To be submitted to Rev. Sci. Instr.

stemOrchestrator: Enabling Seamless Hardware Control and High-Throughput Workflows on Electron Microscopes.

Utkarsh Pratiush[1*], Austin Houston[1], Paolo Longo[2], Remco Geurts[2], Sergei V Kalinin[1,3], and Gerd Duscher[1*]

[1]Department of Materials Science and Engineering, University of Tennessee, Knoxville, TN 37996, USA
[2]ThermoFisher Scientific, Eindhoven, the Netherlands
[3]Pacific Northwest National Laboratory, Richland, WA 99254, USA

*Corresponding author's: upratius@vols.utk.edu and gduscher@utk.edu

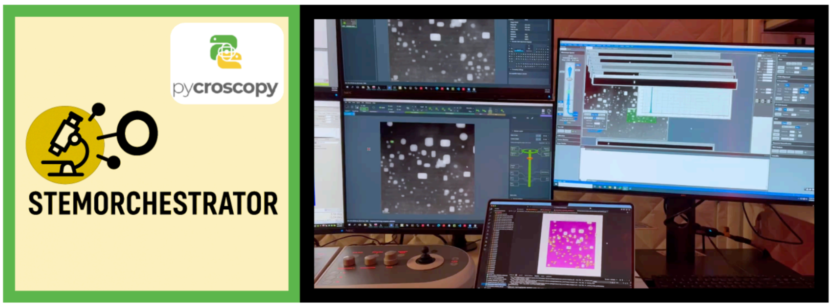

Abstract
Scanning Transmission Electron Microscopy (STEM) is one of the most powerful tools for materials characterization, providing access to atomic-scale structure via direct imaging, chemical composition, orbital populations, vibrational properties and quasiparticles via spectral methods, and crystallographic information through diffraction. However, these diverse functionalities are often supported by hardware components from different manufacturers, creating challenges in seamless operation and integration of detectors, holders, and cameras on a single STEM column. As the field moves toward machine learning (ML) enabled experiments and autonomous discovery, the need for combined control across these hardware interfaces becomes critical. This paper develops stemOrchestrator, a software framework which combines detector- and camera specific Application Programming Interfaces (APIs) in a cohesive platform for controlling various STEM hardware modules and developing sophisticated automated workflows. We illustrate the performance of stemOrchestrator using several model workflows including high-throughput particle characterization, hardware tuning using Bayesian Optimization (BO), cross correlation-based drift correction with informative logging of hardware status. Importantly, these workflows are presented not as isolated technical advances, but to highlight how stemOrchestrator renders their implementation almost trivial by abstracting hardware heterogeneity and execution logic. This framework also enables seamless integration with LLM (Large language model) agents to suggest and run complex automated workflows and opens pathway for orchestration of self-driving labs and geographically distributed instrumentation networks. The codes are available at this link for trying and contributing: https://github.com/pycroscopy/pyAutoMic/tree/main/TEM/stemOrchestrator

I. Introduction

Scanning Transmission Electron Microscopy (STEM) is a cornerstone technique in contemporary materials science, bringing exceptional imaging, spectroscopy, and diffraction capabilities(Pennycook, 2011; Williams & Carter, 1996; Crewe, 1974). By focusing the electron beam into an angstrom-sized probe, STEM enables the direct visualization of crystal defects, interfaces, and grain boundaries with picometer precision (Batson et al., 2002). When coupled with aberration correction, it facilitates simultaneous Z-contrast imaging and atomic-scale chemical mapping, allowing researchers to quantify elemental composition and oxidation states with single-atom sensitivity using Electron Energy Loss Spectroscopy (EELS) and Energy Dispersive X-ray Spectroscopy (EDS) (Krivanek et al., 2010; Muller et al., 2008). Furthermore, recent advancements in high-speed pixelated detectors have unlocked four-dimensional (4D) STEM modalities, permitting the reconstruction of internal electromagnetic fields, strain tensors, and ptychographic phase images for observing light elements (Ophus, 2019; Shibata et al., 2012).

STEM and other microscopy experiments are almost always interactive in nature. The human operator dynamically interacts with the instrument to optimize the alignment, redefining regions of interest, switching modalities, and triggering targeted spectra acquisition as hypotheses evolve. Effectively, the discovery proceeds via on-the-fly workflow construction, enabling the experiment to zoom, select regions of interest, tune, or acquire spectral data within the same session. These traditional manual or semi-automated methods can be labor-intensive, prone to operator bias, and ill-suited for high-throughput experimentation(Kalinin et al., 2022, 2023). In addition, human decision times are well below intrinsic latencies of EM operation, leading to fairly low efficiencies of human-led instrumentation.

These considerations resulted in tremendous interest to ML-enabled automated electron microscopy. These developments span two complementary aspects. The first is the principles of the decision making. Here, areas such as cryo-EM or semiconductor characterization often require fixed policy workflows that can be defined before the experiment or tuned via small number of parameters. More complex reward driven tasks include instrument optimization or structure-property relationship discovery. Finally, multistep workflows will balance the duration of the optimization and data acquisition tasks, as well as drive the physics-based analytics and upstream sample selection(Kalinin et al., 2025).

The second aspect is the instrument control. A number of groups demonstrated real-time data analytics where instrument data streams are analyzed in real time and piped to ML and subsequently MD/DFT environment. For example, Barakati et al. demonstrated the use of custom reward-based segmentation workflows(Barakati et al., 2025) directly on the instrument, but their approach excludes spectroscopic and diffraction acquisition. Cherukara(Cherukara, 2024) et al. and Welborn(Welborn et al., 2024) et al. have made strides in streaming data analysis for real-time insights during microscopy experiments. However, these systems mainly focus on managing data flow rather than implementing closed-loop instrument control. However, in this case the decisions are still made and executed by human operator. Many more efforts have concentrated on post-acquisition data management and analysis(Abebe et al., 2025; Meirovitch et al., 2023; Archit et al., 2025) but do not integrate with real-time hardware control.

More complex examples are real time controls, when ML agents directly control the instrument. Here, Pratiush et al. introduced PyAUTOMIC(Pycroscopy, 2025; Pratiush, Houston, et al., 2024), a platform designed for automating STEM experiments. However, its control capabilities are primarily restricted to Gatan image filters, limiting compatibility with other hardware. Similarly, Pattison et al. and Ma et al. developed BEACON(Pattison et al., 2024) and emittance minimization approach(Ma et al., 2025a, 2025b) respectively to automate aberration correctors' operations. Although effective in optical alignment and tuning, these work do not integrate with spectroscopic detectors like EELS or EDX. Human-in-the-loop frameworks like hAE(Pratiush, Duscher, et al., 2024) support interactive decision-making during microscopy sessions. Additionally, active learning approaches using Bayesian optimization have been explored for planning microscopy experiments(utkarshp1161, 2024; Ziatdinov, Liu, et al., 2021; Ziatdinov, Ghosh, et al., 2021). Despite being effective in theory, many of these methods have only been tested on pre-acquired datasets(Pratiush, Roccapriore, et al., 2024; Zhu et al., 2022; Shi et al., 2022; Ge et al., 2022) and lack the real-time hardware interfacing needed for dynamic experimental control. Furtherore, they are limited in their ability to simultaneously query multiple spectroscopic modalities, which restricts their applicability for complex multi-modal workflows.

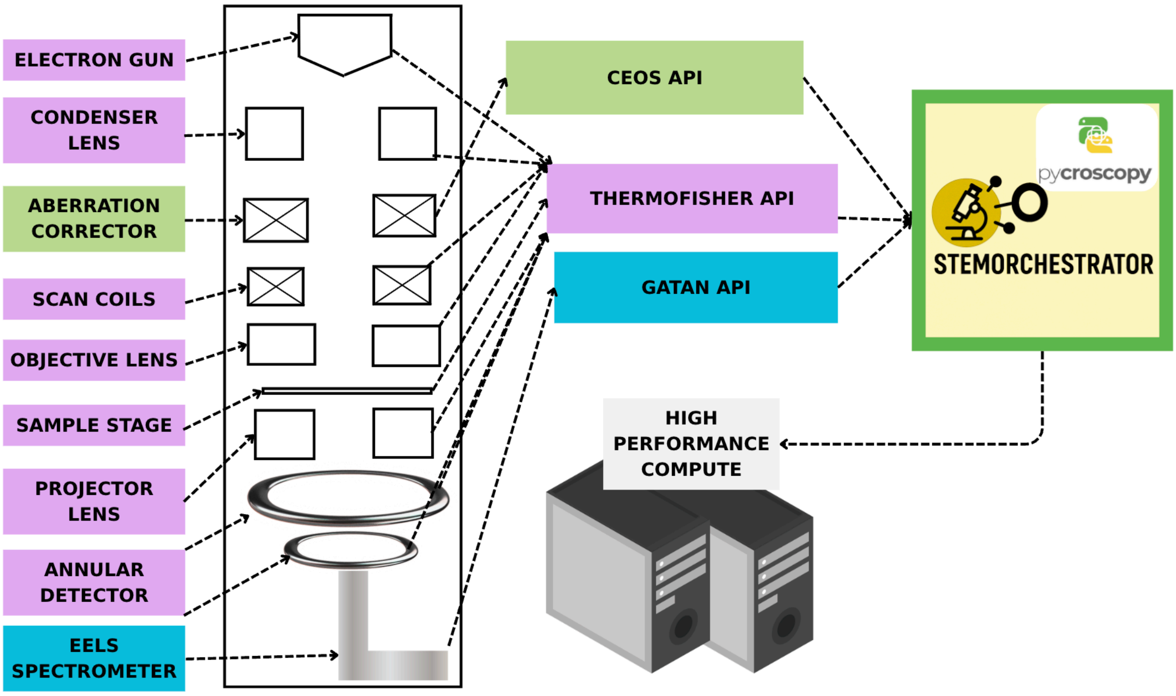

Figure 1: Illustration of the heterogeneous API control in a modern advanced STEM(Schematic). The schematic highlights how different hardware components of the microscope are controlled by disparate manufacturer APIs (ThermoFisher, CEOS, and Gatan).

Overall, while artificial intelligence (AI) and machine learning (ML) have demonstrated significant potential to improve microscopy workflows(Spurgeon et al., 2021; Kalinin et al., 2022, 2023; Botifoll et al., 2022) the increasing complexity of STEM experiments necessitates advanced automation solutions for efficient hardware integration, data acquisition, and analysis. Realizing the potential of machine learning and automated characterization in microscopy requires unified control over fragmented hardware APIs (**Figure 1**).

To address this, here we developed stemOrchestrator: a modular framework that integrates multiple STEM interfaces to facilitate sophisticated, automated workflows. We demonstrate the flexibility of the framework via three representative workflows: high-throughput particle characterization, hardware tuning using BO, and cross-correlation–based drift correction with informative logging of hardware status. Beyond these use cases, the framework is designed to support future integration of large language model (LLM) agents that can intervene, suggest, and execute complex automated workflows. The codebase is openly available for use and community contribution at:
https://github.com/pycroscopy/pyAutoMic/tree/main/TEM/stemOrchestrator

II. Software Implementation and Network Architecture

II. 1. Design Philosophy and Internal Logic

The internal logic of stemOrchestrator is designed to serve as a unified gateway to the microscope, bridging the critical gap between physical hardware and computational logic. While the architecture is necessitated by the sequential data flow of a typical STEM experiment, spanning Acquisition, Processing, Simulation, and Decision Making, the framework does not merely mirror this sequence; rather, it provides the essential "window" through which these distinct stages can interact seamlessly.

In a modern analysis workflow, a researcher must continually cycle through diverse tasks: configuring lenses and detectors, standardizing high-dimensional data streams, and comparing experimental outputs with simulations to make real-time decisions. stemOrchestrator acts as the programmable entrance to this environment. By abstracting the fragmented manufacturer APIs, it transforms the microscope into a cohesive entity accessible to software agents. This "window" allows complex, automated workflows, for example high-throughput particle characterization, active learning for hardware tuning, or adaptive drift correction, to drive the experiment directly, treating the instrument as a responsive component within a larger computational loop.

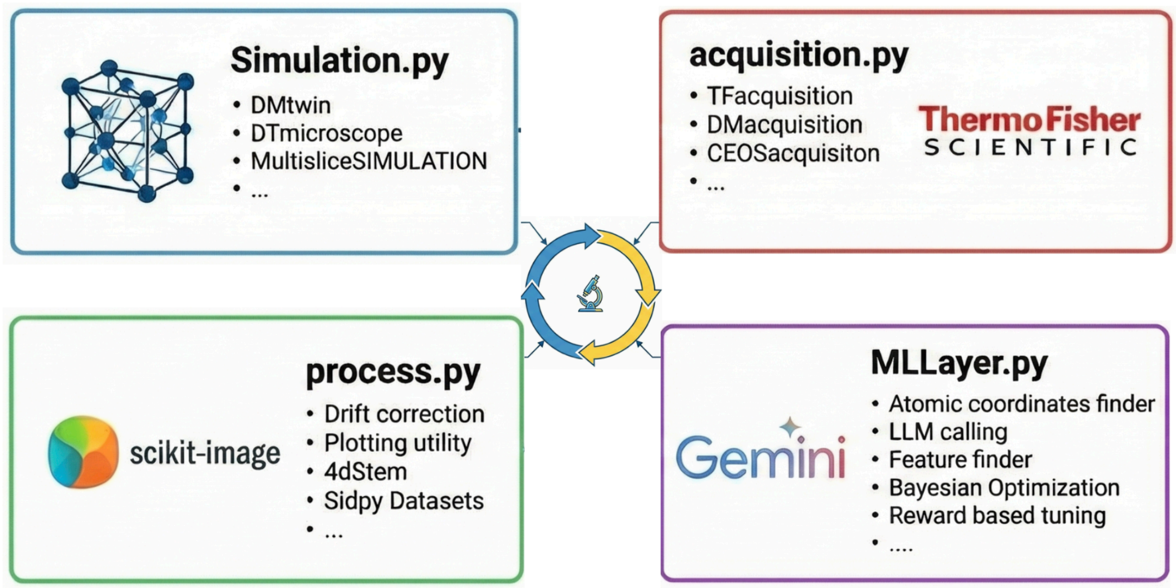

Figure 2. Visual block layout of stemOrchestrator showing the modularized design philosophy. The *acquisition.py* module is the central focus of this work (hardware control).

The codebase is structured into interoperable components to ensure extensibility and ease of debugging.

1. Hardware Abstraction and Protocol Management:
The foundational communication layer (hardware.py) abstracts the complexities of asynchronous network programming by utilizing the Twisted framework. We implemented the _CEOSProtocol class, which inherits from NetstringReceiver, to ensure data integrity by encapsulating JSON payloads within netstrings (length-prefixed byte strings). This design prevents TCP fragmentation issues common in high-frequency instrument polling. The CEOSacquisition class employs a deferred execution model where commands generate unique transaction IDs (self._next_id) and register Deferred objects in a pending dictionary,

allowing the orchestrator to remain responsive to user interrupts while waiting for hardware acknowledgment. For high-speed operations where the overhead of the Twisted event loop is unnecessary, the CEOSacquisitionTCP class provides an alternative blocking socket connection with a 5-minute timeout, optimized for rapid, sequential command execution.

2. Acquisition Layer and Instrument State tracking:
The acquisition.py module serves as the primary interface for instrument state management and automated control routines. The TFacquisition class initializes interactions by verifying detector presence via the TemMicroscopeClient and handles the critical task of normalizing disparate hardware units, converting relative stage translations to nanometers and absolute positions to meters, while tilt angles are transformed between radians and millidegrees using helper methods like rad2miliDegree_stage. To ensure downstream models receive physically calibrated data, image acquisition routines return "Adorned" objects containing essential metadata such as pixel size and defocus rather than simple raw arrays. Beyond basic control, this layer implements intelligent workflows; for example, the jog_and_capture_with_metrics method executes step-wise stage movements while calculating image statistics in real-time, utilizing an early stopping mechanism if the metric declines consecutively to automate dataset creation.

3. Data Processing and Metadata Parsing:
The process.py module provides a unified toolkit for image normalization, metadata extraction, and registration. To ensure data provenance, the tiff_metadata function utilizes tifffile to parse standard tags and extract embedded instrument XML data (specifically under the "FEI_TITAN" tag) into nested Python dictionaries using xml.etree.ElementTree. This metadata supports accurate visualization, such as the HAADF_tiff_to_png function, which uses matplotlib_scalebar to draw physically accurate scale bars based on the extracted pixel size. The module also implements sub-pixel drift correction via compute_drift, which uses phase cross-correlation (skimage.registration.phase_cross_correlation), and verifies the results using plot_drift_comparison to generate composite views of the reference and shifted images.

4. Machine Learning Integration (MLlayer):
The Machine Learning Layer (MLlayer) is divided into specialized sub-modules for optimization and segmentation. The MLlayerBO.py module wraps packages such as BoTorch and GPyTorch to implement Gaussian Process (GP) regression, using classes like Tune1d, Tune2d, and Tune3d to manage the optimization loop. This loop iteratively maximizes the Expected Improvement (EI) acquisition function to select candidate hardware parameters for the next experimental step. Parallel to this, MLlayerSAM.py integrates the Segment Anything Model (SAM) for particle detection, preprocessing grayscale images into RGB format and extracting object boundaries using OpenCV's cv2.findContours. Crucially for automated targeting, the sample_particle_positions function applies a stratified sampling strategy to extract both centroids and a user-defined percentage of boundary points, facilitating precise beam positioning for spectroscopic analysis.

5. Simulation and Offline Development:
Finally, to ensure code reproducibility and facilitate offline development, the simulation.py module includes a DMtwin class. This digital twin mimics the API signatures of the real Gatan spectrometer, generating synthetic EELS spectra that include physically realistic features such as Zero Loss Peaks (ZLP), plasmon resonances, and background noise. This

abstraction allows the entire orchestration pipeline to be tested and debugged without requiring immediate access to the physical microscope.

II. 2. Network Architecture and Communication Protocols

To bridge the gap between legacy instrument control PCs (often running Windows) and modern High-Performance Computing (HPC) environments where ML orchestration occurs, we implemented a multi-protocol communication layer. This design allows the main orchestration script to "import" the microscope as if it were a local library, completely hiding the complexity of the underlying network transport.

We utilized three distinct connection modalities, each optimized for specific hardware constraints. For the CEOS aberration corrector, we implemented an asynchronous, event-driven messaging layer built on the Twisted framework. This architecture structures the communication payload using JSON-RPC over Twisted Netstrings, a method chosen specifically to define rigid message boundaries and prevent packet fragmentation. To ensure experimental reliability, the `_CEOSFactory` class within the hardware layer manages intrinsic reconnection logic, making long-duration data acquisition campaigns resilient to transient network drops or jitter.

In scenarios requiring high-frequency feedback, where the overhead of JSON parsing and event-loop scheduling is undesirable, we developed a lightweight direct control interface. This client communicates via raw TCP sockets, bypassing higher-level abstraction layers to interact directly with ports exposed by the server hardware. This "bare-metal" approach minimizes computational overhead, offering the lowest possible latency for critical control loops where millisecond-level timing is essential.

Finally, for instruments that are restricted to local Python APIs, specifically Thermo Fisher AutoScript and Gatan Digital Micrograph, we employed a remote object proxying strategy using wrappers such as Pyro5. Because these systems do not natively support network control, we wrap their local API objects to expose them over the network. This allows the central client (`acquisition.py`) to interact with a local proxy object that transparently handles the serialization and transmission of calls to the remote instrument PC, unifying disparate interfaces under a single control architecture.

II. 3. Security and Reproducibility

Given that instrument PCs often run older operating systems required by vendor software, the framework allows heavy computational logic to be offloaded to secure HPC environments. Users deploying this framework must ensure that the TCP/RPC ports used by *hardware.py* are firewalled and accessible only via a secure local VLAN or SSH tunnel.
To reproduce this work, researchers can use the modular design; specifically, the simulation.py module allows for the full execution of the *MLlayer.py* and *process.py* workflows using synthetic data, enabling development without immediate access to physical microscope hardware.

III. Example workflow enabled

In this section, we use stemOrchestrator to illustrate some of the most common workflows, demonstrate their implementation, and propose benchmark examples for evaluating performance and reproducibility. Also, highlighting how stemOrchestrator makes it possible that these complex workflows almost trivially are realized on the Real microscope.

IIIA. SAM-EDX-EELS-Diffraction

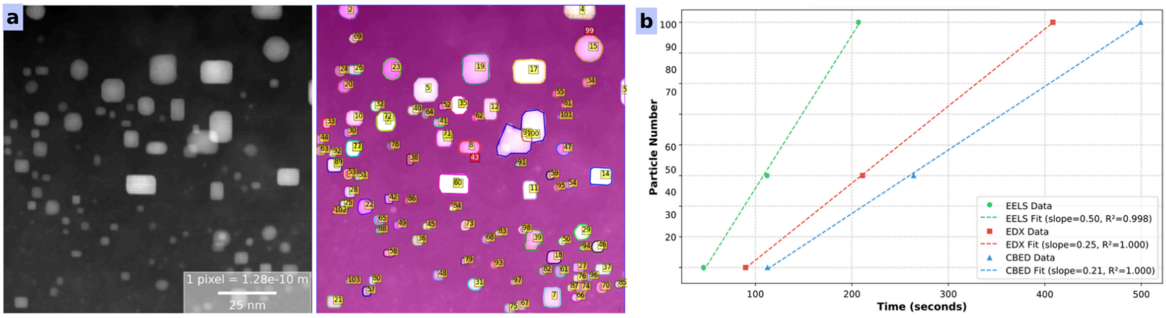

Figure 3. Automated particle analysis using HAADF-STEM imaging and multimodal spectroscopy. (a) Example segmentation maps showing over 100 particles, detected using the Segment Anything Model (SAM). (b) Acquisition times for Electron Energy Loss Spectroscopy (EELS), Energy-Dispersive X-ray Spectroscopy (EDX), and Diffraction across nine experimental runs at different particle counts. The near-linear scaling trend demonstrates the pipeline's efficiency and suitability for high-throughput characterization.

As a first example of common STEM workflows, we demonstrate the high-throughput particle library characterization, a common task modern materials research, especially in studies involving heterogeneous nanoparticles and catalytic systems. Traditional Scanning Transmission Electron Microscopy (STEM)-based Electron Energy Loss Spectroscopy (EELS) workflows can take upwards of 10–12 hours to acquire spectrum images for just ~100 particles, often requiring custom scripts to selectively sample only at particle boundaries or centers to reduce overhead. stemOrchestrator allows native integration between the instrument data stream and the established Python ML/AI ecosystem including foundational vision model.

Here we demonstrate an accelerated, automated STEM workflow powered by the Segment Anything Model (SAM), enabling rapid, multimodal characterization of particles using Diffraction, Energy-Dispersive X-ray Spectroscopy (EDX), and EELS. Here, the SAM model is used to segment images to discover individual particles, and the policies are defined as measurement of spectral signal for each particle. The workflow includes pre-processing to enhance contrast, SAM-based segmentation for robust particle identification, and post-processing to extract morphological and statistical descriptors. This approach bypasses the bottleneck of manual segmentation and spectrum image acquisition, allowing automated analysis with minimal human intervention.

The workflow is realized on a ThermoFisher Scientific (TFS) STEM instrument interfaced with the ISAACS supercomputer at the University of Tennessee, Knoxville (UTK). Hardware control and acquisitions (HAADF, EDX, Diffraction) were managed through ThermoFisher AutoScript, while EELS acquisition was automated via the Gatan DigitalMicrograph (DM) Server software using the Gatan Image Filter (GIF). We demonstrate this capability on a TiO_2–Au sample, where Au nanoparticles are embedded in a TiO_2 matrix. As shown in **Figure 3 a**, our pipeline processes datasets with 20+, 50+, and 100+ particles across three magnifications. Benchmarking against manually curated datasets confirms that we maintain high accuracy while significantly improving throughput.

Furthermore, the use of the ML-workflow allows timing individual operations, providing access to individual latencies and costs. Specifically, our system acquires EELS data for 100 particles in just ~3 minutes, EDX in ~7 minutes, and CBED Diffraction in ~10 minutes, scaling linearly with particle count, as illustrated in **Figure 3**. Acquisition times per particle are as low as 0.1 seconds for CBED, 20 milliseconds for EELS, and 2 seconds for EDX.

III B. Drift correction

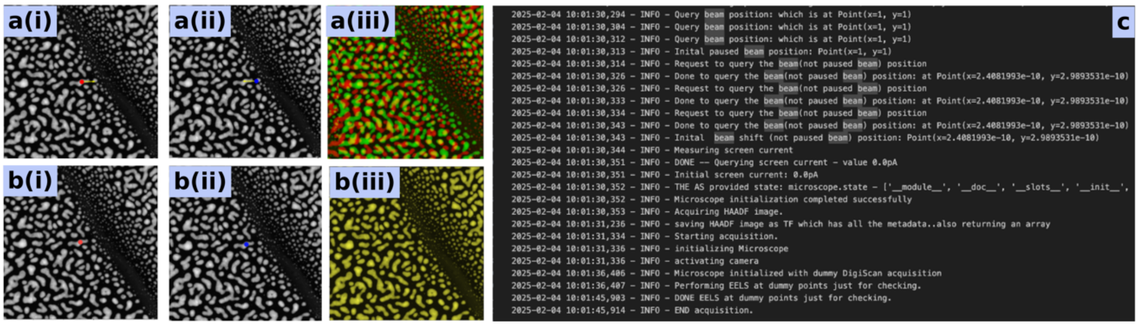

Figure 4. Panels (a) and (b) illustrate an example of the real-time drift correction workflow applied to High-Angle Annular Dark-Field (HAADF) images within the STEM Orchestrator. In (a), the stage was intentionally shifted to compute drift, resulting in a measured shift of 37.14 pixels in X and -0.43 pixels in Y. In (b), the verification step shows no drift (Image 1 = Image 2) with a measured shift of 0 pixels in both X and Y. Panel (c) displays an example of the live logging output, highlighting how various hardware components are actively queried during the drift correction process.

High-precision microscopy workflows often require dynamic interaction between an ML agent and the hardware modules, querying an image, identifying features of interest, determining their coordinates, and then triggering spectroscopic acquisitions at those target locations. However, sample drift during this processing loop introduces spatial misalignments, which can significantly affect the accuracy of data collection, particularly in tasks requiring nanoscale or even sub-angstrom precision. To address this, we implemented an automated drift correction module within the STEM workflow **(Figure 4)**. The core mechanism of our solution relies on phase cross-correlation between a stable "reference" image acquired initially and a "current" image taken immediately prior to spectral acquisition. Rather than simple intensity matching in the spatial domain, this method operates in the frequency domain via Fast Fourier Transforms (FFTs). By calculating the phase difference between the Fourier transforms of the two images, the algorithm determines the relative translational shift robustly against noise and slight illumination variations. To achieve the necessary nanoscale precision beyond the pixel grid, we utilized an efficient sub-pixel registration algorithm, specifically implemented using the phase_cross_correlation function within the scikit-image Python library with up sampling, to estimate the precise drift vector $(\Delta x, \Delta y)$. As demonstrated in Figure 4, this vector is then immediately inverted and applied to the microscope hardware, either by introducing an offsetting beam shift or physically repositioning the stage, ensuring that subsequent spectroscopic measurements are aligned with the originally selected target points, a result confirmed by the verification step shown in Figure 4(b). Example of Logging of this process is shown in Figure 4(c) which is pretty handy for debugging.

Drift can also be handled using predictive models that learn the underlying behavior of the stage or sample over time, which may be beneficial in high-drift or ultra-precise applications such as atomic-scale manipulation. While these approaches are part of our future roadmap, the current cross-correlation method offers a robust and generalizable solution for most use cases. As shown in Figure 3, this correction module has been integrated seamlessly into the real-time feedback loop of our automated STEM control stack.

III C. Bayesian Optimization and Active Learning for Experimental Automation using Rewards

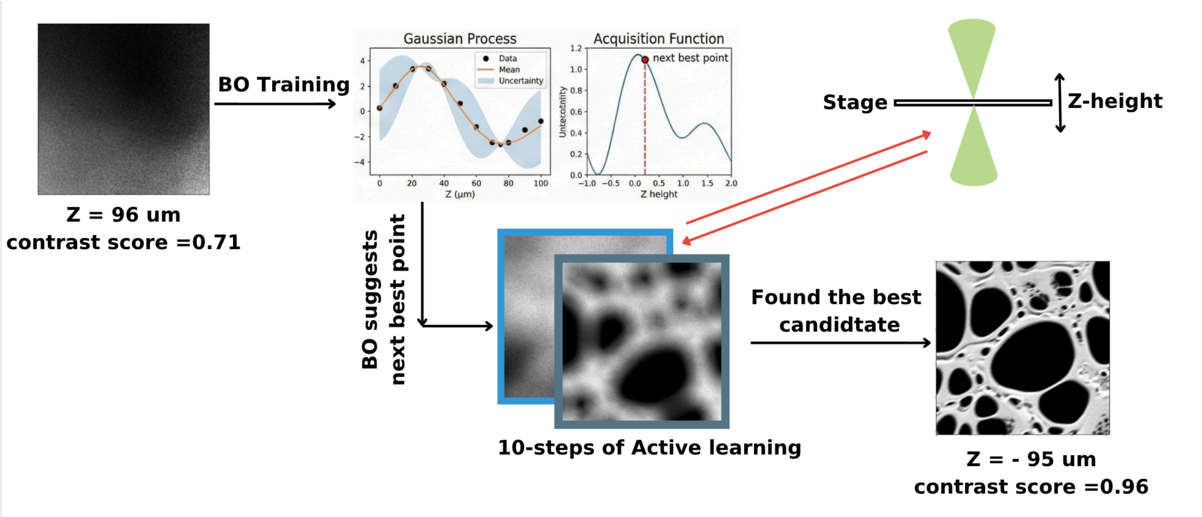

Figure 4. Using BO to tune the sample height of the STEM.

Experimental design in electron microscopy often involves selecting the next measurement point based on prior knowledge, intuition, or visual inspection. However, human decision-making is inherently sequential, potentially biased, and difficult to reproduce. Active learning provides a data-driven alternative by systematically suggesting where to measure next—balancing exploration and exploitation strategies. This can be formalized as a black-box optimization problem:

$$x^{\cdot} = arg\ \boldsymbol{max}_{x \in X}\ f\,(\,\mathbf{x})\qquad (1)$$

where $f(x)$ is an expensive-to-evaluate function (e.g., quality of an image or spectrum), and x represents the experimental parameters (e.g., defocus, aberration coefficients, sample height). Since $f\ (x)$ is not known a priori and costly to sample, Bayesian Optimization (BO) offers an efficient framework to model it using a surrogate (typically a Gaussian Process, GP) and choose the next point via an acquisition function like Expected Improvement (EI), Upper Confidence Bound (UCB), or Entropy Search.

In microscopy, BO has proven useful in automating parameter tuning tasks. For example:, BO can help find optimal corrector settings that minimize image distortion or maximize image sharpness based on sample height, stigmator coefficients or aberration coefficients values, as demonstrated in prior works like BEACON(Pattison et al., 2024). Alternatively, BO can guide the microscope to regions with higher likelihood of exhibiting desired features such as elemental peaks, edge plasmons and other desired features in the EELS spectrum(Pratiush, Roccapriore, et al., 2024).

The optimization algorithms including multiple versions of BO with different surrogate models, reinforcement learning, or multiple versions of stochastic optimizations/dynamic programming are well established in ML/optimization community. Deploying these methods as a part of automated experiments requires the definition of appropriate policies and/or reward or probabilistic reward (roll-out) functions and building the bridge between the ML agent and the instrument. stemOrchestartor addresses the latter aspect, allowing for interfacing the optimization algorithm and full set of instrument controls.

To illustrate this approach, we demonstrate the sample height tuning as shown in **figure 4.** Here, The Optimization problem is same as shown in eqn (1), where f(x) is the contrast measure at the different sample height. The sample height "x" can vary for -100 micrometer to +100 micrometer. The contrast measure is defined as

$$f(\mathbf{x}) = \frac{\sigma(\bar{I}(x))}{\mu(\bar{I}(x))} \tag{2}$$

Where $\bar{I}(x)$ is the normalized HAADF image acquired at sample height "x", $\sigma(\bar{I}(x))$ is the calculated std deviation of the normalized image and $\mu(\bar{I}(x))$ is the mean of the normalized image. Eq. 2 defines the objective function, which is being optimized shown in equation 1. Using the simple Gaussian Process based optimization, the stemOrchestartor allows optimizing sample height.

III. D LLM assisted image understanding and orchestration.

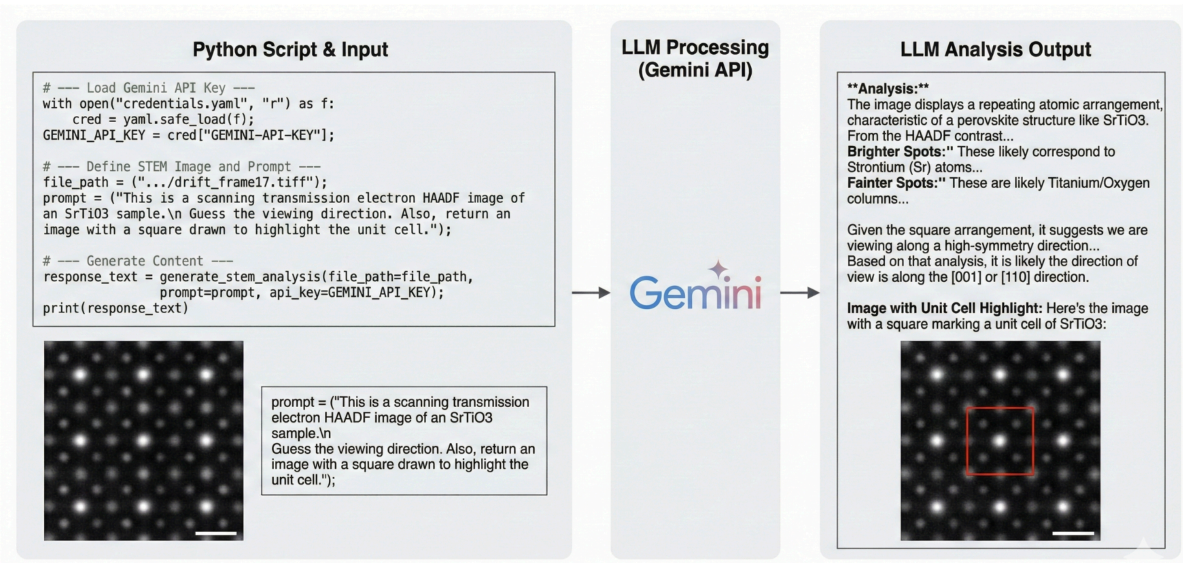

Figure 5. Using LLM in the experimental loop for image understanding.

The hallmark of our time are rapidly emerging LLM- and reasoning-based agentic autonomous scientists. (Sapkota et al., 2026) Since 2020, AI Feynman,(Udrescu & Tegmark, 2020) EUREKA,(Ma et al., 2023) and AI Descartes(Cornelio et al., 2023) brought forth the concept of a closed-loop, self-directed systems capable of simplifying complex relationships into interpretable equations through the iterative logical and physical reasoning and pattern discovery. Over the last 3 years, the rapid advancement of LLMs has led to the development of several frameworks(CrewAI Inc., 2024; Wu et al., 2023, Anon, 2025) that enable the realization of multi-LLM agentic systems capable of independent reasoning, tool use, and inter-agent collaborations.(Li et al., 2023; Yang et al., 2023; Nakajima, 2025; M. Bran et al., 2024; Boiko et al., 2023; Panapitiya et al., 2025; Song et al., 2025) For materials characterization, SciLink (a multi-agent, theory-in-the-loop workflow linking experiment, novelty detection, and simulation),(Yao et al., 2025) AILA (AFM automation via LLM-driven agents),(Mandal et al., 2025) and HoneyComb (an agentic framework integrating a knowledge base and tool hub for autonomous experimentation) has emerged.(Zhang et al., 2024) These are complemented by general-purpose research frameworks AI Scientist,(Lu et al., 2024; Yamada et al., 2025) AI Researcher,(Tang et al., 2025) Curie(Kon et al., 2025) and KOSMOS.(Mitchener et al., 2025) These rapidly emerging systems allow the integration of multiple agents into hierarchical structures for performing cross-discipline and multi-layer scientific investigations. Even turnkey Large Language Models (LLMs) when included in Python workflows offer a powerful new interface between microscopy images and actionable experimental insights. Multiple examples of such integration can be found in the outcomes of recent ML for Microsocpy Hackathons(Pratiush et al., 2025).

stemOrchestrator offers a native bridge allowing to directly connect such reasoning agents to real instrument tos. **Figure 5** shows an example validating zone axis alignment between literature and acquired image and deciding to tilt the stage to another desired zone axis. Beyond interpretation, LLMs act as orchestration assistants(to be added in future), capable of interfacing across the three layers (Hardware, ML and simulation and experiment logic) of the microscope control stack: Layer 1 (Hardware): LLMs can parse equipment manuals and provide quick guidance on commands, syntax, and expected behavior for hardware-level operations such as stage tilt or beam alignment. Layer 2 (Software & ML): They can recommend analysis workflows, pre-processing techniques, or model selection strategies tailored to the sample and imaging goals. Layer 3 (Experiment Logic): LLMs can suggest adaptive experiment plans, including candidate reward functions for optimization, e.g., maximizing image contrast, minimizing drift, or targeting edge-on interfaces in tomography.

IV. Conclusion

In summary, STEM-Orchestrator is a modular, open-source framework that unifies control of diverse STEM hardware, decoupling acquisition, simulation, processing, and ML layers into interoperable modules to enable rapid, high-throughput workflows. We demonstrated its effectiveness through particle mapping an AI-driven segmentation and multimodal spectroscopy pipeline that reduces acquisition times, and a real-time drift-correction module that maintains nanoscale alignment during extended experiments. By seamlessly integrating Bayesian optimization into live feedback loops. Its cohesive interface to beams, stages, detectors, and existing machine learning and electron microscopy codes not only accelerates materials discovery but also fosters collaboration and reproducibility across Machine learning, experiment and simulation communities. The framework's extensible design invites community contributions—whether adding new instrument backends, simulation engines, or advanced AI models and paves the way for next-generation, LLM-assisted agentic workflows.

Building on the modular and extensible design of our current workflow, future efforts will focus on integrating this framework into the broader BlueSky ecosystem(Allan et al., 2019), enabling real-time feedback loops and multithreaded experimental orchestration in electron microscopy. We also envision human in the loop and incorporation of Large Language Model (LLM)-based agentic workflows, where intelligent agents can autonomously plan, execute, and adapt microscopy experiments based on literature insights, past data, and predefined scientific goals.

The source code is available for testing and extension at https://github.com/pycroscopy/pyAutoMic/tree/main/TEM/stemOrchestrator.

Acknowledgement

The work was partially supported (UP) AI Tennessee Initiative at University of Tennessee Knoxville (UTK). (U. P, A.H, G. D, S.V.K) acknowledges support from high performance computing facility, ISAAC and Microscopy Core Facility, of The University of Tennessee, Knoxville (UTK). This work was partially supported (AH, GD) by the U.S. DOE, Office of Science, Materials Sciences and Engineering Division and the Center for Nanophase Materials Sciences, which is a DOE Office of Science User Facility. The work was partially supported (SVK) by the Center for Advanced Materials and Manufacturing (CAMM), the NSF MRSEC center. The preprint of this is article is available at https://doi.org/10.31224/4645

References

ABEBE, W., STRUBE, J., GUO, L., TALLENT, N. R., BEL, O., SPURGEON, S., DOTY, C. & JANNESARI, A. (2025). SAM-I-Am: Semantic boosting for zero-shot atomic-scale electron micrograph segmentation. *Computational materials science* 246, 113400.

ALLAN, D., CASWELL, T., CAMPBELL, S. & RAKITIN, M. (2019). Bluesky's Ahead: A Multi-Facility Collaboration for an *a la Carte* Software Project for Data Acquisition and Management. *Synchrotron Radiation News* 32, 19–22.

ARCHIT, A., FRECKMANN, L., NAIR, S., KHALID, N., HILT, P., RAJASHEKAR, V., FREITAG, M., TEUBER, C., BUCKLEY, G., VON HAAREN, S., GUPTA, S., DENGEL, A., AHMED, S. & PAPE, C. (2025). Segment Anything for Microscopy. *Nature methods* 22, 579–591.

BARAKATI, K., LIU, Y., PRATIUSH, U., SLAUTIN, B. N. & KALININ, S. V (2025). Rewards-based image analysis in microscopy. https://arxiv.org/abs/2502.18522.

BATSON, P. E., DELLBY, N. & KRIVANEK, O. L. (2002). Sub-ångstrom resolution using aberration corrected electron optics. *Nature* 418, 617–620.

BOIKO, D. A., MACKNIGHT, R., KLINE, B. & GOMES, G. (2023). Autonomous chemical research with large language models. *Nature* 624, 570–578. https://www.nature.com/articles/s41586-023-06792-0.

BOTIFOLL, M., PINTO-HUGUET, I. & ARBIOL, J. (2022). Machine learning in electron microscopy for advanced nanocharacterization: current developments, available tools and future outlook. *Nanoscale Horizons* 7, 1427–1477.

CHERUKARA, M. J. (2024). High Performance Computing and Artificial Intelligence Enabled Materials Characterization and Experimental Automation. *Microscopy and Microanalysis* 30, ozae044.1006. https://doi.org/10.1093/mam/ozae044.1006.

CORNELIO, C., DASH, S., AUSTEL, V., JOSEPHSON, T. R., GONCALVES, J., CLARKSON, K. L., MEGIDDO, N., EL KHADIR, B. & HORESH, L. (2023). Combining data and theory for derivable scientific discovery with AI-Descartes. *Nature Communications* 14, 1777. https://www.nature.com/articles/s41467-023-37236-y.

CREWAI INC. (2024). crewAI. https://github.com/crewAIInc/crewAI.

CREWE, A. V. (1974). Scanning transmission electron microscopy*. *Journal of Microscopy* 100, 247–259.

GE, M., LIU, X., ZHAO, Z., SU, F., GU, L. & SU, D. (2022). Ensemble Machine-Learning-Based Analysis for In Situ Electron Diffraction. *Advanced Theory and Simulations* 5.

KALININ, S. V., MUKHERJEE, D., ROCCAPRIORE, K., BLAISZIK, B. J., GHOSH, A., ZIATDINOV, M. A., AL-NAJJAR, A., DOTY, C., AKERS, S., RAO, N. S., AGAR, J. C. & SPURGEON, S. R. (2023). Machine learning for automated experimentation in scanning transmission electron microscopy. *npj Computational Materials* 9, 227.

KALININ, S. V, OPHUS, C., VOYLES, P. M., ERNI, R., KEPAPTSOGLOU, D., GRILLO, V., LUPINI, A. R., OXLEY, M. P., SCHWENKER, E., CHAN, M. K. Y., ETHERIDGE, J., LI, X., HAN, G. G. D., ZIATDINOV, M., SHIBATA, N. & PENNYCOOK, S. J. (2022). Machine learning in scanning transmission electron microscopy. *Nature Reviews Methods Primers* 2.

KALININ, S. V, SPURGEON, S. R. & DRAVID, V. (2025). Self-Driving Microscopy for AI/ML-Enabled Physics Discovery and Materials Optimization. *The Bridge* 55.

KON, P. T. J., LIU, J., DING, Q., QIU, Y., YANG, Z., HUANG, Y., SRINIVASA, J., LEE, M., CHOWDHURY, M. & CHEN, A. (2025). Curie: Toward Rigorous and Automated Scientific Experimentation with AI Agents. https://arxiv.org/abs/2502.16069.

KRIVANEK, O. L., CHISHOLM, M. F., NICOLOSI, V., PENNYCOOK, T. J., CORBIN, G. J., DELLBY, N., MURFITT, M. F., OWN, C. S., SZILAGYI, Z. S., OXLEY, M. P., PANTELIDES, S. T. & PENNYCOOK, S. J. (2010). Atom-by-atom structural and chemical analysis by annular dark-field electron microscopy. *Nature* 464, 571–574.

LangChain (GitHub Repository) (2025). https://github.com/langchain-ai/langchain.

LI, G., HAMMOUD, H., ITANI, H., KHIZBULLIN, D. & GHANEM, B. (2023). CAMEL: Communicative Agents for 'Mind' Exploration of Large Language Model Society. In *Advances in Neural Information Processing Systems 36*. https://proceedings.neurips.cc/paper_files/paper/2023/hash/a3621ee907def47c1b952ade25c67698-Abstract-Conference.html.

LU, CHRIS, LU, CONG, LANGE, R. T., FOERSTER, J., CLUNE, J. & HA, D. (2024). The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery. https://arxiv.org/abs/2408.06292.

M. BRAN, A., COX, S., SCHILTER, O., BALDASSARI, C., WHITE, A. D. & SCHWALLER, P. (2024). Augmenting large language models with chemistry tools. *Nature Machine Intelligence* 6, 525–535.

MA, D., ZELTMANN, S. E., ZHANG, C., BARAISSOV, Z., SHAO, Y.-T., DUNCAN, C., MAXSON, J., EDELEN, A. & MULLER, D. A. (2025a). Emittance minimization for aberration correction II: Physics-informed Bayesian optimization of an electron microscope. *Ultramicroscopy* 273, 114138.

———— (2025b). Emittance minimization for aberration correction I: Aberration correction of an electron microscope without knowing the aberration coefficients. *Ultramicroscopy* 273, 114137.

MA, Y. J., LIANG, W., WANG, G., HUANG, D.-A., BASTANI, O., JAYARAMAN, D., ZHU, Y., FAN, L. & ANANDKUMAR, A. (2023). Eureka: Human-Level Reward Design via Coding Large Language Models. https://arxiv.org/abs/2310.12931.

MANDAL, I., SONI, J., ZAKI, M., SMEDSKJAER, M. M., WONDRACZEK, K., WONDRACZEK, L., GOSVAMI, N. N. & KRISHNAN, N. M. A. (2025). Autonomous Microscopy Experiments through Large Language Model Agents. https://arxiv.org/abs/2501.10385.

MEIROVITCH, Y., PARK, C. F., MI, L., POTOCEK, P., SAWMYA, S., LI, Y., WU, Y., SCHALEK, R., PFISTER, H., SCHOENMAKERS, R., PEEMEN, M., LICHTMAN, J. W., SAMUEL, A. D. T. & SHAVIT, N. (2023). SmartEM: machine-learning guided electron microscopy. *bioRxiv*. https://www.biorxiv.org/content/early/2023/10/08/2023.10.05.561103.

MITCHENER, L., YIU, A., CHANG, B., BOURDENX, M., NADOLSKI, T., SULOVARI, A., LANDSNESS, E. C., BARABASI, D. L., NARAYANAN, S., EVANS, N., REDDY, S., FOIANI, M., KAMAL, A., SHRIVER, L. P., CAO, F., WASSIE, A. T., LAURENT, J. M., MELVILLE-GREEN, E., CALDAS, M., BOU, A., ROBERTS, K. F., ZAGORAC, S., ORR, T. C., ORR, M. E., ZWEZDARYK, K. J., GHAREEB, A. E., MCCOY, L., GOMES, B., ASHLEY, E. A., DUFF, K. E., BUONASSISI, T., RAINFORTH, T., BATEMAN, R. J., SKARLINSKI, M., RODRIQUES, S. G., HINKS, M. M. & WHITE, A. D. (2025). Kosmos: An AI Scientist for Autonomous Discovery. https://arxiv.org/abs/2511.02824.

MULLER, D. A., KOURKOUTIS, L. F., MURFITT, M., SONG, J. H., HWANG, H. Y., SILCOX, J., DELLBY, N. & KRIVANEK, O. L. (2008). Atomic-Scale Chemical Imaging of Composition and Bonding by Aberration-Corrected Microscopy. *Science* 319, 1073–1076.

NAKAJIMA, Y. (2025). BabyAGI (Github Repository). https://github.com/yoheinakajima/babyagi.

OPHUS, C. (2019). Four-Dimensional Scanning Transmission Electron Microscopy (4D-STEM): From Scanning Nanodiffraction to Ptychography and Beyond. *Microscopy and Microanalysis* 25, 563–582.

PANAPITIYA, G., SALDANHA, E., JOB, H. & HESS, O. (2025). AutoLabs: Cognitive Multi-Agent Systems with Self-Correction for Autonomous Chemical Experimentation. https://arxiv.org/abs/2509.25651.

PATTISON, A. J., RIBET, S. M., NOACK, M. M., VARNAVIDES, G., PARK, K., KIRKLAND, E., PARK, J., OPHUS, C. & ERCIUS, P. (2024). BEACON – Automated Aberration Correction for Scanning Transmission Electron Microscopy using Bayesian Optimization. https://arxiv.org/abs/2410.14873.

PENNYCOOK, S. J. , N. P. D. (2011). *Scanning Transmission Electron Microscopy*. Pennycook, S. J. & Nellist, P. D. (Eds.). New York, NY: Springer New York.

PRATIUSH, U., DUSCHER, G. & KALININ, S. (2024). Human-in-the-loop interface for Automated experiments in Electron Microscopy, Automated characterization. In *AI for Accelerated Materials Design - NeurIPS 2024*. https://openreview.net/forum?id=AB4EuiXkAr.

PRATIUSH, U., HOUSTON, A., BARAKATI, K., RAGHAVAN, A., BULANADI, R., YIN, X., WELBORN, S. S., YOON, D., K P, H., BARAISSOV, Z., MA, D., JAKOWSKI, M., BARHORST, S.-P., PATTISON, A. J., MANGANARIS, P., MADUGULA, S. S., GAYATHRI AYYAGARI, S. V., KENNEDY, V., WANG, M., PANG, K. J., ADDISON-SMITH, I., MENACHO, W., GUZMAN, H. V, KIEFER, A., FURTH, N., KOLEV, N. L., PETROV, M., LIU, V., ILYEV, S., RAIRAO, S., RODANI, T., PINTO-HUGUET, I., CHEN, X., CRUAÑES, J., TORRENS, M., POMAR, J., SU, F., VEDANTI, P., LYU, Z., WANG, X., YAO, L., TAQIEDDIN, A., LASKOWSKI, F., SHAO, Y.-T., FEIN-ASHLEY, B., JIANG, Y., KUMAR, V., MISHRA, H., PAUL, Y., BAZGIR, A., PRANEETH MADUGULA, R. C., ZHANG, Y., OMPRAKASH, P., HUANG, JIAN, MONTUFAR-MORALES, E., CHAWLA, V., SETHI, H., HUANG, JIE, KURKI, L., GUINAN, G., SALVADOR, A., TER-PETROSYAN, A., VAN WINKLE, M., SPURGEON, S. R., NARASIMHA, G., WU, Z., LIU, R., LIU, Y., SLAUTIN, B., LUPINI, A. R., VASUDEVAN, R., DUSCHER, G. & KALININ, S. V (2025). Mic-hackathon 2024: hackathon on machine learning for electron and scanning probe microscopy. *Machine Learning: Science and Technology* 6, 040701.

PRATIUSH, U., HOUSTON, A., KALININ, S. V. & DUSCHER, G. (2024). Realizing smart scanning transmission electron microscopy using high performance computing. *Review of Scientific Instruments* 95.

PRATIUSH, U., ROCCAPRIORE, K. M., LIU, Y., DUSCHER, G., ZIATDINOV, M. & KALININ, S. V. (2024). Building Workflows for Interactive Human in the Loop Automated Experiment (hAE) in STEM-EELS.

PYCROSCOPY (2025). pyAutoMic: DigitalMicrograph TEM Scripts.

SAPKOTA, R., ROUMELIOTIS, K. I. & KARKEE, M. (2026). AI Agents vs. Agentic AI: A Conceptual taxonomy, applications and challenges. *Information Fusion* 126, 103599. https://linkinghub.elsevier.com/retrieve/pii/S1566253525006712.

SHI, C., CAO, M. C., REHN, S. M., BAE, S.-H., KIM, J., JONES, M. R., MULLER, D. A. & HAN, Y. (2022). Uncovering material deformations via machine learning combined with four-dimensional scanning transmission electron microscopy. *npj Computational Materials* 8, 114.

SHIBATA, N., FINDLAY, S. D., KOHNO, Y., SAWADA, H., KONDO, Y. & IKUHARA, Y. (2012). Differential phase-contrast microscopy at atomic resolution. *Nature Physics* 8, 611–615.

SONG, T., LUO, M., ZHANG, X., CHEN, L., HUANG, Y., CAO, J., ZHU, Q., LIU, D., ZHANG, B., ZOU, G., ZHANG, G., ZHANG, F., SHANG, W., FU, Y., JIANG, J. & LUO, Y. (2025). A Multiagent-Driven Robotic AI Chemist Enabling Autonomous Chemical Research On Demand. *Journal of the American Chemical Society* 147, 12534–12545. https://pubs.acs.org/doi/10.1021/jacs.4c17738.

SPURGEON, S. R., OPHUS, C., JONES, L., PETFORD-LONG, A., KALININ, S. V., OLSZTA, M. J., DUNIN-BORKOWSKI, R. E., SALMON, N., HATTAR, K., YANG, W.-C. D., SHARMA, R., DU, Y., CHIARAMONTI, A., ZHENG, H., BUCK, E. C., KOVARIK, L., PENN, R. L., LI, D., ZHANG, X., MURAYAMA, M. & TAHERI, M. L. (2021). Towards data-driven next-generation transmission electron microscopy. *Nature Materials* 20, 274–279.

TANG, J., XIA, L., LI, Z. & HUANG, C. (2025). AI-Researcher: Autonomous Scientific Innovation. https://arxiv.org/abs/2505.18705.

UDRESCU, S.-M. & TEGMARK, M. (2020). AI Feynman: A physics-inspired method for symbolic regression. *Science Advances* 6, eaay2631. https://www.science.org/doi/10.1126/sciadv.aay2631.

UTKARSHP1161 (2024). utkarshp1161/Active-learning-in-microscopy: v1.0.0. https://doi.org/10.5281/zenodo.14562670.

WELBORN, S. S., HARRIS, C., RIBET, S. M., VARNAVIDES, G., OPHUS, C., ENDERS, B. & ERCIUS, P. (2024). Streaming Large-Scale Microscopy Data to a Supercomputing Facility. *Microscopy and Microanalysis* 31, ozae109. https://doi.org/10.1093/mam/ozae109.

WILLIAMS, D. B. & CARTER, C. B. (1996). The Transmission Electron Microscope. In *Transmission Electron Microscopy*, pp. 3–17. Boston, MA: Springer US.

WU, Q., BANSAL, G., ZHANG, J., WU, Y., LI, B., ZHU, E., JIANG, L., ZHANG, X., ZHANG, S., LIU, J., AWADALLAH, A. H., WHITE, R. W., BURGER, D. & WANG, C. (2023). AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation. https://arxiv.org/abs/2308.08155.

YAMADA, Y., LANGE, R. T., LU, CONG, HU, S., LU, CHRIS, FOERSTER, J., CLUNE, J. & HA, D. (2025). The AI Scientist-v2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search. https://arxiv.org/abs/2504.08066.

YANG, H., YUE, S. & HE, Y. (2023). Auto-GPT for Online Decision Making: Benchmarks and Additional Opinions. https://arxiv.org/abs/2306.02224.

YAO, L., SAMANTRAY, S., GHOSH, A., ROCCAPRIORE, K., KOVARIK, L., ALLEC, S. & ZIATDINOV, M. (2025). Operationalizing Serendipity: Multi-Agent AI Workflows for Enhanced Materials Characterization with Theory-in-the-Loop. https://arxiv.org/abs/2508.06569.

ZHANG, H., SONG, Y., HOU, Z., MIRET, S. & LIU, B. (2024). HoneyComb: A Flexible LLM-Based Agent System for Materials Science. https://arxiv.org/abs/2409.00135.

ZHU, L., TANG, J., LI, B., HOU, T., ZHU, Y., ZHOU, J., WANG, Z., ZHU, X., YAO, Z., CUI, X., WATANABE, K., TANIGUCHI, T., LI, Y., HAN, Z. V., ZHOU, W., HUANG, Y., LIU, Z., HONE, J. C. & HAO, Y. (2022). Artificial Neuron Networks Enabled Identification and Characterizations of 2D Materials and van der Waals Heterostructures. *ACS Nano* 16, 2721–2729.

ZIATDINOV, M., GHOSH, A. & KALININ, S. V. (2021). Physics makes the difference: Bayesian optimization and active learning via augmented Gaussian process.

ZIATDINOV, M., LIU, Y., MOROZOVSKA, A. N., ELISEEV, E. A., ZHANG, X., TAKEUCHI, I. & KALININ, S. V. (2021). Hypothesis Learning in Automated Experiment: Application to Combinatorial Materials Libraries.