

**Title:** Autonomous Experimentation Systems for Materials Development: A Community Perspective

**Authors:**

Eric Stach,<sup>1,2</sup> Brian DeCost,<sup>3</sup> A. Gilad Kusne,<sup>3</sup> Jason Hattrick-Simpers,<sup>3</sup> Keith A. Brown,<sup>4</sup> Kristofer G. Reyes,<sup>5</sup> Joshua Schrier,<sup>6</sup> Simon Billinge,<sup>7,8</sup> Tonio Buonassisi,<sup>9</sup> Ian Foster,<sup>10,11,12</sup> Carla P. Gomes,<sup>13</sup> John M. Gregoire,<sup>14</sup> Apurva Mehta,<sup>15</sup> Joseph Montoya,<sup>16</sup> Elsa Olivetti,<sup>17</sup> Chiwoo Park,<sup>18</sup> Eli Rotenberg,<sup>19</sup> Semion K. Saikin,<sup>20</sup> Sylvia Smullin,<sup>21</sup> Valentin Stanev,<sup>22</sup> Benji Maruyama<sup>23\*</sup>

1. Department of Materials Science and Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA
2. Laboratory for Research on the Structure of Matter, University of Pennsylvania, Philadelphia, PA 19104, USA
3. National Institute of Standards and Technology, Gaithersburg, MD 20899, USA
4. Department of Mechanical Engineering, Boston University, Boston, MA 02215, USA
5. Department of Materials Design and Innovation, University of Buffalo, Buffalo, NY 14260, USA
6. Department of Chemistry, Fordham University, The Bronx, NY 10458, USA
7. Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY 10027, USA
8. Condensed Matter Physics and Materials Science Department, Brookhaven National Laboratory, Upton, NY 11973, USA
9. Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
10. Joint Center for Energy Storage Research (JCESR), Argonne National Laboratory, Lemont, IL 60439, USA
11. Data Science and Learning Division, Argonne National Laboratory, Lemont, IL 60439, USA
12. Department of Computer Science, University of Chicago, Chicago, IL 60637, USA
13. Department of Computer Science, Cornell University, Ithaca, NY 14850, USA
14. Joint Center for Artificial Photosynthesis, California Institute of Technology, Pasadena, CA 91125, USA
15. Stanford Synchrotron Radiation Lightsource, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA
16. Toyota Research Institute, Los Altos, CA 94022, USA
17. Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
18. Department of Industrial and Manufacturing Engineering, Florida State University, Tallahassee, FL 32306, United States
19. Advanced Light Source, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA
20. Kebotix, Inc., Cambridge, MA 02139, USA
21. X, The Moonshot Factory, Mountain View, CA 94043 USA
22. Department of Materials Science and Engineering, University of Maryland, College Park, MD 20742, USA

23. *Air Force Research Laboratory, Materials and Manufacturing Directorate, Air Force Research Laboratory, Wright-Patterson Air Force Base, OH 45433, USA*

\* Corresponding Author: [benji.maruyama@us.af.mil](mailto:benji.maruyama@us.af.mil)

### **Progress and Potential**

*As human researchers, we are trained to reduce the number of variables to make experiments manageable. This limits the depth and kinds of phenomena we can study. High-dimensional iterative search empowers us to investigate richer, more complex materials phenomena.*

*Importantly we envisage network effects for the globally integrated autonomous experimentation systems, where beyond the tipping point, the size and degree of interconnectedness greatly multiply the impact of each research robot's contribution to the network.*

*However, to truly exploit the potential of autonomous research, we must build substantial programmatic investments to develop a workforce comfortable working with artificial intelligence.*

### **Summary**

Solutions to many of the world's problems depend upon materials research and development. However, advanced materials can take decades to discover and decades more to fully deploy. Humans and robots have begun to partner to advance science and technology orders-of-magnitude faster than humans do today through the development and exploitation of closed-loop, autonomous experimentation systems. This review discusses the specific challenges and opportunities related to materials discovery and development that will emerge from this new paradigm. Our perspective incorporates input from stakeholders in academia, industry, government laboratories, and funding agencies. We outline the current status, barriers, and needed investments, culminating with a vision for the path forward. We intend the article to spark interest in this emerging research area and to motivate potential practitioners by illustrating early successes. We also aspire to encourage a creative reimagining of the next generation of materials science infrastructure. To this end, we frame future investments in materials science and technology, hardware and software infrastructure, artificial intelligence and autonomy methods, and critical workforce development for autonomous research.

**Keywords:** autonomy, artificial intelligence, machine learning, algorithmic development, research methods, human–machine teaming, workforce development, materials discovery, carbon nanotubes, additive manufacturing

## 1. INTRODUCTION

Materials science and technology are at the core of society, and the development of new materials defines our history. Indeed, specific materials technologies give appellation to the Stone, Bronze, and Iron Ages, to the Industrial Revolution (steel), and to our modern Information Age (silicon).<sup>1</sup> Future advances in quantum computation<sup>2</sup> and synthetic biology<sup>3</sup> will similarly arise from advancements in materials research. However, while the pace of technological advancement is ever increasing, the rate of materials development remains slow, with decades typically needed to transition a new material from discovery to commercial use.<sup>4, 5</sup> This slow development directly impedes humanity's ability to solve existential problems such as climate change, and to generate new technologies that fuel economic growth.<sup>6</sup> Indeed, the futurist Hiroaki Kitano has said, "Scientific discovery is at pre-industry revolution level."<sup>7</sup> The importance of artificial intelligence (AI) in augmenting research and autonomous experimentation (AE) is becoming recognized as a solution to these needs. Former US Secretary of Defense M. Esper recently remarked, "...AI is advancing automated chemistry... These advances free up time for our scientists and researchers to focus on next-generation innovation, rather than countless tests and experiments."<sup>8</sup>

### 1.1 Distinguishing AE from Previous Attempts to Increase the Rate of Research

The materials community has worked for many years to improve the rate of research progress. High-throughput and combinatorial (HT/Combi) approaches<sup>9</sup> allow many experiments to be conducted simultaneously, covering large libraries in composition and processing. Integrated computational materials engineering (ICME)<sup>10</sup> is a longstanding effort to reduce the time and costs needed for materials development by substituting modeling and simulation for experiments, with the Materials Genome Initiative (MGI)<sup>4, 11, 12</sup> as the flagship effort. Recently, using the power of AI and machine learning (ML) to segment materials characteristics in high-dimensional parameter space, new potential compounds and processes have been identified using existing databases of simulations and experiments.<sup>13-16</sup> This paper will make the case that, although powerful, these methods have not and will not speed the rate of research as effectively as AE.

AE uses advanced decision algorithms to plan and execute a series of materials experiments iteratively towards human-directed outcomes.<sup>6, 17-27</sup> More precisely, an iterative research loop of planning, experiment, and analysis is carried out autonomously (see Figure 1). Once human researchers have provided the necessary information (e.g., campaign objectives; constraints; relevant data from previous experiments, also referred to as prior knowledge), the AE campaign is initialized, and the AE system *plans* the first group of experiments. These *experiments*—which are broadly defined to encompass physical tests, modeling/simulation, or data mining—are conducted via automation without human intervention to generate experimental outputs from supplied inputs. Next, the results are *analyzed* automatically and incorporated into an updated understanding of the series of experiments in the framework of a knowledge representation.<sup>28</sup> Finally, a decision algorithm employing AI once again *plans* the next experiment phase and generates a new set of experimental inputs by considering the research campaign objectives and the value of a particular next experiment towards furthering the objective. The system autonomously advances through the iterations of planning, experiment, and analysis. Iterations

continue until the campaign objective is achieved or other exit criteria are met, concluding the AE campaign. The hundreds or perhaps thousands of iterations that may comprise an experimental campaign form the powerful core of AE systems.

In previous efforts to speed research, HT/Combi, ICME, and AI methods typically comprise only a few iterations. While they are powerful for exploring materials parameter spaces and producing and analyzing large amounts of data, they have low iteration rates (related to the “Analysis Bottleneck”<sup>29</sup>), where interpreting results and planning further iterations are the rate-limiting factor. AE systems can execute tens or hundreds of iterations without human intervention. It is this power of an iterative, targeted experimental search that enables AE systems to achieve exceptional speed and fidelity of research results.

## 1.2 Value Proposition of AE

The value proposition of AE lies in its ability to advance research progress much faster than current methods, to make better use of human researcher time and effort, to allow for novel unanticipated findings, and to arrive at a better understanding of a system—all while expending fewer resources. Repetitive manual human labor can be done by research robots more efficiently and reproducibly without human confirmation bias,<sup>30</sup> freeing up human researchers to do more creative work. Cognitive labor is achieved through AI, and research robots are able to analyze data in high-dimensional parameter space beyond human capabilities. Highly autonomous systems also facilitate experiments to be performed remotely, including inaccessible environments such as in the Mars missions, or simply over the internet,<sup>31</sup> making AE highly accessible to the broad community. Thus, in general, we can view AE as an opportunity to rethink the division of labor between human and robot researchers.

## 1.3 Organization of the Review

The intent of this review is to inform the broad materials community about the current status and future directions of AE from researchers active in the area. After presenting the basics of AE and considerations in the design of AE campaigns, we will set the background with a brief overview of previous attempts to speed research—through HT/Combi research and by applying AI to materials data sets non-iteratively. We will then illustrate the state of the art of AE systems for materials, describing their achievements in discovery and development, speeding research, and AI learning methods. Next, we will set out a future vision for how to expand and exploit AE, and we will identify needed investments in hardware, data management, software infrastructure, and algorithm development. Additionally, we will present the challenges and needed investments in workforce development, while considering the broad implications of AE on society, including the role of humans in human–robot research teams and the democratization of science.

## 2. WHAT IS AUTONOMOUS EXPERIMENTATION (AE)?

We first clarify our terminology, as many of the terms have multiple interpretations. We use **automation** to refer to a system that can execute experimental actions without human intervention. An example is using robotics to mix chemicals and measure results. In contrast,

**autonomy** involves the independence of action, integration of delegated decision-making, and complexity of operations. An **autonomous** experimentation (**AE**) system uses automation to execute experiments; it critically has the additional capability to act without human intervention (but in partnership with humans) to incorporate new knowledge derived from these experiments and to reason over and make decisions on subsequent iterations. AE systems can incorporate new knowledge and design appropriate experiments towards the research objective using AI and ML. While artificial intelligence (**AI**) and machine learning (**ML**) are often used interchangeably, we will use the broader term AI to emphasize algorithms used for decision-making in experiments; ML will refer to a subset of methods that include interpolation, classification, and statistical inference. In summary, an *AE* system comprises *AI* for choosing conditions of the next experiment or simulation, *automated* physical and computational experiments for generating new information, *automated* perception or sensing, and a knowledge representation or analysis for incorporating the latest information into an updated understanding. In doing so, AE takes advantage of iterative, sequential experimentation to rapidly progress towards a research objective.

## 2.1 The AE Campaign of ARES

To illustrate the AE campaign, we review ARES (Autonomous Research System), which was developed by Nikolaev et al.<sup>18</sup> as the first AE system for materials development. ARES was able to learn to grow carbon nanotubes (CNTs) at controlled rates over a six-dimensional processing parameter space, ultimately delivering an improved understanding of CNT growth phenomena.

Before the first cycle of planning, experiment, and analysis (see Figure 2), a human researcher enabled initialization of the campaign by generating a database to seed the AI planner algorithm of 140 input synthesis conditions with resultant growth rates. A random-forest model was trained on the database to determine the reaction conditions to be pursued, enabling the AE system to *plan* the first experimental phase. Sequential *experiments* were performed on an automated chemical vapor deposition (CVD) system with computer control of the input synthesis conditions, including laser heating conditions, reactor pressure, and mixture ratios of the four process gases. The resultant growth rates were measured using *in-situ* Raman spectroscopy.<sup>32</sup> For *analysis*, the input conditions and output growth rate were updated in the database after each experiment; following each update, a random-forest model was trained on the database. Using the most recent database and random-forest model as well as a genetic algorithm that sampled suggested experimental input conditions, the AI *planner* determined the next experiment most likely to achieve the target growth rate.

After approximately 600 autonomous iterations, ARES concluded the campaign; it had converged on its target and was able to supply input growth conditions that led to the targeted growth rate. Starting from no prior knowledge of nanotube growth physics, ARES taught itself to grow CNTs at controlled rates via optimized sampling of a vast, complex, six-dimensional parameter space that was much too large to sample using grid-based<sup>33</sup> or open-loop DOE (design of experiments)<sup>34</sup> methods. For example, a full factorial span of 6 parameters with 5 conditions per parameter yields 7776 experiments, which would take over 500 years!

These results provide a direct demonstration that an AE system can control complex chemical reactions. It is essential to point out that the AI-driven iterative search over high-dimensional experimental parameter space is a fundamentally new way to conduct experimental research. It represents a marked departure from the conventional reductionist approach, where reduced-order relations are extracted from pre-determined input conditions subject to specific constraints (e.g., by changing one variable at a time).

## 2.2 Considerations in the Design of an AE Campaign

### 2.2.1 Campaign Objective

In the first step in developing an AE system, human researchers design the campaign objective, which will be the goal of the iterative search process. In its most basic form, the objective can be an accurate prediction of a result. The objective in the first ARES campaign was to closely predict the growth rate of CNTs from prior experiments.<sup>18</sup> Other examples of campaign objectives could be to maximize a property,<sup>35</sup> or to test a hypothesis.<sup>36</sup>

### 2.2.2 Knowledge Representation

Considering once again the closed research loop where the AE system plans, experiments, and analyzes (Figure 1), we enter where new results or outputs are generated from experimental inputs or controlled conditions. From output signals or spectra (*experiment*), the next step in the loop is to incorporate these results into a machine-interpretable representation of the knowledge gained from past experiments, including the mapping from inputs to outputs (*analyze*). The representation can also employ AI and further statistical analyses to identify anomalies in data, categorize regions where experiments are prone to fail, detect fundamentally different system responses, or build beliefs into hyperparameters to models. The difference between the results of the experiment and the expected results from the model can be thought of as a feedback signal for the control loop; for example, it can be used as the basis for training subsequent models. As the campaign advances, an understanding of the response of the material from a user-defined knowledge representation develops and improves its fidelity.

### 2.2.3 Decisions and the Exploration–Exploitation Dilemma

The AE system *plans* the next experiment based on the current understanding from the knowledge base and the campaign objective (Figure 1). Design of the AI planner requires careful consideration of the design policy, which is directly related to the field of optimal experimental design.<sup>37</sup> Throughout the execution of a campaign, the task of achieving the research objective (such as minimizing a response) is often in tension with resolving the uncertainties inherent in the autonomous system's knowledge. This tension, also known as the exploration–exploitation dilemma in the AI community,<sup>15, 38, 39</sup> fundamentally arises from the limited, imperfect, and uncertain knowledge the autonomous system has about the physical system under study. The system may choose to perform experiments that are more tailored to reducing overall uncertainties and to searching for new minima of the response function (exploration), or it may choose to perform experiments near minima predicted based on current knowledge, uncertainties

in that knowledge notwithstanding (exploitation). A balance between these two modes, in which the response function is learned globally prior to optimization, is often more efficient than the decoupled alternative. Understanding the optimal balance between exploration and exploitation is an active area of research and will be discussed further (see *Current and Related AI Technologies for AE*).

#### *2.2.4 Attributes of Autonomy*

With respect to AE systems, autonomy has several attributes that need consideration when designing the research loop and its AI algorithms. The interaction between the AE system and the human researcher can be thought of as a collaboration—so called, “human–machine teaming”—so communication is necessary. The attributes for communication include interpretability, explainability, and interrogability (e.g., Can the human researcher understand and exploit the results?). Trust between humans and autonomous systems is also necessary; it is an active area of study<sup>40</sup> and is broadly characterized by predictable behavior, the ability to achieve intended behavior, and the expectation of two-way communication of well-defined and achievable objectives. Another attribute of autonomy is the ability to integrate uncertainty and contextual information into the decision-making of the planning phase. Examples of uncertainty could be an intrinsic variability in the materials phenomena themselves, noise from the feedback characterization tools, or the influence of exogenous parameters we do not control/measure. Additionally, autonomy would not be possible without a depth of intelligence; while a simple home thermostat can act on its own, its degree of intelligence is limited. Profound AI for AE can include logical reasoning, independent hypothesis generation and testing, understanding by analogy, ability to extrapolate concepts, and ability to design experiments to efficiently and effectively discern complex relationships among myriad possible outcomes. A final attribute of autonomy to be considered in loop and algorithm design, is the delegation of decision authority (e.g., When is it important for the AE system to ask permission from the human researcher? How do ethics and policy come into consideration?).

#### *2.2.5 Deciding on the Decision-Maker: AE System or Human?*

As human researchers delegate decision-making to an AE system, one must consider the desired degree of its independence and responsibility to act in experimental campaigns. Delegation to an AE system is most appropriate when an algorithm is likely to make a better decision. Examples of cognitive labor better and more easily done by machine include scenarios where decisions require: (i) a faster pace than that of human cognitive and/or manual ability; (ii) holistic and detailed understanding of every preceding experiment; and (iii) interpolation/extrapolation in multi-dimensional spaces, with requisite tracking uncertainties, variances, and covariances; and lastly, (iv) when decisions are easy but numerous, and tedious such that they will tire or bore a human, potentially leading to errors.

On the other hand, a completely autonomous system is inappropriate in some cases. When there are clear issues of safety and/or ethics involved in the next experiment, it is imperative that humans oversee AI experiments to ensure relevant safety and ethical practices are observed, e.g., that dangerous reactions are avoided. This is an active area of study for autonomous systems in

general,<sup>41, 42</sup> much of which is appropriate to materials AE systems. Human decision-makers are also better than AE systems when new insights or inferences beyond the supplied physical rules are required to understand a phenomenon, when difficult-to-define objective functions are involved, and in general when information beyond the context of what has been supplied to the AE system becomes relevant.

At the campaign level, autonomous workflows work only when the subset of experimental modules to be used are predefined, already automated, and configured to push–pull data in a consistent manner, and when new experiments do not require frequent workflow modifications to incorporate new analytical tools or processes. In all cases, AE campaigns are best pursued with good teaming between human and robot researchers.

### 3. FOUNDATIONS AND STATE OF THE ART OF AE

#### 3.1 Foundations of Autonomous Research and their Contributions to AE

The development of AE builds upon prior investments in technologies developed to accelerate the research process, and AE integrates them in new ways. These technologies include: (i) HT/Combi experimentation as a method to increase the rate at which new experiments are performed; (ii) modeling and simulation as a substitute for slow and costly experiments; and (iii) data science methods to extract information from simulation and experimental data. The groundwork for these developments was laid in part by the MGI,<sup>12</sup> as well as by similar initiatives worldwide.<sup>43, 44</sup> We will briefly review the state of each of these technology areas to clarify their contribution to AE.

##### 3.1.1 *From HT/Combi Experimentation to AE*

Traditional HT/Combi experiments expedite materials science discovery by parallelizing materials synthesis, processing, and characterization.<sup>45</sup> A typical HT/Combi experiment starts with the automated synthesis of a set of  $10^1$ – $10^2$  samples, in which some combination of composition, microstructure, and processing have been systematically varied to cover the entire parameter space of interest. This library of samples is then screened either in parallel or serially using a set of automated measurement tools. HT/Combi experimental campaigns are typically limited to one or a few iterations of libraries. Some representative recent examples<sup>46, 47</sup> of this for materials are reviewed by Green et al.<sup>45</sup>

Historically, HT experimentation (HTE) hardware development efforts have focused on increasing the number of experimental results per unit time and decreasing the cost per experiment. This makes sense in a *non-autonomous* (“open-loop”) scenario, where the desired goal is either to obtain the desired material by a brute-force experimental search or to generate a sufficiently large dataset that can be used *post hoc* to determine the relative composition–processing–structure–property linkages and provide information on regions of optimal performance for subsequent study. The shift towards autonomous approaches may eliminate the need for many experiments and instead favor faster turnaround for smaller batches of targeted experiments, as new results can be incorporated into the experiment planning. In a non-iterative

(“open loop”) system, AI can be used to intelligently guide an automated characterization tool to subsample a pre-deposited compositional spread library, realizing a 2 $\times$  to 10 $\times$  decrease in the number of samples required to extract information from the system.<sup>45</sup> Some relevant examples of the trend towards lower-throughput, low-latency, small-batch laboratory automation include 3D-printed carousels for performing iterative syntheses of gold nanoparticles to obtain a desired spectrum,<sup>24</sup> dexterous, free-roaming robot chemists that synthesize and characterize small batches of photocatalysts,<sup>48</sup> one-at-a-time synthesis and optoelectronic characterization of perovskite thin films,<sup>22</sup> and the iterative synthesis of perovskite nanocrystals.<sup>49</sup> Microfluidic flow chemistry targeting nanocrystalline materials are especially amenable to this type of approach, as the products can be observed in iteratively changed conditions.<sup>50-52</sup>

AE takes this trend further to autonomous systems with fully closed loops, combining automated on-demand synthesis of AI- (or AI–human-) selected samples with automated characterization on a per sample or per few-sample basis. Unlike the traditional combinatorial experimental stack, autonomous systems are capable of adapting sampling as needed. Replicates are made where experimental uncertainties are high. Redundant information is minimized, and regions of optimal material properties are densely mapped.

### *3.1.2 From Modeling and Simulation to AE*

The use of physics-based modeling simulations in materials sciences is a mature field, and it is now widely accepted that physics-based simulations can identify possible materials of interest.<sup>53</sup> This is exemplified in national efforts, such as the MGI,<sup>12</sup> as well as by large-scale computational materials database/repositories,<sup>54</sup> such as the Materials Project,<sup>11</sup> AflowLib,<sup>55</sup> Open Quantum Materials Database,<sup>56</sup> the Harvard Clean Energy Project (for solar materials),<sup>57</sup> and the NOMAD repository.<sup>58</sup> Rich toolsets have been developed for facilitating large-scale computation and data archiving, such as ChemML<sup>59</sup> and Atomate.<sup>60</sup> Whereas past efforts have focused on making predictions that are subsequently tested in the laboratory, autonomy enables the incorporation of this information into the ongoing experimental process. That is, simulations are used to select better experiments, and simultaneously incoming experimental data are used to select more informative simulations, in a closed-loop process. A notable recent example of this idea is in the use of density functional theory (DFT) alloy thermodynamics as a probabilistic constraint in the (experimental) Bayesian optimization (BO) of perovskite alloys for structure and stability.<sup>61</sup>

### *3.1.3 From Data Science Methodologies to AE*

The use of ML and AI methods for materials applications is now well established and is the topic of recent reviews.<sup>14, 30, 62-66</sup> Their use in accelerating tasks in materials research can be broadly classified as learning to “see” (e.g., spectral interpretation), learning to “estimate” (e.g., surrogate models for predicting outcomes), and learning to “search” (e.g., optimization).<sup>67</sup> Many ML predictions of new materials and properties have been confirmed experimentally.<sup>68, 69</sup> In addition to the use of these methods on simulation and experimental data, they have been used to process other sources of information, such as the natural language text descriptions of synthesis conditions and properties in published papers<sup>61</sup> and structured data showing the relationships

between known materials.<sup>70</sup> In addition to mere prediction, ML approaches can play a role in facilitating human understanding. Relevant examples include the use of machine-learned natural language models to provide automated summarization of material properties,<sup>71</sup> collaborative human–algorithm optimization approaches,<sup>72</sup> and explainable AI (XAI) methods.<sup>73, 74</sup> With such versatility, the ML and AI methods should be able to handle the challenges of the analysis and decision phases in the AE process.

### **3.2 State of the Art through a Selection of AE Examples**

AE is a quickly developing field with new systems coming online with increasing frequency. In order to separate the abstract capabilities of the continually evolving robotic systems from the discrete achievements, we will view this progress through the lens of a selection of completed AE research campaigns (see Table 1). One overarching theme to note is that reports of fully autonomous systems are often closely preceded by related advances in hardware automation, in ML-driven experimental planning, or in both, but without the iterative experimentation and learning that is characteristic of an autonomous researcher. The related non-autonomous advances are included to better illustrate the current state of AE development. For instance, in the previously described example of ARES (Table 1, Study A),<sup>18</sup> the fully autonomous implementation was preceded by the realization of an automated system to map reaction conditions.<sup>75</sup> The more recent inclusion of scanning probe lithography to introduce compositional variations for screening CNT catalysts highlights one path of future development.<sup>76</sup>

#### *3.2.1 AE for Solution Chemistry*

One of the main focal points for the development of automation has been platforms for studying solution-phase chemistry in a broad sense. For instance, Bédard et al. developed a plug-and-play continuous flow AE system for performing synthesis and analysis in an automated fashion (Table 1, study B);<sup>77</sup> user-specified reactions were automatically optimized through the exploration of three discrete reactions using a black-box optimization tool known as SNOBFIT.<sup>78</sup> Building on this platform, Coley et al.<sup>79</sup> integrated both a robotic system to dynamically reconfigure the flow chemistry platform and a pipeline to search the literature and predict synthetic pathways. This highly versatile system was used to discover the optimal synthetic pathways for six sample drug substances; however, the process did not comprise experimental feedback nor optimization of reaction conditions due to the complexity inherent to multi-step reaction chemistry.

An alternate approach to general chemistry has recently been shown by Burger et al., wherein a mobile robot can move around a room to access a variety of distinct stand-alone, commercial instruments, reducing the need for instrument customization. This AE system was used to optimize the hydrogen evolution reaction (Table 1, study G).<sup>80</sup>

#### *3.2.2 AE for Emulsions: Algorithmic versus Random Sampling*

Following the initial demonstrations of AE systems, an important trend started in the literature with systems explicitly testing the acceleration inherent to the confluence of automation and algorithmic planning. Specifically, DropFactory was constructed as an automated system that dispenses reagents to form oil-in-water droplets, which exhibit a wide range of behaviors from locomotion to self-dividing.<sup>81</sup> Incorporating it into an AE campaign, Grizou et al. explored the behavioral range resulting from a four-dimensional parameter space (Table 1, study C). One important contribution from this work was the comparison of experimental campaigns run by random sampling versus those in which each subsequent experiment was chosen algorithmically. When given a budget of 1000 experiments, the algorithmically driven system explored 73% of the parameter space while random sampling only explored 22%. Perhaps more importantly, the algorithmic sampling achieved the same performance in 128 experiments as the random sampling achieved in 1000.

### *3.2.3 AE for Additive Manufacturing: Bayesian Optimization versus Grid-Based Exploration*

Building on the trend of introducing new categories of experiments in an autonomous context while benchmarking against traditional techniques, Gongora et al.<sup>35</sup> developed BEAR, a robotic manufacturing and testing system to autonomously optimize the toughness of additively manufactured components (BEAR = Bayesian Experimental Autonomous Researcher; see Figure 3 and Table 1, Study D). As part of the initial demonstration to study components defined by four geometric parameters, the authors included an explicit comparison between experimental campaigns guided by BO and those guided by grid-based exploration, revealing the time- and cost-efficiency of AE. What the grid-based system achieved in about a month, the Bayesian system accomplished in just 12 h; after 24 h, the Bayesian method produced a higher toughness performance than that achieved by the month-long grid-based search. They have now extended their work to include finite-element modeling of the physical response, successfully increasing the toughness by another 30% (see Figure 3, Conclude panel).<sup>82</sup>

### *3.2.4 AE for Thin Films*

There has been a sustained effort by multiple research groups to develop AE to synthesize and study functional thin films for energy applications. Once again, examples in automation and HTE came first. In 2019, Sun et al. developed a HT process that allowed the synthesis and characterization of 75 unique compositions of perovskite-inspired inorganic films over a span of two months.<sup>83</sup> Following these results, Langner et al. developed a robotic system to synthesize polymer blends for organic photovoltaics and to study degradation in a totally automated fashion, at ~300 samples per day. The resulting large dataset in a four-dimensional parameter space of compositional blends was used to simulate autonomous campaigns, which suggested that a self-driving laboratory could achieve equivalent performance in this space with 32 times fewer experiments.<sup>84</sup> A fully autonomous realization of functional films was published shortly thereafter by MacLeod et al., in which they reported a robotic system moving between synthesis, processing, and multiple characterization stations (Table 1, Study E). By guiding this system with BO through two 35-sample experimental campaigns, they optimized the hole mobility of an organic semiconductor film. Significantly, they also identified a region that exhibits a previously unknown local maximum in mobility.<sup>22</sup>

### 3.2.5 AE for Quantum Dots

In addition to films, quantum dots (QDs) have been the subject of advances in both automation and, recently, autonomy. As far back as 2010, HT synthesis had been applied to map the synthetic parameter space corresponding to QDs.<sup>85</sup> Efforts to screen QDs continue with recent reports on metal–halide QDs.<sup>86</sup> Recently, the concept of automated QD synthesis was combined with a ML-guided experimental planner to realize an artificial chemist for optimizing QD synthesis (Table 1, Study F).<sup>51</sup> This system utilized flow reactors to study a variety of decision-making policies in a BO framework. Further, they showed that learning can be accelerated by at least two-fold when the knowledge of one set of precursors was transferred to a different set of precursors.

### 3.2.6 Developments in Characterization and Analytical Methods in Efforts towards AE

In some cases, efforts towards autonomy in the study of complex properties involve innovative approaches to assess properties. Kirman et al. employed optical observation of crystallization to identify novel perovskites.<sup>87</sup> HT experiments were made possible by using instrumentation developed for protein crystallography studies. ML was applied to both optically analyze samples to evaluate crystallization and to build a predictive model of whether samples would crystallize. Independently, Li et al. also combined robotic synthesis with ML-based experimental selection for perovskite synthetic studies.<sup>88</sup> While their analysis involved a number of manual steps including visual inspection, their experimental selection leveraged a previously developed experimental planner termed ESCALATE (Experiment Specification, Capture And Laboratory Automation Technology).<sup>89</sup>

Efforts towards materials AE need not originate from a synthetic viewpoint; the active guidance of analytical systems can itself accelerate the characterization process. For instance, Noack et al. demonstrated how a kriging-based approach could accelerate X-ray scattering experiments by selecting the parameters of subsequent experiments.<sup>90</sup> This approach was experimentally validated through a set of campaigns, each with 600 experiments, on a sample composed of nanoparticles; a reduction in error was observed when the system was guided by active learning (AL), where the ML model’s uncertainty and expected value are used to select new data points. This study highlights a challenge inherent to benchmarking experimental-learning-based studies; comparisons can only be made to previously reported experiments. More recently, real-time control over X-ray measurements was combined with synthetic capabilities by Rakita et al. to dynamically adjust the redox state of compounds in solution.<sup>91</sup> While this approach only featured a single dimension of control (the presence of reducing or oxidizing agents), it is a promising example of how synthesis and characterization can be combined in an autonomous fashion.

### 3.2.7 AE and Materials Discovery

While there are additional examples at various stages of exploration, we end with examples showing the state of the art with respect to materials discovery. Combining HTE and ML, Ren et al. discovered a new metallic glass using an iterative approach and an ML model for

experimental selection.<sup>92</sup> Many important materials properties are intimately tied to the structure. As such, learning the relationship between the structure of a material and how it is formed—i.e., phase map—can serve as a blueprint for guiding materials discovery and optimization. Kusne et al.<sup>20</sup> developed CAMEO (Closed-loop, Autonomous system for Materials Exploration and Optimization), an AE system that maximizes overall knowledge of the composition–structure relationship. By controlling synchrotron X-ray diffraction measurements and exploiting phase-map knowledge, they identified a novel phase-change material, which has recently attracted attention in the electronics industry.<sup>93</sup> Further, recent reports of AE systems using first-principles simulation provide more evidence that this approach is amenable to the rapid discovery of novel materials formulations.<sup>94</sup>

### 3.3 Current and Related AI Technologies for AE

At this nascent stage, many of the existing AE systems offer proof-of-concept demonstrations, opting to use previously developed ML and AI methods. While these generic methods are attractive for their broad and perhaps immediate applicability, they do not necessarily capture aspects of the material system, experimental apparatus, and campaign constraints. The ability to properly model such factors ultimately impacts the effectiveness of the closed-loop search, the agency of the AE system, and the scientific insight gained from such a campaign. To build more robust, intelligent platforms with greater autonomy, these factors should be included. In this section, we highlight a few examples of such problem-specific features and consider how—if at all—current state-of-the-art methods address them.

#### 3.3.1 Bayesian Optimization (BO) and Gaussian Process (GP) Models

Many general methodologies—such as BO,<sup>95</sup> AL,<sup>96</sup> and statistical DOE<sup>97, 98</sup>—suggest a model of the relevant quantities-of-interest to learn sequentially as well as the decision-making policies that can select a set of experimental actions to execute towards a research objective, making them useful to closed-loop techniques. For example, in BEAR (Figure 3 and Table 1, Study D),<sup>35, 82</sup> the mechanical performance of a manufactured structure is viewed as an experimental response function over such structures and is modeled as a random function using a Gaussian process (GP) model.<sup>99</sup> Used with the expected improvement (EI) policy, in which sampling is pursued at the point most likely to maximize improvement of a value, this GP model is used to select the next structure to test.<sup>100</sup> GP models with the EI policy or similar modeling and policy choices are attractive because of the modeling and computational ease. The GP model allows the specification of the assumed structure, such as smoothness, of the response function without being overly restrictive. However, in many materials systems, such assumptions are not globally accurate. The archetypal example of this are critical phenomena. Critical regions of experiment space (e.g., delineating regimes of pressure or temperature) result in responses that change rapidly or discontinuously, which cannot be properly modeled using off-the-shelf GP models. This is not isolated to the use of GP models in BO. Many AL and DOE methods ultimately rely on similar types of generic models. For example, uncertainty-based methods<sup>101-103</sup> often rely on GP or linear models to model responses.

Another feature not immediately captured with off-the-shelf methods is the fact that experiments often yield several types of responses. Among others, this could be in the form of various characterizations, experiment failure, experimental time or cost, or an uncontrolled factor, such as laboratory humidity. More complex models are needed to properly capture the relationships between the different responses, as well as the uncertainties between these relationships. A joint description capturing a variety of measurable responses and phenomena may not be easy to work with. An alternative direction is to utilize an ensemble of more traditional models, each offering simple estimates of the functions of interest; however, the lack of formalism makes inference and predictions more difficult. For example, Powell and Reyes and co-workers<sup>104, 105</sup> describe methods for using an ensemble of physics-based kinetic models to represent beliefs on experimental responses. Other models such as ensembles of neural networks<sup>51</sup> can directly offer multi-variate predictions for several types of responses, in which correlations between outputs are emergent rather than having to explicitly couple them statistically. Such networks have already been used in experimental science and control settings.<sup>106</sup> In a broader context, ensemble-based methods could allow us to use a variety of different types of models in a single decision-making framework. Here, methods such as Bayesian hypothesis testing,<sup>107</sup> model averaging,<sup>108</sup> multi-fidelity modeling,<sup>109, 110</sup> strategies for multi-fidelity optimization with variable dimensional hierarchical models,<sup>109, 111</sup> and multi-information source optimization (MISO),<sup>112</sup> offer potential avenues for more robust modeling and decision-making.

### *3.3.2 Reinforcement Learning (RL)*

Closely related to closed-loop techniques, such as BO, are reinforcement learning (RL)<sup>113</sup> and optimal control.<sup>114</sup> Markov decision processes (MDPs), a core RL framework, models generic states of a closed-loop campaign, stochastic transitions between states upon taking experimental actions, and rewards or costs incurred when making such transitions,<sup>115</sup> offering a more fluent way of modeling many aspects of materials research. Through RL, MDPs allow an agent to make more operational considerations. RL decisions are obtained by estimating expected future cumulative rewards incurred when pursuing a particular branch of an experimental campaign. Many such techniques do so by approximating a value function (i.e., a measure of how “good” states are) or a policy function (i.e., the expected best action we can take in an attempt to transition to high-value states). As with BO, learning such functions can be done with generic black-box models or with more problem-specific models that use probabilistic beliefs on response functions, experimental failure, costs, or rewards obtained.

### *3.3.3 Deep Learning (DL)*

Regardless of the type of modeling, approximating the functions needed to execute decision-making in RL generally requires a significant computational investment. The coupling of deep learning (DL)<sup>116</sup> with RL—so-called deep reinforcement learning (DRL),<sup>117</sup>—to calculate DL-model surrogates of value or policy functions may prove useful here. DL models are trained against a large number of states/value pairs. This can be done offline, by considering a large number of potential states a campaign can be in and assuming that a representative set of potential states can be simulated. While this methodology proved successful in the case of

AlphaGo<sup>118</sup> and other cases,<sup>119</sup> it remains to be seen whether something similar can be applied in the context of AE.

In general, DL methods are also proving useful outside the context of predicting value or policy functions. They work well by self-discovering latent and predictive features from raw, often high-dimensional data.<sup>120</sup> Despite impressive results in many problems, the direct use of DL in materials AE is limited due to the high data requirements needed to train models. Requiring large sets of representative data is somewhat antithetical to the intelligent and nuanced exploration of experiment space discussed above. There are, however, opportunities for this powerful technique inside the closed loop when simulations and physical models are used to generate synthetic data for offline pre-training of the DL model. DL can also be used to autonomously analyze rich characterization data, such as microscopy or tomography data, and possibly map such data into signals that the autonomous agent can use to close the loop. Current examples of this use in non-autonomous settings include DL for optimal microscopy,<sup>121</sup> cryo-electron microscopy,<sup>122</sup> and atom probe tomography.<sup>123</sup>

### 3.3.4 Transfer Learning (TL)

The lack of data is frequently encountered in autonomous research and generally prohibits the use of larger DL outright. To mitigate this, transfer learning (TL) can be used to leverage existing data of previously studied, related materials systems. One way to do this is with deep transfer learning (DTL).<sup>124</sup> Above, we discussed pre-training a DL model in a way similar to what would be encountered during the online execution of the closed loop. In DTL, a DL model is trained using data obtained from a separate task, often in an unsupervised manner, resulting in a learned latent representation of some material in general. Then, within the closed loop, the model is trained from latent representation features to a material property of interest. Pre-training the mapping from material to latent features reduces the data requirements needed to learn the mapping to the property of interest. Alternatively, one can use adjacent data to build more informative priors for BO models used in closed-loop design. This is the perspective taken by Roy and Kaelbling<sup>125</sup> and applied, for example, to building Bayesian priors for the tribological properties in two-dimensional TMD (transition metal dichalcogenide) materials using adjacent materials descriptors.<sup>126</sup>

## 4. POTENTIAL IMPACT AND FUTURE DIRECTIONS OF AE

With these early demonstrations of AE, we can begin to assess the potential impact of AE on the research process. As demonstrated for specific research tasks (ARES,<sup>18</sup> BEAR,<sup>35</sup> CAMEO<sup>20</sup>), AE has enabled materials research to be orders-of-magnitude faster and has been successful in the discovery of novel functional materials.<sup>20</sup> AE can also achieve better research outcomes than current processes in terms of parameters, such as materials performance or fidelity of characterization.<sup>20, 35</sup>

### 4.1 Economics of Research

Maruyama has proposed a "Moore's Law for the Speed of Research".<sup>127</sup> That is, over time, we may see an exponential increase in the speed of research progress akin to Moore's Law for semiconductors. Others expect an initial burst in research speed with incoming investment, followed by slower progress as we surpass the more easily overcome barriers. For either outcome, AE promises to disrupt the current research enterprise and investment structure. While it is difficult to quantify the rate of research progress, it is straightforward to expect large increases over time as experimental hardware is automated and as closed-loop iterative algorithms are implemented and improved. Already, we have seen a  $100\times$  increase in experimental iteration rates in ARES<sup>18</sup> and a  $10\times$  increase in convergence rates in BEAR.<sup>35</sup>

Market economics may be helpful in understanding the implications of the exponential growth of AE on society and the research community. A corollary to Moore's Law for the Speed of Research is that research becomes exponentially cheaper to execute (if fixed costs for equipment and labor are constant). Since labor dominates the cost of research, AE can effectively multiply the effort of an individual researcher. That is over time, hundreds of experimental iterations can be done with the time and labor it previously took to do one, making the marginal cost of subsequent experiments extremely low. Moreover, the progress towards the research objective is faster because of the iterative search, and so progress (as contrasted to more experiments) is exponentially cheaper. Faster research progress will bring materials to market faster, and it will enable agile, efficient, and effective response, better-tailored materials solutions, and greater scientific understanding in society. Furthermore, as research becomes more affordable, we expect it to become more accessible, just as computing power became more accessible with low-cost processors. Greater access to research can spur a rise in the number of people engaging in scientific research, potentially leading to a revolution in citizen science. In *Culture of Growth*,<sup>128</sup> Joel Mokyr describes how a "Marketplace of Ideas" spurred the Industrial Revolution, increasing the rate of production and transmission of useful knowledge. Outlining impediments to scientific progress (e.g., barriers to entry, transaction costs), he noted that "Economists think that knowledge tends to be chronically underproduced..." By making scientific research faster and more accessible, AE may similarly revolutionize the way we do research.

Figure 4 depicts our projections for increasing the speed of research (per researcher), as well as expectations for the associated access to research robots, which increases the number of researchers. We see three phases of AE development stemming from their degree of interconnectedness. Current AE systems are stand-alone and self-contained. In 3–5 years, we anticipate a transition to locally connected systems, where multiple robots can perform mutually dependent research. In 15–20 years, we expect a network of AE systems to be globally integrated, much like the internet is today. Importantly, we envisage network effects for the globally integrated AE systems, where beyond the tipping point, the size and degree of interconnectedness greatly multiply the impact of each new research robot's contribution to the network. We can thus expect solutions to currently intractable problems, as a result of leveraging network effects from data sharing and interpretation and a community-driven approach to scientific investigation.

## 4.2 Impact on Research Strategies

AE has been successfully applied to a broad range of materials, processes, and characterization modalities using complex campaigns. With the potential for increased complexity, we must consider the implications of AE on the design of campaign objectives and search strategies. AI algorithms are expected to explore experimental spaces differently from humans. That is, human researchers design experimental campaigns to balance the likelihood of success, potential benefits of success, and explainability of outcomes. Often this takes the form of starting from known experiments<sup>129</sup> and making modifications one variable at a time.<sup>130</sup> This strategy can be effective for local optimizations, but it has difficulty in multiparameter problems and results in biased datasets.<sup>131</sup> The speed and reduced human effort of AE enable a greater diversity of experiments, and since AI/ML algorithms excel at high-dimensional search problems, they are holistic rather than reductionist.

The further development of AE has strategic implications for the risk appetite per experiment. Human researchers tailor experimental campaigns to balance the likelihood of success with the potential benefit from success. That is, experiments that have a reasonable chance of success using the available timeframe and resources will be pursued. With AE, we can perform many more experimental iterations, and therefore increase the overall likelihood of success. The failure of one or even several experiments does not doom a campaign. In fact, "failed experiments" can serve to inform where experiments do not work and further improve the ML model.<sup>132</sup> Previously intractable problems become more likely to succeed. Using AE, we can pursue more challenging, high-dimensional problems.

We have described the scope of a general AE system as one that encompasses multiple materials and processes, multimodal characterization, constraint optimization, and—distinctively—an iterative search via sequential decision-making for experimental design. Next, we consider advanced strategies for closing the autonomous research loop.

#### 4.3 Hypothesis-Driven AI

AE offers an opportunity to tightly integrate the scientific method's hypothesis generation and testing into the iterative experimentation loop. Advanced strategies to close the loop may exploit the full range of AI, reasoning, decision science, optimal experimental design, and the convergence of the scientific method with research robots. While autonomous systems promise to more reliably perform optimal experiments towards an objective, some have expressed concern that robots will ignore results that are outside the objective but are nonetheless interesting and that they will miss serendipitous and synergistic unanticipated results that a human would naturally recognize.<sup>133</sup> "In the fields of observation chance favors only the prepared mind," said Louis Pasteur.<sup>134</sup> How can we imbue AE with curiosity, creativity and insight? In the future, we may be able to incorporate serendipity-awareness into autonomous research algorithms.<sup>81</sup>

Iterative experimentation can be used to explicitly test physics-based models, which are effectively scientific hypotheses.<sup>135</sup> Thus, it is possible to confirm or negate hypotheses inline, leading to a hypothesis-driven search. Another goal of a campaign could be to search over multiple hypothesized models, known in the AI community as model selection, where the

models would be potentially operative physical and chemical models, enabling physics-informed sequential decision-making. This approach is particularly appealing in that these models incorporate physics-based phenomena, e.g., Arrhenius behavior, into their knowledge representations rather than being naïve or purely statistical representations.<sup>104</sup> King et al.<sup>36</sup> built one of the earliest robot scientists called "ADAM," which produced yeast enzymes by generating hypotheses of biological synthesis routes that it evaluated in a closed loop.<sup>136</sup> Similar to the symbolic regression work of Schmidt and Lipson,<sup>137</sup> we envisage AE campaigns where the objective is to select and parameterize from a broad range of materials phenomena (in place of mathematical symbols) using iterative experimental search strategies that are designed to regress quickly to the operative physics. AE output as physical models is clearly superior to output as naïve or black-box statistical models—which, while they may be predictive, are scientifically uninformative. Additionally, while statistical techniques, such as ML, are appropriate for interpolation, they do not excel at extrapolation, which is where hypothesis-driven research has the advantage. And so, we feel this broader approach to hypothesis-driven, physics-based models can capture unanticipated results and ultimately be more comprehensive than strictly human-supplied hypotheses.

Hypothesis-driven exploration enables extrapolation to unexplored and unanticipated compositions, properties and phenomena, promising to revolutionize research by providing expert-level analyses and super-human capabilities with respect to speed and the complexity of the research problems. Such hypothesis-driven exploration is a ripe opportunity for investment in initiatives that integrate the AI and materials communities, establishing mutual challenge problems and workforce development programs that bridge the two communities.

#### 4.4 Bringing the AI and Materials Communities Together

The distinguishing component of closed-loop AE systems from merely automated experimentation are the sequential, iterative decisions made by an AI/ML planner. The choice of planning approach is currently the subject of intense exploration, with established approaches (such as statistical experimental design, BO, and RL) in use. The exploration of advanced AI/ML methods is also an active area of research.<sup>20, 138-140</sup> Notably, AE systems offer a unique opportunity to the AI community as platforms for the development and testing of their models and algorithms. The value proposition of AE to the AI community is the iterative nature of the platform over unknown search spaces, that nonetheless have a ground truth in materials phenomena because they originate from fundamental chemistry and physics. There is no direct mechanistic analogue in social media response or static voluminous databases, where advanced AI methods are often applied.

A key challenge at the intersection of materials and AI research is the integration of ML with AI reasoning in the context of scientific knowledge, e.g., in data interpretation tasks that humans tackle with a phenomenological approach. AI reasoning comprises the ability to infer new facts via the consideration of various information sources, making it complementary to statistical ML and critical to the emulation of human scientific exploration in automated systems. The broad AI area of knowledge representation and reasoning encompasses various sub-areas, such as search, logic and probabilistic reasoning, knowledge representation, planning, and sequential decision-making. As such, coupled AI-reasoning/ML is a pillar of the so-called third wave of AI,<sup>141-143</sup>

and the phenomenological nature of the physical sciences make it particularly well-suited for the development and demonstration of AI reasoning systems. AI reasoning will enable search, reasoning, and inference over hypothesis space rather than over raw experimental parameter space, elevating autonomous systems from the current emphasis on black-box statistical models to a future emphasis on generation and validation of scientific hypotheses.

SARA (Scientific Autonomous Reasoning Agent) is an AE system under development through a multidisciplinary initiative led by Cornell University.<sup>144</sup> It will use combinatorial samples with laser spike annealing to generate time, temperature, and transformation (TTT) diagrams for novel functional oxides. SARA uses AL strategies with the distinguishing feature that complex reasoning is required to guide experiment selection. For example, one operational mode for SARA is the identification of all unique synthesis routes for a given structure. This requires AI algorithms that are deeply aware of the structure of phase diagrams, non-equilibrium processes, properties of all known phases, and how to recognize new phases, among other issues. Realizing this breadth of expertise in an AI system requires learning-and-reasoning-based algorithms well beyond the purview of ML algorithms demonstrated in materials research to date. The need for revolutionary materials-aware AI combined with the substantial automation complexity of the composition–processing–structure instrumentation makes the SARA project emblematic of the grand challenges in autonomous materials science.

Materials research is clearly benefiting from AI, while the AI community sees AE as a platform to rapidly develop their approaches in what are effectively AL robots. Moreover, unlike social networks, which occupy a significant fraction of AI research and are less subject to controlled scientific experimentation, AE systems can be iteratively probed and critically have physical and chemical phenomena as ground truth. Indeed, an opportunity exists for the autonomous materials and AI communities to pursue the Nobel–Turing Challenge set out by futurist Hiroaki Kitano.<sup>145</sup>

## 5. CHALLENGES AND INVESTMENTS

In order to fully benefit from AE, the community must overcome significant challenges by investing in key areas. Fundamental research funding typically focuses on addressing specific foundational questions. However, investments in AE will establish an infrastructure that will broadly enable faster research towards *many* scientific questions as well as industry-relevant results. Here we summarize challenges and needed investments in experimental hardware, data, software, and workforce development.

### 5.1 Investments in Experimental Hardware

Designing an automated system to increase the speed, agility, and reproducibility of experiments is a challenge when using existing commercial, off-the-shelf equipment, which are primarily designed for human use and not for robotic or automated sample handling. Although it is possible to design mobile robotic systems that can work in existing laboratories,<sup>80</sup> this is not an ideal long-term solution compared to standardized sample exchange interfaces, which will reduce complexity and design or robot-path planning time. The redesign of microscopes, synchrotron beamlines, and other sophisticated instrumentation to be compatible with robotic sample handling—akin to the multi-plate-handling robots in the bio-community—is an essential

area needing investment. It will also be necessary to integrate *in-situ*/inline and real-time metrologies with automated data processing pipelines for various material data formats.<sup>80</sup> For example, *in-situ* microscopes could generate a massive amount of image frames at a microsecond frame rate,<sup>146</sup> and inline image analysis carried out as fast as the frame rate is crucial for accelerating multistage experimental campaigns.<sup>147</sup>

Innovative new technologies for on-demand sample fabrication and *in-situ* characterization are also needed to translate these early wins to the full spectrum of materials science applications. Many material properties and desired functionalities are emergent from multiple phenomena. In thermoelectrics, the interplay of electronic, magnetic, and vibrational effects over length scales from the atomic to mesoscopic, leads to the desired property. Full characterization of these effects requires local transport, optical, electronic, and scanning probe measurements. Such *in-situ* characterization generates multimodal data sets that may span many dimensions of spatial, temporal, and spectral information. This improves the rate and degree of convergence towards campaign goals. However, it is not generally tractable to measure every spot of a heterogeneous sample with every tool, particularly when destructive characterization techniques are necessary. Optimizing the information obtained by multiple techniques requires a judicious subsampling of this measurement space to uncover the relevant descriptors.<sup>148</sup>

Thus for hardware, we encourage non-proprietary interfaces to enable facile sample exchange across multiple commercial tools for synthesis and characterization.

## 5.2 Investments in Data Management and Sharing

By their very nature, AE systems will generate much larger data sets than current laboratory practice, and this data is inherently machine-readable. This creates an opportunity to encourage the organized collection, sharing, and reuse of data at much larger scales than at present. Such accumulated and well-curated databases resulting from many experiments can be reused by a distributed network of AE systems. This in turn provides a dataset from which prior knowledge of related scientific domains<sup>92, 149, 150</sup> can be extracted and then used to supplement RL<sup>151</sup> and TL,<sup>126, 152, 153</sup> algorithms.

Such data repositories will only be successful if they contain comprehensive experimental data and metadata. More attention must be paid to the collection of data from failed experiments<sup>132</sup> and to the automated labeling of data as they are collected. These needs have been highlighted in the MGI<sup>12</sup> and other efforts, and they are captured in the FAIR data principles of findability, accessibility, interoperability, and reusability.<sup>154</sup> Newly developed AE systems present a unique opportunity to satisfy FAIR principles at the point of data generation, rather than at the time of publication. The distinction for AE is that research robots can immediately probe the system and generate new responses with immediate feedback in the form of materials experiments, which is not possible using static databases. This provides new opportunities for identifying and rectifying data anomalies.

We encourage increased investment into open-source/open-standards data file formats and application programming interfaces (APIs), and discourage proprietary software and data

formats for experimental hardware. Investments will be needed for the large-scale structured repositories of both data (e.g., Materials Data Facility,<sup>155</sup> Materials Project,<sup>11</sup> PRISMS<sup>156</sup>) and trained AI/ML models (e.g., DLHub<sup>157</sup>)—designed for use by machines and people—as well as the automated tools for constructing and curating these databases.<sup>11</sup> Additionally, efforts to develop uniform metadata descriptions, such as tracking material sources and workflow methodologies will be needed.

### 5.3 Investments in Software Infrastructure

Irrespective of whether a system is fully automated, the algorithms used to direct experimental decision-making need to be both robust and flexible enough to be used on a variety of different experimental platforms. Investment is needed in the software infrastructure for AE. ChemOS,<sup>158</sup> ESCALATE,<sup>89</sup> LabMate.ML,<sup>159</sup> MAOS,<sup>160</sup> BlueSky,<sup>161</sup> and AREST™ OS<sup>31</sup> are examples of such efforts in progress. However, the broader range of materials, modeling software, and experimental hardware will require further investment into software. Commercial hardware often uses software and data formats that are proprietary and difficult to access or modify for incorporation into AE systems. To build automated and AE systems that can incorporate multiple commercial systems for synthesis and characterization, more open-source software, data standards, and APIs are needed.

### 5.4 Challenges of Decision Algorithms in Novel Environments

In addition to information-theoretic considerations, physical experimental campaigns often have real-life operational considerations or constraints (e.g., time, cost, available inventories) that may need to be incorporated into the decision algorithm. Decisions could also include a "cost–benefit" analysis that utilizes multiple information sources.<sup>112, 162</sup> These information sources could include a mix of experiments that exploit various characterization techniques and *in-silico* simulations of multiple levels of fidelity. Experimental decisions that are most valuable in an information-theoretic sense may not be so easy to execute in practice, and an AE system should be able to capture such practical constraints in its decision-making. A variety of similar, established fields, such as decision theory, optimal design, AL, and RL, provide various methods and algorithms for striking a balance between exploration, exploitation, and operational constraints.

Advanced materials development frequently involves expensive methods and instrumentation or limited-access, highly competitive facilities (e.g., synchrotron X-ray or neutron sources), making the ability to reliably select the right experiment especially impactful. Often-used electron and ion microscopy methods are time-consuming and costly; molecular beam epitaxy and CVD can yield extraordinary control over synthesis at the atomic scale, but they also require expensive instrumentation. New algorithms should be developed for efficient exploration of high-dimensional parameter space in a time-constrained environment, thus reducing the number of required experiments. In the CAMEO system of Kusne et al.,<sup>20</sup> phase maps of composition spreads were identified using synchrotron X-ray diffraction, generating a combinatorial library with an effectively infinite number of compositions to characterize. By combining graph segmentation techniques, BO methods, and physical constraints in closed-loop iteration, they

were able to demonstrate maximal knowledge with minimal experiments and obtain an accurate phase map of the material system. Exploiting the phase map autonomously, they discovered an optimal phase-change memory material using only one-tenth the number of measurements required by the standard grid-based approach. The new material also outperforms the current industry standard, underscoring the effectiveness of implementing AE on a combinatorial platform.

## **5.5 Bringing Together Industry, Academia, and Government – The Challenges and Investments**

Beyond academic research, many materials and chemical corporations (such as BASF, Bayer AG, Merck Group, and Dow) have internal HT/Combi units within their technology innovation divisions. These units actively explore applications of AI tools in research and are directly benefitting from the network of AE labs in academia through licensing the technology, subcontracting research-and-development projects to academic teams, and also involving these teams in joint development processes. There is also an ongoing effort in knowledge transfer from academia to industry, where the research teams directly commercialize technologies of AE materials discovery. Existing examples include ML tools for materials data analysis (Citrine Informatics), a closed-loop discovery platform for molecular materials that combines AI control with robotic synthesis and characterization (Kebotix), software for upgrading existing HT/Combi robots for automation (Atinary), and automation for chemical manufacture (Snapdragon Chemistry).

National-scale scientific user facilities (SUFs) can also play a critical role in encouraging the transition from small-team independent research to cooperative scientific networks. On the one hand, SUFs offer the highest-caliber materials science tools available, concentrated at a few large-scale facilities, which are coupled to high-performance computing resources and are staffed by technique specialists. On the other hand, these tools are typically used by single-investigator teams after a lengthy, competitive proposal process; data formats are diverse, and there is little—if any—coordination between complementary experiments at different SUFs. The centralized nature of these facilities offers an opportunity to establish common data formats, data sharing policies, and new access paradigms such as multi-facility proposals. The national laboratory-scale engineering resources can be leveraged to enhance automation, develop hardware and software standards around which large community-scale AE programs can nucleate.

We expect that forming partnerships and consortia between academic, government, and industrial partners will be one of the pathways for future development. However, the main barrier to effective partnerships between academic and industrial teams is in the ownership of the co-developed intellectual property. With laboratory capabilities distributed across different entities and open-sourced ML algorithms trained with proprietary data, the questions about product ownership and the contributions of involved parties will be a persistent concern.

We have recommended a variety of investments to help establish a new infrastructure for AE to accelerate research progress. We suggest that investments that parallel equipment or capability modernization as a potential path forward. In the longer term, we envision AE systems linked

together over networks where experimental, simulation, and information processing nodes combine with human direction to form autonomous "collaboratories," which generate scientific knowledge at rates barely imaginable today (see Figure 4).<sup>163</sup>

## 5.6 Workforce Development and Human–Robot Teaming

We also see workforce development for AE systems as a critical need going forward. Our existing workforce does not have the skillset to do both materials and autonomy research, and universities are just beginning to develop curricula to address computer science and AI for materials research. However, to truly exploit the potential of AE, we need substantial programmatic investments to develop a workforce of "AI natives." They must be as comfortable doing closed-loop AE as we are doing conventional materials research today to enable effective human–robot research teaming.<sup>164</sup>

### 5.6.1 Human–Robot Research Teaming

Designing effective human–AI teaming is an emerging area in autonomy and user-experience research.<sup>165</sup> Teams of humans and chess-playing computers outperform either humans alone or computers alone.<sup>166</sup> In the 2005 Freestyle Chess Tournament, a team of chess masters and a supercomputer were defeated by a team of amateur humans and desktop computers with superior teaming. Nascent efforts at teaming humans with AI exist for inorganic materials,<sup>72</sup> and such efforts must be accelerated.

To clarify a common misconception, it should be noted that AE is not meant to remove humans from the research process, but to enhance their efforts. Robots excel at performing repetitive work with precision, so it is easy to imagine that manual labor can be done by research robots more quickly, more reproducibly, and more cost-effectively. Research robots can also analyze data in high-dimensional parameter space in ways that are beyond the capabilities of human researchers; they can make principled, more effective decisions towards a set of research goals and without human confirmation bias.<sup>30</sup> With AE relieving human researchers of the tedious manual labor and high-dimensional cognitive efforts, they are freed to focus on creativity, insight, oversight, and high-level goal definition. AE should multiply the effort of human researchers in the same way that a tractor multiplies manual labor or a computer multiplies cognitive effort. To this end, there are on-going efforts to incorporate human expertise, judgement, and prior knowledge into search and decision-making algorithms.<sup>167</sup> Future AE systems would greatly benefit from investments to human–robot research teaming.

### 5.6.2 Maintaining the Human Workforce

The materials research community faces a persistent challenge in attracting the STEM (science, technology, engineering, and mathematics) workforce. Young potential scientists are often discouraged from materials research as they tend to start in the laboratory performing mundane, repetitive tasks as required in the current research workflow; they are the present-day robots. The associated slow pace of research and lack of immediate feedback can also be discouraging. To recruit and retain future researchers, and to maintain the STEM workforce, the materials

community must make the research process more intellectually rewarding. We believe that teams of human and robot researchers, far from displacing human researchers, will make research more attractive and fulfilling.

While studies have indicated that workforce development and curricular innovation is needed at all levels,<sup>164</sup> one particularly pressing need is for technicians to manage the hybrid mechanical–electrical–chemical systems. Because of similarities to workforce needs in advanced manufacturing, there may be opportunities to extend the existing efforts of community colleges.<sup>168</sup> As a result, there is a critical need for pedagogical material that can form the basis of new courses or be incorporated into existing courses.

In addition to education, we must address systemic bias in recruiting young people to science. Lei et al. found that "Children lose confidence in their potential to 'be scientists,' but not in their capacity to 'do science'... This pernicious decline is especially evident among underrepresented groups, including girls, members of some racial and ethnic minorities, and children from lower socioeconomic backgrounds."<sup>169</sup> To partly rectify this problem, we need to invest in making scientific research more accessible to everyone, especially those at risk.

Currently, only those with access to large, well-resourced laboratories are able to participate in materials research at the highest level. A potential outcome of AE is a rise in citizen science where—as in the astronomy and high-energy physics communities—contributions to the field can be made by enthusiasts with access to data or instruments. Today in contrast, most materials science research is conducted by small teams, typically consisting of a principal investigator, students, and post-doctoral researchers. Despite efforts to make data more widely available (e.g., Materials Data Facility<sup>170</sup>), most data are kept proprietary and are not used fully even by the team that produced them. With minimal future reuse, the potential value of the data is not captured.

In the future, greater access to AE will provide more people the opportunity to access research robots and be able to do meaningful research. This may take the form of remote-access "cloud labs" (e.g., Emerald Cloud Lab, Strateos), which may include low-cost, relatively self-contained benchtop equipment analogous to 3D printers (e.g., Molecule Maker Lab, Martin Burke's synthesis machine),<sup>171</sup> or open-access challenges where participants can propose new experiments based on collected datasets (e.g., the DARPA SD2 Perovskites Synthesis challenge, performed on the RAPID system).<sup>88</sup>

## 6. CONCLUSION

We hope that this paper informs, sparks interest, and potentially inspires the larger community for AE systems. The first research robots are already making an impact in materials research and development. From optimizing the growth of CNTs to accelerating the understanding of composition–structure–property maps, they are revolutionizing the way scientific research is conducted. Disrupting conventional research methods, AE has demonstrated an increased rate of knowledge generation by orders of magnitude and has resulted in the discovery of new compounds. Broad deployment of AE will require substantial investment in hardware, software,

and data infrastructure, as well as in education and recruitment to overcome technological and workforce challenges. Integrated, online AE systems need to be made cheaper and exponentially more accessible. Upon the demonstration of a sufficient number of AE platforms, funding of large-scale multi-institutional "collaboratories" will enable researchers to attack civilization's most pressing topics.

AE has led to the collaboration between the materials and autonomy communities. Early collaborators in the autonomy community used materials AE systems as a platform to develop and test advanced AI and autonomy approaches, where materials research problems are both more accessible and representative of complex real-world environments. From the autonomy community, the materials community has learned about algorithm development, application, and search strategies, and notably of the importance of human-machine integration and teaming. Teaming of human and robot researchers shifts mundane manual and computationally intensive cognitive labor to machines. This is critical in research, which, unlike advanced manufacturing, is dominated by labor costs.<sup>172</sup>

Overcoming the challenges identified in the paper has the potential to reshape science and particularly the roles of human researchers, freeing us to engage with science more meaningfully and interactively. This will lower the barrier to entry for asking and answering scientific questions, generating a new breed of scientists who focus on insight and creativity and lowering the barrier to entry for citizen scientists. These tools will bring together artificial and human intelligence in efficient and effective efforts to accelerate technological and fundamental scientific progress transforming the world around us.

### **Acknowledgements:**

The authors gratefully acknowledge the contributions of the participants of the Autonomous Systems for Materials Development Workshop, held at the University of Pennsylvania Singh Center for Nanotechnology, 4–6 September, 2019, with support from the University of Pennsylvania, the Air Force Research Laboratory, Materials & Manufacturing Directorate, and *ACS Combinatorial Science*. We also gratefully acknowledge John Russell of the Singh Center for his assistance in organizing this workshop. This research used resources of the Advanced Light Source, which is a DOE Office of Science User Facility under contract no. DE-AC02-05CH11231. Support from the Air Force Research Laboratory Air Force Office of Scientific Research (AFOSR) LRIR No. 19RXCOR322 is gratefully acknowledged.

### **REFERENCES** to be moved here after removing Endnote links

1. Sass, S.L. (1998). *The substance of civilization: Materials and human history from the stone age to the age of silicon* (Arcade Publishing).
2. USDOE Office of Science (SC)(United States). (2017). Basic Energy Sciences Roundtable: Opportunities for Quantum Computing in Chemical and Materials Sciences (USDOE Office of Science (SC)(United States)).
3. Ball, P. (2018). Synthetic biology—Engineering nature to make materials. *MRS Bulletin* *43*, 477-484.
4. Holdren, J. (2014). National Science and Technology Council, Committee on Technology, Subcommittee on the Materials Genome Initiative, Materials Genome Initiative Strategic Plan.
5. Gordon, R.J. (2016). Perspectives on the rise and fall of American growth. *American Economic Review* *106*, 72-76.
6. Tabor, D.P., Roch, L.M., Saikin, S.K., Kreisbeck, C., Sheberla, D., Montoya, J.H., Dwaraknath, S., Aykol, M., Ortiz, C., Tribukait, H., *et al.* (2018). Accelerating the discovery of materials for clean energy in the era of smart automation. *Nature Reviews Materials* *3*, 5-20. 10.1038/s41578-018-0005-z.
7. Kitano, H. Nobel-Turing Challenge. <http://www.oecd.org/going-digital/ai-intelligent-machines-smart-policies/conference-agenda/ai-intelligent-machines-smart-policies-kitano.pdf>.
8. Esper, M. (2020). Secretary of Defense Remarks for DOD Artificial Intelligence Symposium and Exposition.
9. Green, M.L., Takeuchi, I., and Hattrick-Simpers, J.R. (2013). Applications of high throughput (combinatorial) methodologies to electronic, magnetic, optical, and energy-related materials. *Journal of Applied Physics* *113*, 9\_1.
10. Allison, J., Backman, D., and Christodoulou, L. (2006). Integrated computational materials engineering: a new paradigm for the global materials profession. *Jom* *58*, 25-27.
11. Jain, A., Ong, S.P., Hautier, G., Chen, W., Richards, W.D., Dacek, S., Cholia, S., Gunter, D., Skinner, D., and Ceder, G. (2013). Commentary: The Materials Project: A materials genome approach to accelerating materials innovation. *Apl Materials* *1*, 011002.
12. Materials Genome Initiative. <https://www.mgi.gov/>.
13. Butler, K.T., Davies, D.W., Cartwright, H., Isayev, O., and Walsh, A. (2018). Machine learning for molecular and materials science. *Nature* *559*, 547-555.
14. Schmidt, J., Marques, M.R.G., Botti, S., and Marques, M.A.L. (2019). Recent advances and applications of machine learning in solid-state materials science. *npj Computational Materials* *5*, 83. 10.1038/s41524-019-0221-0.
15. Ramprasad, R., Batra, R., Pilania, G., Mannodi-Kanakkithodi, A., and Kim, C. (2017). Machine learning in materials informatics: recent applications and prospects. *npj Computational Materials* *3*, 1-13.
16. Wang, A.Y.-T., Murdock, R.J., Kauwe, S.K., Oliynyk, A.O., Gurlo, A., Brgoch, J., Persson, K.A., and Sparks, T.D. (2020). Machine Learning for Materials Scientists: An introductory guide towards best practices. *Chemistry of Materials*.
17. Sparkes, A., Aubrey, W., Byrne, E., Clare, A., Khan, M.N., Liakata, M., Markham, M., Rowland, J., Soldatova, L.N., and Whelan, K.E. (2010). Towards Robot Scientists for autonomous scientific discovery. *Automated Experimentation* *2*, 1.
18. Nikolaev, P., Hooper, D., Webber, F., Rao, R., Decker, K., Krein, M., Poleski, J., Barto, R., and Maruyama, B. (2016). Autonomy in materials research: a case study in carbon nanotube growth. *npj Computational Materials* *2*, 16031.

19. Häse, F., Roch, L.M., and Aspuru-Guzik, A. (2019). Next-generation experimentation with self-driving laboratories. *Trends in Chemistry* *1*, 282-291.

20. Kusne, A.G., Yu, H., Wu, C., Zhang, H., Hattrick-Simpers, J., DeCost, B., Sarker, S., Oses, C., Toher, C., Curtarolo, S., *et al.* (2020). On-the-fly closed-loop materials discovery via Bayesian active learning. *Nature Communications* *11*, 5966. 10.1038/s41467-020-19597-w.

21. Flores-Leonar, M.M., Mejía-Mendoza, L.M., Aguilar-Granda, A., Sanchez-Lengeling, B., Tribukait, H., Amador-Bedolla, C., and Aspuru-Guzik, A. (2020). Materials Acceleration Platforms: on the way to autonomous experimentation. *Current Opinion in Green and Sustainable Chemistry*, 100370.

22. MacLeod, B.P., Parlante, F.G., Morrissey, T.D., Häse, F., Roch, L.M., Dettelbach, K.E., Moreira, R., Yunker, L.P., Rooney, M.B., and Deeth, J.R. (2020). Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances* *6*, eaaz8867.

23. Correa-Baena, J.-P., Hippalgaonkar, K., van Duren, J., Jaffer, S., Chandrasekhar, V.R., Stevanovic, V., Wadia, C., Guha, S., and Buonassisi, T. (2018). Accelerating materials development via automation, machine learning, and high-performance computing. *Joule* *2*, 1410-1420.

24. Salley, D., Keenan, G., Grizou, J., Sharma, A., Martín, S., and Cronin, L. (2020). A nanomaterials discovery robot for the Darwinian evolution of shape programmable gold nanoparticles. *Nature Communications* *11*, 1-7.

25. Chang, J., Nikolaev, P., Carpena-Núñez, J., Rao, R., Decker, K., Islam, A.E., Kim, J., Pitt, M.A., Myung, J.I., and Maruyama, B. (2020). Efficient Closed-loop Maximization of Carbon Nanotube Growth Rate using Bayesian Optimization. *Scientific Reports* *10*, 1-9.

26. Montoya, J.H., Winther, K., Flores, R.A., Bligaard, T., Hummelshøj, J.S., and Aykol, M. (2020). Autonomous intelligent agents for accelerated materials discovery. *Chemical Science* *11*, 8517-8535.

27. Stein, H.S., and Gregoire, J.M. (2019). Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chemical Science* *10*, 9640-9649.

28. Davis, R., Shrobe, H., and Szolovits, P. (1993). What is a knowledge representation? *AI magazine* *14*, 17-17.

29. Bessa, M., Bostanabad, R., Liu, Z., Hu, A., Apley, D.W., Brinson, C., Chen, W., and Liu, W.K. (2017). A framework for data-driven analysis of materials under uncertainty: Countering the curse of dimensionality. *Computer Methods in Applied Mechanics and Engineering* *320*, 633-667.

30. DeCost, B., Hattrick-Simpers, J.R., Trautt, Z., Kusne, A.G., Campo, E., and Green, M.L. (2020). Scientific AI in materials science: a path to a sustainable and scalable paradigm. *Machine Learning: Science and Technology*.

31. Deneault, J.R., Chang, J., Myung, J., Hooper, D., Armstrong, A., Pitt, M., and Mauryama, B. (2021). Toward Autonomous Additive Manufacturing: Bayesian Optimization on a 3D Printer. *MRS Bulletin*.

32. Rao, R., Liptak, D., Cherukuri, T., Yakobson, B.I., and Maruyama, B. (2012). In situ evidence for chirality-dependent growth rates of individual carbon nanotubes. *Nature materials* *11*, 213-216.

33. Tanoto, A., and Rückert, U. (2012). Local navigation strategies for multi-robot exploration: from simulation to experimentation with mini-robots. *Procedia Engineering* *41*, 1197-1203.

34. Kuehl, R.O., and Kuehl, R. (2000). Design of experiments: statistical principles of research design and analysis.

35. Gongora, A.E., Xu, B., Perry, W., Okoye, C., Riley, P., Reyes, K.G., Morgan, E.F., and Brown, K.A. (2020). A Bayesian experimental autonomous researcher for mechanical design. *Science Advances* 6, eaaz1708. 10.1126/sciadv.aaz1708.

36. King, R.D., Rowland, J., Oliver, S.G., Young, M., Aubrey, W., Byrne, E., Liakata, M., Markham, M., Pir, P., and Soldatova, L.N. (2009). The automation of science. *Science* 324, 85-89.

37. Box, G.E. (1982). Choice of response surface design and alphabetic optimality. Paper presented at: Proceedings of the Conference on the Design of Experiments in Army Research, Development and Testing (US Army Research Office).

38. Lookman, T., Balachandran, P.V., Xue, D., Hogden, J., and Theiler, J. (2017). Statistical inference and adaptive design for materials discovery. *Current Opinion in Solid State and Materials Science* 21, 121-128.

39. Powell, W.B., and Ryzhov, I.O. (2012). Optimal learning, Vol 841 (John Wiley & Sons).

40. Ososky, S., Schuster, D., Phillips, E., and Jentsch, F.G. (2013). Building appropriate trust in human-robot teams. Paper presented at: 2013 AAAI spring symposium series.

41. Winfield, A.F., Michael, K., Pitt, J., and Evers, V. (2019). Machine Ethics: The Design and Governance of Ethical AI and Autonomous Systems [Scanning the Issue]. *Proceedings of the IEEE* 107, 509-517. 10.1109/JPROC.2019.2900622.

42. Shahriari, K., and Shahriari, M. (2017). IEEE standard review — Ethically aligned design: A vision for prioritizing human wellbeing with artificial intelligence and autonomous systems. Paper presented at: 2017 IEEE Canada International Humanitarian Technology Conference (IHTC). 10.1109/IHTC.2017.8058187.

43. de Pablo, J.J., Jones, B., Kovacs, C.L., Ozolins, V., and Ramirez, A.P. (2014). The materials genome initiative, the interplay of experiment, theory and computation. *Current Opinion in Solid State and Materials Science* 18, 99-117.

44. de Pablo, J.J., Jackson, N.E., Webb, M.A., Chen, L.-Q., Moore, J.E., Morgan, D., Jacobs, R., Pollock, T., Schlam, D.G., and Toberer, E.S. (2019). New frontiers for the materials genome initiative. *npj Computational Materials* 5, 41.

45. Green, M.L., Choi, C., Hattrick-Simpers, J., Joshi, A., Takeuchi, I., Barron, S., Campo, E., Chiang, T., Empedocles, S., and Gregoire, J. (2017). Fulfilling the promise of the materials genome initiative with high-throughput experimental methodologies. *Applied Physics Reviews* 4, 011105.

46. Cui, J., Chu, Y.S., Famodu, O.O., Furuya, Y., Hattrick-Simpers, J., James, R.D., Ludwig, A., Thienhaus, S., Wuttig, M., and Zhang, Z. (2006). Combinatorial search of thermoelastic shape-memory alloys with extremely small hysteresis width. *Nature materials* 5, 286-290.

47. Fujino, S., Murakami, M., Anbusathaiah, V., Lim, S.-H., Nagarajan, V., Fennie, C., Wuttig, M., Salamanca-Riba, L., and Takeuchi, I. (2008). Combinatorial discovery of a lead-free morphotropic phase boundary in a thin-film piezoelectric perovskite. *Applied Physics Letters* 92, 202904.

48. Burger, B., Maffettone, P.M., Gusev, V.V., Aitchison, C.M., Bai, Y., Wang, X., Li, X., Alston, B.M., Li, B., and Clowes, R. (2020). A mobile robotic chemist. *Nature* 583, 237-241.

49. Li, J., Lu, Y., Xu, Y., Liu, C., Tu, Y., Ye, S., Liu, H., Xie, Y., Qian, H., and Zhu, X. (2018). AIR-Chem: Authentic intelligent robotics for chemistry. *The Journal of Physical Chemistry A* 122, 9142-9148.

50. Lignos, I., Morad, V., Shynkarenko, Y., Bernasconi, C., Maceiczyk, R.M., Protesescu, L., Bertolotti, F., Kumar, S., Ochsenbein, S.T., and Masciocchi, N. (2018). Exploration of near-infrared-emissive colloidal multinary lead halide perovskite nanocrystals using an automated microfluidic platform. *ACS nano* *12*, 5504-5517.

51. Epps, R.W., Bowen, M.S., Volk, A.A., Abdel-Latif, K., Han, S., Reyes, K.G., Amassian, A., and Abolhasani, M. (2020). Artificial Chemist: An Autonomous Quantum Dot Synthesis Bot. *Advanced Materials*, 2001626.

52. Mekki-Berrada, F., Ren, Z., Huang, T., Wong, W.K., Zheng, F., Xie, J., Tian, I.P.S., Jayavelu, S., Mahfoud, Z., and Bash, D. (2020). Two-Step Machine Learning Enables Optimized Nanoparticle Synthesis.

53. Curtarolo, S., Hart, G.L., Nardelli, M.B., Mingo, N., Sanvito, S., and Levy, O. (2013). The high-throughput highway to computational materials design. *Nature materials* *12*, 191-201.

54. Thygesen, K.S., and Jacobsen, K.W. (2016). Making the most of materials computations. *Science* *354*, 180-181.

55. Curtarolo, S., Setyawan, W., Wang, S., Xue, J., Yang, K., Taylor, R.H., Nelson, L.J., Hart, G.L., Sanvito, S., and Buongiorno-Nardelli, M. (2012). AFLOWLIB. ORG: A distributed materials properties repository from high-throughput ab initio calculations. *Computational Materials Science* *58*, 227-235.

56. Choudhary, K., Garrity, K.F., Reid, A.C., DeCost, B., Biacchi, A.J., Walker, A.R.H., Trautt, Z., Hattrick-Simpers, J., Kusne, A.G., and Centrone, A. (2020). The joint automated repository for various integrated simulations (JARVIS) for data-driven materials design. *npj Computational Materials* *6*, 1-13.

57. Hachmann, J., Olivares-Amaya, R., Atahan-Evrenk, S., Amador-Bedolla, C., Sánchez-Carrera, R.S., Gold-Parker, A., Vogt, L., Brockway, A.M., and Aspuru-Guzik, A. (2011). The Harvard clean energy project: large-scale computational screening and design of organic photovoltaics on the world community grid. *The Journal of Physical Chemistry Letters* *2*, 2241-2251.

58. Draxl, C., and Scheffler, M. (2019). The NOMAD laboratory: from data sharing to artificial intelligence. *Journal of Physics: Materials* *2*, 036001.

59. Haghaghatali, M., Vishwakarma, G., Altarawy, D., Subramanian, R., Kota, B.U., Sonpal, A., Setlur, S., and Hachmann, J. (2019). Chemml: A machine learning and informatics program package for the analysis, mining, and modeling of chemical and materials data. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, e1458.

60. Mathew, K., Montoya, J.H., Faghaninia, A., Dwarakanath, S., Aykol, M., Tang, H., Chu, I.-h., Smidt, T., Bocklund, B., and Horton, M. (2017). Atomate: A high-level interface to generate, execute, and analyze computational materials science workflows. *Computational Materials Science* *139*, 140-152.

61. Sun, S., Tiihonen, A., Oviedo, F., Liu, Z., Thapa, J., Hartono, N.T., Goyal, A., Batali, C., Encinas, A., Yoo, J., *et al.* (2020). A Physical Data Fusion Approach to Optimize Compositional Stability of Halide Perovskites.

62. Morgan, D., and Jacobs, R. (2020). Opportunities and challenges for machine learning in materials science. *Annual Review of Materials Research* *50*.

63. Schleider, G.R., Padilha, A.C., Acosta, C.M., Costa, M., and Fazzio, A. (2019). From DFT to machine learning: recent approaches to materials science—a review. *Journal of Physics: Materials* *2*, 032001.

64. Vasudevan, R.K., Choudhary, K., Mehta, A., Smith, R., Kusne, G., Tavazza, F., Vlcek, L., Ziatdinov, M., Kalinin, S.V., and Hattrick-Simpers, J. (2019). Materials science in the artificial intelligence age: high-throughput library generation, machine learning, and a pathway from correlations to the underpinning physics. *MRS Communications* *9*, 821-838.

65. Hong, Y., Hou, B., Jiang, H., and Zhang, J. (2020). Machine learning and artificial neural network accelerated computational discoveries in materials science. *Wiley Interdisciplinary Reviews: Computational Molecular Science* *10*, e1450.

66. Chibani, S., and Coudert, F.-X. (2020). Machine learning approaches for the prediction of materials properties. *APL Materials* *8*, 080701.

67. Suh, C., Fare, C., Warren, J.A., and Pyzer-Knapp, E.O. (2020). Evolving the Materials Genome: How Machine Learning Is Fueling the Next Generation of Materials Discovery. *Annual Review of Materials Research* *50*.

68. Saal, J.E., Oliynyk, A.O., and Meredig, B. (2020). Machine Learning in Materials Discovery: Confirmed Predictions and Their Underlying Approaches. *Annual Review of Materials Research* *50*.

69. Sparks, T.D., Kauwe, S.K., Parry, M.E., Tehrani, A.M., and Brgoch, J. (2020). Machine Learning for Structural Materials. *Annual Review of Materials Research* *50*.

70. Hegde, V.I., Aykol, M., Kirklin, S., and Wolverton, C. (2020). The phase stability network of all inorganic materials. *Science Advances* *6*, eaay5606.

71. Ganose, A.M., and Jain, A. (2019). Robocrystallographer: automated crystal structure text descriptions and analysis. *MRS Communications* *9*, 874-881.

72. Duros, V., Grizou, J., Sharma, A., Mehr, S.H.M., Bubliauskas, A., Frei, P., Miras, H.N., and Cronin, L. (2019). Intuition-enabled machine learning beats the competition when joint human-robot teams perform inorganic chemical experiments. *Journal of chemical information and modeling* *59*, 2664-2671.

73. Yeung, C., Tsai, J.-M., King, B., Kawagoe, Y., Ho, D., Knight, M.W., and Raman, A.P. (2020). Elucidating the Behavior of Nanophotonic Structures through Explainable Machine Learning Algorithms. *ACS Photonics* *7*, 2309-2318.

74. Feng, J., Lansford, J.L., Katsoulakis, M.A., and Vlachos, D.G. (2020). Explainable and trustworthy artificial intelligence for correctable modeling in chemical sciences. *Science advances* *6*, eabc3204.

75. Nikolaev, P., Hooper, D., Perea-López, N., Terrones, M., and Maruyama, B. (2014). Discovery of Wall-Selective Carbon Nanotube Growth Conditions via Automated Experimentation. *ACS Nano* *8*, 10214-10222. 10.1021/nn503347a.

76. Kluender, E.J., Hedrick, J.L., Brown, K.A., Rao, R., Meckes, B., Du, J., Moreau, L., Maruyama, B., and Mirkin, C.A. (2019). Catalyst discovery through megalibraries of nanomaterials. *Proceedings of the National Academy of Sciences of the United States of America* *116* 40-45.

77. Bédard, A.-C., Adamo, A., Aroh, K.C., Russell, M.G., Bedermann, A.A., Torosian, J., Yue, B., Jensen, K.F., and Jamison, T.F. (2018). Reconfigurable system for automated optimization of diverse chemical reactions. *Science* *361*, 1220-1225.

78. Huyer, W., and Neumaier, A. (2008). SNOBFIT--stable noisy optimization by branch and fit. *ACM Transactions on Mathematical Software (TOMS)* *35*, 1-25.

79. Coley, C.W., Thomas, D.A., Lummiss, J.A., Jaworski, J.N., Breen, C.P., Schultz, V., Hart, T., Fishman, J.S., Rogers, L., and Gao, H. (2019). A robotic platform for flow synthesis of organic compounds informed by AI planning. *Science* *365*, eaax1566.

80. Burger, B., Maffettone, P.M., Gusev, V.V., Aitchison, C.M., Bai, Y., Wang, X., Li, X., Alston, B.M., Li, B., Clowes, R., *et al.* (2020). A mobile robotic chemist. *Nature* *583*, 237-241. 10.1038/s41586-020-2442-2.

81. Grizou, J., Points, L.J., Sharma, A., and Cronin, L. (2020). A curious formulation robot enables the discovery of a novel protocell behavior. *Science advances* *6*, eaay4237.

82. Gongora, A.E., Snapp, K.L., Whiting, E., Riley, P., Reyes, K.G., Morgan, E.F., and Brown, K.A. (2020). Using Simulation to Accelerate Autonomous Experimentation (AE): A Case Study Using Mechanics. [10.2139/ssrn.3751791](https://doi.org/10.2139/ssrn.3751791).

83. Sun, S., Hartono, N.T.P., Ren, Z.D., Oviedo, F., Buscemi, A.M., Layurova, M., Chen, D.X., Ogunfunmi, T., Thapa, J., Ramasamy, S., *et al.* (2019). Accelerated Development of Perovskite-Inspired Materials via High-Throughput Synthesis and Machine-Learning Diagnosis. *Joule* *3*, 1437-1451. <https://doi.org/10.1016/j.joule.2019.05.014>.

84. Langner, S., Häse, F., Perea, J.D., Stubhan, T., Hauch, J., Roch, L.M., Heumueller, T., Aspuru-Guzik, A., and Brabec, C.J. (2020). Beyond Ternary OPV: High-Throughput Experimentation and Self-Driving Laboratories Optimize Multicomponent Systems. *Advanced Materials* *32*, 1907801.

85. Chan, E.M., Xu, C., Mao, A.W., Han, G., Owen, J.S., Cohen, B.E., and Milliron, D.J. (2010). Reproducible, high-throughput synthesis of colloidal nanocrystals for optimization in multidimensional parameter space. *Nano letters* *10*, 1874-1885.

86. Higgins, K., Valletti, S.M., Ziatdinov, M., Kalinin, S.V., and Ahmadi, M. (2020). Chemical Robotics Enabled Exploration of Stability in Multicomponent Lead Halide Perovskites via Machine Learning. *ACS Energy Letters* *5*, 3426-3436.

87. Kirman, J., Johnston, A., Kuntz, D.A., Askerka, M., Gao, Y., Todorović, P., Ma, D., Privé, G.G., and Sargent, E.H. (2020). Machine-Learning-Accelerated Perovskite Crystallization. *Matter*.

88. Li, Z., Najeeb, M.A., Alves, L., Sherman, A.Z., Shekar, V., Cruz Parrilla, P., Pendleton, I.M., Wang, W., Nega, P.W., Zeller, M., *et al.* (2020). Robot-Accelerated Perovskite Investigation and Discovery. *Chemistry of Materials* *32*, 5650–5663.

89. Pendleton, I.M., Cattabriga, G., Li, Z., Najeeb, M.A., Friedler, S.A., Norquist, A.J., Chan, E.M., and Schrier, J. (2019). Experiment Specification, Capture and Laboratory Automation Technology (ESCALATE): a software pipeline for automated chemical experimentation and data management. *MRS Communications* *9*, 846-859.

90. Noack, M.M., Yager, K.G., Fukuto, M., Doerk, G.S., Li, R., and Sethian, J.A. (2019). A kriging-based approach to autonomous experimentation with applications to x-ray scattering. *Scientific reports* *9*, 1-19.

91. Rakita, Y., O’Nolan, D., McAuliffe, R.D., Veith, G.M., Chupas, P.J., Billinge, S.J., and Chapman, K.W. (2020). Active Reaction Control of Cu Redox State Based on Real-Time Feedback from In Situ Synchrotron Measurements. *Journal of the American Chemical Society*.

92. Ren, F., Ward, L., Williams, T., Laws, K.J., Wolverton, C., Hattrick-Simpers, J., and Mehta, A. (2018). Accelerated discovery of metallic glasses through iteration of machine learning and high-throughput experiments. *Science advances* *4*, eaaq1566.

93. Kusne, A.G. U.S. Provisional Patent Application No. 63/151,263.

94. Montoya, J.H., Winther, K.T., Flores, R.A., Bligaard, T., Hummelshøj, J.S., and Aykol, M. (2020). Autonomous intelligent agents for accelerated materials discovery. *Chemical Science* *11*, 8517-8532. 10.1039/D0SC01101K.

95. Frazier, P.I., and Wang, J. (2016). Bayesian optimization for materials design. In *Information Science for Materials Discovery and Design* (Springer), pp. 45-75.

96. University of Wisconsin-Madison Department of Computer Sciences. (2009). Active learning literature survey (University of Wisconsin-Madison Department of Computer Sciences).

97. Fisher, R.A. (1936). Design of experiments. *Br Med J 1*, 554-554.

98. Montgomery, D.C. (2017). *Design and analysis of experiments* (John wiley & sons).

99. Williams, C.K., and Rasmussen, C.E. (2006). *Gaussian processes for machine learning*, Vol 2 (MIT press Cambridge, MA).

100. Bull, A.D. (2011). Convergence Rates of Efficient Global Optimization Algorithms. *Journal of Machine Learning Research 12*.

101. Settles, B., Craven, M., and Ray, S. (2008). Multiple-instance active learning. Paper presented at: *Advances in neural information processing systems*.

102. Cohn, D.A., Ghahramani, Z., and Jordan, M.I. (1996). Active learning with statistical models. *Journal of artificial intelligence research 4*, 129-145.

103. Cressie, N. (1990). The origins of kriging. *Mathematical geology 22*, 239-252.

104. Chen, S., Reyes, K.-R.G., Gupta, M.K., McAlpine, M.C., and Powell, W.B. (2015). Optimal learning in experimental design using the knowledge gradient policy with application to characterizing nanoemulsion stability. *SIAM/ASA Journal on Uncertainty Quantification 3*, 320-345.

105. He, X., Reyes, K.G., and Powell, W.B. (2020). Optimal Learning with Local Nonlinear Parametric Models over Continuous Designs. *SIAM Journal on Scientific Computing 42*, A2134-A2157.

106. Raza, S.E.A., Cheung, L., Epstein, D., Pelengaris, S., Khan, M., and Rajpoot, N.M. (2017). Mimo-net: A multi-input multi-output convolutional neural network for cell segmentation in fluorescence microscopy images. Paper presented at: *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)* (IEEE).

107. Bernardo, J.M., and Rueda, R. (2002). Bayesian hypothesis testing: A reference approach. *International Statistical Review 70*, 351-372.

108. Hoeting, J.A., Madigan, D., Raftery, A.E., and Volinsky, C.T. (1999). Bayesian model averaging: a tutorial. *Statistical science*, 382-401.

109. Robinson, T., Eldred, M., Willcox, K., and Haimes, R. (2006). Strategies for multifidelity optimization with variable dimensional hierarchical models. Paper presented at: *47th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference 14th AIAA/ASME/AHS Adaptive Structures Conference 7th*. 10.2514/6.2006-1819.

110. Pilania, G., Gubernatis, J.E., and Lookman, T. (2017). Multi-fidelity machine learning models for accurate bandgap predictions of solids. *Computational Materials Science 129*, 156-163.

111. Sen, R., Kandasamy, K., and Shakkottai, S. (2018). Multi-Fidelity Black-Box Optimization with Hierarchical Partitions. Paper presented at: *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research: PMLR)*.

112. Poloczek, M., Wang, J., and Frazier, P. (2017). Multi-information source optimization. Paper presented at: *Advances in Neural Information Processing Systems*.

113. Sutton, R.S., and Barto, A.G. (2018). *Reinforcement learning: An introduction* (MIT press).

114. Bertsekas, D.P. (1995). Dynamic programming and optimal control, Vol 1 (Athena scientific Belmont, MA).

115. Puterman, M.L. (2014). Markov decision processes: discrete stochastic dynamic programming (John Wiley & Sons).

116. Goodfellow, I., Bengio, Y., Courville, A., and Bengio, Y. (2016). Deep learning, Vol 1 (MIT press Cambridge).

117. Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. Paper presented at: International conference on machine learning.

118. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., Lanctot, M., Sifre, L., Kumaran, D., and Graepel, T. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* *362*, 1140-1144.

119. Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing atari with deep reinforcement learning. arXiv preprint arXiv:13125602.

120. LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature* *521*, 436-444.

121. Rivenson, Y., Göröcs, Z., Günaydin, H., Zhang, Y., Wang, H., and Ozcan, A. (2017). Deep learning microscopy. *Optica* *4*, 1437-1443.

122. Zhu, Y., Ouyang, Q., and Mao, Y. (2017). A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy. *BMC bioinformatics* *18*, 1-10.

123. Madireddy, S., Chung, D.-W., Loeffler, T., Sankaranarayanan, S.K., Seidman, D.N., Balaprakash, P., and Heinonen, O. (2019). Phase Segmentation in Atom-Probe Tomography Using Deep Learning-Based Edge Detection. *Scientific reports* *9*, 1-10.

124. Tan, C., Sun, F., Kong, T., Zhang, W., Yang, C., and Liu, C. (2018). A survey on deep transfer learning. Paper presented at: International conference on artificial neural networks (Springer).

125. Roy, D.M., and Kaelbling, L.P. (2007). Efficient Bayesian Task-Level Transfer Learning. Paper presented at: IJCAI.

126. Sattari Baboukani, B., Ye, Z., G Reyes, K., and Nalam, P.C. (2020). Prediction of Nanoscale Friction for Two-Dimensional Materials Using a Machine Learning Approach. *Tribology Letters* *68*, 1-14.

127. Seffers, G.I. (2017). Scientists Pick AI for Lab Partner Signals. <https://www.afcea.org/content/scientists-pick-ai-lab-partner>.

128. Mokyr, J. (2018). A Culture of Growth (Princeton, NJ: Princeton University Press).

129. Aykol, M., Hegde, V.I., Hung, L., Suram, S., Herring, P., Wolverton, C., and Hummelshøj, J.S. (2019). Network analysis of synthesizable materials discovery. *Nature communications* *10*, 1-7.

130. Cao, B., Adutwum, L.A., Oliynyk, A.O., Luber, E.J., Olsen, B.C., Mar, A., and Buriak, J.M. (2018). How to optimize materials and devices via design of experiments and machine learning: Demonstration using organic photovoltaics. *ACS nano* *12*, 7434-7444.

131. Javdekar, C., Watson, E., Kapilow, V., Bograd, M., Boyer, P., Zeid, I., and Duggan, C. (2016). Closing the advanced manufacturing talent gap. *Procedia Manufacturing* *5*, 1197-1207.

132. Raccuglia, P., Elbert, K.C., Adler, P.D., Falk, C., Wenny, M.B., Mollo, A., Zeller, M., Friedler, S.A., Schrier, J., and Norquist, A.J. (2016). Machine-learning-assisted materials discovery using failed experiments. *Nature* *533*, 73-76. 10.1038/nature17439.

133. Yaqub, O. (2018). Serendipity: Towards a taxonomy and a theory. *Research Policy* *47*, 169-179.

134. Wikipedia: Louis Pasteur. [https://en.wikiquote.org/wiki/Louis\\_Pasteur](https://en.wikiquote.org/wiki/Louis_Pasteur).

135. Sun, S., Tiihonen, A., Oviedo, F., Liu, Z., Thapa, J., Hartono, N.T.P., Goyal, A., Batali, C., Encinas, A., and Yoo, J. (2020). A Physical Data Fusion Approach to Optimize Compositional Stability of Halide Perovskites. *Proceedings of the National Academy of Sciences* *116*, 18142-18147. 10.1073/pnas.1900548116.

137. Schmidt, M., and Lipson, H. (2010). Symbolic regression of implicit equations. In *Genetic Programming Theory and Practice VII* (Springer), pp. 73-85.

138. Park, C., Rao, R., Nikolaev, P., and Maruyama, B. (2019). Gaussian process metamodeling for experiments with manipulating factors. *arXiv preprint arXiv:191100955*.

139. Segler, M.H.S., Preuss, M., and Waller, M.P. (2018). Planning chemical syntheses with deep neural networks and symbolic AI. *Nature* *555*, 604-610. 10.1038/nature25978.

140. Korovina, K., Xu, S., Kandasamy, K., Neiswanger, W., Poczos, B., Schneider, J., and Xing, E. (2020). Chembo: Bayesian optimization of small organic molecules with synthesizable recommendations. Paper presented at: International Conference on Artificial Intelligence and Statistics.

141. National Academies of Sciences, E., and Medicine (2018). Artificial Intelligence and Machine Learning to Accelerate Translational Research: Proceedings of a Workshop—in Brief.

142. Haenlein, M., and Kaplan, A. (2019). A brief history of artificial intelligence: On the past, present, and future of artificial intelligence. *California management review* *61*, 5-14.

143. Gomes, C.P., Selman, B., and Gregoire, J.M. (2019). Artificial intelligence for materials discovery. *MRS Bulletin* *44*, 538-544.

144. Scientific Autonomous Reasoning Agent (SARA): Integrating Materials Theory, Experiment and Computation. <https://community.apan.org/wg/afosr/w/researchareas/22949/scientific-autonomous-reasoning-agent-sara-integrating-materials-theory-experiment-and-computation/>.

145. Kitano, H. (2016). Artificial intelligence to win the nobel prize and beyond: Creating the engine for scientific discovery. *AI magazine* *37*, 39-49.

146. Horwath, J.P., Zakharov, D.N., Mégret, R., and Stach, E.A. (2020). Understanding important features of deep learning models for segmentation of high-resolution transmission electron microscopy images. *npj Computational Materials* *6*, 1-9.

147. Park, C., and Ding, Y. (2019). Automating material image analysis for material discovery. *MRS Communications* *9*, 545-555.

148. Melton, C.N., Noack, M.M., Ohta, T., Beechem, T.E., Robinson, J., Xiaotian Zhang, Bostwick, A., Jozwiak, C., Koch, R., Zwart, P.H., *et al.* (2020). K-Means-Driven Gaussian Process Data Collection for Angle-Resolved Photoemission Spectroscopy Machine Learning Science and Technology *in review*.

149. Lake, B.M., Salakhutdinov, R., and Tenenbaum, J.B. (2015). Human-level concept learning through probabilistic program induction. *Science* *350*, 1332-1338.

150. Olson, M., Wyner, A., and Berk, R. (2018). Modern neural networks generalize on small data sets. Paper presented at: *Advances in Neural Information Processing Systems*.

151. Arulkumaran, K., Deisenroth, M.P., Brundage, M., and Bharath, A.A. (2017). A brief survey of deep reinforcement learning. *arXiv preprint arXiv:170805866*.

152. Weiss, K., Khoshgoftaar, T.M., and Wang, D. (2016). A survey of transfer learning. *Journal of Big data* 3, 9.

153. Chen, D., Bai, Y., Zhao, W., Ament, S., Gregoire, J.M., and Gomes, C.P. Deep Reasoning Networks for Unsupervised Pattern De-mixing with Constraint Reasoning.

154. Wilkinson, M.D., Dumontier, M., Aalbersberg, I.J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J.-W., da Silva Santos, L.B., and Bourne, P.E. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data* 3, 1-9.

155. Blaiszik, B., Chard, K., Pruyne, J., Ananthakrishnan, R., Tuecke, S., and Foster, I. (2016). The Materials Data Facility: Data services to advance materials science research. *JOM* 68, 2045-2052.

156. Aagesen, L., Adams, J., Allison, J., Andrews, W., Araullo-Peters, V., Berman, T., Chen, Z., Daly, S., Das, S., and DeWitt, S. (2018). Prisms: An integrated, open-source framework for accelerating predictive structural materials science. *JOM* 70, 2298-2314.

157. Chard, R., Li, Z., Chard, K., Ward, L., Babuji, Y., Woodard, A., Tuecke, S., Blaiszik, B., Franklin, M., and Foster, I. (2019). DLHub: Model and data serving for science. Paper presented at: 2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS) (IEEE).

158. Roch, L.M., Häse, F., Kreisbeck, C., Tamayo-Mendoza, T., Yunker, L.P., Hein, J.E., and Aspuru-Guzik, A. (2020). ChemOS: An orchestration software to democratize autonomous discovery. *PLoS One* 15, e0229862.

159. Reker, D., Hoyt, E.A., Bernardes, G.J., and Rodrigues, T. (2020). Adaptive Optimization of Chemical Reactions with Minimal Experimental Information. *Cell Reports Physical Science* 1, 100247.

160. Li, J., Tu, Y., Liu, R., Lu, Y., and Zhu, X. (2020). Toward “On-Demand” Materials Synthesis and Scientific Discovery through Intelligent Robots. *Advanced Science* 7, 1901957.

161. Bluesky. <https://blueskyproject.io>

162. Reyes, K., Chen, S., Li, Y., and Powell, W.B. (2015). Quantifying experimental characterization choices in Optimal Learning and materials design. Paper presented at: TMS 2015 144th Annual Meeting & Exhibition (Springer).

163. Wulf, W.A. (1993). The collaborative opportunity. *Science* 261, 854-856.

164. The Minerals, Metals & Materials Society (TMS). (2019). Creating the next-generation Materials Genome Initiative workforce (Pittsburgh, PA: The Minerals, Metals & Materials Society (TMS)). DOI: 10.7449/mgiworkforce\_1. [www.tms.org/mgiworkforce](http://www.tms.org/mgiworkforce).

165. Amershi, S., Weld, D., Vorvoreanu, M., Journey, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., *et al.* (2019). Guidelines for Human-AI Interaction. Paper presented at: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK: Association for Computing Machinery). 10.1145/3290605.3300233.

166. Carter, S., and Nielsen, M. (2017). Using Artificial Intelligence to Augment Human Intelligence. *Distill*.

167. Nikolaidis, S., and Shah, J. (2012). Human-robot teaming using shared mental models. *ACM/IEEE HRI*.

168. Reynolds, E.B., and Uygun, Y. (2018). Strengthening advanced manufacturing innovation ecosystems: The case of Massachusetts. *Technological Forecasting and Social Change* 136, 178-191.

169. Lei, R.F., Green, E.R., Leslie, S.J., and Rhodes, M. (2019). Children lose confidence in their potential to “be scientists,” but not in their capacity to “do science”. *Developmental science* 22, e12837.

170. Blaiszik, B., Ward, L., Schwarting, M., Gaff, J., Chard, R., Pike, D., Chard, K., and Foster, I. (2019). A data ecosystem to support machine learning in materials science. *MRS Communications* 9, 1125-1133.

171. Li, J., Ballmer, S.G., Gillis, E.P., Fujii, S., Schmidt, M.J., Palazzolo, A.M., Lehmann, J.W., Morehouse, G.F., and Burke, M.D. (2015). Synthesis of many different types of organic small molecules using one automated process. *Science* 347, 1221-1226.

172. Bloom, N., Jones, C.I., Van Reenen, J., and Webb, M. (2020). Are ideas getting harder to find? *American Economic Review* 110, 1104-1144.

## Figure Captions

### Figure 1. Schematic showing the different parts of an autonomous experimentation (AE) campaign.

An AE campaign comprises an iterative research loop that is carried out autonomously towards a research objective. *Initialize*: Before the first research cycle, the AE system incorporates information provided by human researchers to initialize the campaign. *Plan*: In this first part of the research loop, the system considers the predefined campaign objective and the most recent knowledge base and plans the next experiments to be pursued. *Experiment*: The experiments—which are broadly defined to encompass physical tests, modeling/simulation, or data mining—are carried out without human intervention. *Analyze*: Finally, the AE system uses the output data of the experiments to update the knowledge base, which will then be used in the planning of the next loop. *Conclusion*: Once the campaign objective is reached or some other criteria is met, the system completes the AE campaign and discontinues the iterations of planning, experiment, and analysis. Further development of this new research process is expected to significantly increase the efficiency of scientific investigations and completely shift the way research is carried out.

### Figure 2. Schematic showing the autonomous experimentation (AE) campaign of ARES (Autonomous Research System), the first reported AE system.<sup>18</sup>

*Initialize*: ARES was provided with a database to seed its AI planner algorithm of 140 input synthesis condition with resultant growth rates. *Plan*: ARES used the subsequent database to train a random-forest model, which it used to determine the reaction conditions of the first experimental phase, beginning the first cycle of planning, experiment, and analysis. *Experiment*: Using automated apparatus, carbon nanotubes (CNTs) were synthesized via chemical vapor deposition (CVD), and CNT growth was tracked via *in-situ* Raman spectroscopy. Left inset shows the experimental set-up: an array of pillars for experiments with a laser heating one pillar. Right inset shows the time series of spectra (waterfall plot), revealing CNT growth via the increasing intensity of the G peak with time. *Analyze*: The growth rate,  $v$ , of each experiment was extracted as shown by plotting G-peak intensity versus time. Along with the results of previous experiments (input conditions and output results), the results were analyzed and used to update the random-forest knowledge representation. *Plan*: Considering the latest knowledge representation, the AI planner once again decides on new experimental input conditions to target growth rates using a genetic algorithm. *Conclude*: After hundreds of iterations, the system converged on the maximum growth rate, demonstrating that ARES taught itself to grow CNTs at controlled rates.

### Figure 3. Schematic showing features of autonomous experimentation (AE) campaigns of BEAR (Bayesian Experimental Autonomous Researcher).<sup>35, 82</sup>

BEAR is an AE system for producing and mechanically testing additively manufactured components. *Initialize*: The diagrams define the strut thickness ( $t$ ), strut radius ( $r$ ), number of struts ( $n$ ), and twist ( $\theta$ ) of the components. These four parameters were varied to optimize toughness, the campaign objective. Performance was measured during uniaxial compression in which the structure was compressed by displacement  $D$ . *Plan*: The plot is an example of how parameter selection in one of BEAR's AE campaigns progressed with campaign time. Planning involved Bayesian optimization (BO) in all AE campaigns; in a set of campaigns, finite-element

modeling of prior physical data was also included through transfer learning to evaluate the inclusion of this data into the AE campaign. *Experiment*: The image shows the automated experimental apparatus of BEAR, where components were manufactured and tested for toughness ( $U$ ). *Analyze*: Plot showing how  $U$  was obtained by measuring the force ( $F$ ) as  $D$  was varied, adding to the knowledge base. *Conclude*: BEAR was used to benchmark the performance of AE by exploring the mechanical toughness of components that were either chosen from a grid or by an active-learning algorithm. Plotted are the median performance of the grid-based exploration and the AE campaigns. Even after 60 experiments, AE (blue diamonds) outperformed the 1800 experiments chosen from a grid (black squares). Providing the system with prior information about physical response (orange triangles) led to a +30% improvement in median performance.

**Figure 4. Schematic showing the expected exponential increase of the speed of research as autonomous experimentation (AE) is further developed.** We see a progression from connected AE systems to locally integrated systems, and finally to globally integrated systems. At a critical (or tipping) point, integration will create network effects that multiply the contribution of individual research nodes, greatly increasing research speed. Global integration and reduced cost will exponentially impact the access of researchers to AE systems. By leveraging network effects from data sharing and interpretation, and from the community-driven approach to scientific investigation, we anticipate solutions to currently intractable problems.

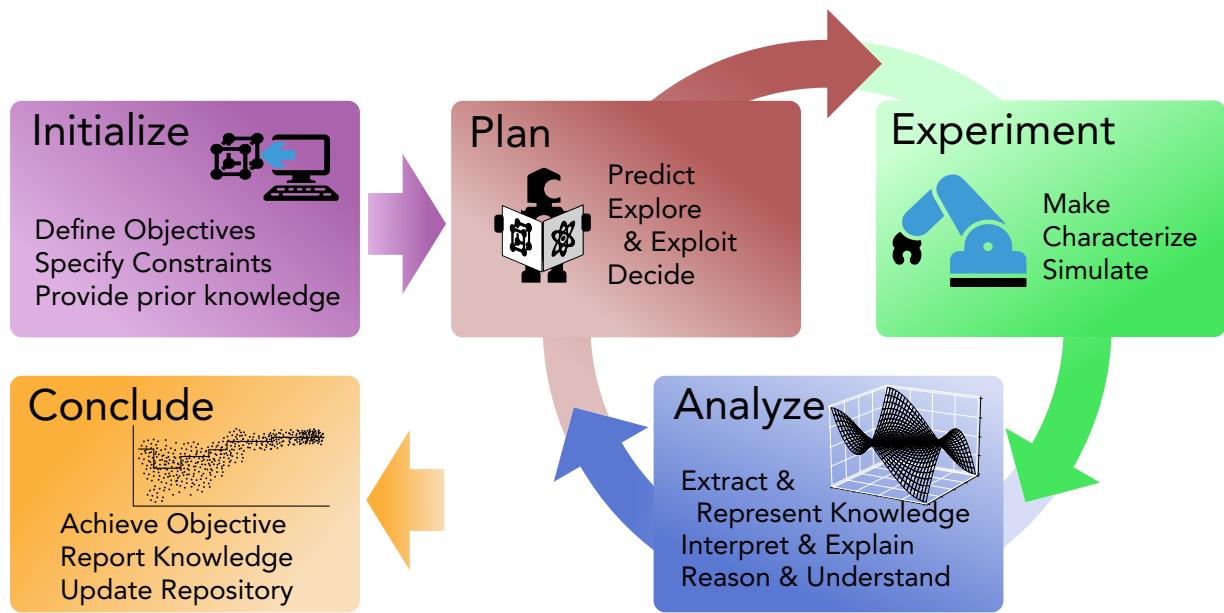
## Table

**Table 1. Selected Autonomous Experimentation (AE) Campaigns for Materials**

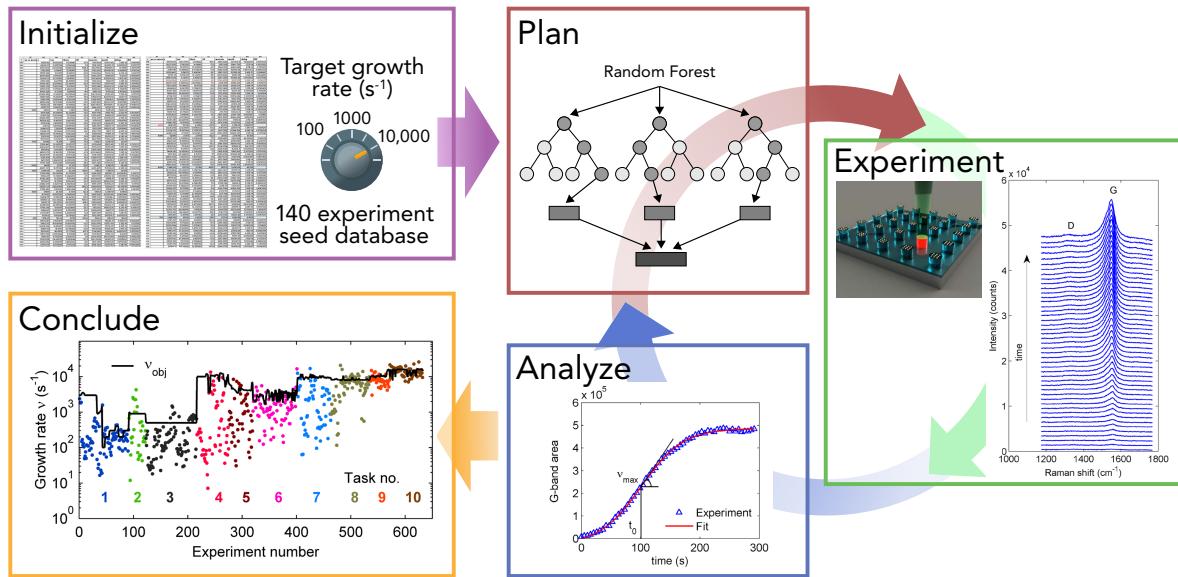
**Development.** The breadth of materials classes, synthesis methods, and characterization methods reveals the versatility of AE, and the benchmarked examples show that AE has successfully accelerated the research process. (UV, vis, and NIR represent ultraviolet, visible, and near infrared, respectively.)

Study (including publication date)	Material class and synthesis method	Characterization method	Planning and learning algorithm	Experimental campaigns and objectives	Metric of acceleration and benchmarking
(A) Autonomy in materials research: a case study in carbon nanotube growth (Oct 2016) <sup>18</sup> see Figure 2 and the ARES example in the text	Chemical vapor deposition of nano-materials	<i>in situ</i> Raman spectroscopy	Random-forest model	600 experiments to obtain a controlled growth rate	none
(B) Reconfigurable system for automated optimization of diverse chemical reactions (Sept 2018) <sup>77</sup>	Flow-based chemistry of soluble molecules	High-performance liquid chromatography	Blackbox optimization software (SNOBFIT)	112 experiments to optimize three chemical reactions	none
(C) A curious formulation robot enables the discovery of a novel protocell behavior (Jan 2020) <sup>81</sup>	Syringe-based liquid handling of oil-in-water emulsions	Optical imaging	Random goal exploration on a support vector regressor	1000 experiments to explore temperature response of emulsions	8× reduction in the number of experiments needed to match the performance of 1000 random experiments
(D) A Bayesian experimental autonomous researcher for mechanics (April 2020) <sup>35</sup> see Figure 3 and the BEAR example in text	Additive manufacturing of structural polymers	Mechanical uniaxial compression, weight measurement, optical imaging	Bayesian optimization	6 repetitions of 100 to maximize component toughness	55× reduction in number of experiments needed to match 1800 experiments on a grid
(E) Self-driving laboratory for accelerated discovery of thin-film materials (May 2020) <sup>22</sup>	Spin-coating of mixtures of photoactive chemicals	Dark-field photography, UV-vis–NIR spectroscopy, 4-point probe	Bayesian optimization	2 campaigns, each with 35 experiments to maximize hole mobility	none
(F) Artificial chemist: an autonomous quantum dot synthesis (June 2020) <sup>51</sup>	Solution-phase quantum-dot halide exchange reaction in flow	<i>in situ</i> UV-vis absorption and photoluminescence spectroscopy	Neural network ensemble, Bayesian optimization	X campaigns of 25 experiments to obtain a target emission energy with maximized brightness	Comparison of decision-making policies and role of pre-training
(G) A mobile robotic chemist (July 2020) <sup>80</sup>	Vial-based solution chemistry	Gas chromatography	Batched, constrained, Bayesian optimization algorithm	688 experiments performed over 8 days	none

**Figure 1**



**Figure 2**



**Figure 3**

