

Digital Discovery

Accepted Manuscript

This article can be cited before page numbers have been issued, to do this please use: S. Bagchi, A. Biswas, P. V. Balachandran, A. Ghosh and P. Ganesh, *Digital Discovery*, 2026, DOI: 10.1039/D6DD00049E.



This is an Accepted Manuscript, which has been through the Royal Society of Chemistry peer review process and has been accepted for publication.

Accepted Manuscripts are published online shortly after acceptance, before technical editing, formatting and proof reading. Using this free service, authors can make their results available to the community, in citable form, before we publish the edited article. We will replace this Accepted Manuscript with the edited and formatted Advance Article as soon as it is available.

You can find more information about Accepted Manuscripts in the [Information for Authors](#).

Please note that technical editing may introduce minor changes to the text and/or graphics, which may alter content. The journal's standard [Terms & Conditions](#) and the [Ethical guidelines](#) still apply. In no event shall the Royal Society of Chemistry be held responsible for any errors or omissions in this Accepted Manuscript or any consequences arising from the use of any information it contains.

Towards "on-demand" van der Waals epitaxy with adaptive ensemble sampling atomistic workflows

Soumendu Bagchi,^{*,†} Ankita Biswas,[‡] Prasanna V. Balachandran,[‡] Ayana Ghosh,^{*,¶} and P. Ganesh^{*,†}

[†]*Center for Nanophase Materials Sciences, Oak Ridge National Laboratory, Oak Ridge, TN, 37831, USA*

[‡]*Department of Materials Science and Engineering, University of Virginia, Charlottesville, Virginia, 22902, United States*

[¶]*Department of Physics, Indian Institute of Technology, Madras, Chennai, 600036, Tamil Nadu, India*

E-mail: bagchis@ornl.gov; ghosha@iitm.ac.in; ganeshp@ornl.gov

Abstract

Traditional approaches to achieve targeted epitaxial growth involves exploring a vast parameter space of thermodynamical and kinetic drivers (e.g., temperature, pressure, chemical potential etc). This tedious and time-consuming approach becomes particularly cumbersome to accelerate synthesis and characterization of novel materials with complex dependencies on local chemical environment, temperature and lattice-strains, specifically for nanoscale heterostructures of layered 2D materials. We combine the strength of next generation supercomputers at the extreme scale, machine learning and classical molecular dynamics simulations within an adaptive real time closed-loop



virtual environment steered by Bayesian algorithms to enable asynchronous ensemble sampling of the synthesis space, and apply it to the recrystallization phenomena of amorphous transition-metal dichalcogenide (TMDC) bilayer to form stacked moiré superstructures under various growth parameters. We show that such batch parallel Bayesian optimization-based online ensemble sampling frameworks for materials simulations can be promising towards achieving and accelerating on-demand epitaxy of van der Waals stacked moiré devices, paving the way towards a robust autonomous materials synthesis pipeline to enable discovery of unprecedented functionalities.

Introduction

Accelerating conventional synthesis efforts in van der Waals (vdW) heterostructures, is an important aspect to discover novel pathways which can harness unprecedented functionalities in 2D materials.¹ Recent surge in efforts towards realizing autonomous synthesis^{2,3} and characterization^{4,5} frameworks is expediting such discoveries. However, narrowing down the optimal thermodynamic and kinetic parameter space is a nontrivial automation challenge, especially when gaining finer control over tuning these emerging functionality *on demand* is the ultimate goal. For instance, vertically stacked van der Waals (vdW) heterostructures⁶ with optimal interlayer twists are extremely promising for manifesting tunable and novel functionalities at the nanoscale. These include emerging quantum phenomena, e.g., correlated electronic phases,⁷ unconventional superconductivity⁸ and interlayer excitons⁹ etc. which can enable unprecedented optical and electronic properties in nanoscale devices. Traditionally, top-down methods such as mechanical manipulations of as-grown epitaxial layers (via release, transfer and stamping¹⁰ etc.) have been a preferred mode to produce heterostructures with tunable geometries. However, achieving fine control over interlayer orientation as well as ensuring high quality lattice and desired functionality at the wafer-scale can be challenging,^{1,11} especially in the context of automated platforms.

To overcome the bottlenecks associated with conventional top-down approaches, alterna-



tive synthesis routes often exploit thermodynamically and kinetically driven pathways^{4,12} of rotational self-organization^{13,14} in vdW epitaxial systems.^{4,15–17} It has been indeed demonstrated that bottom-up growth strategies can be leveraged to spontaneously evolve a range of rotationally aligned homo and heterostructures by tuning the driving forces e.g., substrate interactions,^{13,14} epitaxial strain,^{18,19} electron irradiation^{5,20} and/or thermal annealing^{21,22} etc. Harnessing such growth drivers either individually or combinatorially can therefore pave the way towards scalable syntheses recipes for vdW layered structures with targeted twists or moiré patterns.

However, efficacy of such bottom-up techniques to design heterostructures with desired functionality, often hinges upon the ability to navigate the complex and intertwined parameter space of synthesizing vdW layered structures.²³ In addition to nucleation and growth, the prominence of solid-state phase transitions²⁴ and out-of-equilibrium pathways^{25,26} involving the effects of defects, precursors and metastable phases¹⁴ only adds toward the complexity of the design space. To this end, the recent rise of high-throughput experimentation techniques^{27,28} integrated in closed-loop fashion with advanced artificial intelligence (AI)/machine learning (ML) platforms have shown to be promising in rapid and autonomous exploration² of vast synthesis-structure-property spaces.³

While active learning (AL) driven robotic discovery platforms^{3,29,30} are in general tailored to facilitate synthesis aided by autonomous characterizations of target materials through optimal design of experiments, several challenges³¹ still remain to achieve accelerated discovery of "new" functional materials. Many of these campaigns can benefit from deeper and accessible physical insights through the robust integration of theory, modeling and simulation capabilities^{32,33} in the autonomous experimentation (AE) loop. Moreover, it has been recently demonstrated how virtual environments³⁴ can serve as useful testbeds to rapid refinement and systematic benchmarking of algorithms and rigorous uncertainty quantification shielding the prohibitively expensive experimental data acquisition loops.

In the context of accelerating epitaxial growth of desired twisted vdW layered structures,



existing self-driving thin-film synthesis platforms³⁵ could naturally exploit the predictive insights from atomistic simulation methods. In past, purely simulation driven investigations^{36–39} have been able to propose novel strategies to design twisted metastable states in graphene and transition metal dichalcogenides (TMDCs) layers by means of lattice strain and thermally activated self-assembly pathways, ultimately predicting and motivating subsequent experimental efforts.^{12,22,40}

To bridge and guide autonomous nanoscale experimental (AE) platforms with insightful molecular dynamics simulations, Bayesian active learning workflows⁴¹ exploiting the asynchronous throughput accessible at the Exascale can act as catalyst towards further acceleration in autonomous materials discovery. Moreover, the recent advances⁴² in high performance computing (HPC) has pushed the boundaries of state-of-the art extreme scale computational frameworks⁴³ yielding unprecedented rates of floating point operations.

However, current ML and UQ-driven workflows for atomistics, especially dynamical trajectory datasets have not yet fully and efficiently leveraged the extreme scale computational resources, and are mostly limited in their current role to serve as initial motivation and/or offline knowledge-base to implicitly to kick start an AE platform. Few key challenges associated with these efforts include **(a)** orchestration and scheduling bottlenecks to handle large-volume of resource (node-hours), **(b)** navigating asynchronous on-line inference frameworks along with **(c)** on-the-fly data reduction—all three of which are paramount to towards achieving *theory-experiment integrated closed-loop hpc-enabled AI workflows* incorporating predictive materials theory and simulation tools in automated experiments to make them truly autonomous.

Inspired by recent experimental efforts to recrystallize MoS₂ crystals from amorphous precursors,⁴⁴ here we computationally sample optimal growth parameters to recrystallize (i.e. to evolve and reorganize long range order) a variety of target moiré heterostructures on-demand from a partially disordered amorphous system. We achieve this by demonstrating a resource-adaptive, on-the-fly decision making and parallelizable asynchronous Bayesian en-



semble learning framework for dynamical atomistic simulations, we computationally sample optimal growth parameters to recrystallize (i.e. to evolve and reorganize long range order) a variety of target moiré heterostructures *on-demand* from a partially disordered amorphous system. We apply active learning strategies using both conventional acquisition function (e.g. expected improvement) and direct posterior sampling (Thompson sampling) based Bayesian optimization (BO), and study the acceleration for asynchronous batch parallel acquisition in our on-the-fly ensemble molecular dynamics workflow environment. Excellent advantages over parallel random sampling as well as efficacy in identifying optimal growth space for evolving "on-demand" target structural phases are then shown.

Results

A real-time hpc-driven asynchronous and autonomous workflow for accelerated computational synthesis

Our hpc resource-driven workflow to enable accelerated computational autonomous synthesis is demonstrated in Figure 1. Choosing recrystallization of an amorphous Molybdenum disulfide (MoS_2) layer on top of a monolayer crystalline substrate (Figure 1a) as a case study (amorph+crystal), we aim to predict the parameter space (Figure 1c) which could transform into crystalline bi-layers with target interlayer twist angles (Figure 1e). We hypothesize, that for a given stoichiometry, the parameters space for the recrystallization process involves annealing temperature (T), in-plane lattice mismatch strain (ϵ_{lat} , assuming it to be tensile for the thin-film) and shear (ϵ_{sh}), and we could identify the appropriate growth regime leading to a specific target moiré structure (Figure 2). Choice of these parameters are motivated from the experimentally reported investigations of a variety of transformation and reconstruction pathways involving van der Waals layered epitaxial phases.^{12,24,26,45} While several other kinetic and thermodynamic driving forces e.g., chemical potential, electron beam effects, effect of edge crystallography and pressure etc. will play crucial roles in the emergence



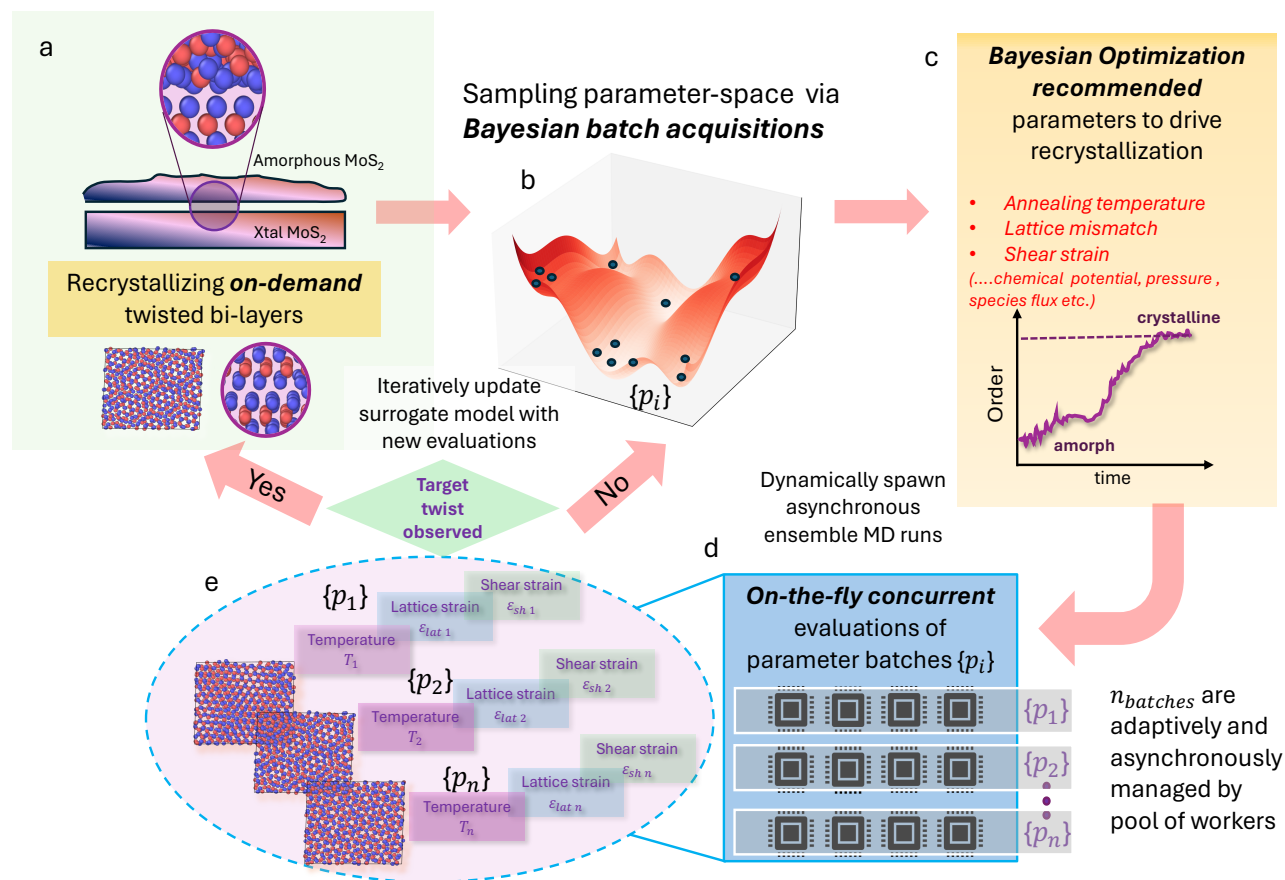



Figure 1: **Asynchronous active batch sampling strategy for optimal driving parameters for target twist angles.** Goal is to transform a given initial structure with amorphous MoS₂ layer on top of a crystalline layer (amorph+Xtal), into crystallized twisted bi-layer with a predefined target (a). Asynchronous ensemble optimization (b) strategies using Bayesian acquisition of candidate recrystallization parameters (c). *On the fly* concurrent evaluations of atomistic simulation via a resource-adaptive dynamic orchestration under a parameter space consisting of annealing temperature, lattice and shear strains. Additionally, this space can easily be expanded (as indicated in smaller fonts) to consider other synthesis relevant parameters like chemical potential, pressure, ion/species flux etc. can naturally be (d) to identify the driving parameters ($\{p_i\}$) to achieve recrystallization with the target twist angle (e).

of long-range order during recrystallization and as well dictating the final orientation of phases, for practicality, we restrict ourselves to implicitly incorporate such effects through annealing temperature, hydrostatic (e.g. uniform lattice mismatch strains) as well as symmetry distortions i.e. shear components in the mechanical strain tensor to demonstrate the potential of such computational investigations in the autonomous discovery campaigns. We note that by considering this parameter space, we only demonstrate a representative work-



flow where a diverse set of local environments resulting from temperature and strain effects are covered during the amorphous to crystalline phase transformation dynamics; further consideration of drivers like chemical potential and pressure could also be similarly used to explore the structural evolution pathways relevant for typical growth conditions within our current framework.

In principle, this involves exploration of materials dynamics over a nontrivial landscape sensitive to small perturbations in temperature and strain. Accelerating such search will require adaptive sampling techniques. We take a multi-level iterative acceleration approach by adopting i) Bayesian optimization (BO) to predict along with ii) an on-the-fly asynchronous high-throughput sampling orchestrated on a leadership class single batch computational allocation (c.f. methods) bypassing scheduling bottlenecks. Such workflows, hence, ensure an integrated gain in terms of optimal resource usage and computational throughput through uncertainty-aware dynamic task execution based sampling strategies.⁴⁶

Active learning of molecular dynamics (MD) trajectories performed at a candidate parameter combination set ($p_i \equiv \{T_i, \epsilon_{lat\ i}, \epsilon_{sh\ i}\}$, where $i \in \{1, 2, \dots, N_{batch}\}$) also requires automated parsing through the large volume simulation output as well as analytics related to the quantity of interest (q.o.i) which in this current case is the interlayer twist angle. Our solution to address this problem is to analyze simulation data at a regular time interval on-the-fly as soon as the data becomes available through an in-memory message passing and data-parallel processing approach (c.f. Methods). Hyperparameters of this adaptive workflow, e.g., batch size of candidate parameters (n_{batch}) to be evaluated with MD and the available allocation are correlated (for instance, with the allocation required MD evaluation) and are managed by adaptive scheduling so as to maintain a uniform number of tasks running during most part of the wallclock limit. The capability (Figure 1) to combine on-the-fly analytics of individual simulations with a concurrent asynchronous evaluation framework makes it possible to expand the global search horizon while exploiting uncertainty of the data through BO.



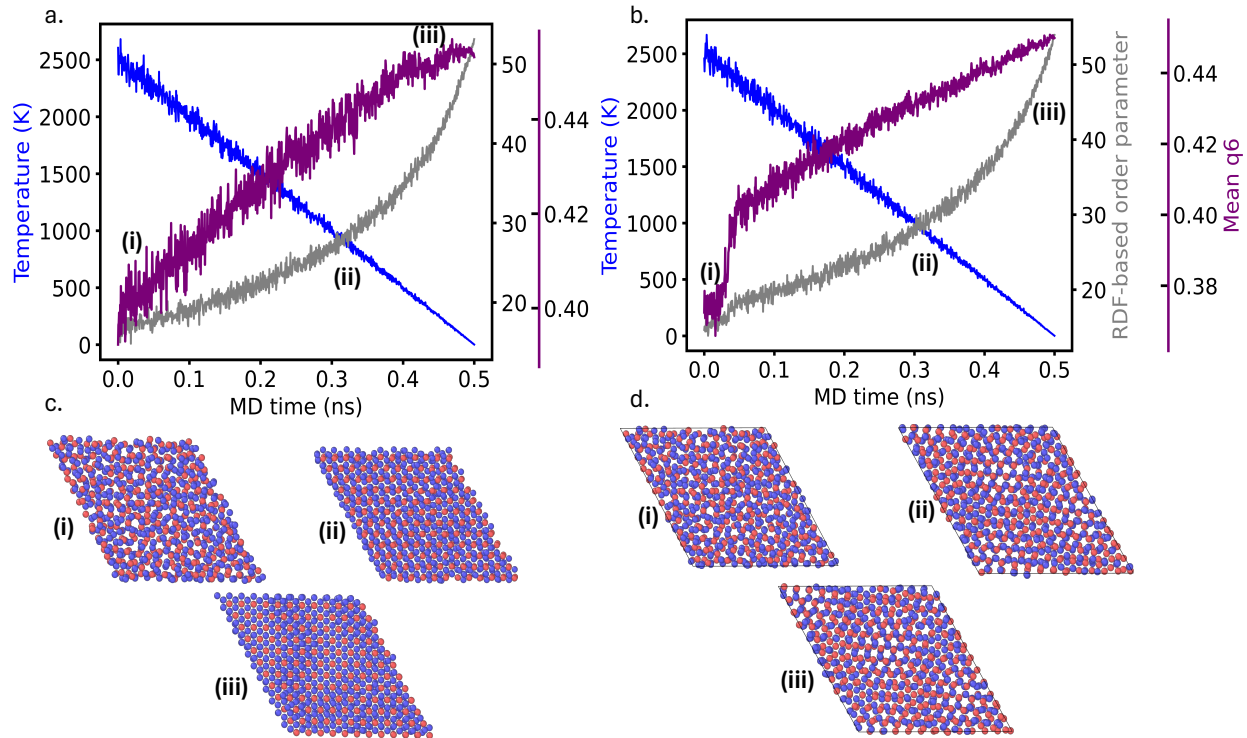


Figure 2: **Recrystallization under temperature and strain effects** As temperature is quenched near 0 K (**a, b**), overall order starts to emerge in the top layer as shown by both the mean-per atom q_6 (purple) and pair-count based order parameters (grey) in the bottom panel (**c,d**) for different snapshot in the MD trajectory. While only temperature driven (i.e., heating at 2500 K) trajectories lead to almost no significant interlayer twist—AB stacked 2H phase (**c-iii**), applying a lattice strain (ϵ_{lat}) with a supercell shear of (ϵ_{sh}) 1.5 % each, can induce a recrystallized bilayer with a twist angle (θ) of around 8.7° (**d-iii**).

Tunable twisted MoS₂ bilayers through recrystallization

We start from an initially disordered top layer (amorph-MoS₂) generated with melt-quench MD runs (c.f. Methods) placed on top of 2H crystalline layer (Figure 2c,d). Upon rapid heating (e.g., $T=2500$ K in Figure 2) followed by a slow quench until ~ 0 K, recrystallized MoS₂ top layer forms an AB stacked 2H bi-layer phase in Figure 2a with gradually emerging order (c.f. Methods details on order quantification using a pair-distance count based simplistic order parameter (gray lines in Figure 2) which compares well with a conventional hexatic symmetry identifier like Q_6 shown in purple in Figure 2). The final ordered structure is drastically altered as 1.5% strain in terms of both the lattice strain (ϵ_{lat}) and shear (ϵ_{sh}) is induced (Figure 2b). A moiré bilayer with a final interlayer twist angle of 8.7° with a supercell wave-



length of 20.8Å appears as a result of this change in the recrystallization conditions. This indicates that there exists a promising approach to bias recrystallization pathways towards tunable twisted bi-layers as final phases with the help of controlling temperature and strain induced parameters. Based on such observations as well as guided by past computational investigations of melt-quench crystallization analysis⁴⁷ of MoS₂ involving reactive force-fields (c.f. modeling details) we pick a range of 1500K-3000K for annealing temperature and lattice strains upto 5% for our autonomous search space. We note that using such approaches to control solid-solid phase transitions via self-organization of disordered 2D materials can emerge as cleaner and tunable alternates over traditional CVD driven growth of twisted-bilayer⁴⁸ as they involve complex interplay of nucleation and growth kinetics dominated by chemically sensitive driving factors.

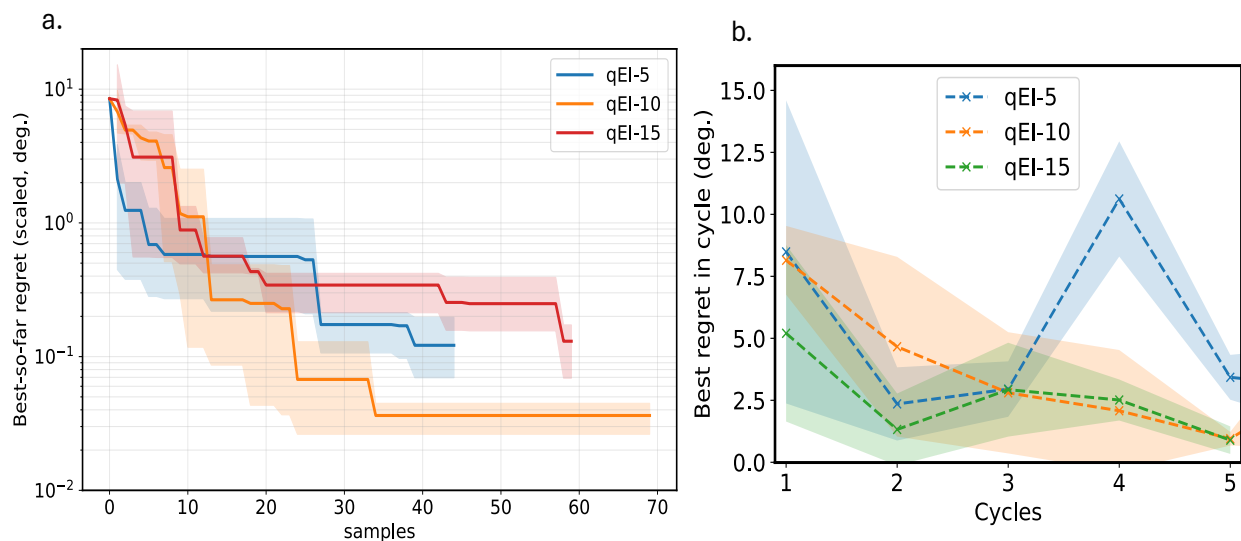


Figure 3: **Sampling a target twist with batch acquisition based on q-Expected Improvement** Regret (absolute error) between sampled twist angle and the predefined target for q-EI based acquisitions with different batch sizes (n_{batch}) in (a). Solid lines correspond to the cumulative minimum of regret over the iterations. Mean regret evolution in (b) for the best sample in each batch over BO cycle iterations for varying batch sizes. The confidence interval shows standard deviation over 3 repetitions



Bayesian batch acquisitions to sample a target twist

Given that we have a tunable epitaxial phenomenon to generate twisted homobilayer which can be captured through atomistic simulations, we now turn towards an autonomous approach to expedite the process of discovering optimal parameter combinations leading to a target twisted final state, given *on demand*. While sequential acquisition based Bayesian optimization loops have been commonly deployed in self-driving laboratories,^{3,30,41,49} in the context of building virtual environments, throughput in terms of both sampling and decision making can be increased through extreme scale computational resources. Furthermore, many of the state-of-the-art materials simulation capabilities are highly optimized into applications exploiting leadership class supercomputers.⁴²

Our approach therefore is to accelerate the discovery of recrystallization parameters through asynchronous iterative evaluations of a batch of candidates (p_i , $i = 1, 2 \dots n_{batch}$) recommended by BO. Several algorithms have explored the idea of parallel optimization of the acquisition function with various approximation^{50,51} to account for the uncertainties of the pending evaluations. First, we start with q-expected improvement (EI) (c.f. Methods) as our acquisition function to sample a target twist angle, $\theta_{target} = 8.5^\circ$. In Figure 3a, we show the variation in predictions with different batch sizes (n_{batch}). The regrets from the target for each of the acquisition events are shown in Figures 3a and 5a. These acquisitions are driven by Bayesian uncertainty quantification metrics as each of the MD trajectory is spawned with a set of annealing temperature, lattice and shear strains under recommendation of algorithms as defined in the methods section. To better follow the progress of this regret minimization process, we added cumulative minimum lines for each case (c.f. Figures 3a and 5a). From this analysis we can see that the cumulative minimum gradually drops discretely as sample trajectories are evaluated. For qEI driven acquisition, with a batch size=5, the best regret (0.1°) in Figure 3a leads a final average recrystallized moiré twist of 8.4° , i.e. Target (8.5°) – regret (0.1°) as achieved on the 38th recommendation (5-th cycle) at a temperature of 1386 K, 1.2% lattice strain and 5.53% shear (Supplementary Figure 5).



Whereas, starting with an identical set of initial observation data, with a larger batch size ($n_{batch}=15$), the workflow predicts a twisted phase closer to the target, $\theta = 8.46^\circ$ (Figure 3a and Supplementary Figure 5c,d) for temperature=2077 K, 2.3% lattice strain and 5.76% shear at 4-th cycle (58-th sample).

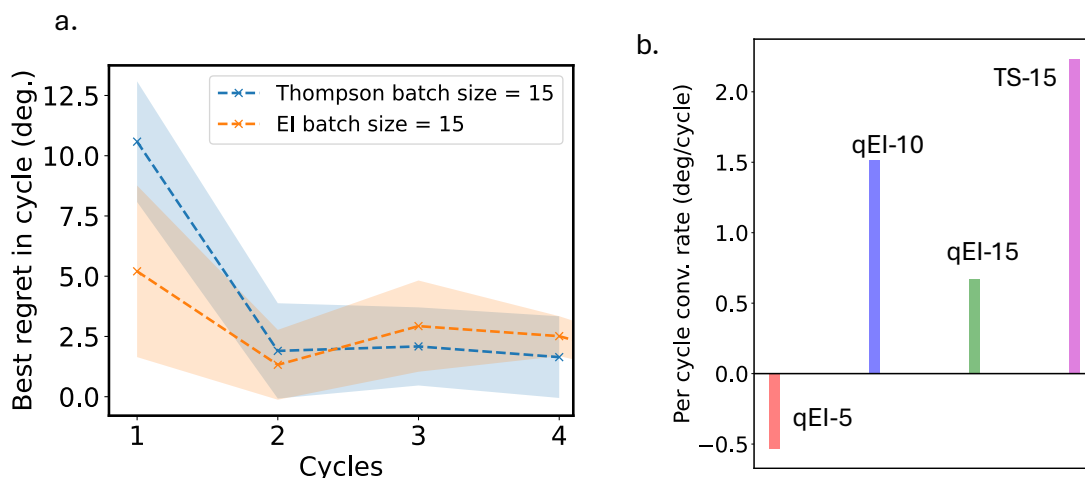


Figure 4: **Effect of batch size and sampling strategies** Initialized with identical set of observations, performance of two different acquisition strategies with batch evaluation—Thompson sampling and Expected Improvement in (a). Overall efficiency of various batch sampling strategies with respect to size and acquisitions in (b). Inherent stochastic nature of Thompson sampling outperforms EI based acquisition which can become explorative beyond an optimal batch size

While it is obvious that larger batches can be harnessed to cover a larger search space (Supplementary Figure 5b,d) in faster wallclock time with scalable resources, the indirect exploitation of the posterior, as it is the case for EI as an acquisition function (⁵²), might not always lead to efficient searches (Figure 4a) with a larger batch size. We could define a per-cycle sampling efficiency metric as an average rate (computed normalizing the change of mean best regret per cycle over 4 cycles) describing how fast BO iteration cycles generate an MD trajectory with a target twisted phase. As seen in Figure 4(b), this is driven by the fact that beyond a point larger batches tend to make the search more explorative rather than efficiently exploiting model uncertainties accessed through an implicit acquisition criteria like EI.⁵³



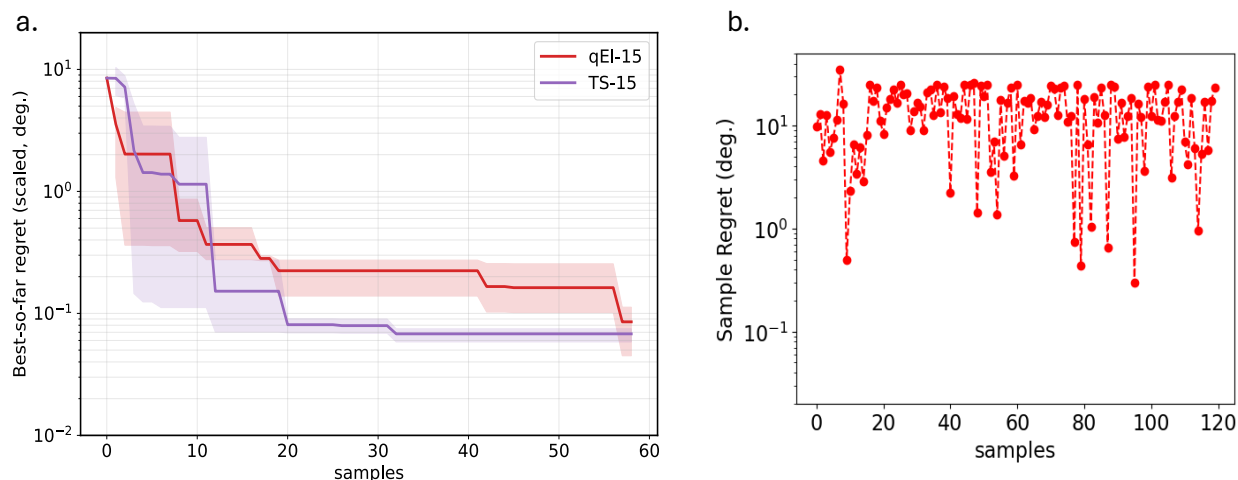


Figure 5: **Comparison of adaptive batch sampling via BO with random sampling** Performance of two different acquisition strategies with batch evaluation—Thompson sampling and Expected Improvement (a). Solid lines correspond to the cumulative minimum of regret over the iterations. Initialized with identical set of observations, random sampling (b) underperforms within a comparable sampling duration.

To probe the effect of acquisition strategies, we adopt batch Thompson sampling (c.f. methods), shown in Figure 5a, for the same target angle ($\theta = 8.5^\circ$) and identical prior observation set. It is interesting to note that Thompson sampling with a comparable batch size (i.e., $n_{batch} = 15$) leads to a more consistent sampling near the target yielding three twist angles (sample index) e.g., $8.39^\circ(18)$, $8.57^\circ(47)$ at different parameter combinations ($T_{18} = 1362\text{K}$, $\epsilon_{lat\ 18} = 0.4\%$, $\epsilon_{sh\ 18} = 5.69\%$), ($T_{47} = 3195\text{K}$, $\epsilon_{lat\ 47} = 0.2\%$, $\epsilon_{sh\ 47} = 5.5\%$) than EI. This ability to scan through a diverse range of recrystallization parameters could be due to the probabilistic nature of Thompson sampling from the posterior leads to a better exploration-exploitation trade-off in batch mode⁵² than primarily indirect exploitation (of posterior properties) driven acquisitions through EI. Furthermore, we assess the advantage of our active learning-driven sampling approach over vanilla random sampling (Figure 5b). Starting with the identical initial observation set like all the previous sampling demonstrations, random sampling with $n_{batch} = 15$ is unable to sample twist angles close to the target (8.5°) even upto 120 samples, with the best value being 8.2° .



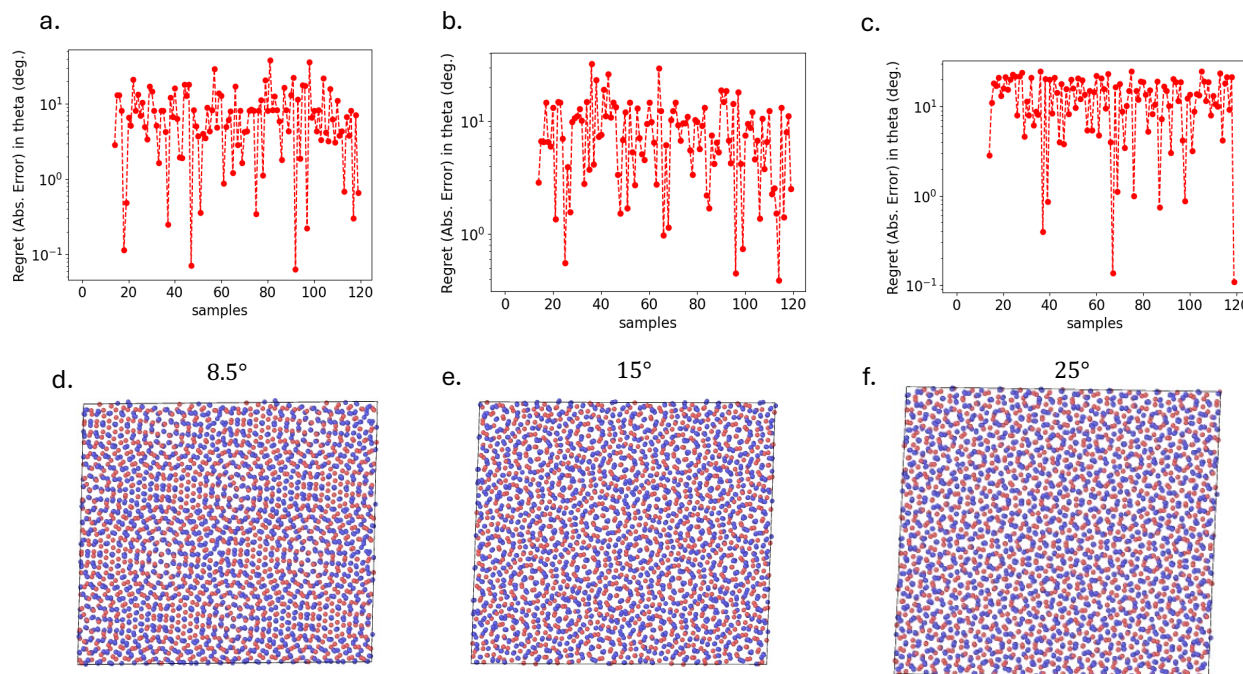


Figure 6: **Achieving recrystallization with *on demand* final twist angles** General flexibility of the workflow is shown using batch Thompson sampling. Starting from the same set of initial samples, the asynchronous Thompson sampling could search three diverse set of moiré sublattices with target interlayer twists— 8.5° (a,d), 15° (b,e) and 25° (c,f).

Discussions

Designing twisted bilayer interfaces *on-demand*

We thus far have demonstrated a capability to identify target recrystallization conditions through our Bayesian optimization workflow coupled with asynchronous ensemble evaluation strategies driven by scalable computational resources. While constrained in the scope of parameter space (e.g., more realistic conditions might involve more drivers like pressure, chemical potential, rate of heating etc.), these results show the potential of accelerating targeted solid state phase transitions in vdW 2D materials interfaces. To further demonstrate the efficacy of our workflow we provide quantitative evidence towards enabling *on-demand* design of twisted moiré interfaces of MoS₂ bilayers. Using batch Thompson sampling as a mode to evaluate the parameter space comprising temperature, lattice strain and shear, we focus our workflow towards sampling twist angles 15° and 25° in addition to 8.5° —thereby



covering a wide range of moiré wavelengths and superlattices. As shown in Figure 6, the current approach is able to achieve recrystallization into moiré lattice with predefined target twist angles, especially for structures with nontrivial target twist angles (e.g. 25°) (c.f. Supplementary Figure 6 for more details).

It will be further intriguing to probe the effects of edge reactions, stacking orders, perturbed local environments and defect-pinning/doping that results from atomic rearrangement during the melt-quench recrystallization as well as precursor-driven growth,⁵⁴ on their electronic structure via first principles. The advantage of proposed online workflow orchestration strategy is that such effects can now be more readily investigated leveraging HPC platforms.

Moreover, we successfully identify multiple combinations of parameter corresponding to for each of these diversely specified target twists (Figure 6). For instance, interlayer twists close to 15° can be recrystallized at i) 2725 K, 2.7% lattice strain and 5.5% shear as well as at ii) 2107 K, 1.3% lattice strain with 5.69% shear. Likewise, we identify more than one set of optimal parameters for $\theta \sim 25^\circ$ —i) {2797 K, 0.9% lattice strain and 5.57% shear} and ii) {1379 K, 3.2% lattice strain and 5.51% shear}. We note that very small changes to the recrystallization parameters, especially shearing the supercell leads to drastically different final moiré superlattices. For this reason, we keep the search window for ϵ_{sh} to vary from 5.5% to 5.75% strain levels which corresponds to absolute distances of 1\AA to 1.5\AA tilts of in-plane supercell dimensions. Furthermore, the wide variations (Supplementary Figure 4) in optimal temperature and lattice strain conditions associated with a specific final twisted bi-layer phase, leads us to speculate that the recrystallization landscape could be rough having multiple locally varying optimal cases. Traditional grid search or sequential active learning protocols hence might not be adequate to exhaustively identify favorable recrystallization conditions for a given target twist (Supplementary Figure 1). Furthermore, current approach of high-throughput classical force-field driven simulations being linear scaling $O(N)$ in comparison to cubic scaling ab initio/DFT simulations $O(N^3)$, making them much more competitive towards accessing sufficient timespan for the *in-silico* recrystallization simula-



tion recipe describe here, within a matter of few hours of wallclock time with hpc-scalable batch parallel BO as long as enough compute is available. However, we note that bridging the timescales of state of the art atomistic simulations with what is relevant for experiments is a long standing multiscale challenge, for which we remark that our end-to-end developments of asynchronous on the fly active learning atomistics are ideally suited to upscale the dynamics with coarse grained models e.g., kinetic monte carlo, phase field simulations etc. The presented online ensemble sampling workflow for identifying recrystallization regimes of MoS2 bi-layers could be seamlessly extended towards various other TMDC systems (.e.g, W/Mo-S/Se/Te etc.), multilayers or even heterostructures are of increasing interest due to their highly tunable quantum properties. In practice, such adoptions are not limited by any of the workflows features which are rather general and mostly take atomistic systems and force-field as inputs. If appropriate force-field descriptions are available for a TMDC system of choice, the current framework should be fully transferable. With the rise of highly accurate foundation MLIPs, the bottlenecks of force descriptions are rapidly disappearing. Adapting these methods to more stochastic or continuum simulation approaches is also quite straightforward.

Conclusion

We have demonstrated that combining extreme-scale computations, Bayesian optimisation and asynchronous automated computational workflows, a powerful building block to enable *on-demand* synthesis can be achieved, and that it allows time and resource efficient discovery of synthesis pathways for complex materials. The targeted problem of achieving ‘on-demand’ van der Waals epitaxy is a major technological challenge, specifically when controlling both the phase and the interlayer stacking orientations. We demonstrate that within 20 iterations the autonomous computational synthesis platform achieves the targeted moiré quantum heterostructure. The global sampling opportunities boosted by extreme-scale



resource-driven batch acquisition workflows and on-the-fly analysis and UQ capabilities are ideally suited to pin down more sophisticated interplay of thermodynamic and kinetic factor driving self-organized and strain-engineered interfacial evolution during vdW epitaxy. Furthermore, we note that integrating the autonomous computational synthesis workflows with real experimental synthesis efforts can quickly narrow down the number of experimental attempts to grow a targeted quantum heterostructure. The current ensemble simulation framework can also be used to generate digital-twins on experimental measurements, such as X-ray diffraction (XRD) and reflective high-energy electron diffraction (RHEED). As such, integration of our framework with experimental molecular beam-epitaxy (MBE) or pulsed-laser-deposition (PLD) systems will allow us to perform on-the-fly comparison with experimental XRD/RHEED measurements, so as to guide experiments in real-time, allowing a real autonomous thin-film synthesis platform with theory-in-the-loop.

Methods

MatEnsemble: Adaptive real-time ensemble task management environment for extreme scale computing

Within a single HPC batch allocation, ensemble evaluations/jobs which require MPI parallel resources at the individual level, often suffer from scheduling challenges with typical SLURM-based workload managers.⁵⁵ This is primarily due to the fact that a hard ceiling (e.g., 100 for Frontier, Perlmutter/NERSC) are often imposed toward chaining together multiple job submissions in batch mode (e.f. `srun` for SLURM). Such setups are highly unsuitable when **(i)** size of the ensemble exceeds the ceiling of a scheduler and **(ii)** dynamic heterogeneous tasks are at the core of the workflow e.g., based on outcomes of the ensemble of MD simulations active learning or UQ algorithms (e.g., BO) has to operate and further spawn next generation of ensembles.

To mitigate the above challenges, we implement all our active learning MD simulations



within a recently developed adaptive and dynamic orchestrator framework `matEnsemble`⁴⁶ which has an executor back-end of `Flux`, an Exascale friendly HPC resource manager with hierarchical graph-based scheduler allowing for a generalized flexible *custom* operations within a single batch allocation. `MatEnsemble`⁴⁶ benefits from the native python executor-interface of `Flux`,⁵⁶ and the concurrent asynchronous programming model of core python through `Future` objects.⁵⁷ A continuous throughput is maintained via dynamically spawning and monitoring task. Furthermore, to enable real-time streaming of post-processed data (e.g., interlayer twist angle evolution) from large-scale atomistic trajectories an *in-memory* data analysis protocol is used by exploiting the heterogeneous (GPU+CPU) architecture of Exascale systems (e.g., Frontier) via a round-robin MPI-communicator splitting approach (c.f.⁴⁶). As explained in the following sections, such an online adaptive framework enables efficiently coupling between available computing resource chunks for ensemble evaluations guided by adaptive sampling methods.

Parallel batch sampling with Bayesian optimization

Traditionally optimal design of experiments via active learning based on Bayesian optimization⁵⁸ proceeds through sequential acquisitions. With the adaptive asynchronous capability to produce scalable ensemble evaluations in our workflow infrastructure, batch mode acquisitions algorithms are natural choices to effectively utilize available hpc resources and accelerate the search towards target twist angles in the spirit of efficient global optimization⁵⁹ with the following acquisition functions.

q-Expected Improvement: With a goal to accelerate the solution of the minimization problem for an expensive black-box function f , we have used parallel/batch expected im-



provement^{60–63} in this study. It is computed by Monte Carlo sampling and is given by:

$$qEI_{y^*}(\mathbf{X}) = \mathbb{E}_{f(\mathbf{x})} \left[\max_{j=1, \dots, q} \{ [f(\mathbf{x}_j) - y^*]_+ \} \right] \quad (1)$$

$$\approx \sum_{i=1}^N \max_{j=1, \dots, q} \{ [\xi^i(\mathbf{x}_j) - y^*]_+ \} \quad (2)$$

where $\xi^i(\mathbf{x}) \sim f(\mathbf{x})$ are samples drawn randomly from the joint posterior distribution of the model evaluated at batch points \mathbf{x} , before optimizing the expectation which dictates the acquisition strategy rather than inferences which are directly based on the posterior, and top performing n_{batch} number of candidates from the optimization are finally obtained. Relevant to the last point, we also explored more direct posterior sampling strategies e.g. Thompson sampling as described in the following.

Batch Thompson sampling: Unlike several acquisition functions which are mostly driven by exploitative strategies, Thompson sampling, due to its stochastic nature, is able to handle the exploration-exploitation trade-off in a better fashion. In our simulation we adopt the batch parallel version of Thompson sampling first proposed by Kandasamy et al.⁵¹ We choose a quasi-Monte Carlo sampler which uses Sobol sequences⁶⁴ to sample from the joint-posterior of over q (n_{batch}) batch and optimizes over simple regret,⁵²

$$\text{Simple Regret}(\mathbf{X}) = f^* - \min_{i=1, \dots, n_{\text{batch}}} f(\mathbf{x}_i) \quad (3)$$

The optimization leads to n_{batch} number of top candidates to be evaluated in the following cycle. Upon finishing the evaluation of all candidates in the batch the next cycle of the BO is resumed.

We implement all our BO algorithms using the BOTorch library⁶⁵ for single task Gaussian process approximated with Matérn kernel.⁶⁶



Molecular dynamics simulations of recrystallization in MoS₂ bilayers

We perform all our MD simulations by wrapping the shared library utilities of the open source package LAMMPS⁶⁷ through `matEnsemble` drivers. Starting with a 2H bilayer MoS₂ phase (rhombohedral supercell with 864 atoms), we first induce disorder in the top crystalline layer by selectively heating up the layer at a high temperature of 5000K for 50ps with a timestep of 0.5fs. We use a classical reactive force field ReaxFF, specifically re-parametrized to accurately capture order-disorder transitions in MoS₂ during multiple melt-quench cycles.⁴⁷ After the top layer undergoes melting, the whole bilayer system is equilibrated at 300K in atmospheric pressure using a Berendsen NPT thermostat for 50ps to achieve the disordered top layer placed on top of a crystalline bottom layer—this served as the initial sample (amorph+Xtal) for the recrystallization simulations.

Finally to simulate the recrystallization process under various candidate input parameter combination of $p_i \equiv \{T_i, \epsilon_{lat}, \epsilon_{sh}\}$, the system is rapidly heated for 50ps at a temperature T_i along with in-plane lattice rescaling (x, y) of $(1 + \epsilon_{lat})$ (to model epitaxial strain) and an in-plane shear strain (xy) of ϵ_{sh} measured on the substrate crystal layer (c.f. Figure 1a). This is then followed by a slow quench of the bilayer system to reduce the temperature level from T_i to 0 K over 0.5 ns i.e. 1×10^6 MD steps. The final step of rapid melt-slow quench simulations are performed using Langevin thermostat (NVT). Periodic boundaries are maintained along in-plane (i.e. x and y) directions. Lennard-Jones reflective walls are implemented near the bottom (0Å) and top (12Å) edges. We use Ovito for all the visualizations and analysis of MD trajectories.

Important to note that the annealing approach does not guarantee lowest energy configurations which is intentional for the present study to let the workflow explore high-energy metastable basins and discover a diverse range of on-demand Moire motifs which might not necessarily be equilibrium states. In the context of synthesis and materials discovery, where tracking down non-equilibrium phases are of particular interest, such melt-quench dynamics approach could specifically be useful. While it is possible for atoms to undergo rare diffusion



event which are path dependent mostly in the long-time slow quench limit, our 0.5ns quench protocol could still be relatively fast to favor such path dependent quenching behavior in the first order. This could be explored in more detail by studying long-time dynamics limits which is beyond the current scope.

Constructing an order parameter—To quantify the emerging order in the MoS₂ bilayers (Figure 2), we define a simple pair correlation based custom order parameter which estimates sharpness of the first nearest neighbor peak distribution. First, the pair correlation function is computed with cutoff of 3Å based on first nearest neighbors. Then we count the number of pairs within a pair distance window ($[2.4\text{Å}, 2.6\text{Å}]$), where first nearest neighbor peak/mode of the distribution for pristine MoS₂ is around 2.47Å.

Computing interlayer twist-angles on-the-fly

Adopting a streaming approach to analyze MD trajectories as they become available through a distributed in-memory message passing framework, we accelerate and perform on-the-fly data reduction. As described in more detail in,⁴⁶ apart from being scalable and ideally suited for integrated active learning workflows, such frameworks serve as general testbeds to a variety of atomistic analysis and visualization algorithms. For this study, we wrap through a simple and computationally light-weight iterative algorithm to compute the interlayer twists for atomistic snapshots along the recrystallization trajectory.



Algorithm 1 Iterative estimation of interlayer twist angles

Input: atom positions $x_i^j, i \in \{1, 2, 3\}$ and $j \in \{1, 2, \dots, N_{\text{atoms}}\}$ and atom types $t_j =$ either "Mo" or "Se"

Output: Twist (θ_{12}) between upper and lower layers

- 1: For each layer l randomly select an atom (j) of a particular type (say, "Mo")
 - 2: Find its second nearest neighbors $\{x^k\}$ (e.g., for crystalline layers $size(\{x^k\}) \sim 6$ and $type(\{x^k\}) \equiv$ "Mo")
 - 3: Sort $\{x^k\}$ based on their in-plane polar coordinates $\{r, \theta\}^k$ with central atom j position (x_i^j) as the origin.
 - 4: $ic \leftarrow 0, \theta_{ic} \leftarrow \min\{\theta\}^k, atom_{ic} \leftarrow atom(\text{argmin}\{\theta\}^k)$
 - 5: **while** $ic \leq$ max iterations **do**
 - 6: shift central atom at $atom_{ic}$ repeat steps 2 to 3
 - 7: **if** the neighbor selection fails due to a vacancy or other variations **then**
 - 8: Add a ghost atom at the guess position by replicating and shifting the central atom along r_{ic} in θ_{ic}
 - 9: **end if**
 - 10: $\theta_{ic} \leftarrow \min\{\theta\}^{ic}, atom_{ic} \leftarrow atom(\text{argmin}\{\theta\}^{ic})$
 - 11: $ic \leftarrow (ic + 1)$
 - 12: **end while**
 - 13: Fit a straight line ($y = m^l x + c^l$) over the atom positions stored in $\{atoms\}_{ic}, ic \in \{1, 2, \dots, \text{max iterations}\}$, where l is layer number $\in \{1, 2\}$
 - 14: Finally calculate the twist $\theta_{12} \leftarrow \arctan(\frac{|m^1 - m^2|}{|1 + m^1 m^2|})$
-



Analysis of polycrystallinity and validity of the twist computation algorithm

To test the validity of our approach, we performed a more rigorous per-atom bond-orientation based interlayer twist computation approach. As can be seen (Supplementary Figure 7) from the randomly considered final quenched phases for a series of recrystallization simulations under various temperature, lattice and shear strain conditions, the twist distributions mostly tend to be **unimodal** leading to no speculation of multiple grains/polycrystalline nature. Barely some small clusters of other angles are observed which are likely due to disorder and vacancy debris after the quench. Furthermore, the **mode** of per-atom twist-histograms are with excellent match with our much simpler single twist results for most of the cases. Note that due to hexagonal symmetry, angles that are numerically close to 0 deg. and 60 deg. are physically identical—hence the bimodal distributions are still single crystal phases.

Acknowledgement

This research is sponsored by the INTERSECT Initiative as part of the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy under contract DE-AC05-00OR22725 via the QCAD project. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, via the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. Resources of the Compute and Data Environment for Science (CADES) was used via a user project at the Center for Nanophase Materials Sciences (CNMS), which is a US Department of Energy, Office of Science User Facility at Oak Ridge National Laboratory



Author Contributions

SB designed the core real-time adaptive learning workflow and on-the-fly MD simulations of recrystallization, implemented batch TS sampling and wrote the manuscript draft. AG co-wrote and designed the BO approach. AB co-wrote, implemented BO and analyzed results. PVB discussed overall results and co-wrote. PG wrote, supervised and conceived the project.

References

- (1) Li, J.; Yang, X.; Zhang, Z.; Yang, W.; Duan, X.; Duan, X. Towards the scalable synthesis of two-dimensional heterostructures and superlattices beyond exfoliation and restacking. *Nature Materials* **2024**, *23*, 1326–1338.
- (2) Szymanski, N. J.; Rendy, B.; Fei, Y.; Kumar, R. E.; He, T.; Milsted, D.; McDermott, M. J.; Gallant, M.; Cubuk, E. D.; Merchant, A.; others An autonomous laboratory for the accelerated synthesis of novel materials. *Nature* **2023**, *624*, 86–91.
- (3) Ament, S.; Amsler, M.; Sutherland, D. R.; Chang, M.-C.; Guevarra, D.; Connolly, A. B.; Gregoire, J. M.; Thompson, M. O.; Gomes, C. P.; Van Dover, R. B. Autonomous materials synthesis via hierarchical active learning of nonequilibrium phase diagrams. *Science Advances* **2021**, *7*, eabg4930.
- (4) Liu, C.; Lin, Y.-C.; Yoon, M.; Yu, Y.; Puretzky, A. A.; Rouleau, C. M.; Chisholm, M. F.; Xiao, K.; Eres, G.; Duscher, G.; others Understanding substrate-guided assembly in van der Waals epitaxy by in situ laser crystallization within a transmission electron microscope. *ACS nano* **2021**, *15*, 8638–8652.
- (5) Roccapiore, K. M.; Boebinger, M. G.; Dyck, O.; Ghosh, A.; Unocic, R. R.; Kalinin, S. V.; Ziatdinov, M. Probing electron beam induced transformations on a




- single-defect level via automated scanning transmission electron microscopy. *ACS nano* **2022**, *16*, 17116–17127.
- (6) Geim, A. K.; Grigorieva, I. V. Van der Waals heterostructures. *Nature* **2013**, *499*, 419–425.
- (7) Cao, Y.; Fatemi, V.; Demir, A.; Fang, S.; Tomarken, S. L.; Luo, J. Y.; Sanchez-Yamagishi, J. D.; Watanabe, K.; Taniguchi, T.; Kaxiras, E.; others Correlated insulator behaviour at half-filling in magic-angle graphene superlattices. *Nature* **2018**, *556*, 80–84.
- (8) Cao, Y.; Fatemi, V.; Fang, S.; Watanabe, K.; Taniguchi, T.; Kaxiras, E.; Jarillo-Herrero, P. Unconventional superconductivity in magic-angle graphene superlattices. *Nature* **2018**, *556*, 43–50.
- (9) Rivera, P.; Yu, H.; Seyler, K. L.; Wilson, N. P.; Yao, W.; Xu, X. Interlayer valley excitons in heterobilayers of transition metal dichalcogenides. *Nature nanotechnology* **2018**, *13*, 1004–1015.
- (10) Kim, K.; Yankowitz, M.; Fallahazad, B.; Kang, S.; Movva, H. C.; Huang, S.; Larentis, S.; Corbet, C. M.; Taniguchi, T.; Watanabe, K.; others van der Waals heterostructures with high accuracy rotational alignment. *Nano letters* **2016**, *16*, 1989–1995.
- (11) Xu, X.; Pan, Y.; Liu, S.; Han, B.; Gu, P.; Li, S.; Xu, W.; Peng, Y.; Han, Z.; Chen, J.; others Seeded 2D epitaxy of large-area single-crystal films of the van der Waals semiconductor 2H MoTe₂. *Science* **2021**, *372*, 195–200.
- (12) Baek, J.-H.; Kim, H. G.; Lim, S. Y.; Hong, S. C.; Chang, Y.; Ryu, H.; Jung, Y.; Jang, H.; Kim, J.; Zhang, Y.; others Thermally induced atomic reconstruction into fully commensurate structures of transition metal dichalcogenide layers. *Nature materials* **2023**, *22*, 1463–1469.



- (13) Sutter, P.; Wimer, S.; Sutter, E. Chiral twisted van der Waals nanowires. *Nature* **2019**, *570*, 354–357.
- (14) Sutter, P.; Ibragimova, R.; Komsa, H.-P.; Parkinson, B. A.; Sutter, E. Self-organized twist-heterostructures via aligned van der Waals epitaxy and solid-state transformations. *Nature Communications* **2019**, *10*, 5528.
- (15) Gehring, P.; Gao, B. F.; Burghard, M.; Kern, K. Growth of high-mobility Bi₂Te₂Se nanoplatelets on hBN sheets by van der Waals epitaxy. *Nano letters* **2012**, *12*, 5137–5142.
- (16) Kim, J.; Bayram, C.; Park, H.; Cheng, C.-W.; Dimitrakopoulos, C.; Ott, J. A.; Reuter, K. B.; Bedell, S. W.; Sadana, D. K. Principle of direct van der Waals epitaxy of single-crystalline films on epitaxial graphene. *Nature communications* **2014**, *5*, 4836.
- (17) Lin, Y.-C.; Lu, N.; Perea-Lopez, N.; Li, J.; Lin, Z.; Peng, X.; Lee, C. H.; Sun, C.; Calderin, L.; Browning, P. N.; others Direct synthesis of van der Waals solids. *ACS Nano* **2014**, *8*, 3715–3723.
- (18) Wan, Y.; Huang, J.-K.; Chuu, C.-P.; Hsu, W.-T.; Lee, C.-J.; Aljarb, A.; Huang, C.-W.; Chiu, M.-H.; Tang, H.-L.; Lin, C.; others Strain-directed layer-by-layer epitaxy toward van der Waals homo-and heterostructures. *ACS Materials Letters* **2021**, *3*, 442–453.
- (19) Lu, Y.; Chen, J.; Coupin, M. J.; Sinha, S.; Warner, J. H. Lattice-Mismatch-Driven Small-Angle Moiré Twists in Epitaxially Grown 2D Vertical Layered Heterostructures. *Advanced Materials* **2022**, *34*, 2205403.
- (20) Maksov, A.; Dyck, O.; Wang, K.; Xiao, K.; Geohegan, D. B.; Sumpter, B. G.; Vasudevan, R. K.; Jesse, S.; Kalinin, S. V.; Ziatdinov, M. Deep learning analysis of defect and phase evolution during electron beam-induced transformations in WS₂. *npj Computational Materials* **2019**, *5*, 12.



- (21) Omambac, K. M.; Hattab, H.; Brand, C.; Jnawali, G.; N'Diaye, A. T.; Coraux, J.; van Gastel, R.; Poelsema, B.; Michely, T.; Meyer zu Heringdorf, F.-J.; others Temperature-controlled rotational epitaxy of graphene. *Nano letters* **2019**, *19*, 4594–4600.
- (22) de Jong, T. A.; Benschop, T.; Chen, X.; Krasovskii, E. E.; de Dood, M. J.; Tromp, R. M.; Allan, M. P.; Van der Molen, S. J. Imaging moiré deformation and dynamics in twisted bilayer graphene. *Nature Communications* **2022**, *13*, 70.
- (23) Jenks, C.; Lee, H. N.; Lewis, J.; Kagan, C. R.; Nealey, P.; Braun, P.; Holladay, J. E.; Gao, Y.; Sholl, D. S.; Helms, B. A.; others *Basic Research Needs for Transformative Manufacturing*; 2020.
- (24) Li, W.; Qian, X.; Li, J. Phase transitions in 2D materials. *Nature Reviews Materials* **2021**, *6*, 829–846.
- (25) Woods, J. M.; Jung, Y.; Xie, Y.; Liu, W.; Liu, Y.; Wang, H.; Cha, J. J. One-step synthesis of MoS₂/WS₂ layered heterostructures and catalytic activity of defective transition metal dichalcogenide films. *ACS nano* **2016**, *10*, 2004–2009.
- (26) Bianchini, M.; Wang, J.; Clément, R. J.; Ouyang, B.; Xiao, P.; Kitchaev, D.; Shi, T.; Zhang, Y.; Wang, Y.; Kim, H.; others The interplay between thermodynamics and kinetics in the solid-state synthesis of layered oxides. *Nature materials* **2020**, *19*, 1088–1095.
- (27) Kusne, A. G.; Gao, T.; Mehta, A.; Ke, L.; Nguyen, M. C.; Ho, K.-M.; Antropov, V.; Wang, C.-Z.; Kramer, M. J.; Long, C.; others On-the-fly machine-learning for high-throughput experiments: search for rare-earth-free permanent magnets. *Scientific reports* **2014**, *4*, 6367.
- (28) Hattrick-Simpers, J. R.; Gregoire, J. M.; Kusne, A. G. Perspective: composition–structure–property mapping in high-throughput experiments: turning data into knowledge. *APL Materials* **2016**, *4*.



- (29) Chen, J.; Cross, S. R.; Miara, L. J.; Cho, J.-J.; Wang, Y.; Sun, W. Navigating phase diagram complexity to guide robotic inorganic materials synthesis. *Nature Synthesis* **2024**, 1–9.
- (30) Sheng, H.; Sun, J.; Rodríguez, O.; Hoar, B. B.; Zhang, W.; Xiang, D.; Tang, T.; Hazra, A.; Min, D. S.; Doyle, A. G.; others Autonomous closed-loop mechanistic investigation of molecular electrochemistry via automation. *Nature Communications* **2024**, *15*, 2781.
- (31) Leeman, J.; Liu, Y.; Stiles, J.; Lee, S. B.; Bhatt, P.; Schoop, L. M.; Palgrave, R. G. Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis. *PRX Energy* **2024**, *3*, 011002.
- (32) Sumpter, B. G.; Vasudevan, R. K.; Potok, T.; Kalinin, S. V. A bridge for accelerating materials by design. *NPJ Computational Materials* **2015**, *1*, 1–11.
- (33) Ghosh, A.; Ziatdinov, M.; Dyck, O.; Sumpter, B. G.; Kalinin, S. V. Bridging microscopy with molecular dynamics and quantum simulations: An atomAI based pipeline. *npj Computational Materials* **2022**, *8*, 74.
- (34) Kanarik, K. J.; Osowiecki, W. T.; Lu, Y.; Talukder, D.; Roschewsky, N.; Park, S. N.; Kamon, M.; Fried, D. M.; Gottscho, R. A. Human–machine collaboration for improving semiconductor process development. *Nature* **2023**, *616*, 707–711.
- (35) MacLeod, B. P.; Parlane, F. G.; Morrissey, T. D.; Häse, F.; Roch, L. M.; Dettelbach, K. E.; Moreira, R.; Yunker, L. P.; Rooney, M. B.; Deeth, J. R.; others Self-driving laboratory for accelerated discovery of thin-film materials. *Science Advances* **2020**, *6*, eaaz8867.
- (36) Zhu, S.; Pochet, P.; Johnson, H. T. Controlling rotation of two-dimensional material flakes. *ACS nano* **2019**, *13*, 6925–6931.



- (37) Bagchi, S.; Johnson, H.; Chew, H. B. Rotational stability of twisted bilayer graphene. *Physical Review B* **2020**, *101*, 054109.
- (38) Dai, S.; Xiang, Y.; Srolovitz, D. J. Twisted bilayer graphene: Moiré with a twist. *Nano letters* **2016**, *16*, 5923–5927.
- (39) Bagchi, S.; Johnson, H. T.; Chew, H. B. Strain-controlled dynamic rotation of twisted 2D atomic layers for tunable nanomechanical systems. *ACS Applied Nano Materials* **2020**, *3*, 10878–10884.
- (40) Kazmierczak, N. P.; Van Winkle, M.; Ophus, C.; Bustillo, K. C.; Carr, S.; Brown, H. G.; Ciston, J.; Taniguchi, T.; Watanabe, K.; Bediako, D. K. Strain fields in twisted bilayer graphene. *Nature materials* **2021**, *20*, 956–963.
- (41) Kusne, A. G.; Yu, H.; Wu, C.; Zhang, H.; Hattrick-Simpers, J.; DeCost, B.; Sarker, S.; Oses, C.; Toher, C.; Curtarolo, S.; others On-the-fly closed-loop materials discovery via Bayesian active learning. *Nature communications* **2020**, *11*, 5966.
- (42) Alexander, F.; Almgren, A.; Bell, J.; Bhattacharjee, A.; Chen, J.; Colella, P.; Daniel, D.; DeSlippe, J.; Diachin, L.; Draeger, E.; others Exascale applications: skin in the game. *Philosophical Transactions of the Royal Society A* **2020**, *378*, 20190056.
- (43) Mniszewski, S. M.; Belak, J.; Fattebert, J.-L.; Negre, C. F.; Slattery, S. R.; Adedoyin, A. A.; Bird, R. F.; Chang, C.; Chen, G.; Ethier, S.; others Enabling particle applications for exascale computing platforms. *The International Journal of High Performance Computing Applications* **2021**, *35*, 572–597.
- (44) Wang, R.; Koch, N.; Martin, J.; Sadofev, S. Recrystallization of MBE-Grown MoS₂ Monolayers Induced by Annealing in a Chemical Vapor Deposition Furnace. *physica status solidi (RRL)–Rapid Research Letters* **2023**, *17*, 2200476.



- (45) Weston, A.; Zou, Y.; Enaldiev, V.; Summerfield, A.; Clark, N.; Zólyomi, V.; Graham, A.; Yelgel, C.; Magorrian, S.; Zhou, M.; others Atomic reconstruction in twisted bilayers of transition metal dichalcogenides. *Nature nanotechnology* **2020**, *15*, 592–597.
- (46) Bagchi, S.; others MatEnsemble: An exascale compatible asynchronous integrated on-the-fly active learning framework for materials ensemble simulations. <https://github.com/Q-CAD/MatEnsemble>, 2025.
- (47) Ponomarev, I.; Polcar, T.; Nicolini, P. New reactive force field for simulations of MoS₂ crystallization. *The Journal of Physical Chemistry C* **2022**, *126*, 9475–9481.
- (48) Xu, M.; Ji, H.; Zheng, L.; Li, W.; Wang, J.; Wang, H.; Luo, L.; Lu, Q.; Gan, X.; Liu, Z.; others Reconfiguring nucleation for CVD growth of twisted bilayer MoS₂ with a wide range of twist angles. *Nature Communications* **2024**, *15*, 562.
- (49) Shields, B. J.; Stevens, J.; Li, J.; Parasram, M.; Damani, F.; Alvarado, J. I. M.; Janey, J. M.; Adams, R. P.; Doyle, A. G. Bayesian reaction optimization as a tool for chemical synthesis. *Nature* **2021**, *590*, 89–96.
- (50) Hernández-Lobato, J. M.; Requeima, J.; Pyzer-Knapp, E. O.; Aspuru-Guzik, A. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. International conference on machine learning. 2017; pp 1470–1479.
- (51) Kandasamy, K.; Krishnamurthy, A.; Schneider, J.; Póczos, B. Parallelised Bayesian optimisation via Thompson sampling. International conference on artificial intelligence and statistics. 2018; pp 133–142.
- (52) González, J.; Dai, Z.; Hennig, P.; Lawrence, N. Batch Bayesian optimization via local penalization. Artificial intelligence and statistics. 2016; pp 648–657.
- (53) De Palma, A.; Mender-Dünner, C.; Parnell, T.; Anghel, A.; Pozidis, H. Sampling



- acquisition functions for batch Bayesian optimization. *arXiv preprint arXiv:1903.09434* **2019**,
- (54) Filho, M. A. M.; Farmer, W.; Hsiao, C.-L.; Dos Santos, R. B.; Hultman, L.; Birch, J.; Ankit, K.; Gueorguiev, G. K. Density functional theory-fed phase field model for semiconductor nanostructures: The case of self-induced core-shell InAlN nanorods. *Crystal growth & design* **2024**, *24*, 4717–4727.
- (55) Yoo, A. B.; Jette, M. A.; Grondona, M. Slurm: Simple linux utility for resource management. Workshop on job scheduling strategies for parallel processing. 2003; pp 44–60.
- (56) Ahn, D. H.; Bass, N.; Chu, A.; Garlick, J.; Grondona, M.; Herbein, S.; Ingólfsson, H. I.; Koning, J.; Patki, T.; Scogland, T. R.; others Flux: Overcoming scheduling challenges for exascale workflows. *Future Generation Computer Systems* **2020**, *110*, 202–213.
- (57) Quinlan, B. Futures - Execute computations asynchronously. PEP 3148, 2009; [Online]. Available: <https://www.python.org/dev/peps/pep-3148/>.
- (58) Močkus, J. On Bayesian methods for seeking the extremum. Optimization techniques IFIP technical conference: Novosibirsk, July 1–7, 1974. 1975; pp 400–404.
- (59) Jones, D. R.; Schonlau, M.; Welch, W. J. Efficient global optimization of expensive black-box functions. *Journal of Global optimization* **1998**, *13*, 455–492.
- (60) Daulton, S.; Ament, S.; Eriksson, D.; Balandat, M.; Bakshy, E. Unexpected improvements to expected improvement for Bayesian optimization. Proceedings of the 37th International Conference on Neural Information Processing Systems. Red Hook, NY, USA, 2023.
- (61) Ginsbourger, D.; Le Riche, R.; Carraro, L. *Computational intelligence in expensive optimization problems*; Springer, 2010; pp 131–162.



- (62) Kushner, H. J. A New Method of Locating the Maximum Point of an Arbitrary Multi-peak Curve in the Presence of Noise. *Journal of Basic Engineering* **1964**, *86*, 97–106.
- (63) Ginsbourger, D.; Riche, R. L.; Carraro, L. A Multi-points Criterion for Deterministic Parallel Global Optimization based on Gaussian Processes. 2008.
- (64) Snoek, J.; Larochelle, H.; Adams, R. P. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems* **2012**, *25*.
- (65) Balandat, M.; Karrer, B.; Jiang, D.; Daulton, S.; Letham, B.; Wilson, A. G.; Bakshy, E. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. *Advances in neural information processing systems* **2020**, *33*, 21524–21538.
- (66) Williams, C. K.; Rasmussen, C. E. *Gaussian processes for machine learning*; MIT press Cambridge, MA, 2006; Vol. 2.
- (67) Thompson, A. P.; Aktulga, H. M.; Berger, R.; Bolintineanu, D. S.; Brown, W. M.; Crozier, P. S.; In't Veld, P. J.; Kohlmeyer, A.; Moore, S. G.; Nguyen, T. D.; others LAMMPS-a flexible simulation tool for particle-based materials modeling at the atomic, meso, and continuum scales. *Computer Physics Communications* **2022**, *271*, 108171.



Data and Code Availability

Data and code for this article, including example workflows and analysis notebooks are available at Zenodo at <https://doi.org/10.5281/zenodo.18444064>.

