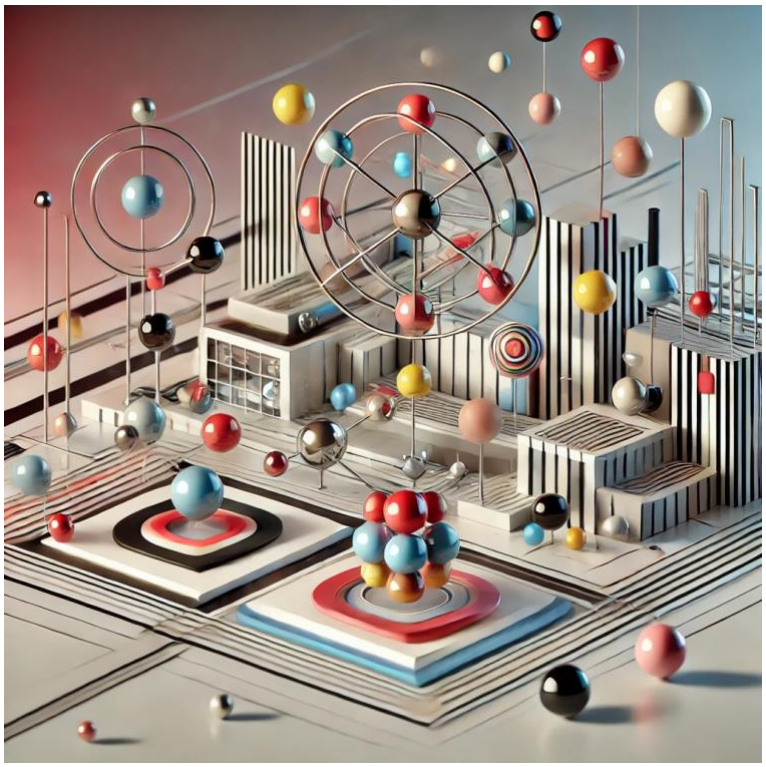
The atomic simulation environment ecosystem: Present and perspectives



June 23 - 27, 2025
CECAM-HQ-EPFL, BCH2103, Lausanne, Switzerland

Michele Ceriotti
EPFL, Switzerland

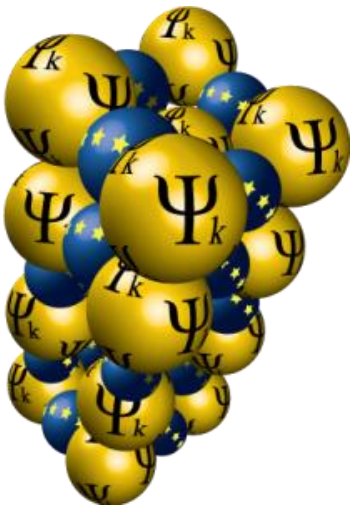
Alin Elena
Science and Technology Facilities Council, United Kingdom

James Kermode
University of Warwick, United Kingdom

Ask Hjorth Larsen
Technical University of Denmark, Denmark

Karsten Wedel Jacobsen
Technical University of Denmark, Denmark

With the support of:



Psi-k

1. Description

Atomistic simulations are increasingly widespread in modern condensed matter physics, computational chemistry and materials science. They can provide insight into experimental data, elucidate mechanisms underlying real-world processes and even help to design new materials with improved properties. A wide range of approaches from quantum chemistry to machine learning interatomic potentials (MLIPs) are used, with correspondingly wide-ranging software implementations. The diversity of software packages leads to different levels of user-friendliness, interoperability, and ease of automation.

It can be challenging to trace the origin of a calculation or ensure that the results used in a study are consistent, and it can be costly to move from one research group to another with an incompatible set of in-house tools. Without shared collaborative tools, high-level algorithms such as geometry optimisation or transition state search which are in principle code independent may instead be implemented in one particular package, leaving the rest of the community to create in-house variants for their own work. This inefficiency results both in needless duplication of effort and code and the increased potential for bugs.

The Atomic Simulation Environment (ASE) is a community-driven Python package that mitigates the N^2 problem of maintaining pairwise interfaces between codes by providing standard data structures – principally for atomic structures (the Atoms object) and calculation methods (the Calculator object) – as well as interfaces to ca. 100 file and ca. 30 simulation codes, acting as useful "glue" for work spanning multiple packages. A 2017 paper describing ASE has attracted over 500 citations every year for the past 5 years, demonstrating the broad adoption of ASE [1].

ASE integrates methods from ab-initio and classical molecular dynamics codes. In addition, it can manipulate structures and run calculations, providing a range of generic dynamics and geometry-optimisation routines with a toolkit for the further development of these methods. New high-level algorithms are written once and immediately available to users of established atomistic codes. ASE is increasingly used as a *lingua Franca* for MLIP fitting and evaluation: for example, the GAP, ACE, MACE, CHGNET, M3GNET, ALIGNFF and NEQUIP codes all use ASE to name but few, and the OpenKIM knowledge base of interatomic models is strongly interfaced to ASE. The existence of ASE frees the developers of new packages to focus on novel aspects, making use of existing tools for structure manipulation and dynamics.

We are proposing to organise a CECAM workshop on ASE and its broader community of dependent packages. We think this will be a good opportunity to bring together developers and users of core ASE and other packages in its ecosystem. This workshop is timely because ASE has grown rapidly in recent years, meaning that maintaining it is becoming a challenge for the core development team. At the same time, ASE has effectively become a critical piece of research infrastructure for an increasingly broad community of users and developers. An effective way ahead has been identified through the establishment, promotion and curation of an interoperable ecosystem of packages rather than a monolithic code base. 24 such packages have already been identified; their authors are represented amongst our proposed invited speakers. It is now critical to assemble a team of developers and end users to consider how to make this transition smoothly and sustainably.

Challenges to be addressed at the workshop will include:

- What are the dominant current use cases of ASE? How will this change in the next 5 years?
- What functionality should the core ASE include?
- Are the current Atoms and Calculator interfaces sufficient, or do they need to be generalized? How should they be serialized to/from files?
- What can be moved from the core code to plugins?
- How can the promotion of a distributed ecosystem of plugin file-formats and calculators be catalyzed?
- How can the development of the ASE ecosystem remain sustainable?
- Does a community governance model for ASE need to be established?

Key References

*[1] A. Hjorth Larsen, J. Jørgen Mortensen, J. Blomqvist, I. Castelli, R. Christensen, M. Dułak, J. Friis, M. Groves, B. Hammer, C. Hargus, E. Hermes, P. Jennings, P. Bjerre Jensen, J. Kermode, J. Kitchin, E. Leonhard Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. Bergmann Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, K. Jacobsen, J. Phys.: Condens. Matter, **29**, 273002 (2017)*

2. Program

Day 1 - Monday June 23rd 2025

- 13:00 to 13:30 – Registration
- 13:30 to 13:45 - Welcome & Introduction

Use-case talks session 1 - Methods

- 13:45 to 14:15 - **Bjørk Hammer**
  Atomic global optimization X (AGOX)
- 14:15 to 14:30 – Discussion
- 14:30 to 15:00 - **Lars Pastewka**
  matscipy: materials science at the atomic scale with Python
- 15:00 to 15:15 – Discussion
- 15:15 to 15:45 - Coffee break
- 15:45 to 16:15 - **Elena Gelzinyte**
  The `wfl` Python toolkit for creating machine learning interatomic potentials and related atomistic simulation workflows.
- 16:15 to 16:30 – Discussion
- 16:30 to 17:00 - **Xing Wang**
  Integrating ASE into reproducible and automated workflows with AiiDA
- 17:00 to 17:15 – Discussion
- 17:30 to 19:30 - Poster session & aperitif

Day 2 - Tuesday June 24th 2025

Use-case talks session 2 - Materials

- 09:00 to 09:30 - **Gabor Csanyi**
  Title TBC
- 09:30 to 09:45 – Discussion
- 09:45 to 10:15 - **Adam Jackson**
  Big things are coming slowly
- 10:15 to 10:30 – Discussion
- 10:30 to 11:00 - Coffee break

- 11:00 to 11:30 - **Paul Erhart**
  Predicting scattering experiments from first principles: understanding dynamics via correlation functions

- 11:30 to 11:45 – Discussion

- 11:45 to 12:15 - **Yuji Ikeda**
  Achieving coupled-cluster accuracy for periodic van der Waals materials with ASE-aided Δ-machine-learning interatomic potentials

- 12:15 to 12:30 – Discussion

- 12:30 to 14:00 - Lunch

Use-case talks session 3- Reactions

- 14:00 to 14:30 - **Hanna Türk**
  Unlocking the potential of lithium thiophosphate: atomistic in- sights on its surface reactivity

- 14:30 to 14:45 – Discussion

- 14:45 to 15:15 - **Andrew Peterson**
  Atomistic simulations into electrochemistry

- 15:15 to 15:30 – Discussion

- 15:30 to 16:00 - Coffee break

- 16:00 to 16:30 - **Michael Herbst**
  The Julia atomistic modelling ecosystem: an outside perspective on ASE

- 16:30 to 16:45 – Discussion

- 16:45 to 17:15 - **Andrew Rosen**
  The quantum accelerator: accessible and scalable materials science workflows

- 17:15 to 17:30 – Discussion

- 17:30 to 18:00 - General discussion 1

- 18:00 to 19:30 - Informal discussions with snacks

Day 3 - Wednesday June 25th 2025

- 09:00 to 10:30 - Discussion: The Future Purpose of ASE

- 10:30 to 11:00 - Coffee break

- 11:00 to 12:30 - Challenge Formulation

- 12:30 to 14:00 – Lunch

- 14:00 to 16:00 - Design work in Breakout Groups
- 16:00 to 16:30 - Coffee break
- 16:30 to 17:30 - Design work in Breakout Groups
- 17:30 to 18:00 - Feedback from Breakout Groups
- 19:30 to 22:00 - Social dinner

Day 4 - Thursday June 26th 2025

- 09:00 to 10:30 - Design work in Breakout Groups
- 10:30 to 11:00 - Coffee break
- 11:00 to 12:30 - Design work in Breakout Groups
- 12:30 to 14:00 – Lunch
- 14:00 to 16:00 - Design work in Breakout Groups
- 16:00 to 16:30 - Coffee break
- 16:30 to 17:30 - Design work in Breakout Groups
- 17:30 to 18:00 - Feedback from Breakout Groups

Day 5 - Friday June 27th 2025

- 09:00 to 09:30 - Presentation from Group 1
- 09:30 to 09:45 – Discussion
- 09:45 to 10:15 - Presentation from Group 2
- 10:15 to 10:30 – Discussion
- 10:30 to 11:00 - Coffee break
- 11:00 to 11:30 - Presentation from Group 3
- 11:30 to 11:45 – Discussion
- 11:45 to 12:30 - Closing discussion and road-map formulation

3. Abstracts

Achieving coupled-cluster accuracy for periodic van der Waals materials with ASE-aided Δ-machine-learning interatomic potentials

Yuji Ikeda, Blazej Grabowski
University of Stuttgart, Germany

Coupled-cluster theory at the CCSD(T) level constitutes the benchmark for electronic-structure accuracy, providing a rigorously non-empirical description of van der Waals (vdW) interactions. Yet its N⁷ scaling and the scarcity of periodic implementations preclude routine application to extended systems. This work introduces an ASE-aided Δ-learning workflow in which a machine-learning interatomic potential (MLIP) is calibrated against a concise set of molecular CCSD(T) reference data. By learning the energy differences relative to a computationally inexpensive tight-binding model, the MLIP attains CCSD(T) fidelity with far fewer training configurations than direct energy fitting. During production runs, the Δ-ML correction is combined on-the-fly with the underlying tight-binding Hamiltonian through ASE's "SumCalculator", delivering coupled-cluster accuracy at tight-binding cost. The protocol is demonstrated for a prototypical covalent organic framework (COF)—a quasi-two-dimensional material whose covalently bonded layers are held together by vdW forces. The composite potential reproduces CCSD(T) energies within the chemical accuracy (1 kcal/mol) for systems containing a few tens of atoms and yields converged predictions of symmetry, inter-layer binding energies, and hydrogen adsorption enthalpies. These findings show that CCSD(T) quality can now be achieved for periodic vdW materials at a fraction of the customary computational expense, enabling routine high-level benchmarking and materials discovery within the ASE ecosystem.

Atomic global optimization X (AGOX)

Bjørk Hammer, Mads-Peter Verner Christiansen
Aarhus University, Denmark

We present the Python package AGOX [1], that extends ASE with various optimization algorithms, including Basin Hopping, the Bayesian method GOFEE, and Replica Exchange-type optimization. AGOX leverages the on-the-fly construction of a surrogate energy landscape, an MLIP, that is used for local relaxation of new structural candidates before their evaluation with DFT. The MLIP may take a foundational model as a prior [2]. Various descriptors are supported, such as global pair-distribution functions, SOAP, and complex embedding vectors from foundational graph neural networks [3]. It will be described how AGOX extends the Atoms class of ASE in order to store information such as descriptors.

*[1] M. Christiansen, N. Rønne, B. Hammer, The Journal of Chemical Physics, **157**, (2022)*
*[2] J. Pitfield, F. Brix, Z. Tang, A. Slavensky, N. Rønne, M. Christiansen, B. Hammer, Phys. Rev. Lett., **134**, 056201 (2025)*
*[3] M. Christiansen, B. Hammer, The Journal of Chemical Physics, **162**, (2025)*

Atomistic simulations into electrochemistry

Andrew Peterson
Brown University, United States

In this talk, I'll highlight how we employ atomistic simulations in simulations of electrified interfaces, including how open-source software in the ASE ecosystem has enabled us to study a myriad of phenomena, including strain effects, coverage, electrification, surrogate models, global optimization, solvent effects, novel barrier searches, and many related phenomena. I'll describe some key insights we have gained that allow us to have improved understandings of such systems, which are of tremendous importance in the energy transition.

Big things are coming slowly

Adam Jackson
STFC Rutherford Appleton Laboratory, United Kingdom

The current ASE codebase starts in 2007 and has over 500 contributors. Python 2 was dropped in 2019; a lot of the code predates some useful "modern" tools. The number of Calculators has expanded and some are periodically orphaned or adopted by contributors from the ASE community.
In this time certain "big ideas" have emerged and been played with, but prove difficult to apply. Here we look at a few changes that seem "stuck in the pipeline", and outline the current maintainer position and projects to progress them.
A plug-in system for Calculators (and more) solves several maintenance problems: with a more consistent experience across internal/external implementations, the barrier to use is lowered and becomes more attractive to extend ASE using external dependencies on an independent development cycle.
Recent years have seen great progress in tools to assure Python code quality: static linting and type annotations allow standards and correctness to be enforced without runtime overhead. However, there are some nuances to consider when applying such tools to a 200,000-line legacy codebase!
The strong coupling between Atoms and Calculators is a fundamental part of ASE's design that serves some workflows very well while obstructing others. We show some emerging ideas about how to loosen this coupling, but community discussion will be needed in order to decide the ultimate roadmap.

Integrating ASE into reproducible and automated workflows with AiiDA

Xing Wang
Paul Scherrer Institute, Switzerland

The Atomic Simulation Environment (ASE) plays a central role in scripting atomistic simulations, but integrating it into reproducible, automated workflows remains a challenge for many users. In this talk, I will introduce AiiDA, an open-source infrastructure for managing computational workflows with full data provenance [1]. I will highlight recent developments in the AiiDA ecosystem, including WorkGraph (a flexible graph-based workflow definition model) and aiida-pythonjob (a plugin designed to execute arbitrary Python functions remotely with provenance tracking), which is ideal for integrating ASE calculators into structured workflows.
I will also outline our ongoing efforts to extend FAIR principles beyond data to include workflows and analysis tools. If time allows, I will propose enhancements to ASE to improve its integration into workflow engines, such as support for restarts and structured exit codes.

[1] https://www.aiida.net/

Matscipy: materials science at the atomic scale with python

Lars Pastewka[1], Petr Grigorev[2], Lucas Frérot[3], Fraser Birks[4], Adrien Gola[5], Jacek Golebiowski[6], Jan Grießer[1], Johannes Hörmann[1], Andreas Klemenz[7], Gianpietro Moras[7], Wolfram Nöhring[1], Jonas Oldenstaedt[8], Punit Patel[4], Thomas Reichenbach[7], Thomas Rocke[4], Lakshmi Shenoy[4], Michael Walter[1], Simon Wengert[9], Lei Zhang[10], James Kermode[4]
[1]University of Freiburg, Germany
[2]INSA Lyon, France
[3]Sorbonne Université, France
[4]Warwick University, United Kingdom
[5]Karlsruhe Institute of Technology, Germany
[6]Imperial College London, United Kingdom
[7]Fraunhofer IWM, Germany
[8]Technical University of Munich, Germany
[9]Fritz Haber Institute of the Max Planck Society, Germany
[10]University of Groningen, Netherlands

Behaviour of materials is governed by physical phenomena that occur at an extreme range of length and time scales. Computational modelling requires multiscale approaches. Simulation techniques operating on the atomic scale serve as a foundation for such approaches, providing necessary parameters for upper-scale models. The

physical models employed for atomic simulations can vary from electronic structure calculations to empirical force fields. However, construction, manipulation and analysis of atomic systems are independent of the given physical model but dependent on the specific application. matscipy implements such tools for applications in materials science, including fracture, plasticity, tribology and electrochemistry. This talk will present architecture and functionality of matscipy, with a focus on data structures for fast neighbor lists.

[1] P. Grigorev, L. Frérot, F. Birks, A. Gola, J. Golebiowski, J. Grießer, J. Hörmann, A. Klemenz, G. Moras, W. Nöhring, J. Oldenstaedt, P. Patel, T. Reichenbach, T. Rocke, L. Shenoy, M. Walter, S. Wengert, L. Zhang, J. Kermode, L. Pastewka, JOSS., 9, 5668 (2024)

Predicting scattering experiments from first principles: understanding dynamics via correlation functions

Paul Erhart
Chalmers University of Technology, Sweden

Machine learning has emerged as a powerful tool in materials discovery, enabling the rapid design of novel materials with tailored properties for countless applications, including in the context of energy and sustainability. To ensure the reliability of these methods, however, rigorous validation against experimental data is essential. Scattering techniques –using neutrons, X-rays, or electrons– offer a direct way to probe atomic-scale structure and dynamics, making them ideal for this purpose. Here, I will describe a computational workflow that bridges machine learning– based simulations with experimental validation. The workflow combines density functional theory calculations, machine-learned interatomic potentials, molecular dynamics, and autocorrelation function analysis to simulate, e.g., diffraction and quasi-elastic scattering of neutrons, X-rays, and electrons but also infrared and Raman spectra. The good agreement between simulated and experimental results highlights the potential of this approach for guiding and interpreting experiments, while also pointing out areas for further improvement.

The `wfl` Python toolkit for creating machine learning interatomic potentials and related atomistic simulation workflows.

Elena Gelzinyte
Fritz-Haber Institut der Max Planck Gesellschaft, Germany

The Atomic Simulation Environment (ASE) offers a suite of foundational functions that can be assembled into workflows of varying complexity, supported by numerous packages. Conventionally, these packages have been tailored for computationally intensive ab initio calculations, primarily in materials science applications.
The requirements in more recently emerging applications, such as the development of machine learning interatomic potentials (MLIPs), operate in a different computational regime. For example, application of MLIPs may involve a larger number of less computationally demanding tasks, requiring a different parallelisation strategy to make the best use of high-performance computing resources. Additionally, the ease of adaptation and extension is crucial, allowing users to modify and iterate over workflows before scaling them up, rather than merely re-executing predefined workflows.
The `wfl` package for atomistic simulation and MLIP fitting workflow management addresses these evolving requirements. It is designed to support a wide range of ASE-derived workflows, built on a low-level, developer-oriented framework that enables the creation of high-level, user-friendly functionalities. In this talk, I will explain how the `wfl` package operates, emphasising key design choices and the ASE features it relies on. Practical examples will illustrate its application and benefits in real-world scenarios.

The Julia atomistic modelling ecosystem: an outside perspective on ASE

Michael Herbst
EPFL, Switzerland

Over the past years a number of atomistic modelling efforts have been started, which employ the Julia programming language. Very early one these efforts aggregated around the https://juliamolsim.org/ GitHub organisation, which by now features Julia codes to perform DFT simulation (https://dftk.org), build ACE-based interatomic potentials (https://acesuit.github.io/), or run molecular dynamics (https://juliamolsim.github.io/Molly.jl). From the start, integration of our software with efforts outside the Julia programming language has been an important consideration to us and the design of ASE in itself served as a source of inspiration for some of our own decisions.
Focusing on the density-functional toolkit developed in my group I will provide a quick overview of some of Julia's unique ability to lead to high-level, efficient code with added benefits such as differentiability. I will explain what consequences Julia's language design and features has when it comes to interfacing codes and to what extend this lead us to make deviating design decisions when comparing to ASE. I will further express opportunities I see from our perspective to more deeply integrate with ASE and what obstacles I see to make this happen.

The quantum accelerator: accessible and scalable materials science workflows

Andrew Rosen
Princeton University, United States

In this talk, I will briefly introduce the Quantum Accelerator (https://github.com/Quantum-Accelerators/quacc). Built around ASE, the Quantum Accelerator is a materials science workflow library with an interface to many different underlying workflow engines and numerous simulation codes. I will place specific emphasis on the critical role ASE plays in the Quantum Accelerator as well as areas where future enhancements to ASE could prove beneficial.

Unlocking the potential of lithium thiophosphate: atomistic in- sights on its surface reactivity

Hanna Türk, Davide Tisi, Michele Ceriotti
École polytechnique fédérale de Lausanne, Switzerland

Lithium thiophosphate (LPS) emerges as a promising solid electrolyte for the next-generation of lithium-ion batteries. However, the high reactivity of LPS with common contaminants such as atmospheric water, solvents used for preparation, and common electrode materials hinders commercialization of the cells. The unknown structure, morphology, and chemical behavior of LPS's surface prevents strategic solutions to this problem. Here, we employ a machine learning interatomic potential in form of an Atomic Simulation Environment (ASE) calculator to gain fundamental, atomistic understanding on the mechanical and chemical properties of an LPS surface in contact with water. Our focus lies first on the identification of relevant surface complexions formed by surface reconstructions, which define the form of the reactive sites and thus the degradation mechanism of the material with relevant contaminants. Metadynamics and nudged-elastic band simulations then reveal the complex reactivity of the LPS surface with water. This allows us to identify the surface moteties that play a relevant role in the material's degradation.

4. Participant list

Organizers

Ceriotti, Michele
EPFL, Switzerland
Elena, Alin
STFC, United Kingdom
Jacobsen, Karsten Wedel
Technical University of Denmark, Denmark
Kermode, James
University of Warwick, United Kingdom
Larsen, Ask Hjorth
Technical University of Denmark, Denmark

On-site participants

Azar, Said - Assistant Professor of Physics, Jordan
Berger, Esmée - Chalmers University of Technology, Sweden
Chehel Amirani, Morteza - Svante, Canada
Csanyi, Gabor - University of Cambridge, United Kingdom
Erhart, Paul - Chalmers University of Technology, Sweden
Fraux, Guillaume - EPFL, Switzerland
Gelzinyte, Elena - Fritz-Haber Institut der Max Planck Gesellschaft, Germany
Grigorev, Petr - CNRS, France
Hammer, Bjørk - Aarhus University, Denmark
Herbst, Michael - EPFL, Switzerland
Ikeda, Yuji - University of Stuttgart, Germany
Jackson, Adam - STFC, United Kingdom
Janke, Svenja - University of Warwick, United Kingdom
Kasoar, Elliott - STFC, United Kingdom
Keller, Elisabeth - University of Bayreuth, Germany
Langer, Marcel, EPFL, Switzerland
Li, Yutao - EPFL Valais Wallis, Switzerland
Lysgaard, Steen - Technical University of Denmark, Denmark
Maurya, Manoj Kumar - University of Freiburg, Germany
Pastewka, Lars - University of Freiburg, Germany
Peterson, Andrew - Brown University, United States
Qaisrani, Muhammad Nawaz - Ilmenau University of Technology, Germany
Tetenoire, Auguste - EPFL, Switzerland
Tsirkin, Stepan - EPFL, Switzerland
Türk, Hanna EPFL, Switzerland
Verner Christiansen, Mads-Peter - Aarhus University, Denmark
Wang, Xing - Paul Scherrer Institute, Switzerland
Wang, Yuandong - Fritz-Haber-Institut der Max-Planck-Gesellschaft, Germany

On-line participants

Arvelos, Graciele - Institute of Theoretical Physics, Brazil
Bayer, Kutand - İzmir Institute of Technology, Turkey
Bernart, Sarah - Technical University of Eindhoven, Netherlands
Egeolu, Ifeanyichukwu - Jan Evangelista Purkyně University, Czech Republic

Gong, Yilun - University of Oxford, United Kingdom
Gula, Igor - Continental GbmH, Germany
Heidari Pebdani, Zeinab - Helmholtz-Zentrum Hereon, Germany
Jena, Nityasagar - IFM, Linköping University, Sweden
Khusinov, Dmitriy - No, No affiliation
Liberati, Diego - National Research Council of Italy, Italy
Makhdoom, Madiha Masood - University of Padova, Italy
Rosen, Andrew - Princeton University, United States
Sarkar, Arka Prava - Iit Kharagpur, India
Verma, Rahul - NC State University, United States
Yadav, Saanya - Indian Institute of Technology, Hyderabad, India
Zhang, Ruiqi - Binghamton University, United States