrating the AiiDALab platform with ELNs. We discuss this in detail in Section 3.



In the following sections, we elaborate on the challenges we encountered in developing and deploying the AiiDALab platform, on the solutions we designed to address them, and on its successes in and beyond materials science. Our discussion considers organizational, technical, and usability aspects, emphasizing the complex nature of building scientific platforms. We present our approaches to lowering the entry barrier to the platform, allowing users to onboard more easily and facilitating their research. We also discuss the impact and use of AiiDALab beyond just a graphical interface for AiiDA.

2 Beyond computational materials science

Despite its origins in materials science, AiiDA was designed from the outset as a domain-agnostic infrastructure. Its core provides a standardized data and provenance model, workflow engine, and execution layer, while domain-specific functionality is introduced *via* a well-established plugin mechanism. In practice, extending AiiDA to a new scientific field primarily involves defining new data types and workflows as plugins, without modifying the core system. The same architectural principle applies to AiiDALab. AiiDALab applications are built as modular, plug-and-play apps on top of a common framework and make extensive use of reusable UI components provided by the *aiidalab-widget-base* library. Many of these widgets (*e.g.*, for workflow configuration, job submission, monitoring, and data visualization) are domain-independent and are already reused across multiple applications. Naturally, moving to a new scientific domain requires the implementation of domain-specific workflows and data types, and in some cases, additional UI elements. However, the underlying data pipelines, execution logic, and large parts of the user interface are shared and reused. This design significantly reduces the amount of domain-specific redevelopment required and supports AiiDALab's role as a general and extensible platform rather than a collection of bespoke, isolated applications.

With that in mind, and building on the successes of the platform in computational materials science, the AiiDALab team, in collaboration with partners in the research community, explored the application of the AiiDALab platform in other scientific domains. In this section, we present three representative AiiDALab apps: (i) FLEXPART, which aims to simplify and automate inverse modeling of greenhouse gas emissions; (ii) AtmoSpec, which enables *ab initio* prediction of ultraviolet-visible spectra for organic molecules; and (iii) Aurora, which focuses on orchestrating charge-discharge cycling experiments for large sets of battery cells. For each project, we will highlight the unique challenges met, the tailored solutions developed, and the lessons learned along the way.

2.1 Inverse modeling of greenhouse gas emissions

Monitoring greenhouse gas (GHG) fluxes to and from the atmosphere is critical, as their accumulation drives global warming and climate change. The United Nations Framework Convention on Climate Change (UNFCCC) addresses this by

requiring participating nations to regularly report their GHG fluxes as part of ongoing mitigation efforts. These reports are typically based on “bottom-up” approaches, which rely on statistical data about emission sources, their activity levels, and corresponding emission factors.⁴⁸ An alternative method involves “top-down” estimates, where direct atmospheric GHG concentration measurements are combined with atmospheric-transport simulations to infer GHG fluxes in a more integrated manner.⁴⁹ Two prominent examples of how atmospheric observations have led to tightening of emission control mechanisms and revealed shortcomings in abatement strategies are given by Park *et al.*⁵⁰ and Rust *et al.*⁵¹ for the ozone depleting and greenhouse gas CFC-11 and the very potent greenhouse gas HFC-23, respectively. However, atmospheric transport and inverse modeling involve multiple data processing and modeling steps, which are often treated as separate tasks and require significant manual intervention.

For example, one common task is to launch simulations that generate concentration footprints based on a Lagrangian Particle Dispersion Model (LPDM) implemented in the FLEXPART code.⁵² These footprints must be computed for every atmospheric concentration observation to be assimilated in the subsequent inverse modeling step. Currently, greenhouse gases are monitored at a growing number of sites across Europe (currently 40–50), where *in situ* gas analyzers provide high-temporal-resolution measurements. When these data are aggregated to hourly averages for assimilation, the total number of required footprint calculations for a one-year inversion can reach approximately $O(10^5)$. Although this number can be reduced by grouping observations into combined simulations, a substantial number—on the order of $O(10^3)$ —of individual LPDM simulations typically remains. Managing this volume of computations becomes increasingly complex without a robust and reliable workflow management system. Looking ahead, the volume of observational data is expected to grow significantly, particularly with the expansion of both *in situ* monitoring networks and satellite-based remote sensing platforms.

Within the *Explore AiiDA for Regional Inverse modeling of Greenhouse Gases* (ExAIRIM) Open Research Data (ORD) project of the Federal Institute of Technology (ETH) domain, the AiiDALab team developed an AiiDA plugin for the FLEXPART code⁵³ to manage a robust inverse modeling workflow focused on previously used regional inversion domains, such as Europe and Switzerland. The underlying AiiDA workflow triggers the inverse modeling calculations, which ingest both the monthly concentration footprints and the corresponding atmospheric concentration observations. The current implementation supports a regional inversion system developed at the Swiss Federal Laboratories for Materials Science and Technology (Empa): Empa Lagrangian Regional Inversion System (ELRIS).⁵⁴ Using this setup, AiiDA interfaces with ELRIS to retrieve flux estimates, which can then be analyzed to quantify greenhouse gas emissions from individual countries, thus providing a scientific basis to track progress toward national GHG reduction targets. For example, ELRIS contributes annually to Switzerland's reporting to the UNFCCC,⁵⁵ delivering national estimates of nitrous oxide and methane fluxes. Beyond enabling



the submission of individual transport model simulations to the HPC infrastructure and retrieving their outputs, the AiiDA-based implementation also automates the concatenation of simulation results over monthly periods. These aggregated outputs are provided in a standardized, interchangeable footprint data format, facilitating seamless sharing and reuse by other research groups.

A major technical challenge in this context was the handling of large volumes of meteorological input and output data. In contrast to many materials-science applications, atmospheric transport simulations often rely on substantial external datasets (usually three-dimensional fields of initial and boundary conditions) and produce comparably large outputs. Due to their size, these data cannot be permanently stored on the AiiDALab server and instead must be retrieved from dedicated long-term storage systems before execution. Similarly, simulation outputs need to be transferred from temporary directories (typically subject to periodic cleanup policies) to persistent storage locations to ensure long-term availability. Addressing these requirements led to the implementation of dedicated data-transfer workflows and further motivated the extension and refinement of AiiDA's stashing mechanisms. In practice, users can request the transfer of generated data to a secure storage location, thereby minimizing the risk of accidental data loss while maintaining reproducibility and provenance tracking.

The accompanying FLEXPART⁵⁶ app, fully developed by the ExAiRIM team, provides an interface to interact with the two core components of the workflow: the atmospheric transport model and the inversion model. The app enables new users to quickly submit large numbers of atmospheric simulations. Users first select the geographic location and time window for which simulations should be performed. Once complete, the app visualizes simulation results along with selected parameters, assuring that model settings were used consistently across individual simulations. Most settings are hidden from the lay user, largely abstracting the level of complexity, but can still be accessed through the GUI for more advanced control. An example of the FLEXPART app interface, along with the retrieved concentration footprints for the Swiss tall-tower site Beromünster, is shown in Fig. 1.

2.2 *Ab initio* UV/vis spectroscopy for volatile organic compounds

Greenhouse gases are only one class of compounds in the chemically diverse mixture that makes up the Earth's atmosphere. Volatile Organic Compounds (VOCs), which originate from both anthropogenic and biogenic sources, play a central role in atmospheric chemistry and participate in a complex network of reactions.⁵⁷ In order to accurately model the evolving composition of the atmosphere, *e.g.*, to estimate the effects of new industrial chemicals on air pollution,⁵⁸ it is essential to understand the reactivity of these compounds. While the existing chemical kinetics models, such as the Master Chemical Mechanism (MCM),⁵⁹ include a large amount of thermally activated chemical reactions, they largely omit photochemical pathways, *i.e.*, reactions initiated by the absorption of visible

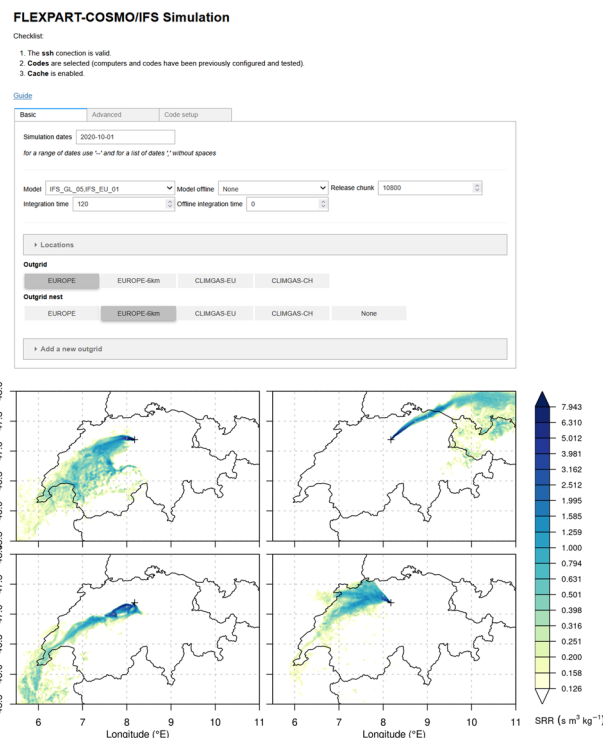


Fig. 1 Top: the FLEXPART app interface to run FLEXPART simulations for greenhouse gas monitoring sites. Bottom: example concentration footprints for the Beromünster tall-tower site in Switzerland, given for four different hourly time intervals. The concentration footprints indicate where the air sampled at the monitoring site was previously in contact with the surface and was potentially impacted by greenhouse gas emissions.

and UV light.⁶⁰ The potential importance of these pathways is largely unexplored and may contribute to some of the observed discrepancies between the models and field measurements.

A crucial quantity for predicting a molecule's photochemical reactivity is its absorption cross-section, which measures the probability of the molecule absorbing a photon at a specific wavelength. While UV/vis spectroscopy can be used to experimentally measure absorption cross-sections, obtaining precise data for gas-phase molecules is generally challenging and often impractical for transient reactive intermediates. Therefore, accurately predicting absorption cross-sections using quantum chemistry tools is highly desirable.⁶⁰ Accurate quantum-chemistry modeling of UV/vis spectra is, however, a challenging task. The simplest workflow, based on a fixed minimum molecular geometry, can only provide positions of spectral peaks, not their shape and widths. More advanced techniques, such as the Nuclear Ensemble Approach (NEA),⁶¹ require execution of many individual quantum-chemistry calculations and are usually performed only by experts in the field of theoretical photochemistry.

The AiiDALab platform presented a unique opportunity to make these simulations more accessible. The AtmoSpec app (see Fig. 2) provides (i) a robust, AiiDA workflow based on NEA, and (ii) a GUI for workflow setup, visualization, and analysis of the calculated UV/vis spectra. In the GUI, the user begins by



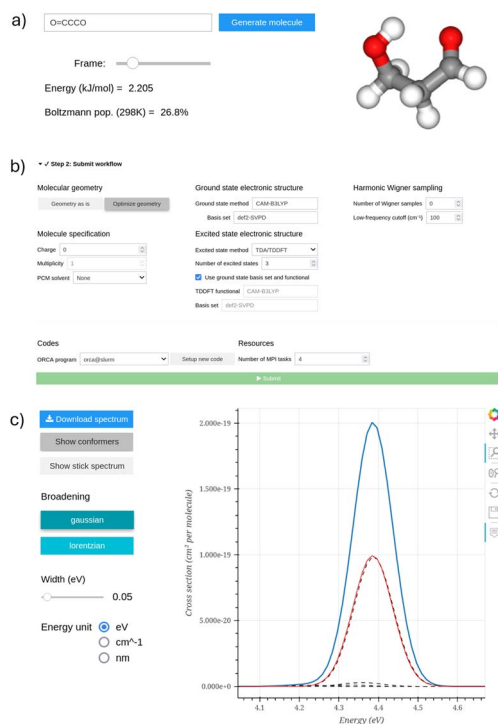


Fig. 2 Example workflow in the AtmoSpec app for a 3-hydroxypropanol molecule. First, (a) molecular conformers are autogenerated from a SMILES string. In the next step, (b) the simulation parameters (such as a DFT functional and basis set) are specified, with pre-determined default values provided. Once the calculation is complete, (c) the UV/vis spectrum is plotted in the spectrum widget, with the option to decompose the spectrum into the contributions of individual conformers.

specifying the molecule of interest by entering a SMILES string (a simple text notation specifying a molecule⁶²) in the *SmilesWidget*, which then automatically generates a corresponding 3D molecular structure, including different molecular conformers. Alternatively, the initial molecular structure can be uploaded from a file. Next, the user enters several computational parameters, such as the DFT functional and basis set (we provide well-tested default settings).

On confirmation, the app launches an AiiDA workflow to perform several quantum chemistry calculations. First, each conformer is optimized at the DFT level. Next, the UV/vis spectrum is modeled using NEA, where the ground-state geometry ensemble is randomly sampled from a harmonic vibrational distribution. Quantum chemistry calculations are performed with the ORCA program,⁶³ interfaced through the *aiida-orca* plugin.⁶⁴ Once all the necessary calculations are finished, the app presents the user with a final UV/vis spectrum through a dedicated *SpectrumWidget*. The spectrum is modeled as a Boltzmann-weighted sum over all molecular conformers, and the widget enables visualization of the individual contributions from each conformer.

By combining automated AiiDA workflows with an intuitive graphical interface, AtmoSpec broadens access to reliable predictions of molecular absorption cross-sections, encapsulating a technically demanding expert procedure within a streamlined

and reproducible protocol. A more detailed description of the AtmoSpec app is given in the corresponding publication.⁶⁵

2.3 Battery assembly and testing

Developing new battery materials extends beyond identifying promising individual compounds. Battery scientists must evaluate combinations of materials (such as anode, cathode, electrolyte solvents, electrolyte salts, and various electrolyte additives) as well as charge–discharge cycling parameters, including current density, voltage cut-offs, and temperature. The goal is to identify material combinations that deliver batteries with both high energy density and long-term cycling stability. High-throughput experimentation, made possible by recent advances in robotic automation platforms, addresses this complexity.

One such platform is Aurora, an automated robotic platform for battery cell assembly and cycling developed by Empa's Materials for Energy Conversion Laboratory in collaboration with the Swiss company Chemspeed Technologies.⁶⁶ Aurora automates the formulation of electrolytes, assembly of battery cells, and their cycling, thereby enabling a closed-loop, application-relevant workflow. Typically, cells with varying combinations of anode, cathode, and electrolyte components are assembled in batches of up to 36 and cycled over extended periods (often several months) using one of the more than 1500 available potentiostat channels. This approach facilitates the evaluation of energy density, rate capability, cycling stability, and the identification of performance degradation mechanisms from the cycling data (see ref. 66 for details on the evaluation of these quantities).

With Aurora's robotic battery cell assembly and a rapidly growing number of cycling channels, experimentation enters a “high-throughput” regime. This was identified in the BIG-MAP⁶⁷ project, in which members of the Aurora and AiiDA teams collaborated to develop and integrate AiiDA-based tools to manage Aurora's experimental workflows. As discussed in a recent publication,²⁴ the automation of the cycling workflow and the tracking of data provenance within the Aurora platform are enabled by the AiiDA infrastructure. In particular, the AiiDA-Aurora plugin⁶⁸ facilitates the integration between AiiDA and the tomato package,⁶⁹ which manages direct communication with the battery cycling hardware. The tomato package includes a job scheduling daemon, device drivers for interfacing with battery cycling potentiostats, and a command-line user interface, collectively enabling scalable control of high-throughput battery cycling experiments.

The AiiDALab Aurora app,⁷⁰ developed by the AiiDALab team, provides a GUI for launching and managing the battery cycling process (see Fig. 3). The app enables users to easily design new cycling protocols which can be saved within AiiDALab for future reuse and applied to either individual cycling jobs or entire batches. Batch submission is supported *via* the select samples widget, which provides filtering by any battery cell parameter. Additionally, users can configure job monitoring parameters directly within the app, including the monitoring frequency, criteria, and associated thresholds. Once the cycling job(s) are



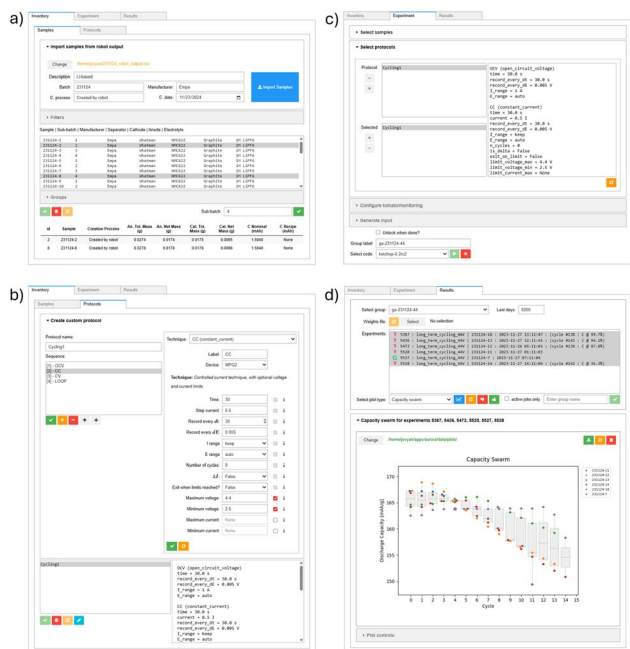


Fig. 3 The Aurora app GUI showing (a) the sample creation widget for importing the output files from the cell assembly robot, (b) the protocol creation widget for defining the cycling conditions, (c) the experiment building wizard, and (d) the results widgets here showing a multi-cell analysis view.

fully defined, AiiDALab transfers all relevant information to AiiDA, which then monitors the cycling process autonomously, without requiring further user intervention. The results are automatically and periodically parsed into a FAIR-compliant data format using the open-source parser *yadg*,^{71,72} storing data together with rich metadata, thus preserving context and meaning. Following parsing, the experimental results are stored in the AiiDA database and can be visualized directly within the app. Attaching monitoring to a process also enables partial data visualization.

The AiiDALab Aurora app was designed in close collaboration with the Aurora team to ensure it best addresses the challenges of daily battery cell cycling. The integrated system provides a single interface in which scientists can design, run, and visualize results from high-throughput battery experiments, accelerating the search for next generation battery materials.

We emphasize that, while AiiDA provides a framework for data and workflow management, interaction with battery cyclers requires dedicated hardware interfaces, data parsers, and corresponding AiiDA plugins. Since there is no widely adopted standard for battery cycler interfaces, separate implementations are needed for each class of cycler hardware, such as *yadg*^{71,72} and *fastnda*⁷³ for data parsing, and *tomato*,⁶⁹ *aurora-neware*,⁷⁴ and *aurora-biologic*⁷⁵ for hardware control.

3 Integration with ELNs and LIMs

Historically, unlike primary data, which might have been acquired using digital systems, experimental notes in scientific

research were primarily recorded using pen and paper. While this traditional approach has served the community for centuries, it presents clear limitations regarding accessibility, reproducibility, and data integration. As a result, more laboratories started adopting ELNs, which offer structured and searchable platforms for managing and sharing experimental information. However, numerical experiments (*i.e.*, simulations) have often remained outside of the scope of ELN integration. Yet, in principle, they follow a similar workflow to laboratory experiments: selecting and preparing a sample, performing the experiment (simulation), and analyzing the results. It is then reasonable to consider both laboratory and numerical experiments as equal candidates for recording in an ELN as part of a scientific study. To streamline the recording of simulations in ELNs, we have developed an integration between AiiDALab and ELNs, allowing users to record AiiDA data and processes conducted in AiiDALab directly in an ELN. The integration positions AiiDALab as a “simulation lab” that enables researchers to analyze experimental and computational work seamlessly across platforms. Such integration facilitates quick and direct comparison of results from experiment and simulation, thus assisting researchers in drawing scientific conclusions.

In Fig. 4, we show an example of AiiDALab interacting with openBIS,⁷⁶ an open-source ELN/LIMS platform developed by the Scientific IT Services (SIS) of ETH Zürich (ETHZ) in close collaboration with scientists at ETHZ and other institutions. In the example, we demonstrate a round-trip between the two platforms, starting from selecting a system of interest in the ELN, setting up and submitting a simulation in AiiDALab's On-Surface Chemistry app, and, upon completion, recording simulation results in the ELN. First, the user establishes a connection between the platforms (not shown). Next, in the ELN, the user selects a molecule from the inventory, and on a click of an “AiiDALab” button (Fig. 4a), sends the molecule to AiiDALab as metadata in the query parameter of the AiiDALab URL. In AiiDALab, the user can leverage dedicated widgets to modify the molecular structure, for example, by adding a substrate (Fig. 4b). The user then selects the executable code to be used and submits the simulation (not shown). After the simulation is complete, the results can be explored within AiiDALab (Fig. 4c). Optionally, the results can be sent back to the ELN for recording in the form of images and/or an AiiDA archive containing the full simulation provenance, explorable *via* AiiDA's tools, *e.g.*, provenance graph visualizer, Materials Cloud Explore section, *etc.* (Fig. 4d). Once data is deposited in the ELN, it may be queried from AiiDALab, for example, *via* a chatbot (Fig. 4e).

The transfer from AiiDALab to openBIS is done *via* the openBIS pyBIS⁷⁷ Python API. The round-trip process enables linking ELN inventory entries for both the molecule and the substrate to the simulation and associating the simulation itself with a specific experiment defined in the ELN. The example also demonstrates more generally the availability of scanning tunneling microscopy (STM) simulations to ELN users *via* AiiDALab and how results can be exported back to an ELN, where they can be compared against experimental data.



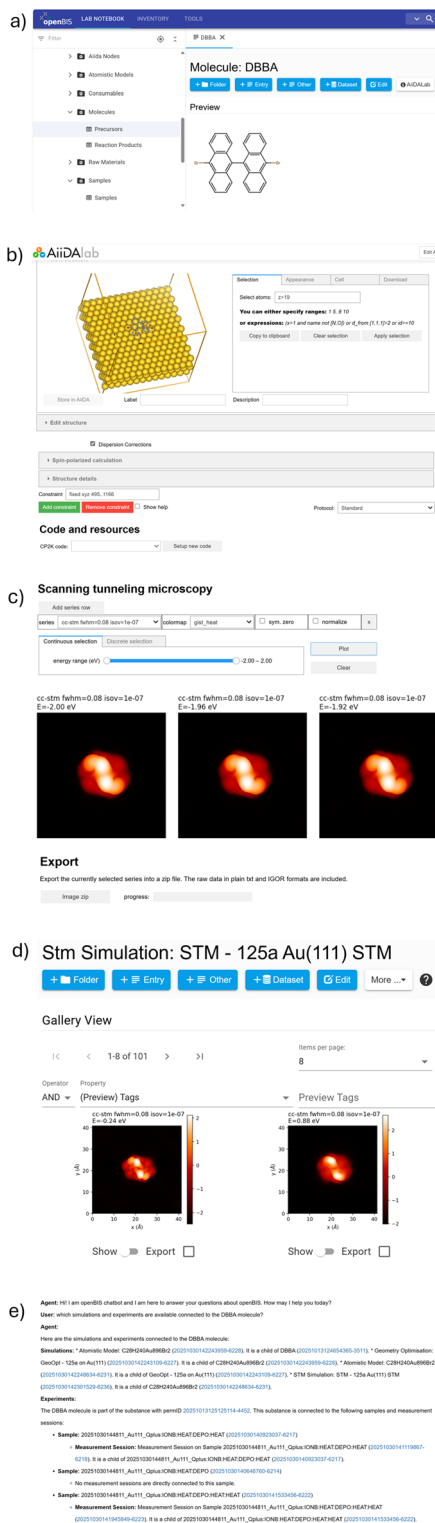


Fig. 4 Screenshots of the integration between AiiDALab and the openBIS ELN-LIMS, for a use case where (a) a molecule is extracted from the ELN inventory and (b) imported to the AiiDALab submission interface for STM simulations. (c) Results of the simulations (the STM images) are shown in AiiDALab. (d) Once exported to the openBIS ELN, results can be displayed directly in the ELN using a dedicated viewer for microscopy images. (e) A chatbot can be used to query the database, e.g., to retrieve available simulations and experiments related to a specific molecule.

Although the above example utilizes openBIS as an ELN, the workflow is not specific to openBIS and can be adapted for use with other ELNs. To support a broader integration, the AiiDALab team developed the *aiidalab-eln* package,⁷⁸ which enables users to add ELN connectors by implementing standardized functions for managing interactions. For detailed instructions on adding a new ELN integration, please refer to the package's GitHub repository. However, as stated above, the integration between AiiDALab and openBIS leverages the specific APIs of the two platforms – the AiiDALab URL and the openBIS pyBIS API. To enable a seamless integration with multiple ELNs without the need to develop a custom connection for each one, one would need to generalize the interoperability layer.

One possible solution is to leverage semantic annotations, which involves linking data elements to shared concepts defined in a common ontology, such as an RDF file or a JSON-LD document. By semantically annotating data and metadata, different ELNs, and even WFMSs beyond AiiDA/AiiDALab, could interpret and process the data consistently, regardless of their internal terminologies. For example, suppose an ELN sends data referring to a “Molecule” along with metadata properly annotated using a shared ontology. The target AiiDALab app could recognize the incoming data as a molecule by mapping it against its own context for a molecule, even in cases where a molecule is internally stored using a different (but semantically equivalent) term, such as “MolecularEntity”. By automatically recognizing the incoming data as a molecule, AiiDALab apps could then recommend compatible simulations for a molecule. Analogously, simulated STM images generated by a WFMS, if appropriately annotated, could be directly displayed using any specific functionality provided in the ELN (if supporting semantic annotations), as shown in Fig. 4, without the need of developing specific code to pair the WFMS and the ELN. This approach facilitates interoperability by ensuring both sides refer to the same standardized concepts, thereby reducing ambiguity and simplifying data exchange.

Research on the application of semantic technologies in this context is being conducted as part of the Open and Reproducible Materials Science Research (PREMISE⁷⁹) project. Additionally, these ideas are being discussed and refined within the Machine-Actionable Data Interoperability for the Chemical Sciences (MADICES⁸⁰) community, which focuses on data interoperability.

4 Integrated data analysis from experiments

The original impetus for the development of AiiDALab was to provide a graphical environment for setting up, submitting, and analyzing computational workflows using AiiDA. A key advantage of such an environment is that all necessary AiiDA components (Python environment, PostgreSQL database, RabbitMQ message broker, and user profile) are pre-configured, significantly reducing the effort required to set up an AiiDA-compatible environment. However, beyond serving as a pre-configured AiiDA environment, AiiDALab can also serve as



a ready-to-use environment for experimental data analysis, especially in scenarios where ease of access, reproducibility, and tight integration with experimental infrastructure are key. In this section, we discuss how AiiDALab can streamline experimental data analysis at large-scale facilities, with a focus on neutron experiments.

At the Paul Scherrer Institute's (PSI) Swiss Spallation Neutron Source (SINQ), the Multiplexing Spectrometer CAMEA⁸¹ has been optimized for energy-resolved detection of in-plane scattered neutrons and is used to map large (q , ω) energy regions in neutron excitation studies. Like many of PSI's facilities, SINQ instruments, including CAMEA, are offered as a user facility, allowing research groups to use its neutron instruments by applying for beamtime. To use CAMEA, visiting researchers must first configure their laptops with the software stack needed for data analysis (using the dedicated in-house code MJOLNIR⁸²). This process often takes place under tight time constraints, as on-the-fly analysis is critical for steering the experiment within the short beamtime window. The post-experiment analysis is also hindered by limited access to raw data, which is often transferred *via* external drives and may no longer be reachable once the user leaves the PSI network, especially due to required VPN or data access policy constraints. Moreover, configuring MJOLNIR on each user's machine is a challenge, as it depends on Python environments, system configurations, and the user's technical expertise. Hardware limitations, operating system incompatibilities, or restricted administrative rights exacerbate the matter further, making installation and use difficult or impossible.

To tackle these challenges, the PSI Laboratory for Materials Simulations (LMS), in collaboration with the PSI Laboratory for Neutron Scattering and Imaging (LNS), developed a bespoke AiiDALab deployment dedicated to streamlining onboarding, authentication, configuration, data access, and data analysis for incoming visiting researchers. The AiiDALab team deployed a local AiiDALab instance using a Micro-Kubernetes (MicroK8s) service (an industry standard for container management),

integrating built-in AiiDALab authentication (*via* JupyterHub) with PSI's institutional authentication infrastructure. Once authenticated *via* institutional credentials, users arrive at an isolated, containerized environment (AiiDALab instance).

Such dedicated AiiDALab instances also support fine-grained data access. In particular, for CAMEA, we developed an access protocol to stream experimental data directly into a shared network file system (NFS) directory accessible from within the PSI network. When a user launches their AiiDALab instance, only specific folders (containing the data of the experiments to which the user has access) are mounted in read-only mode within the user's containerized environment. The authorization policy is queried dynamically *via* an application programming interface (API) from PSI's Digital User Office (DUO) database, which stores a many-to-many map of user accounts to the corresponding proposal IDs of experiments. This ensures live, secure, and granular access to data, fully compliant with PSI's Authentication and Authorization Infrastructure (AAI) and state-of-the-art security standards. The full setup is summarized in Fig. 5.

To streamline data analysis for CAMEA, we also developed the AiiDALab LNS-app,⁸³ which encapsulates MJOLNIR and its dependencies (see Fig. 6). The app's workflow proceeds as follows: (i) while the experiment is performed, data is saved locally on the computer attached to the instrument; (ii) the data is automatically mirrored in real time to an NFS share; (iii) users access the data in real time (read-only and restricted to authorized datasets) *via* the NFS share mounted within their AiiDALab instance; (iv) the analysis is performed using the MJOLNIR code directly within the app.

Upon launching the app, the user can browse for their proposal IDs (or project numbers). The available proposal IDs are limited to only those assigned to the user, as mentioned above. Once a proposal is selected, the user can open a Jupyter notebook connected to the selected project's data and pre-configured with MJOLNIR's analytical tools. The user may freely modify the provided notebook for their specific needs. By combining secure data access and pre-configured analytical tools, the LNS-app effectively reduces setup time to zero, ensuring that novice and advanced users alike can begin analyzing experimental data immediately. After the experiment is completed, the user can install the LNS-app on a local AiiDALab deployment (see Section S4.2) to continue the analysis without reconfiguration, simply by importing the data into the environment. The LNS-app thus ensures consistency and reproducibility across users and beamtimes, both at PSI and on local user machines. In its current pilot phase, the dedicated AiiDALab deployment has already been demonstrated for CAMEA measurements by LNS staff and collaborators, already saving ~ 1 – 2 hours per user by eliminating local software installation and reducing data transfer, environment setup, and plotting configuration through the pre-configured Jupyter notebooks.

A plan has been established to extend the aforementioned data access protocol to other instruments at SINQ, with staff dedicated to this project. For example, AiiDALab is also being adopted at the ICON beamline⁸⁴ for data analysis of cold

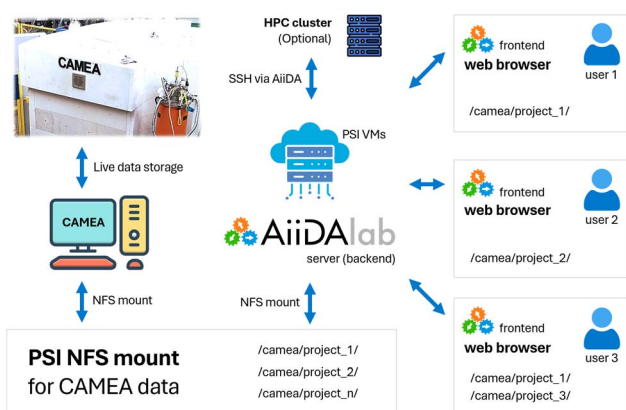


Fig. 5 CAMEA data pipeline. The data of the CAMEA experiment are first pushed from local to a shared NFS storage, then accessed for post-process analysis *via* the AiiDALab PSI deployment and the LNS app. CAMEA users will see experimental data “live” in their AiiDALab deployment, eliminating the need to manually transfer it periodically from the storage attached to the CAMEA hardware.



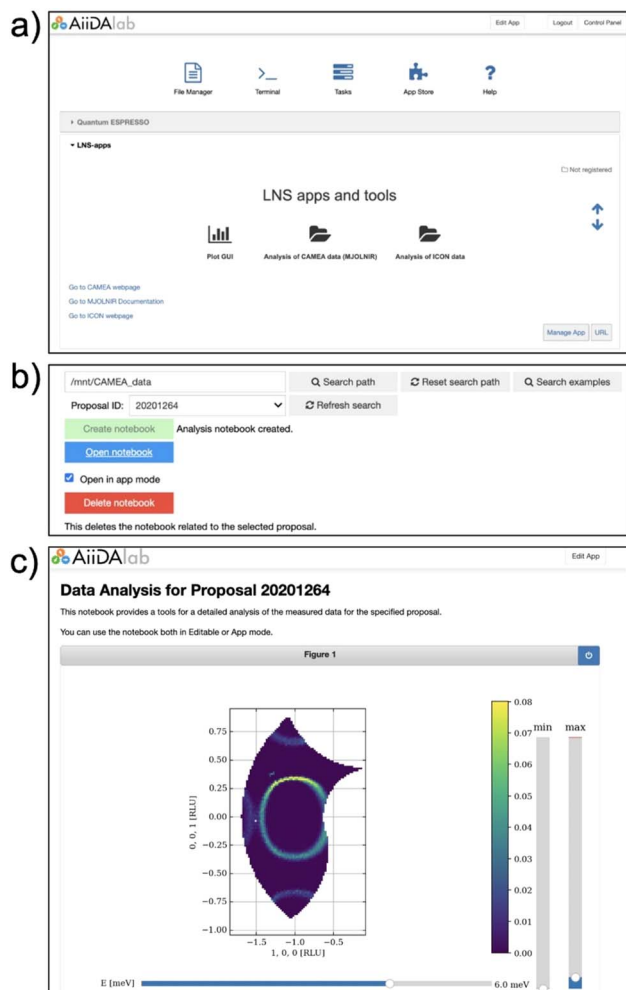


Fig. 6 The LNS app. (a) The app as shown in the AiiDALab home page, once installed. (b) Browsing of proposal folders and buttons to create notebooks for the data associated with the experiment of a given proposal. (c) One of the notebooks created by the buttons in panel (b), showing the MJOLNIR GUI to analyze experimental data.

neutron imaging experiments. These experiments can generate large volumes of complex image data, requiring efficient and reproducible processing workflows. The advantages of using AiiDALab here are similar to the CAMEA example, providing live data inspection without the need to download large amounts of data, centralized reduction, and scalable analysis. Further details on the integration of AiiDALab with the ICON beamline are provided in the Section S1.

In addition to using AiiDALab for secure data access and analysis, LNS personnel leverage the AiiDA engine to submit DFT-based simulation workflows, for example, to model neutron scattering experiments using the QE app (*via* the app's vibroscopy plugin). In the future, even experimental workflows may need to be offloaded to HPC systems, *e.g.*, to enable further advanced reconstructions and machine-learning applications. A natural evolution of the LNS app is therefore to take advantage of AiiDA's native HPC integration to also support remote execution of these workflows. Through AiiDA, remote job

submission and monitoring remain transparent to the user. Moreover, the existing AiiDALab widget ecosystem can be reused to expose HPC capabilities within the same unified UI framework.

The deployment of AiiDALab at PSI illustrates how AiiDALab can support not only computational workflows but also secure, institutionally integrated experimental data analysis, based on standard and robust technology (Kubernetes, JupyterHub). The key advantages include easy integration of HPC into experimental data pipelines and improved data accessibility and security through fine-grained, institution-based authorization policies. Moreover, the PSI deployment offers a unified user environment across experiments and simulations. Together, these features represent a major step forward in creating a truly integrated and reproducible research environment, bridging the gap between experimental facilities, HPC resources, simulation workflows, and user-friendly scientific tools.

5 AiiDALab for teaching

Local or cloud-based installations of AiiDALab provide an ideal platform for exercise sessions in semester-long courses on computational materials science. AiiDALab's Jupyter (Python) environment, powered by AiiDA and a host of apps supporting advanced simulation codes, including Quantum ESPRESSO,⁸⁵ CP2K,⁴⁴ LAMMPS,⁸⁶ and others, allows students to immediately start working with state-of-the-art software without worrying about complex installations. For advanced simulations, AiiDALab's intuitive graphical interfaces and automated workflows allow students to explore sophisticated methods that would otherwise be too time-consuming to set up manually in the scope of an exercise session. The full simulation provenance made available by AiiDA is accessible to students for inspection, allowing them to explore the inputs, processes, and outputs of the simulation. The containerized AiiDALab environment significantly reduces the risk of disrupting the local environment on students' computers by providing a predictable and reproducible learning environment. Furthermore, as AiiDALab is based on Jupyter, it exposes access to many existing education tools, *e.g.*, the *scicode-widgets*⁸⁷ package, a widget library used to teach students how to code and interpret computational experiments.

These features make AiiDALab not only effective for semester-long courses but also particularly well-suited for short, intensive training events, such as those organized by CECAM⁸⁸ or Psi-k,⁸⁹ where time is limited, and a consistent software environment is crucial. Importantly, by introducing students early on to tools that inherently support reproducibility, provenance tracking, and automation, AiiDALab prepares the next generation of scientists for the adoption of FAIR principles. By designing and performing simulations that are transparent, reusable, and reproducible by default, the student learns firsthand how FAIR science is conducted and experiences the benefits of FAIR data *via* AiiDA's built-in querying tools.

A recent example of the use of AiiDALab as a tool for learning was a short-term school on *Electronic-structure simulations for user communities at large-scale facilities* held at PSI in April



2025.⁹⁰ Organized by LMS, the school aimed to provide a gentle introduction to the current capabilities of electronic-structure simulations for materials, with a particular target audience of experimentalists working at large-scale research facilities. The event was conducted in a hybrid format, with approximately 35% of participants attending in person and 65% joining remotely. Participants included both students (~40%, at the Master's and PhD level) and more senior researchers (~60%, including postdocs, tenure-track scientists, and professors). Notably, over 60% of attendees had little or no prior experience with DFT simulations.

The first part of the school covered the theoretical foundations of first-principles modeling, DFT in particular. This was followed by a hands-on session using the AiiDALab platform, specifically the QE app, to run simulations in a user-friendly environment. Exercises covered a range of applications directly relevant to experimental research, including electronic band structures and density of states, inelastic neutron scattering, Raman spectra, muon-stopping sites, as well as X-ray photoelectron (XPS) and absorption (XAS) spectra.

The overall feedback for the school was highly positive, averaging at 8.5 (out of 10) points. This underscores how valuable participants, particularly experimentalists, found the opportunity to gain first-hand experience with electronic-structure simulations in a practical and accessible format. As potential improvements, attendees indicated the need for a slower pace in hands-on sessions and deeper coverage of advanced topics (relevant for their own research field). This will be addressed in future editions through improved time allocation and targeted follow-up materials.

The use of AiiDALab eliminated many technical barriers for newcomers by removing the need to install AiiDA, configure HPC communication, or compile Quantum ESPRESSO executables. Thanks to this setup, even participants with no prior background were able to launch first-principles simulations with just a few clicks. To support the school, a dedicated JupyterHub-based Kubernetes deployment of AiiDALab was made temporarily available on *Microsoft Azure*, with all required apps pre-installed. We discuss these tools in the following section in the scope of deploying the AiiDALab platform. We also discuss in the following section briefly the option of deploying AiiDALab locally, an option that was recommended to students for continued practice.

6 Enabling smooth adoption of AiiDALab solutions

General user feedback revealed that first-time use of AiiDALab often poses challenges due to the number of setup steps required before running the first simulation. We took several steps to alleviate this.

First, we developed *aiidalab-launch*,⁹¹ a command-line interface that automates most steps required to deploy an AiiDALab instance on a local machine. The tool streamlines the creation and configuration of a local AiiDALab container and exposes, as a final step, a URL for the user to open AiiDALab in

their browser. The only requirement left to the user is the prior installation of Docker,⁹² which is well-documented online.⁹³ The instance is empty by default but provides access to the AiiDALab app store, from which the user can install any registered app relevant to their research. As mentioned above, this is particularly useful for students learning how to interact with AiiDALab.

Second, we took steps to simplify, document, and automate the setup of AiiDALab deployments *via* Kubernetes. This allows small-to-medium groups, as well as large-scale institutions, to deploy multi-instance AiiDALab environments for use by their research communities. The new deployment scheme was used to deploy AiiDALab at Empa (production) and PSI (production and school).

Lastly, to better expose the public to AiiDALab's capabilities, we also used the above deployment scheme to deploy the AiiDALab demo server,⁹⁴ a fully-configured AiiDALab instance including the flagship QE app and its dependencies, deployed on the cloud (Microsoft Azure) and accessible to all registered GitHub users. The demo server provides an overview of AiiDALab, deployed with sufficient computational resources to already run small but meaningful simulations, with no required setup. This allows interested parties to explore the features of the platform without having to go through the deployment steps. We believe that the demo server thus provides an effective overview of AiiDALab's capabilities. Once users gain some experience, they can use one of the deployment schemes above to deploy AiiDALab either locally or on a server for multi-instance usage.

We provide further details in the SI on the above deployment techniques and describe more technical, "under-the-hood" developments taken to improve the performance of AiiDALab and the experience of its users.

7 Conclusion and outlook

In recent years, AiiDALab has evolved into a powerful platform to accelerate scientific discovery by making computational science more accessible. The platform transforms complex simulation workflows into structured interactive tasks, implicitly handling much of the lower-level setup. In the background, AiiDA manages a host of powerful community-developed workflows, ensuring that each input, process, and output is tracked, enabling queryability and guaranteeing reproducibility.

Originally developed for materials science, AiiDALab is now expanding into a wider range of scientific domains, from quantum chemistry and atmospheric modeling to experimental battery research, proving to be a flexible and scalable solution for diverse research needs. Built on top of JupyterHub, the platform also provides secure access and a Jupyter (Python) environment for data analysis. This has been leveraged to support experimental labs at the PSI SINQ facility, highlighting AiiDALab's use beyond an AiiDA workflow GUI, also capable of supporting experimental data analysis at large-scale facilities.

Beyond its conceptual and architectural contributions, AiiDA and AiiDALab have demonstrated sustained real-world adoption across multiple scientific domains. The platform is actively used in peer-reviewed research and in operational laboratory



environments, where it supports both computational and experimental workflows. In several research groups, standardized AiiDA workflows exposed through AiiDALab have shifted routine simulation tasks from small teams of computational specialists to a broader community of experimental researchers, lowering the barrier to entry and accelerating iteration cycles. At large-scale facilities, integration with data analysis pipelines has reduced manual software setup, data transfer, and configuration steps, enabling more immediate feedback during experiments. These experiences indicate that AiiDALab not only enhances reproducibility in principle but also tangibly streamlines and democratizes workflow-driven research in practice.

AiiDALab has also proven itself as an education platform, providing a stable and predictable Jupyter environment to instructors and students of both semester-long university courses and short-format training events in computational materials science. Here again, by basing itself on Jupyter, the platform inherently makes available many existing open-source educational tools for instruction, as well as administration (class management, grading, and reporting).

User feedback has always been a key driver of development, guiding the creation of most new features and improvements to provide a platform that enables scientists to focus on the science itself, rather than on the computational infrastructure behind it. We have discussed examples of such improvements, including (but not limited to) simplifying onboarding, streamlining access to HPC resources, and facilitating the handling of large data volumes. Moreover, continuous efforts have been devoted to enhancing service stability, implementing automated testing, and adopting modern software design patterns, which have ultimately resulted in a robust and reliable platform.

So far, adoption of AiiDALab has been mostly happening within Swiss institutions, reflecting the platform's origins and its initial collaborative network. Nevertheless, international uptake already exists. For example, the AtmoSpec application developed at the University of Bristol demonstrates independent use beyond Switzerland. Looking ahead, we see substantial potential for broader geographic and disciplinary expansion and actively welcome engagement from research groups worldwide interested in applying AiiDALab to their own computational or experimental workflows. We are committed to supporting external teams through documentation, onboarding activities, and community workshops, fostering an internationally distributed, community-driven ecosystem.

With these developments, we aim to advance sustainable scientific infrastructure that ensures both scientific data and workflows adhere to FAIR principles, thereby fostering open, reproducible, and collaborative research across various domains. We hope that this approach inspires other scientific software projects, which likely share many common challenges, to adopt similar guidelines to enable FAIR and open science.

Author contributions

Conceptualization: AVY, NM, CAP, GP; methodology: AVY, JY, DH, EB, MB, AOG, XW, CAP, GP; software: AVY, JY, DH, EB, MB,

LFV, GK, JL, FL, DGM, AOG, XW, CAP, GP; validation: AVY, JY, DH, EB, MB, LFV, AOG, XW, CAP, GP; resources: CB, MK, JL, DGM, NM, CAP, GP; writing – original draft: AVY, DH, EB, CB, MB, SH, XW, CAP, GP; writing – review & editing: AVY, DH, EB, JY, CB, MB, LFV, SH, AK, MK, GK, JL, FL, DGM, AOG, XW, NM, CAP, GP; visualization: AVY, JY, DH, EB, MB, LFV, AK, GK, JL, FL, DGM, AOG, XW, CAP, GP; supervision: CB, SH, MK, NM, CAP, GP; project administration: NM, CAP, GP; funding acquisition: CB, SH, MK, NM, CAP, GP.

Conflicts of interest

There are no conflicts of interest to declare.

Data availability

AiiDALab is an open source software ecosystem developed as a collection of interoperable repositories. The source code is publicly available through the AiiDALab GitHub organization at <https://github.com/aiidalab>, which serves as the central entry point to all AiiDALab components described in this work.

The AiiDALab apps presented in this manuscript are also publicly available at the following repositories:

- AiiDALab:
Source <https://github.com/aiidalab/aiidalab>.
DOI: <https://doi.org/10.5281/zenodo.19556336>.
- FLEXPART app:
Source <https://github.com/C2SM/aiidalab-flexpart>.
DOI: <https://doi.org/10.5281/zenodo.19354072>.
- AtmoSpec app:
Source <https://github.com/isp-g-group/aiidalab-isp-g>.
DOI: <https://doi.org/10.5281/zenodo.11075300>.
- Aurora app:
Source <https://github.com/EmpaEconversion/aiidalab-aurora>.
DOI: <https://doi.org/10.5281/zenodo.19354532>.
- ELN app:
Source <https://github.com/aiidalab/aiidalab-eln>.
DOI: <https://doi.org/10.5281/zenodo.19354283>.
- LNS app:
Source <https://github.com/mikibonacci/LNS-app>.
DOI: <https://doi.org/10.5281/zenodo.19348108>.
- aiidalab-launch:
Source <https://github.com/aiidalab/aiidalab-launch>.
DOI: <https://doi.org/10.5281/zenodo.19354408>.
- aiidalab-widgets-base:
Source <https://github.com/aiidalab/aiidalab-widgets-base>.
DOI: <https://doi.org/10.5281/zenodo.19354452>.

Supplementary information (SI): (i) the upcoming integration of AiiDALab with the ICON beamline at PSI; (ii) features in AiiDA and AiiDALab implemented to support large data handling; (iii) the deployment of the AiiDALab demo server, local deployments *via* the aiidalab-launch tool, and general aspects of managing an AiiDALab server; and (iv) under-the-hood improvements to the AiiDALab platform to enhance the user experience. See DOI: <https://doi.org/10.1039/d5dd00567a>.



Acknowledgements

We gratefully thank Carl Simon Adorf for his contributions and developments to the AiiDALab platform. We acknowledge financial support by the NCCR MARVEL, a National Centre of Competence in Research, funded by the Swiss National Science Foundation (grant number 205602). EB, FL, CB, CAP, and GP acknowledge financial support by the Open Research Data Program of the ETH Board (project “PREMISE”: Open and Reproducible Materials Science Research). JY and GP acknowledge financial support by the Swiss National Science Foundation (SNSF) Project Funding (grant 200021E_206190 “FISH4DIET”). JY, XW, and GP acknowledge financial support from the SwissTwins project, funded by the Swiss State Secretariat for Education, Research and Innovation (SERI). JY, NM, and GP acknowledge financial support from the MARKETPLACE project funded by Horizon 2020 under the H2020-NMBP-25-2017 call (Grant No. 760173). DH acknowledges financial support by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 803718 (SINDAM) and the UK Research and Innovation (UKRI) EPSRC grants ref. EP/X026973/1 and EP/V026690/1. CB acknowledges financial support by the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 957189 (BIG-MAP) and 957213 (BATTERY2030PLUS). SH and LFV acknowledge financial support by the Open Research Data Program of the ETH Board (project “ExAiRIM”: Explore AiiDA for Regional Inverse modeling of Greenhouse Gases). We acknowledge access to Piz Daint and Alps at the Swiss National Supercomputing Centre (CSCS), Switzerland, under MARVEL’s share with the IDs mr32 and mr33.

Notes and references

- 1 T. Weinzierl, in *The Pillars of Science*, Springer International Publishing, Cham, 2021, pp. 3–9.
- 2 A. Jain, S. P. Ong, W. Chen, B. Medasani, X. Qu, M. Kocher, M. Brafman, G. Petretto, G.-M. Rignanese, G. Hautier, D. Gunter and K. A. Persson, *Concurrency and Computation: Practice and Experience*, 2015, vol. 27, pp. 5037–5059.
- 3 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, *Comput. Mater. Sci.*, 2012, **58**, 218–226.
- 4 C. S. Adorf, P. M. Dodd, V. Ramasubramani and S. C. Glotzer, *Comput. Mater. Sci.*, 2018, **146**, 220–229.
- 5 J. Goecks, A. Nekrutenko, J. Taylor and T. G. Team, *Genome Biol.*, 2010, **11**, R86.
- 6 J. Janssen, S. Surendralal, Y. Lysogorskiy, M. Todorova, T. Hickel, R. Drautz and J. Neugebauer, *Comput. Mater. Sci.*, 2019, **163**, 24–36.
- 7 A. S. Rosen, M. Gallant, J. George, J. Riebesell, H. Sahasrabudde, J.-X. Shen, M. Wen, M. L. Evans, G. Petretto, D. Waroquiers, G.-M. Rignanese, K. A. Persson, A. Jain and A. M. Ganose, *Open Source Software*, 2024, vol. 9, p. 5995.
- 8 P. Amstutz, M. Mikheev, M. R. Crusoe, N. Tijanić and S. Lampa *et al.*, *Existing Workflow Systems*, 2025, <https://github.com/common-workflow-language/common-workflow-language/wiki/Existing-Workflow-systems>, accessed: 2025-12-08.
- 9 F. Suter, T. Coleman, I. Altıntaş, R. M. Badia, B. Balis, K. Chard, I. Colonnelli, E. Deelman, P. Di Tommaso, T. Fahringer, C. Goble, S. Jha, D. S. Katz, J. Köster, U. Leser, K. Mehta, H. Oliver, J.-L. Peterson, G. Pizzi, L. Pottier, R. Sirvent, E. Suchyta, D. Thain, S. R. Wilkinson, J. M. Wozniak and R. Ferreira da Silva, *Future Gener. Comput. Syst.*, 2026, **174**, 107974.
- 10 G. Pizzi, A. Cepellotti, R. Sabatini, N. Marzari and B. Kozinsky, *Comput. Mater. Sci.*, 2016, **111**, 218–230.
- 11 S. P. Huber, S. Zoupanos, M. Uhrin, L. Talirz, L. Kahle, R. Häuselmann, D. Gresch, T. Müller, A. V. Yakutovich, C. W. Andersen, F. F. Ramirez, C. S. Adorf, F. Gargiulo, S. Kumbhar, E. Passaro, C. Johnston, A. Merkys, A. Cepellotti, N. Mounet, N. Marzari, B. Kozinsky and G. Pizzi, *Sci. Data*, 2020, **7**, 300.
- 12 M. Uhrin, S. P. Huber, J. Yu, N. Marzari and G. Pizzi, *Comput. Mater. Sci.*, 2021, **187**, 110086.
- 13 D. A. Norman, *The Design of Everyday Things: Revised and Expanded Edition*, Basic Books, New York, 2013.
- 14 B. Shneiderman, C. Plaisant, M. Cohen, S. Jacobs, N. Elmquist and N. Diakopoulos, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, Pearson, Boston, 6th edn, 2016.
- 15 A. V. Yakutovich, K. Eimre, O. Schütt, L. Talirz, C. S. Adorf, C. W. Andersen, E. Ditley, D. Du, D. Passerone, B. Smit, N. Marzari, G. Pizzi and C. A. Pignedoli, *Comput. Mater. Sci.*, 2021, **188**, 110165.
- 16 AiiDALab Team, *AiiDALab website*, <https://www.aiidalab.net>, accessed 2026-04-27.
- 17 O. Bentaleb, A. S. Z. Belloum, A. Sebaa and A. El-Maouhab, *J. Supercomput.*, 2022, **78**, 1144–1181.
- 18 B. E. Granger and F. Pérez, *Comput. Sci. Eng.*, 2021, **23**, 7–14.
- 19 K. Eimre, J. I. Urgel, H. Hayashi, M. Di Giovannantonio, P. Ruffieux, S. Sato, S. Otomo, Y. S. Chan, N. Aratani, D. Passerone, O. Gröning, H. Yamada, R. Fasel and C. A. Pignedoli, *Nat. Commun.*, 2022, **13**, 511.
- 20 A. Kinikar, M. Di Giovannantonio, J. I. Urgel, K. Eimre, Z. Qiu, Y. Gu, E. Jin, A. Narita, X.-Y. Wang, K. Müllen, P. Ruffieux, C. A. Pignedoli and R. Fasel, *Nat. Synth.*, 2022, **1**, 289–296.
- 21 A. Kinikar, X.-Y. Wang, M. Di Giovannantonio, J. I. Urgel, P. Liu, K. Eimre, C. A. Pignedoli, S. Stolz, M. Bommert, S. Mishra, Q. Sun, R. Widmer, Z. Qiu, A. Narita, K. Müllen, P. Ruffieux and R. Fasel, *ACS Nanosci. Au*, 2024, **4**, 128–135.
- 22 M. Di Giovannantonio, Z. Qiu, C. A. Pignedoli, S. Asako, P. Ruffieux, K. Müllen, A. Narita and R. Fasel, *Nat. Commun.*, 2024, **15**, 1910.
- 23 C. Zhao, D. D. Bhagwandin, W. Xu, P. Ruffieux, S. I. Khan, C. A. Pignedoli, R. Fasel and Y. Rubin, *J. Am. Chem. Soc.*, 2024, **146**, 2474–2483.
- 24 P. Kraus, E. Bainglass, F. F. Ramirez, E. Svaluto-Ferro, L. Ercole, B. Kunz, S. P. Huber, N. Plainpan, N. Marzari,



- C. Battaglia and G. Pizzi, *J. Mater. Chem. A*, 2024, **12**, 10773–10783.
- 25 X. Xu, A. Kinikar, M. Di Giovannantonio, C. A. Pignedoli, P. Ruffieux, K. Müllen, R. Fasel and A. Narita, *Precis. Chem.*, 2024, **2**, 81–87.
- 26 O. Deniz, C. Sánchez-Sánchez, Q. Chen, P. Shinde, C. A. Pignedoli, J. Liu, X. Feng, A. Narita, K. Müllen, R. Fasel, D. Passerone and P. Ruffieux, *Carbon*, 2025, **243**, 120610.
- 27 S. E. Ammerman, V. Jelic, Y. Wei, V. N. Breslin, M. Hassan, N. Everett, S. Lee, Q. Sun, C. A. Pignedoli, P. Ruffieux, R. Fasel and T. L. Cocker, *Nat. Commun.*, 2021, **12**, 6794.
- 28 S. Mishra, G. Catarina, F. Wu, R. Ortiz, D. Jacob, K. Eimre, J. Ma, C. A. Pignedoli, X. Feng, P. Ruffieux, J. Fernández-Rossier and R. Fasel, *Nature*, 2021, **598**, 287–292.
- 29 G. Catarina, E. Turco, N. Krane, M. Bommert, A. Ortega-Guerrero, O. Gröning, P. Ruffieux, R. Fasel and C. A. Pignedoli, *Nano Lett.*, 2024, **24**, 12536–12544.
- 30 N. Bassi, X. Xu, F. Xiang, N. Krane, C. A. Pignedoli, A. Narita, R. Fasel and P. Ruffieux, *Commun. Chem.*, 2024, **7**, 274.
- 31 F. Xiang, Y. Gu, A. Kinikar, N. Bassi, A. Ortega-Guerrero, Z. Qiu, O. Gröning, P. Ruffieux, C. A. Pignedoli, K. Müllen and R. Fasel, *Nat. Chem.*, 2025, **17**, 1356–1363.
- 32 F. Xiang, A. Kinikar, M. Mühlinghaus, N. Bassi, C. A. Pignedoli, K. Müllen, R. Fasel and P. Ruffieux, *J. Am. Chem. Soc.*, 2025, **147**, 4930–4936.
- 33 L. Cai, T. Gao and A. T. S. Wee, *Nat. Commun.*, 2024, **15**, 3235.
- 34 N. Merino-Diez, R. Amador, S. T. Stolz, D. Passerone, R. Widmer and O. Gröning, *Adv. Sci.*, 2024, **11**, 2309081.
- 35 R. S. Nielsen, A. L. Álvarez, Y. Tømm, G. Gurieva, A. Ortega-Guerrero, J. Breternitz, L. Bastonero, N. Marzari, C. A. Pignedoli, S. Schorr and M. Dimitrievska, *Adv. Opt. Mater.*, 2025, **13**, e00915.
- 36 M. Bommert, B. Schuler, C. A. Pignedoli, R. Widmer and O. Gröning, *Carbon*, 2024, **216**, 118592.
- 37 F. Paschke, L.-A. Lieske, F. Albrecht, C. J. Chen, J. Repp and L. Gross, *ACS Nano*, 2025, **19**, 2641–2650.
- 38 C. Zhao, Q. Huang, L. c. v. Valenta, K. Eimre, L. Yang, A. V. Yakutovich, W. Xu, J. Ma, X. Feng, M. Juriček, R. Fasel, P. Ruffieux and C. A. Pignedoli, *Phys. Rev. Lett.*, 2024, **132**, 046201.
- 39 A. Culich, C. Willing, C. Holdgraf, E. Sundell, R. Lovett, Y. Panda and L. Goderre, *Zero to JupyterHub with Kubernetes*, <https://github.com/jupyterhub/zero-to-jupyterhub-k8s>, 2025, accessed: 2025-12-08.
- 40 G. de Miranda Nascimento, F. J. dos Santos, M. Bercx, D. Grassano, G. Pizzi and N. Marzari, Accurate and efficient protocols for high-throughput first-principles materials simulationss, *npj Comput. Mater.*, 2026, DOI: [10.1038/s41524-026-02097-8](https://doi.org/10.1038/s41524-026-02097-8).
- 41 S. P. Huber, M. Minotakis, M. Bercx, T. Reents, K. Eimre, N. Paulish, N. Hörmann, M. Uhrin, N. Marzari and G. Pizzi, *Digital Discovery*, 2026, **5**, 1114–1131.
- 42 AiiDALab Team, *AiiDALab Widget Base*, GitHub repository: <https://github.com/aiidalab/aiidalab-widgets-base>, 2025, accessed: 2025-12-08.
- 43 K. Eimre, A. V. Yakutovich and C. A. Pignedoli, *AiiDALab on surface chemistry*, <https://github.com/nanotech-empaiidalab-empia-surfaces.git>, 2025, accessed: 2025-12-08.
- 44 T. D. Kühne, M. Iannuzzi, M. Del Ben, V. V. Rybkin, P. Seewald, F. Stein, T. Laino, R. Z. Khaliullin, O. Schütt, F. Schiffmann, D. Golze, J. Wilhelm, S. Chulkov, M. H. Bani-Hashemian, V. Weber, U. Borštnik, M. Taillefumier, A. S. Jakobovits, A. Lazzaro, H. Pabst, T. Müller, R. Schade, M. Guidon, S. Andermatt, N. Holmberg, G. K. Schenter, A. Hehn, A. Bussy, F. Belleflamme, G. Tabacchi, A. Glöß, M. Lass, I. Bethune, C. J. Mundy, C. Plessl, M. Watkins, J. VandeVondele, M. Krack and J. Hutter, *J. Chem. Phys.*, 2020, **152**, 194103.
- 45 N. Oinonen, A. V. Yakutovich, A. Gallardo, M. Ondráček, P. Hapala and O. Krejčí, *Comput. Phys. Commun.*, 2024, **305**, 109341.
- 46 X. Wang, E. Bainglass, M. Bonacci, A. Ortega-Guerrero, L. Bastonero, M. Bercx, P. Bonfà, R. De Renzi, D. Du, P. N. O. Gillespie, M. A. Hernández-Bertrán, D. Hollas, S. P. Huber, E. Molinari, I. J. Onuorah, N. Paulish, D. Prezzi, J. Qiao, T. Reents, C. J. Sewell, I. Timrov, A. V. Yakutovich, J. Yu, N. Marzari, C. A. Pignedoli and G. Pizzi, *npj Comput. Mater.*, 2026, **12**, 72.
- 47 M. D. Wilkinson, M. Dumontier, I. J. Aalbersberg, G. Appleton, M. Axton, A. Baak, N. Blomberg, J.-W. Boiten, L. B. da Silva Santos, P. E. Bourne, *et al.*, *Sci. Data*, 2016, **3**, 1–9.
- 48 IPCC, *Refinement to the 2006 IPCC Guidelines for National Greenhouse Gas Inventories*, 2019, <https://www.ipcc-nggip.iges.or.jp/public/2019rf/index.html>, accessed: 2025-12-08.
- 49 A. Ganesan, A. Manning, O. Barrigar, H. Bovensmann, G. Brailsford, D. Brunner, B. Bukosa, D. Chan, E. Chan, Y. Chen, P. DeCola, B. Dunse, B. Ellerhoff, L. Emmenegger, A. Fix, A. Geddes, C. Gerbig, S. Henne, J. Herner, J. Hong, S. Houweling, L. Hu, H. Imhof, M. Ishizawa, A. Kaiser-Weiss, E. Keller, D. Kennett, R. Kiese, J. Kim, S. Kim, P. Krummel, D. Kubistin, S. E. Mikaloff Fletcher, C. Mielke, J. Miller, S. Moore, J. Müller-Williams, M. Neish, S. Nichol, S. O'Doherty, J. Pitt, A. Redington, S. Reimann, C. Scheer, D. Smale, M. Steinbacher, S. Smyth, K. Stanley, C. Trudinger, A. Vijayan, F. Vogel, M. K. Vollmer, A. Wenger, D. Worthy and D. Young, *Integrated Global Greenhouse Gas Information System Good Practice Guidance for Estimating National-scale Greenhouse Gas Emissions using Atmospheric Observations*, 2025, <https://library.wmo.int/idurl/4/69714>, accessed: 2025-12-08.
- 50 S. Park, L. M. Western, T. Saito, A. L. Redington, S. Henne, X. Fang, R. G. Prinn, A. J. Manning, S. A. Montzka, P. J. Fraser, A. L. Ganesan, C. M. Harth, J. Kim, P. B. Krummel, Q. Liang, J. Mühle, S. O'Doherty, H. Park, M.-K. Park, S. Reimann, P. K. Salameh, R. F. Weiss and M. Rigby, *Nature*, 2021, **590**, 433–437.
- 51 D. Rust, M. K. Vollmer, S. Henne, A. Frumau, P. van den Bulk, A. Hensen, K. M. Stanley, R. Zenobi, L. Emmenegger and S. Reimann, *Nature*, 2024, **633**, 96–100.



- 52 A. Stohl, C. Forster, A. Frank, P. Seibert and G. Wotawa, *Atmos. Chem. Phys.*, 2005, **5**, 2461–2474.
- 53 Atmospheric Modeling and Remote Sensing, *Empa Group, AiiDA-FLEXPART*, GitHub repository: <https://github.com/aiidaplugins/aiida-flexpart>, 2025, accessed: 2025-12-08.
- 54 S. Henne, D. Brunner, B. Oney, M. Leuenberger, W. Eugster, I. Bamberger, F. Meinhardt, M. Steinbacher and L. Emmenegger, *Atmos. Chem. Phys.*, 2016, **16**, 3683–3710.
- 55 FOEN, Switzerland's greenhouse gas inventory 1990–2022. National Inventory Document. Submission of April 2024 under the United Nations Framework Convention on Climate Change and under the Kyoto, *Protocol*, 2024, <https://unfccc.int/documents/637871>.
- 56 *C2SM, AiiDALab-FLEXPART*, GitHub repository, <https://github.com/C2SM/aiidalab-flexpart>, 2025, accessed: 2025-12-08.
- 57 B. Ervens, A. Rickard, B. Aumont, W. P. L. Carter, M. McGillen, A. Mellouki, J. Orlando, B. Picquet-Varrault, P. Seakins, W. R. Stockwell, L. Vereecken and T. J. Wallington, *Atmos. Chem. Phys.*, 2024, **24**, 13317–13339.
- 58 C. Mapelli, J. K. Donnelly, U. E. Hogan, A. R. Rickard, A. T. Robinson, F. Byrne, C. R. McElroy, B. F. E. Curchod, D. Hollas and T. J. Dillon, *Atmos. Chem. Phys.*, 2023, **23**, 7767–7779.
- 59 M. E. Jenkin, J. C. Young and A. R. Rickard, *Atmos. Chem. Phys.*, 2015, **15**, 11433–11459.
- 60 B. F. E. Curchod and A. J. Orr-Ewing, *J. Phys. Chem. A*, 2024, **128**, 6613–6635.
- 61 A. Prlj, E. Marsili, L. Hutton, D. Hollas, D. Shchepanovska, D. R. Glowacki, P. Slavíček and B. F. E. Curchod, *ACS Earth Space Chem.*, 2022, **6**, 207–217.
- 62 D. Weininger, *J. Chem. Inf. Comput. Sci.*, 1988, **28**, 31–36.
- 63 F. Neese, *Wiley Interdiscip. Rev. Comput. Mol. Sci.*, 2012, **2**, 73–78.
- 64 P. Zarabadi-Poor, D. Hollas and S. Huber, *pzarabadip/aiida-orca: Release v0.7.0*, Zenodo, 2023, DOI: [10.5281/zenodo.7955223](https://doi.org/10.5281/zenodo.7955223).
- 65 D. Hollas and B. F. E. Curchod, *J. Phys. Chem. A*, 2024, **128**, 8580–8590.
- 66 E. Svaluto-Ferro, G. Kimbell, Y. Kim, N. Plainpan, B. Kunz, L. Scholz, R. Läubli, M. Becker, D. Reber, P. Kraus, R.-S. Kühnel and C. Battaglia, *Batteries Supercaps*, 2025, **8**, e202500155.
- 67 *BIG-MAP website*, <https://www.big-map.eu/>, accessed 2026-04-27.
- 68 Empa, *Materials for Energy Conversion Laboratory, aiida-aurora*, GitHub repository: <https://github.com/EmpaConversion/aiida-aurora>, accessed 2026-04-27.
- 69 P. Kraus, A. Gbocho, L. Ercole and G. Kimbell, *Tomato: automation without pain!*, GitHub repository: <https://github.com/EmpaConversion/tomato>, 2025, accessed: 2025-12-08.
- 70 Empa, *Materials for Energy Conversion Laboratory, aiidalab-aurora*, GitHub repository: <https://github.com/EmpaConversion/aiidalab-aurora>, 2025, accessed: 2025-12-08.
- 71 P. Kraus, N. Vetsch and C. Terboven, *YADG: yet another datagram*, GitHub repository: <https://github.com/EmpaConversion/yadg>, 2025, accessed: 2025-12-08.
- 72 P. Kraus, N. Vetsch and C. Battaglia, *Open Source Software*, 2022, **7**, 4166.
- 73 *fastnda*, <https://github.com/EmpaConversion/fastnda>, accessed 2026-04-27.
- 74 *aurora-neware*, <https://github.com/EmpaConversion/aurora-neware>, accessed 2026-04-27.
- 75 *aurora-biologic*, <https://github.com/EmpaConversion/aurora-biologic>, accessed 2026-04-27.
- 76 A. Bauch, I. Adamczyk, P. Buczek, F.-J. Elmer, K. Enimanev, P. Glyzowski, M. Kohler, T. Pylak, A. Quandt, C. Ramakrishnan, C. Beisel, L. Malmström, R. Aebersold and B. Rinn, *BMC Bioinform.*, 2011, **12**, 468.
- 77 *openBIS Team, pyBIS*, <https://pypi.org/project/pybis/>, accessed 2026-04-27.
- 78 *AiiDALab Team, AiiDALab-ELN*, GitHub repository, <https://github.com/aiidalab/aiidalab-eln>, 2025, accessed: 2025-12-08.
- 79 PREMISE organization, *PREMISE website*, <https://ord-premise.org/>, accessed 2026-04-27.
- 80 MADICES organization, *MADICES website*, <https://madices.github.io/>.
- 81 J. Lass, H. Jacobsen, K. M. L. Krighaar, D. Graf, F. Groitl, F. Herzog, M. Yamada, C. Kägi, R. A. Müller, R. Bürge, M. Schild, M. S. Lehmann, A. Bollhalder, P. Keller, M. Bartkowiak, U. Filges, U. Greuter, G. Theidel, H. M. Rønnow, C. Niedermayer and D. G. Mazzone, *Rev. Sci. Instrum.*, 2023, **94**, 023302.
- 82 J. Lass, *MJOLNIR: v1.3.3*, Zenodo, doi: DOI: [10.5281/zenodo.13772982](https://doi.org/10.5281/zenodo.13772982).
- 83 M. Bonacci, G. Pizzi, J. Lass and D. G. Mazzone, *LNS-app*, <https://github.com/mikibonacci/LNS-app>, 2025, accessed: 2025-12-08.
- 84 A. Kaestner, S. Hartmann, G. Kühne, G. Frei, C. Grünzweig, L. Josic, F. Schmid and E. Lehmann, *Nucl. Instrum. Methods Phys. Res., Sect. A*, 2011, **659**, 387–393.
- 85 P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. Buongiorno Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A. Dal Corso, S. de Gironcoli, P. Delugas, R. A. DiStasio Jr, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo, R. Gebauer, U. Gerstmann, F. Giustino, T. Gorni, J. Jia, M. Kawamura, H.-Y. Ko, A. Kokalj, E. Küçükbenli, M. Lazzeri, M. Marsili, N. Marzari, F. Mauri, N. L. Nguyen, H.-V. Nguyen, A. Otero-de-la-Roza, L. Paulatto, S. Poncé, D. Rocca, R. Sabatini, B. Santra, M. Schlipf, A. P. Seitsonen, A. Smogunov, I. Timrov, T. Thonhauser, P. Umari, N. Vast, X. Wu and S. Baroni, *J. Phys.: Condens. Matter*, 2017, **29**, 465901.
- 86 A. P. Thompson, H. M. Aktulga, R. Berger, D. S. Bolintineanu, W. M. Brown, P. S. Crozier, P. J. In't Veld, A. Kohlmeyer, S. G. Moore, T. D. Nguyen, et al., *Comput. Phys. Commun.*, 2022, **271**, 108171.
- 87 A. Goscinski, T. J. Baird, D. Du, J. Prado, D. Suman, T.-E. Sodjargal, S. Bonella, G. Pizzi and M. Ceriotti,



- scicode-widgets: Bringing Computational Experiments to the Classroom with Jupyter Widgets, *arXiv*, 2025, preprint, arXiv:2507.05734, DOI: [10.48550/arXiv.2507.05734](https://doi.org/10.48550/arXiv.2507.05734), <https://arxiv.org/abs/2507.05734>.
- 88 CECAM organization, *CECAM website*, <https://www.cecama.org/>, accessed 2026-04-27.
- 89 Psi-k organization, *Psi-k website*, <https://psi-k.net/>, accessed 2026-04-27.
- 90 N. Marzari and G. Pizzi, *Electronic-Structure Simulations for User Communities at Large-Scale Facilities*, <https://indico.psi.ch/event/17436/>, 2025, accessed: 2025-12-08.
- 91 AiiDALab Team, *AiiDALab Launch*, *GitHub repository*, <https://github.com/aiidalab/aiidalab-launch>, 2025, accessed: 2025-12-08.
- 92 D. Merkel, *Linux J.*, 2014, **2014**, 2.
- 93 Docker Inc., *Docker installation*, <https://docs.docker.com/get-started/get-docker/>, accessed 2026-04-27.
- 94 AiiDALab Team, *AiiDALab Demo Server*, <https://demo.aiidalab.io>, 2025, accessed: 2025-12-08.

