Matter

CellPress

Perspective

Toward a composable, modular laboratory ecosystem for autonomous materials research and development

Howie Joress,[1,*] Brian DeCost,[1,2] Katelyn Jones,[1,2] A. Gilad Kusne,[1,2] Austin McDannald,[1,2] Zachary Trautt,[1,2] and Francesca Tavazza[1]

[1]Materials Measurement Science Division, National Institute of Standards and Technology (NIST), Gaithersburg, MD 20899, USA

[2]These authors contributed equally

*Correspondence: howie.joress@nist.gov

https://doi.org/10.1016/j.matt.2026.102756

PROGRESS AND POTENTIAL Autonomous research and development methods, which combine artificial intelligence and automated experimental apparatuses, have the potential to be a powerful tool in materials science and other physical science domains. However, the high level of investment currently needed to adopt these platforms is a major obstacle to their broad adoption. Creating modular components, both software and experimental hardware, which can be easily integrated into a single autonomous research platform, can greatly decrease the time and financial cost needed to engineer these systems. Developing a modular ecosystem requires standardizing the purpose of these components and how they interact with each other. In this paper, we describe a framework for modularizing the various tasks required for an autonomous experimental platform, along with discussing the various types of standards that need to be developed to allow for component interoperability.

SUMMARY

Autonomous and automated experimentation have demonstrated the possibility of accelerating materials research and development by orders of magnitude. However, the investment in both time and money to design and build these platforms has proven to be a major roadblock in their adoption. This cost is driven by the bespoke nature of these platforms, requiring redundant engineering due to the complexity of handling multiple interfaces from multiple vendors. To mitigate these challenges, we suggest the development of a set of community-driven interchange standards for both hardware and software, allowing these platforms to be rapidly constructed from off-the-shelf modular components which can each be designed and built by multiple vendors. This would enable the creation of an autonomous materials research ecosystem.

INTRODUCTION

Filling the materials technology gap

The foundational objective of materials science is to elucidate the relationships between the structure, processing, and properties of materials and how they ultimately determine a material's performance in engineering applications. The work of materials design is to harness the knowledge of those relationships in order to make materials with the needed properties and performance required for critical societal needs.

The space of possible material composition and processing, the resultant structures, and their respective properties is too large to explore and be optimized efficiently with traditional research methods. Most technologically relevant materials have effects across multiple lengthscales that are too intricate for effective theory-driven approaches. Conversely, brute-force trial-and-error approaches suffer from the "curse of dimensionality," the exponential explosion of possibilities with an increasing number of design parameters. Yet, these are still the default approaches to materials discovery and processing optimization problems in many industries. Autonomous laboratories, often referred to as self-driving laboratories (SDLs), combine AI and automated equipment to carry out experimental design, execution, and analysis in a closed loop to efficiently navigate the search space of these materials science challenges. Here, we describe our vision and approach for the creation of a modular SDL ecosystem that can be composed from a set of commercial off-the-shelf components. The overarching objective of the effort described here is to lower the barriers and required investment for the adoption of autonomous labs

for materials science in industrial as well as academic applications.

Autonomous R&D: A new paradigm

To meet the materials needs of the 21st century, it is necessary to change the way materials R&D is done. Research is typically categorized into 4 major paradigms of scientific study: (1) empirical science, (2) theory-driven science, (3) computational science (e.g., simulations), and (4) data-driven science. Various industries are at differing levels of embracing each paradigm. Integrated computational materials engineering (ICME)[1] efforts and the materials genome initiative (MGI)[2] have propelled industry forward in harnessing computational and data-driven methods over the past decade. The 4th paradigm leverages "big data" and AI in the form of machine learning (ML) to make predictions and interpolations from what is already known, often in cases that are too complex for current theory-based models to provide quantitative predictive accuracy.[3] However, many materials problems require traversing unexplored regions of parameter space, where conventional (4th paradigm) "big data" extrapolations will fail. Furthermore, many important materials problems (often because they exist in high-dimensional parameter spaces) are impossible to explore with any of the four existing paradigms of scientific discovery.

Currently, most R&D experiments are carried out using independent instruments with a large amount of manual interaction for operations such as sample synthesis, characterization, data analysis, and knowledge extraction (e.g., finding patterns and trends). SDLs carrying out autonomous experimentation (AE) offer a new operational paradigm with the goal of systematically generating the most informative data toward the end goal of the particular problem at hand, greatly increasing the efficiency of research budgets. AE leverages cutting-edge advances in automated experimentation and AI algorithms, along with theory-based predictive modeling, to rapidly acquiring new materials knowledge. This frees human scientists to focus on imbuing the system with intuition and leveraging their creativity to guide the system toward important work.

The core of an autonomous system is a tightly coupled feedback loop. This loop has been described in multiple places in the literature, including in the work by Stach et al.[4] To start the loop, an experimenter frames the problem with a set of objectives and constraints. Prior knowledge can also be provided to the platform, including databases of known materials and their properties, physicochemical heuristics, and physical laws. The AI-based tool then takes all of this information and uses it to generate a computationally inexpensive representative model of the system; the model can then be used to make predictions across the vast parameter space. Based on the objective and the model, a "decision-making agent" decides which new data will provide the most critical piece of information to increase its knowledge and reach its objective, typically some combinations of exploration and optimization. The automated system then performs an experiment or a theory-based simulation to produce that knowledge. The platform then analyzes, interprets, and stores the resulting data. For complex experimental systems, some form of orchestration is typically required to coordinate the various experimental actions. The scientific AI agent then looks at the updated information, make new predictions, and then decide on the next data point to collect. This loop is repeated until sufficient information is gathered to achieve its objective.

Benefits of autonomous methods

Autonomous methods have several advantages over traditional research methods. The use of active learning, along with the inclusion of theory-based modeling, can greatly increase the value of each experimental data point collected. Conversely, these methods can reduce the amount of low-value data that needs to be collected in order to generate the same knowledge. Particularly in fields like materials science, generating data can be fiscally expensive and resource and labor intensive. Furthermore, AI-based decision-making systems with well-characterized assumptions can reduce the unquantifiable effects of bias imparted by human experimenters, both on decision-making and analysis.

Collecting data at scales sufficient to meet societal needs requires agile experimental capabilities to dynamically explore disparate hypotheses. Robotic, automated platforms achieve this by quickly and accurately executing experiments on demand. These automated platforms bring several distinct advantages: First, data are generated with robotic control, which is generally more precise, traceable, repeatable, and systematic. Second, the digitally native nature of automated labs streamlines the collection and recording of crucial metadata about the experiment. Third, these laboratories can run without the need for downtime required by human lab workers. The time and cost savings associated with these platforms do not simply reflect a difference in degree, but in many cases, a difference in kind—making previously intractably complex problems tractable.

Autonomous methods can solve grand challenges in materials science

Autonomous physical science is still in its infancy. But as this field matures, it can grow to include not only platforms with multiple instruments of the same type but also platforms with multiple types of instruments and multiple synthesis techniques (potentially with different form factors). Adding new techniques and instruments to an ecosystem of tools allows for the scope of the questions to grow. Many materials R&D challenges involve combining and processing different kinds of raw materials in different ways, and there are more possible combinations than stars in the known universe. This near-infinite solution space is the underlying basis of many grand challenges in materials and manufacturing. There are some technical areas where AE platforms may be the only way to make materials design tractable. Some areas such as high entropy alloys—-alloys with multiple elements in large proportions which promise extraordinary properties for applications as varied as hypersonic coatings and water splitting catalysts—-exist in such a large compositional and processing design space that collecting informative data on these alloys through standard experimental approaches is insurmountable. Another place where autonomous methods are crucial is in domains where the problem is constantly mutating. Two examples of this are metal-based additive manufacturing (AM) and circular economy. In the former example, optimization

of AM processing parameters is a complex problem—each unique combination of part geometry, materials, and printing platform requires customized process optimization. Efficiently mapping generalizable relationships between processing, structure, and material properties is the only path for being able to design parts without *de novo* re-optimization for each unique part and process. Similarly, using recycled materials, as part of a circular economy, requires continually optimizing their processing to account for the ever-mutating input material stream, which can have strong impacts on the end-product's properties and performance. Understanding what effects are caused by the changing input stream, which materials can be made, and what adjustments to the processing are needed is an ongoing design challenge. In all of these cases, the materials physics is too complex to be modeled using currently available physics-based simulation approaches; thus, experimentation is a necessary driving force for innovation. Eventually, a suite of autonomous tools could begin to tackle truly open-ended engineering challenges and scientific questions. An AE platform that includes not only synthesis and characterization instruments but also theoretical simulations could begin to test boundaries of current scientific understanding, searching for new operant mechanisms to explain observed behaviors. What causes batteries to degrade, and how can we make them safer? Do room-temperature superconductors exist, and what is their underlying mechanism of operation? The realization of an AE platform capable of addressing this scope of open-ended challenges is predicated on a mature ecosystem containing a large variety of different types of hardware and software components, each with well-described trustworthiness and explainability and the ability to store and handle the required complex data.

Lack of standardized protocols: Major challenge to autonomous implementation

R&D labs that would like to implement an autonomous platform, and taking advantage of the benefits that AE can provide, face a variety of challenges and their associated costs. Many of these challenges have been highlighted in a recent US Federal Government-wide report.[5] Currently, there are no agreed-upon standards for how samples should be transferred between instruments nor standardized protocols for data sharing or instrument control. This leaves researchers with two options: seek an outside equipment "vendor" to engineer and build a bespoke instrument, typically at great cost, or take the time to cobble together instruments in house. Because there is no standardization, these platforms are typically designed for a single purpose, meaning the engineering work invested in them cannot typically be leveraged for future platforms. Further, once the objective of these single purpose platforms has been met or when the scope of inquiry shifts, the platform becomes obsolete, necessitating a major redesign. Similar issues exist for the software and data portion of SDLs. Various vendors produce pieces of software that perform various functions, but there is typically not a simple way to connect these disparate pieces with each other or with the physical instrumentation.

The current rarity of these systems has led to and is caused by a "chicken and egg problem" between vendors and end-users. Because of the large investment currently needed to procure or develop an autonomous system, few end-users are attempting to use one as part of their R&D efforts. As a result, vendors do not believe there is a need in the market to create hardware and software to support these types of systems and, therefore, are reluctant to make the investment to develop these products. In turn, this makes building an autonomous system more challenging and expensive.

A clear example of this was discussed in detail at the "October 2023 Autonomous Methodologies for Accelerating X-ray Measurements" workshop hosted by NIST and the International Center for Diffraction Data.[6] Several manufacturers of common X-ray diffractometers were present, and many described their reluctance to create interfaces necessary to enable autonomous control of their instruments. The clear message was that, at present, the level of customer demand for these features is not significant. The further concern was that the lack of a standard interface meant different customers request different interfaces, requiring repetitive custom engineering. A related risk for vendors is that there could be changes to software that interacts with their interfaces but is out of their control. Such changes are a liability for the vendors as they would adversely affect the functionality of their tool—a problem for which customers would demand fixes. Given these challenges, these equipment manufacturers decided that, for the moment, they would avoid offering this feature set.

The path forward

In order to radically transform the way materials R&D is done, we call for the development and demonstration of an ecosystem of materials R&D infrastructure that enables autonomous materials workflows. The availability of such an ecosystem will drastically reduce the cost and risk to companies developing and acquiring autonomous systems by increasing system modularity and interoperability. In turn, this will enable a reduction of duplicate engineering through assembly of off-the-shelf, ecosystem-enabled components (both hardware and software) available on the open market. Creating this ecosystem will necessitate establishing new documentary standards that will define it.

A centrally coordinated effort is necessary to generate critical mass sufficient to overcome the problems described above, inducing commercial interest in developing autonomous hardware and software, as well as reducing the duplication of efforts to ensure efficient use of research funding. This central coordination additionally facilitates the sharing of methods, tools, and best practices.

For inspiration, we can look at SEMI (Semiconductor Equipment and Materials International), a trade association of manufacturers and their supply chain partners across the semiconductor industry. Since 1973, a major function of SEMI has been the development and dissemination of technical standards, of which there are now over 1,000. Their standards, developed by representatives of over 2,000 companies, greatly simplify the supply chain and increase interoperability of both equipment and raw materials for semiconductor manufacturing, with the average semiconductor industry procurement citing to 25 SEMI standards.[7] These standards are wide ranging, including wafer sizes (the first SEMI standard), fabrication equipment communications and wafer interchange protocols, and safety requirements.

This article provides our prospective on an architecture for a modular SDL ecosystem and the standards needed to define it. We also call for the creation of community groups (e.g., consortia and working groups) to develop these standards. We further call for the creation of a center funded to develop, test, and demonstrate this ecosystem.

There are, of course, several efforts aimed at advancing autonomous platforms for materials science, each working to various degrees toward a scalable, modular ecosystem concept described in this work (this paragraph is by no means a complete listing). There are several large-scale, academic-driven efforts such as the Canadian Accelerate Consortium (https://acceleration.utoronto.ca/) and the American Carnegie Mellon University Cloud Laboratory. Although many tools and software products developed by these types of efforts are reusable, they are predominately built to be internally modular without aiming to be scalable or broadly adoptable. Further, while these efforts have some materials science focus, most of their scope is in the chemistry and bio-pharmaceutical space. The US Department of Energy-based INTERSECT program (https://www.ornl.gov/intersect) and the related Autonomous Interconnected Science Lab Ecosystem (AISLE)[8] have made efforts in developing broad ecosystems, though many of the physical endpoints are based on retrofitting existing tools, including user facilities. In Japan, the Laboratory Automation Suppliers' Association (LASA) has built a large, multifaceted organization focused on the hardware end of the platform, including the Measurement Analysis Instrument Markup Language (MaiML).[9] There are also a variety of efforts focused on developing standardized FAIR (findable, accessible, interoperable, and reusable) data storage,[10] including the Materials Research Data Alliance (MaRDA, https://www.marda-alliance.org/). Standardized data is critical to the development of an AE research ecosystem.

Before we describe our vision in detail, we define two types of stakeholder roles: end-users and vendors. End-users are entities interested in discovering and developing new materials that can be used to meet societal needs. Vendors are organizations that make R&D products to support the work of end-users. In both cases, stakeholders will primarily be the companies who fall into one category or the other but may also be academic groups, government labs, or non-profit labs, and these groups (and in rare cases, companies), in some cases, may act as both vendors and end-users.

It is also worth noting that in this article, we have chosen to focus on the materials science research domain because (1) it is the area we are most familiar with, (2) there is a vast array of types of materials and equipment for processing and characterizing them, (3) materials exist in vast design spaces that can benefit strongly from autonomous R&D, and (4) autonomous methods in materials are under-developed. In many ways, other fields such as biology (particularly pharmaceutical sector) and chemistry are ahead of materials science. These fields benefit from a much-reduced palate of instrument types, and most of the work in automated systems happens in a single-form factor, mL-scale liquids. Microplates (defined by ANSI/SLAS 1–2004 through ANSI/SLAS 4–2004) have been a standard sample holder in many fields of biology and chemistry going back decades,[11] enabling a range of automation to be built around them. Other standards on the software side have also been developed, such as Autoprotocol (http://autoprotocol.org/), for describing experimental workflows. That being said, much of this work focuses on achieving automation but not autonomy. In many cases, the materials community can look to existing standards in the biology and chemistry communities. While these standards can sometimes be adopted *as is*, we will often need to expand upon them or, perhaps in many cases, move beyond them altogether. It is likelier that software interfaces in the "science layer" (defined below) will be more readily usable than software and hardware interfaces related to instruments.

DISCUSSION

A proposed ecosystem

The first step in creating an ecosystem is to define a set of common modular subsystems, from which a wide variety of autonomous platform configurations can be built. Figure 1 illustrates our vision for these functional modules and their required interfaces (i.e., connections between modules), using a self-contained autonomous platform as an example. The possibility of building more complex and extensible systems from these modules is discussed below. Each module in the ecosystem, represented by words in the figure, is a piece of hardware or software and is connected to at least two other modules. In order to create a modular ecosystem, these connections, represented in the figure by arrows, must be standardized to allow various components to be swapped out. A component, a specific piece of software or hardware, could be deemed "ecosystem ready" if it is capable of 1) performing the function of a specific module (or modules) and 2) meeting the requirements of the various connections required of its function or role by the standard. The physical connection standards must define a set of dimensional and other mechanical descriptors that allows for seamless motion of samples across the platform. The virtual connection standards will describe certain information to be conveyed in a specific format between specific module types.

Within this ecosystem, the internal implementation of any particular module is abstracted from the rest of the ecosystem, meaning that each component can be arbitrarily simple or complex and may exhibit additional functionalities or connections not explicitly defined by the standard. A single component may also carry out the role of multiple modules under the constraint that, externally, it performs all the functions of the subsumed modules and has all the same inputs and outputs.

Starting at the top of the figure, we have 3 interconnected virtual resources, referred to collectively as the science layer, which make up the brain of the platform. The rest of the platform is referred to as the "data generation layer," which serves as the rest of the body to schedule and execute the tasks decided on by the science layer. We will walk through the various components of both layers.

Prediction module

The purpose of this module is to use all the information available to it, including platform-generated data, external databases, physical laws, and other materials science knowledge, to make predictions of the relevant properties of study across the range of input space (or some space requested by the decision-making

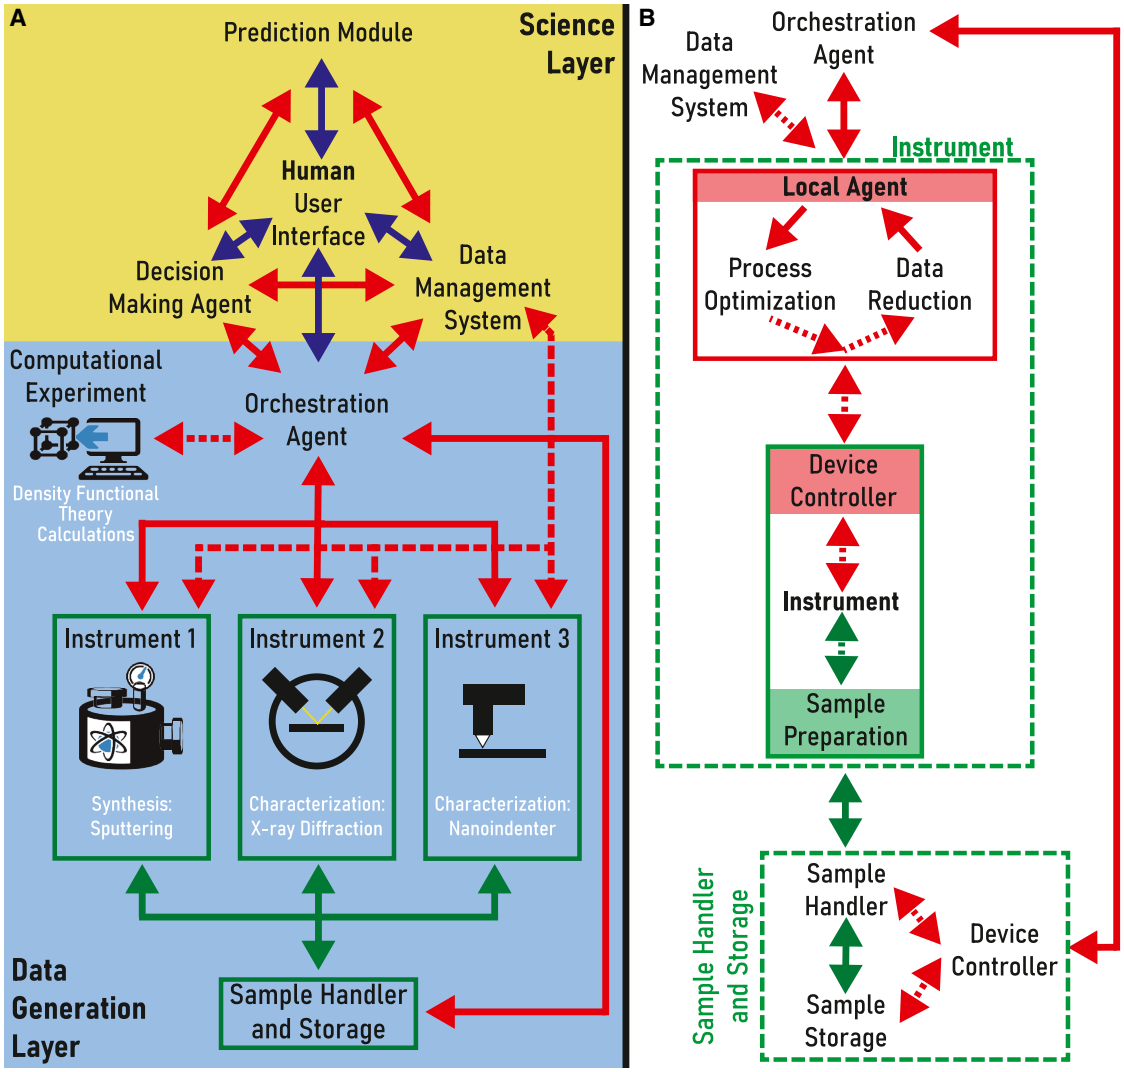

Figure 1. An architecture for a modular autonomous materials research ecosystem

Each module (node in the diagram) represents a major function or component of an autonomous lab. Particular algorithms or components may cover multiple functions as long as they have the same external connections (e.g., one module may handle both the "decision-making agent" and "prediction module" functions). Modularity is achieved through the standardization of interfaces between these modules, indicated in this figure by arrows (dashed arrows may be bespoke or domain-specific interfaces). Red arrows represent virtual interfaces (i.e., an API), green arrows represent sample interchange standards, and blue arrows represent graphical user interfaces (GUIs). This diagram contains 3 representative instruments and a representative computational experiment; a particular platform may have fewer, or likely, more of any type. Individual components may have additional features, such as an instrument may have a local GUI for direct human input. The yellow and blue fields denote the "science layer" and "data generation layer," respectively. (A) illustrates the entirety of the framework, while details of the instrument and sample handler are shown in (B).

agent). This agent can be purely statistical, such as a Gaussian process model, or arbitrarily complex, such as a multi-scale simulation that blends equation-based physical models with ML components.[12] This agent may also perform high-level information fusion by combining data from multiple experimental data streams. The decision-making agent receives these predictions, which can also be logged in the "data management system" for retrospective study and model auditing purposes.

Decision-making agent

The role of the decision-making agent is to identify priority experiments (within the capabilities of the platform) that will produce the data most valuable in achieving the task of platform, be it optimization, exploration, or knowledge generation (i.e., generating and testing scientific theories with experimental data). This agent must also be aware of what experiments have already been performed by accessing the data management system.

Data management system

The data management system is the last part of the science layer. Its primary job is to catalog the data and related metadata generated by the platform and track the history of samples through the various instruments. In addition, it may hold external data available to the platform to help inform predictions and

decisions, and it may track decisions and predictions generated by the decision-making and "prediction module," respectively.

Human user interface

Even in a SDL, there is need, both scientifically and logistically, for humans to interact with the laboratory. Scientifically, this may involve inputing external knowledge (e.g., new physics and external databases), scientific objectives, and experimental campaign parameters and reviewing generated data and model states (intermediate and final). In some cases, input might be needed for "human in the loop" applications, where expert humans will provide analysis, knowledge, or direction as part of the AE loop.[13] Logistically, laboratory staff need to be able to interact with the platform to add consumables and precursors, remove completed experimental materials, and perform maintenance functions.

Human interfaces may take the form of a graphical user interface (GUI), coding interface, or most likely both. Human interfaces may also include scoreboards and other displays of the systems scientific and functional status.

Orchestration agent

In humans, the brain stem, which is just below the brain itself, is responsible for distributing and coordinating the high-level commands of the brain (and also handles some simple logical responses in the form of reflexes). In the case of an ecosystem, there is an "orchestration agent" that carries out these tasks. Its role is to handle the timing and coordination of the experiments requested by the decision-making agent. This includes communicating with individual instruments and coordinating sample movements between instruments through the sample handler and storage system. This agent is also responsible for collecting data and metadata and transmitting them to the data management system.

Instruments

Continuing our analogy, a set of instruments constitute the musculoskeletal system of the platform. These instruments may either carry out synthesis (e.g., sputtering, arc melting, and additive manufacturing), processing (e.g., thermo-mechanical or chemical treatment), characterization (e.g., X-ray diffraction, microscopy and mechanical testing), or some combination of those tasks.

Our representation is shown schematically in Figure 1B, along with the sample management system. In our architecture, we describe two levels of programmatic control. Depending on the specifics of the instrument, there may or may not be separate threads running on one device/computer or they may be operating on entirely separate devices.

The lower-level "device controller" is the electronics native to the device. This may be a complex program with an API operating on an instrument support computer or as simple as an Arduino-based controller that is hardwired into an analog instrument.

The higher layer we call the "local agent." A prime function of this layer is to translate the generic experiment protocol into the language of the instrument. For instance, if we consider a homogeneous sample, like a combinatorial library, the agent translates a request for a measurement in composition space to the necessary coordinates in the instruments' reference frame. Another key function is to collate the data and metadata produced by the actions of the instrument. This function may also involve performing routine data reduction, such as a compositional analysis of an X-ray fluorescence measurement. A more advanced version of this agent might perform real-time measurement optimization independently of the higher-level science agent, such as dynamic X-ray diffraction resolution tuning or automatic focusing of a microscope.

The interface between the "device controller" and the "local agent" will be highly dependent on the details of the instrument and, so, is likely not be able to be defined by a single standard. However, there may be domain-specific standards that can be created. For instance, one or more vendors may work with a group of SEM manufacturers to create a standard SEM interface so that their local agent can work with a variety of SEMs built by various manufacturers. Such interactions already exist to allow third party detectors, such as electron backscatter diffraction (EBSD) or energy dispersive spectroscopy (EDS) detectors, to work with multiple microscopes.

The physical instrument itself may be similar to current equipment, perhaps modified to allow for robotic sample loading using a standardized sample holder (discussed below), or may be designed more specifically for use in an autonomous system. The instrument may be monolithic but, in some cases, may need to consist of two subunits. For many instruments, the instrument will interact with the sample mounted in the holder directly. In others cases, particularly for instruments not designed explicitly to be used in an SDL, the sample may need to be removed from the holder (e.g., tensile testing) or some type of sample preparation may need to be applied (e.g., metallurgical polishing for imaging or creation of a coin cell battery). In the latter case, the main instrument may be mechanically connected to the ecosystem through a sample preparation instrument (e.g., a coin cell battery preparation robot[14]) that can then be connected to the main instrument via an application-specific sample exchanger. Ultimately, tools may become available that can handle both of these functions.

Sample management system

The final component of the system, perhaps the circulatory system in our metaphor, is the sample management system. This system will be controlled by the "orchestration agent" through a device controller (similar to the instrument "device controllers"), which will move samples to the various instruments on the system. Physically, the system will consist of a sample handler to move samples between instruments and a sample storage mechanism.

Computational experiment

Often, there is a desire to combine experimental data with computational modeling. In many senses, there is, by definition, modeling in the science layer of an SDL. However, the role of these predictive models is generally a means to an end—they are not necessarily expected to provide a globally accurate predictive picture of the materials phenomena of interest. For example, Bayesian optimization explicitly trades off accuracy in the less favorable regions of design space in order to more efficiently explore parts of design space with higher probability of optimality. Further, practicality dictates that models used as part of the modeling agent can readily make use of new data generated by the platform and be updated/retrained and

generate predictions of properties across some region of materials design space—all on a timeline sufficiently fast compared with the physical and computational experiments executed on the platform.

There is, of course, a whole range of simulation tools that do not meet this criterion but that researchers may wish to use. These approaches may include density functional theory (DFT), molecular dynamics (MD), or finite element analysis (FEA). The simulations may use purely physics-based material property models or data-driven surrogate models (e.g., ML force fields for MD simulations), as long as they are considered to provide sufficiently accurate representations of reality. As such, they can provide valuable data but typically do so slowly and can only simulate a small number of select materials/conditions. Beyond being slow, in many cases it can be complicated to update these models with new information or materials property data. In these cases, we can treat these computational simulations as *in silico* experiments. Ecosystem-wise, they require the same inputs and generate outputs the same way the experimental equipment does except for the lack of a physical sample.

Ecosystem extensibility

Figure 1 illustrates a minimal platform (excluding the number of instruments). There are many ways that AE systems can have increased complexity. The ecosystem naturally enables a distributed network of experimental hardware, thanks to the "orchestration agent," which separates the tasks of the science layer and the local agents. A specialized federated "orchestration agent" could be constructed to distribute tasks across geographically separate hardware components while still allowing the rest of the ecosystem to function normally. This specialized set of agents would still have the same connections to the science layer (presumably on one side of the pair), and each piece would control the hardware local to it. The connection between systems may have to deal with institutional barriers such as firewalls. If there is sufficient need, a later standard can be established to allow for interactions between these partial "orchestration agents" built by different companies or other vendors. The use of this distributed "orchestration agent" remains abstracted from the "science layer" and the individual instruments.

Alternatively, it may be desirable to have multiple science layers controlling a single (or multiple, as above) experimental platforms. This may allow for more specialized scientific agents to interact[15] or may support a multi-user model, where different users have different scientific objectives. This can similarly be achieved through a specialized "orchestration agent" (and likely a database that allows for communications from multiple modules).

Beyond distributed systems, individual tasks may be complex enough to warrant a nested autonomous system, where one or more measurement or processing sub-tasks require planning and execution of its own. One example is described in the work by McDannald et al.[16], where discovery of a materials for CO_2 absorption requires 1) selection of a candidate test materials, 2) determining the optimal synthesis for each candidate sorbent material, which itself is a complex task, and 3) measurements of these materials across a high-dimensional operational parameter space.

It is also worth noting that, in some cases, only portions of the ecosystem will be used. For instance, one might want to use only the automated portion of the ecosystem—the data generation layer—manually deciding what experiments to carry out. Conversely, one might make use of the science layer as a method of determining key experiments, which are then carried out manually in a traditional or semi-automated lab. A third case might be that the experimental hardware included in the ecosystem are not scoped for a particular task (e.g., equipment inside a nanofabrication facility or a cementitious materials development platform[17]); in this case, the science layer and the "orchestration agent" can be used, but a different set of instruments (which might or might not use the instrument communication standards) and sample handlers will be required.

Community-driven design

The version of the ecosystem described here was designed by the co-authors of this work, using our best efforts to create an ecosystem that would be generalizable across many use cases within materials science. Of course, in order to make any distributed ecosystem functional, there must be broad consensus including both equipment vendors and end-users. As such, this architecture is intended to be viewed as a starting place for conversation. Even as we write this, there are areas of uncertainty. For example, should computational experiments be implemented as virtual instruments or be tied directly to the modeling agent?; should data flow from the instruments (and their respective local agents) through the "orchestration agent" to the "data management system," should there be a separate data broker, or should the instruments move data directly to the data management system? In the end, decisions must be made that work for the community and balance potentially competing goals including minimizing the number of requisite components, decreasing the complexity of the various interchange protocols, and decreasing the required complexity of the individual modules. Decisions must also be made as to the scope of the ecosystem. We discuss pathways toward formalizing these discussions below.

Protocols and standards

Once an initial architecture has been agreed upon, a critical next step will be to create, in concert with relevant stakeholders, a set of protocols, which can be transitioned to formal documentary standards, upon which the autonomous ecosystem will be based. The goal of these standards is to enable modularity. Modular hardware and digital components will greatly reduce integration costs and will allow autonomous systems to seamlessly evolve with changing research scopes, goals, and technologies. Protocols will need to be developed in a number of areas, all related to the interaction between the various modules described above. These protocols may be based on existing standards, which may or may not be broadly adopted currently within the vendor community or may need to be created from scratch.

Sample-management standards

Compared with the many existing automated platforms in the chemistry and bio-pharmaceutical fields, platforms with a high degree of automation are much rarer in materials science. A particular challenge in automated systems for materials science

is the need to move around solid samples (of course, many materials research problems also require liquids to be manipulated[18]), represented by green arrows in Figure 1. Solid materials almost entirely fall into three categories: thin films on substrates, bulk samples, and powders. Each of these has unique sample handling challenges and requirements. The challenge here would be to create a series of universal sample holders that could handle one or more of these sample form factors. Consider the USB-C standard; it defines a variety of aspects (e.g., physical form factor, data transfer protocols and rate, power transfer rate, and alternate use modes), which can be supported (or not supported) at a variety of levels. Similarly, the sample holder interchange standard can describe a variety of aspects, such as sample form factor, size, number of samples, temperature, and sample atmosphere. Each particular instrument and sample holder can be designed to handle each of these aspects to a different degree. For example, one instrument may be designed to support only thin films, while another can be designed to support thin films, bulk samples, and powders. Another example is that some sample holders may be designed to be heated and others can only be used at room temperature. Depending on the nature of a particular instrument, the sample holder may be used natively in that instrument (X-ray diffractometer being a likely instance), while in other cases the sample holder will not be conducive to the tools operation (a likely instance being a device to test a material as part of a device like a battery), and the sample will need to be manipulated by the instrument or a sample preparation tool.[14] Overall, in designing a sample holder, one thing to consider is the price of the holder itself as each platform will likely need dozens or more, one for each individual sample on the platform.

Digitally, these samples are complex to track as well. Unlike liquid samples, many materials are spatially inhomogeneous, which must be managed. These samples may undergo multiple complex processes, each of which may affect the sample overall or only specific portions of the sample.

Inter-module communications

Digital connectivity between AI infrastructure, data infrastructure, and laboratory hardware, represented by red arrows in Figure 1, is critical for the operation of all autonomous systems. For most of the past century, scientific instrumentation has been designed around human operators. The transition from operation via human intelligence to AI has been mired by fragile hacks, rather than robust interfaces. Efforts to address these issues have been made in internet of things (IoT) communication protocols (e.g., MQTT[19]) and within some sectors of scientific instrumentation (e.g., SiLA[20] and EPICS[21]) as it pertains to controlling and monitoring instruments. In developing a standard for the ecosystem, one will need to evaluate existing protocols and identify paths forward that address the unique challenges faced by experimental material hardware and related connected systems.

Data-management standards

In autonomous systems, machine-actionable and AI-ready data are required to inform advanced algorithms and models. Significant progress has been made in elevating the FAIRness[10] of computational data through the establishment of OPTIMADE[22] and related community efforts. Unfortunately, experimental equivalents have lagged behind. Designing an autonomous ecosystem will place an increased emphasis on building consensus on the data interchange formats, including priority instrument types and knowledge graphs.

Module ontology standards

In order to be effective, various parts of the ecosystem, particularly the decision-making agent and the "orchestration agent," need to be aware of all the components that constitute a given system and how to interact with them. These may include what the instruments are, what their function is, what type of data they return, how accurate they are, how long they take to perform a given process, and what the cost associated with their use is. The ability to dynamically and rapidly create a singular platform from disparate parts would be feasible only with a common ontology describing each component. This ontology would have to be flexible enough to encompass a wide variety of component and instrument types.

Algorithm interchange standards

Currently, there is a lack of portability that would enable an algorithm to engage multiple autonomous systems. Some of these challenges are addressed above with communication standards, but an additional layer of abstraction may provide significant benefit for future portability. The power of uniform, high-level abstractions has been clearly demonstrated by scientific modeling software that enables generalized scientific problem statements to be addressed by diverse underlying computational codes. One example is the Atomic Simulation Environment,[23] which provides a generic language for describing atomistic simulations and optimization problem statements that can be seamlessly solved with a variety of different atomistic codes and methodologies. The ecosystem will require a comparable high-level interface for experimental materials science, such as an AE environment, which builds upon precursors such as BlueSky,[24] ChemOS,[25] Hermes,[26] HELAO,[27] and others. Once an open-source environment is established, outside vendors may produce their own proprietary innovation connected via the open standards.

Furthermore, such an AE environment allows for a more rigorous comparison of AI algorithms for scientific tasks, which, in turn, fosters innovation. Some of the most impactful developments in generalist AI have come from open-source libraries, such as PyTorch,[28] Keras,[29] TensorFlow,[30] scikit-learn,[31] and many others. Despite their success, they are neither specialized to the specific needs of materials and manufacturing nor informed by known physics. Unlike the unstructured data from social networks that the generalist AI tools were developed to handle, scientific data are highly contextualized and imbued with specific physical meaning. Because of this, generalist AI is likely to fail to gain meaningful insights into many scientific or engineering problems. Physics-aware ML (aka. scientific AI) is a nascent research area that, in addition to learning from the data, seeks to take advantage of established scientific knowledge. This can include constraints by physical laws, meaning of the parameters, or the context of the data could have specific meaning. However, the nascent scientific AI field lacks any standardization in terms of how it is implemented, documented, or evaluated. This effort will develop an ecosystem of open-source software packages for scientific AI. These packages will be

empowered with known physics and have the capability to discover new physics. These packages, together with the AE environment, will provide broad benefits as well as stimulate the development of proprietary AI within industry.

Community engagement

As discussed above, an ecosystem of this breadth is possible only through broad community buy-in and engagement. A set of consortia will likely need to be established to provide critical guidance and input to a centralized effort. These consortia would encompass a range of stakeholders, from industry and government sectors, as well as academia and other non-governmental organizations, while emphasizing the distinct needs and viewpoints of vendors and end-users. We believe that this approach will increase the adoption of autonomous systems by industry and other end-users by reducing risk and the time and financial investment needed to standup an autonomous platform. Simultaneously, it will increase the number of vendors supporting these platforms by producing a standards-based roadmap for them to develop products around, directly decreasing the associated cost of product development, and increasing their market for these products. A precursor to these consortia will likely be a series of workshops and working groups.

An example: Structural analysis via X-ray diffraction

Structural analysis through X-ray diffraction is a common method in materials science, and, as such, it is a useful example of how an instrument in an ecosystem might work, illustrated in Figure 2.

Let us start with the instrument itself. On the hardware side, there is the instrument with a source, detector, and a set of rotational and linear motion stages to scan the diffraction angles and align and localize the sample with the interaction volume. A sample must be able to be loaded into the instrument; for most lab diffraction, this involves placing the sample, having some flat surface, in a known position/orientation, and then securing the radiation enclosure. The sample might be heated or cooled depending on the experiment. An automated system would need to involve some kind of sample holder that is capable of holding the sample in the goniometer, typically with known positioning, and being installed and removed by mechanical means and able to withstand any environment being imparted on the sample during the measurement. Opening and closing of the radiation enclosure must be automated as well. A standard sample holder and mounting scheme are needed, in addition to any standards for loading the holder and interacting with the enclosure.

At the lowest level, an X-ray diffractometer has a variety of electrical controls that need to be manipulated to make a measurement. These include changing motor positions to scan the sample and diffraction angles and voltage and current settings for the X-ray source. There is also an outgoing data stream from the detector. In most modern diffractometers, these signals and controls are handled by software on a control computer (typically not even the one the user would use, but a server mounted with the instrument) written by the vendor and are highly obscured from the end-user. In our parlance from above, this is the device controller. For most automated measurements, interaction with the instrument on this level is generally not necessary and likely should remain obscured from the general workflow.

The device controller in a diffractometer takes in and generates several types of signals from the higher-level controllers. It provides a set of operations signals such as source ready, enclosure closed, measurement complete, etc. It also needs to be able to accept commands for specific types of measurements—what axes are being scanned, over what range, and how fast—and convert them to the inputs the diffractometer needs to run. This is also where safety checks should be performed, for instance, ensuring that moving components do not collide or that the enclosure door is closed before opening the shutter. Finally, the device controller needs to export the data in a usable format. In many existing systems, the inputs and outputs to the device controller are done through a proprietary GUI. For such a device to be integrated into an automated platform, an API or other scriptable interface must be exposed. Similarly, data should be available in an open and well-documented format (e.g., tiff, csv, json, and HDF5) so that they can be ingested into a variety of software packages. No information about scientific questions or sample information is needed at this level of abstraction.

Above the device controller in this scheme is a local agent. The local agent can be instantiated with a variety of levels of complexity. The simplest would be to have it act only as a translator between other pieces of software in the platform and the device controller, turning requests for data into scan parameters that the device controller can understand. This communication happens over some set of standardized APIs. Let us consider inhomogeneous sample, where one wants to measure a particular point on the sample. It is the job of the local agent to interpret those sample coordinates into stage coordinates for the diffractometer. Once the data have been collected, it is also the job of the local agent to combine the data with metadata about the measurement (e.g., when the measurement was taken and where on the sample the measurement was taken). However, at this level of simplicity, the data would be supplied to the science layer of the platform in a nearly raw form.

The first layer of complexity that one may add to the local agent would be a data reduction and analysis tool. Data reduction reduces the size or dimensionality of the data while maintaining as much information content as possible; data analysis extracts scientific quantities such as phase fraction or grain size and provides them to the science layer (as well as other data analysis tasks being performed on the sample by other local agents). This data analysis can be improved by including known information about the sample (ultimately being read from the database), such as composition or structural simulation results. Data analysis at this level also provides an opportunity to look for errors in the diffraction, from things like misalignment or detector errors.

Another layer of complexity, either in addition to or separately, is the ability to optimize the process. For diffraction, this may include setting a sample-to-detector distance, scan speed, or even scan type. In the most complex version, this allows the science layer to ask specific questions of the diffractometer, e.g., estimate a phase fraction within 1% uncertainty. The process optimizer can tune the measurement to provide that result with the greatest efficiency. This may rely on doing simulation using the analysis code.

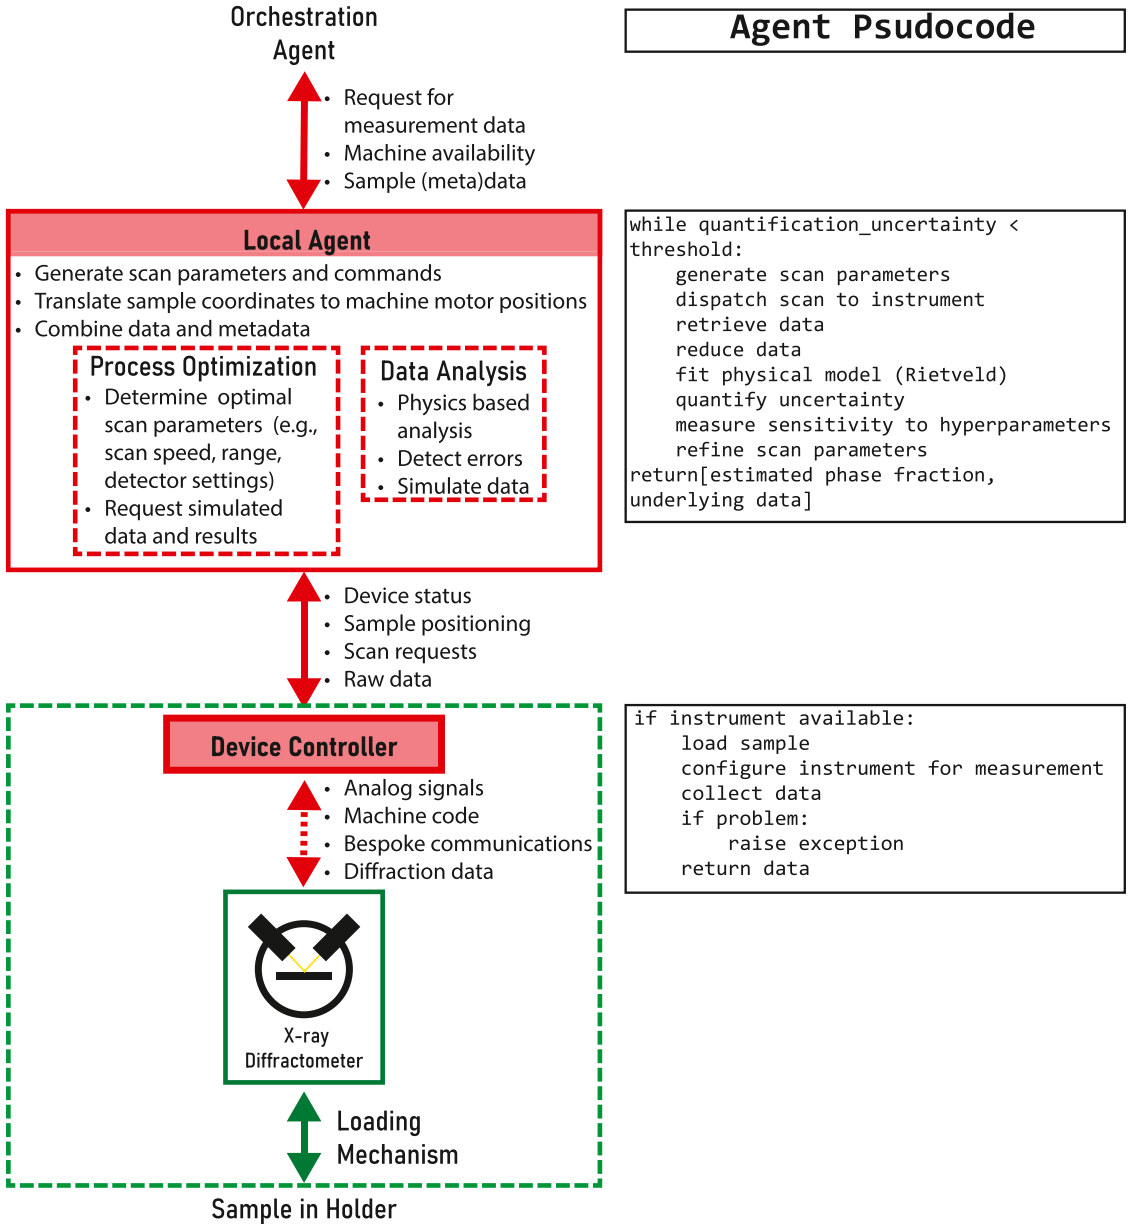

Figure 2. An example of a portion of the ecosystem, using an XRD instrument, showing examples of module functions and data flow; on the right is an example of course pseudocode for the "local agent" and "device controller"

Regardless of how simple or complex the local agent is, we can abstract away specific low-level knowledge of the diffractometer, such as how to move individual motors, signals for manipulating the source voltage and current, or low-level processing of detector output. On the other end, the local agent does not need to know about the broader scientific questions the platform is trying to solve or what else is happening on the platform at any given time.

At the highest level, there are various decisions that need to be made about what to measure and how. The planning agent may also decide which sample needs to be measured, what information is needed, and with what accuracy and precision. The planning agent may be empowered to choose between instruments that can provide similar results. In the case of structural analysis, this may be deciding between two different types of diffractometers or a diffractometer and an EBSD measurement. The orchestration agent manages selecting which diffractometer to use (if there are multiple), decides when the best time to run the measurement is, makes sure the sample arrives at the diffractometer when needed, and ensures the measurement completes without issues.

A center for autonomous materials science

To accelerate the development of an SDL ecosystem, we suggest the creation of a "center." This center, guided by

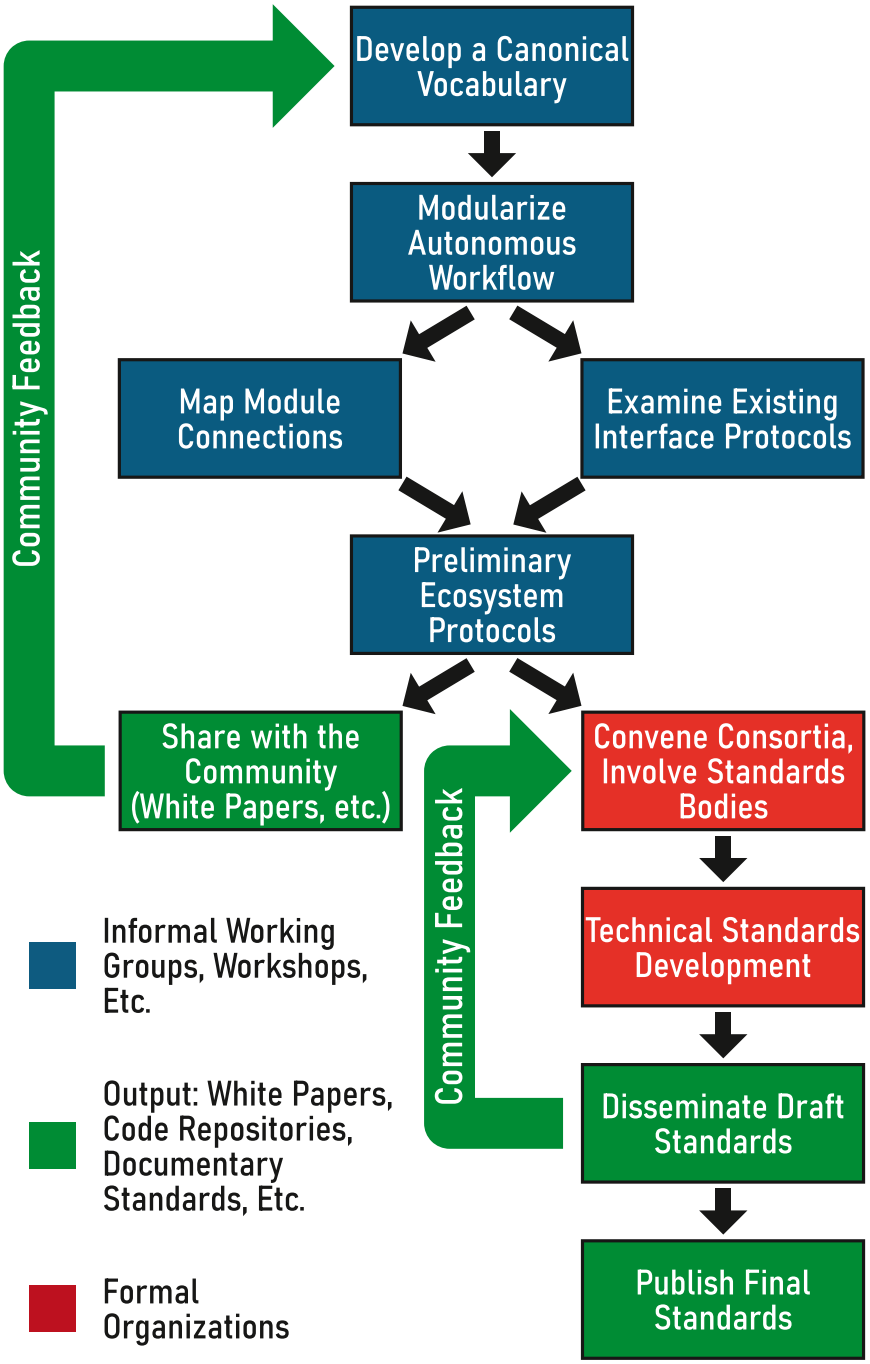

Figure 3. Roadmap for the development of standards around which the autonomous ecosystem can be built

stakeholders, forms the basis for a centralized effort, providing several integrated functions to create and demonstrate an ecosystem of autonomous tools, which industry, in turn, could leverage to accelerate their materials R&D cycles. While not strictly necessary, such a centralized effort would greatly reduce the time needed to develop and gain critical mass for a fully featured ecosystem. Beyond the technical goals, such a center could also serve to train the next-generation workforce, providing hands-on practical training across a range of educational levels, from vocational to doctoral programs. While it is possible that funding could be generated by an industrial consortium, to build this center at scale would almost certainly require the formation of a public-private partnership.

Testbed facility

The core of such a "center" would be a multifunctional autonomous laboratory. This laboratory would comprise various instruments, both for materials synthesis as well as materials characterization, of interest to the center's industrial and governmental partners. These instruments will all be interconnected at both the hardware and software levels, allowing seamless movement of samples and data between instruments with minimal human interaction. The laboratory could serve several purposes and have interaction modes related to each of them, as described below. In addition to the centrally located facility, pipelines for samples and data could be made to other critical materials R&D infrastructure, including X-ray and neutron beamlines and nanofabrication facilities.

Development sandbox

The first role of a testbed facility would be to serve as a sandbox and testbench for developing new hardware, software, and methods. Critically, the building of a testbed will require a set of protocols for the integration of the various instruments and software from a variety of vendors. The development of these protocols, described above, will form the basis of the future ecosystem. As the testbed is built up over time, it will continue to serve as a substrate for refining or expanding these standards. Once established, the testbed will allow for testing and development of the interoperability of new products and features within the ecosystem. Importantly, it can also be used to develop and test new methodologies, both for new characterization methods as well as new AI algorithms and new data analysis methods. For instance, the platform can be used to benchmark the efficiency of new recommendation engines for solving cutting-edge materials science problems.

Exemplar system

The second purpose is to act as an exemplar system. The testbed could showcase the capabilities and features of the

autonomous ecosystem. Vendors could gain an appreciation for how the various components interact and how their instruments/software can interact with the ecosystem. The center will also use the platform to educate end-users on its capabilities as well as how to frame research questions for autonomous laboratories. If they are interested in these types of AE systems, they can then use the platform as a user (see "user facility") or procure the necessary components for their own laboratory.

User facility

Third, the physical laboratory can act as a user facility, similar to a synchrotron light source or a nanofabrication facility. In this mode, end-users would submit proposals to perform experiments on the platform. These proposals may be from, for instance, small companies or academic groups looking to perform a simple optimization that is beyond the capability of their in-house resources, or they may be from larger companies looking to explore, in a deeper way, what autonomous R&D can do for their company, to de-risk investment in their own autonomous platform. By default, the data generated by these user-initiated experiments can be used to populate a publicly available database of experimental results. Alternatively, companies can opt to perform their experiments in a proprietary mode, maintaining their own IP. In concept, this would act as a materials science-focused "cloud lab" (current instantiations of cloud labs are chemistry based due to the simple interchange of samples), often discussed as a potential future for research. Because facility users do not need to purchase their own experimental equipment, these types of facilities can serve to democratize science and innovation.

Data generator

One of the advantages of an autonomous laboratory is the high degree of automation means that the required downtime of the equipment is negligible. The fourth mode of operation of the facility would be as a data-generation platform, adding to the data generated by user-initiated experiments. The platform's planning algorithm can identify knowledge gaps in the testbed facility's database, as well as other publicly available experimental databases, and autonomously perform experiments to fill those gaps. In this mode, the autonomous laboratory will generate high-quality, highly consistent, informative, and publicly accessible datasets. These data can be used to expand the predictive power of models trained on the database, accelerating materials engineering future efforts.

CONCRETE NEXT STEPS

While the end vision is grand, there are concrete steps that can be taken toward this goal. Figure 3 illustrates the necessary steps to develop a standards-based ecosystem. The first step is to reach community consensus (as discussed above) on an architecture that the ecosystem can be built on, such as the one described in this article. A critical first objective in designing this architecture is developing a canonical vocabulary. For instance, "orchestrator" is used in a variety of ways by different groups within the AE community. From there, identifying a modularization of the necessary infrastructure, including identifying module scope and what modules need to be connected, is an essential next step. Simultaneously, existing standards, protocols, and other modular interfaces, many of which are described above, can be examined, identifying their scope, limitations, and general fitness for use in the ecosystem. Existing protocols can then be mapped onto the idealized framework to develop preliminary standardized protocols for the ecosystem. This work can be explored in workshops and disseminated through white papers and other publications. To this point, efforts can be made diffusely within the community. However, to reach full maturity, maximizing adoption and impact, the standardization work will likely need to be focused into more formal and singular bodies, recognized within the community, such as consortia or existing standards bodies. The deeper technical work of finalizing the architecture and then developing documentary standards and protocols for those individual connections can then begin. It is important to note that this step does not have to occur all at once. Any protocol that can be established and broadly adopted can be immediately implemented while others are still being worked out.

CONCLUSION

Creation of an ecosystem for autonomous, enabled hardware and software will produce a wide and long-lasting impact across the entirety of the materials R&D sector through a sea change in the way research is performed. The ecosystem will enable industrial researchers to apply the latest, cutting-edge advancements in automation and AI to their research. The interoperability of this ecosystem will be enabled by the standards and methods developed through community input. The market availability of components for this ecosystem will greatly reduce the investment in time and money that research laboratories in industry, academia, and government need to standup these autonomous facilities, and its modularity will minimize the risk of obsolescence. The creation of a center for developing autonomous methods and platforms will also reduce the (apparent) risk for companies, allowing them to try these platforms prior to investing. As these types of platforms become more prevalent, companies will begin to compete to develop new products to add to these ecosystems, spurring new innovations in experimental hardware, software, and algorithms. Altogether, increasing the availability of this technology will decrease the time-to-market of new materials, many of which are critical for enableing new end-products, which can advance the safety, comfort, and economic wellbeing of us all.

ACKNOWLEDGMENTS

The opinions, recommendations, findings, and conclusions of this work do not necessarily reflect the views or policies of NIST or the United States government. The authors thank D. Holbrook (NIST) and J. Warren (NIST) for helpful discussion.

AUTHOR CONTRIBUTIONS

Conceptualization, all authors; writing – original draft, H.J.; writing – review & editing, all authors.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

1. National Research Council (2008). Division on Engineering and Physical Sciences and National Materials Advisory Board and Committee on Integrated Computational Materials Engineering. Integrated Computational Materials Engineering: A Transformational Discipline for Improved Competitiveness and National Security (National Academies Press).
2. Materials genome initiative for global competitiveness. Tech. Rep. National Science and Technology Council (2011). https://www.mgi.gov/sites/default/files/documents/materials_genome_initiative-final.pdf.
3. Hey, T., Tansley, S., Tolle, K., and Gray, J. The Fourth Paradigm: Data-Intensive Scientific Discovery. Microsoft Research (2009). ISBN 978-0-9825442-0-4. https://www.microsoft.com/en-us/research/publication/fourth-paradigm-data-intensive-scientific-discovery/.
4. Stach, E., DeCost, B., Kusne, A.G., Hattrick-Simpers, J., Brown, K.A., Reyes, K.G., Schrier, J., Billinge, S., Buonassisi, T., Foster, I., et al. (2021). Autonomous experimentation systems for materials development: A community perspective. Matter *4*, 2702–2726.
5. Boswell-Koller, C., Maruyama, B., Warren, J., and Yang, . C. Accelerated materials experimentation enabled by the autonomou s materials innovation infrastructure (AMII): A workshop report. Tech. Rep. Materials Genome Initiative Autonomous Materials Innovation In frastructure Interagency Working Group (2024). https://www.mgi.gov/sites/default/files/documents/MGI_Autonomous_Materials_Innovation_Infrastructure_Workshop_Report.pdf.
6. Trautt, Z.T., DeCost, B.L., Joress, H., McDannald, A.S., Kusne, A.G., Tavazza, F., and Blanton, T.N. (2024). Workshop Report on Autonomous Methodologies for Accelerating X-Ray Measurements, *1500* (NIST Special Publication), p. 25.
7. Davis, S. 1,000th semi industry standard marks 40+ years of microelectronics innovation (2019). https://www.semiconductor-digest.com/1000th-semi-industry-standard-marks-40-years-of-microelectronics-innovation/.
8. Ferreira da Silva, R., Abolhasani, M., Antonopoulos, D.A., Biven, L., Coffee, R., Foster, I.T., Hamilton, L., Jha, S., Mayer, T., Mintz, B., et al. (2025). A grassroots network and community roadmap for interconnected autonomous science laboratories for accelerated discovery. In Workshop Proceedings of the 54th International Conference on Parallel Processing, pp. 142–150.
9. Yoshikawa, N., Asano, Y., Futaba, D.N., Harada, K., Hitosugi, T., Kanda, G.N., Matsuda, S., Nagata, Y., Nagato, K., Naito, M., et al. (2025). Self-driving laboratories in japan. Digital Discovery *4*, 1384–1403.
10. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L.B., Bourne, P.E., et al. (2016). The fair guiding principles for scientific data management and stewardship. Sci. Data *3*, 160018. https://doi.org/10.1038/sdata.2016.18.
11. Banks, P. (2009). The microplate market past, present and future. Drug Discovery *85*.
12. Zhu, S., Sarıtürk, D., and Arróyave, R. (2025). Accelerating calphad-based phase diagram predictions in complex alloys using universal machine learning potentials: Opportunities and challenges. Acta Mater. *286*, 120747.
13. Adams, F., McDannald, A., Takeuchi, I., and Kusne, A.G. (2024). Human-in-the-loop for bayesian autonomous materials phase mapping. Matter *7*, 697–709.
14. Zhang, B., Merker, L., Sanin, A., and Stein, H.S. (2022). Robotic cell assembly to accelerate battery research. Digital Discovery *1*, 755–762.
15. Kusne, A.G., and McDannald, A. (2023). Scalable multi-agent lab framework for lab optimization. Matter *6*, 1880–1893.
16. McDannald, A., Joress, H., DeCost, B., Baumann, A.E., Kusne, A.G., Choudhary, K., Yildirim, T., Siderius, D.W., Wong-Ng, W., Allen, A.J., et al. (2022). Reproducible sorbent materials foundry for carbon capture at scale. Cell Rep. Phys. Sci. *3*, 101063.
17. Joress, H., Cook, R., McDannald, A., Kozdras, M., Hattrick-Simpers, J., Newman, A., and Jones, S. (2024). Autonomous cementitious materials formulation platform for critical infrastructure repair. Digital Discovery *3*, 231–237.
18. DeCost, B., Joress, H., Sarker, S., Mehta, A., and Hattrick-Simpers, J. (2022). Towards automated design of corrosion resistant alloy coatings with an autonomous scanning droplet cell. JOM *74*, 2941–2950.
19. Mqtt: The standard for iot messaging. https://mqtt.org/. Accessed: 2024-06-17.
20. Juchli, D. (2022). SiLA 2: The Next Generation Lab Automation Standard (Springer International Publishing)978-3-031-12582-9, pp. 147–174. https://doi.org/10.1007/10_2022_204.
21. Thuot, M., Dalesio, L., Clausen, M., and Katoh, T. (1996). The success and the future of epics. Tech. Rep. Los Alamos National Lab. https://inis.iaea.org/records/f3rc1-c0t02.
22. Andersen, C.W., Armiento, R., Blokhin, E., Conduit, G.J., Dwaraknath, S., Evans, M.L., Fekete, Á., Gopakumar, A., Gražulis, S., Merkys, A., et al. (2021). Optimade, an api for exchanging materials data. Sci. Data *8*, 217. https://doi.org/10.1038/s41597-021-00974-z.
23. Hjorth Larsen, A., Jørgen Mortensen, J., Blomqvist, J., Castelli, I.E., Christensen, R., Dułak, M., Friis, J., Groves, M.N., Hammer, B., Hargus, C., et al. (2017). The atomic simulation environment—a python library for working with atoms. J. Phys. Condens. Matter *29*, 273002. https://doi.org/10.1088/1361-648X/aa680e.
24. Allan, D., Caswell, T., Campbell, S., and Rakitin, M. (2019). Bluesky's ahead: A multi-facility collaboration for an a la carte software project for data acquisition and management. Synchrotron Radiat. News *32*, 19–22. https://doi.org/10.1080/08940886.2019.1608121.
25. Roch, L.M., Häse, F., Kreisbeck, C., Tamayo-Mendoza, T., Yunker, L.P.E., Hein, J.E., and Aspuru-Guzik, A. (2018). Chemos: Orchestrating autonomous experimentation. Sci. Robot. *3*, eaat5559. https://doi.org/10.1126/scirobotics.aat5559.
26. hermes. https://github.com/usnistgov/hermesAccessed: 2024-06-17.
27. Rahmanian, F., Flowers, J., Guevarra, D., Richter, M., Fichtner, M., Donnely, P., Gregoire, J.M., and Stein, H.S. (2022). Enabling modular autonomous feedback-loops in materials science through hierarchical experimental laboratory automation and orchestration. Adv. Mater. Interfaces *9*, 2101987. https://doi.org/10.1002/admi.202101987.
28. Ansel, J., Yang, E., He, H., Gimelshein, N., Jain, A., Voznesensky, M., Bao, B., Bell, P., Berard, D., Burovski, E., et al. (2024). PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, *2* (ASPLOS '24). ACM), pp. 929–947. https://doi.org/10.1145/3620665.3640366. https://pytorch.org/assets/pytorch2-2.pdf.
29. Keras 3: Deep Learning for Humans. https://github.com/keras-team/kerasAccessed: 2024-06-17.
30. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., et al. (2016). Large-scale Machine Learning on Heterogeneous Systems (2016). In OSDI'16: Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation. https://dl.acm.org/doi/10.5555/3026877.3026899.
31. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in Python. J. Mach. Learn. Res. *12*, 2825–2830.