

Dynamic Processes in Chemical-Vapor Deposition of Transition-Metal Dichalcogenides Revealed by AI-Driven Operando Metrology

Mingwei Feng^{1,2} Xilu Zou³ Haotian Zhang⁴ Caiqi Zou⁴ Silin Chen⁴ Wei Xu⁵ Xiaotian Zhang⁶
Taotao Li^{4,7} Ningmu Zou^{4,7,*}

Abstract— We present an operando metrology framework for AI-enabled high-throughput experimentation and computation in the synthesis of two-dimensional transition-metal dichalcogenides (TMDs). The platform enables operando monitoring of TMD growth in a CVD tube furnace using high-temperature optical microscopy through a reactor viewport. During growth at 900–1000 °C under atmospheric-pressure argon flow, video streams capture crystal nucleation, edge propagation, and domain coalescence in real time. An AI-driven image analysis pipeline processes these data through frame stabilization and region-of-interest detection, precise boundary extraction using an attention-enhanced edge-guided segmentation network (Aegis-Net), spatially resolved growth-velocity mapping, and disentanglement of overlapping domain boundaries. This workflow converts raw operando videos into quantitative growth-kinetics heatmaps with minimal human intervention, supporting scalable, high-throughput experimental analysis. To connect experiment with computation, the extracted growth maps are compared with CFD simulations of the CVD reactor, incorporating compressible flow, heat and mass transport, multiphase VOF modeling, and UDF-based gas–solid reaction kinetics. The simulations reproduce oscillatory growth-ring features consistent with concentric deposition fronts observed experimentally, and the extracted ring periodicity provides a quantitative feedback variable for tuning precursor placement and carrier-gas flow. By coupling AI-based metrology with physics-based high-throughput modeling, this framework reveals otherwise inaccessible spatial–temporal growth dynamics and provides a pathway toward closed-loop optimization, wafer-scale monitoring, and autonomous control of 2D TMD synthesis.

This work was supported by the Natural Science Foundation of China under Grant 62341408 and the Fundamental Research Funds for the Central Universities 2024300421.

¹ College of Engineering and Applied Sciences, Nanjing University, Nanjing 210023, China

² Key Laboratory of Intelligent Optical Sensing and Manipulation, Ministry of Education, Nanjing University, Nanjing 210093, China

³ National Laboratory of Solid-State Microstructures, School of Electronic Science and Engineering and Collaborative Innovation Center of Advanced Microstructures, Nanjing University, Nanjing 210023, China

⁴ School of Integrated Circuits, Nanjing University, Suzhou 215163, China

⁵ School of Materials Science and Engineering, Southeast University, Nanjing 211189, China

⁶ Suzhou Laboratory, Suzhou 215123, China

⁷ Interdisciplinary Research Center for Future Intelligent Chips (Chip-X), Nanjing University, Suzhou 215163, China

*Corresponding author: nzou@nju.edu.cn

I. INTRODUCTION

High-quality two-dimensional TMD films (e.g. MoS₂, WS₂) are critical for next-generation electronics, photonics and catalysts. Chemical vapor deposition (CVD) is the leading method for scalable 2D TMD growth [1], but its complex chemistry and transport make process control challenging. In particular, the interplay of vapor-phase diffusion, nucleation, and edge growth leads to rich spatiotemporal dynamics (e.g. varying edge speeds, dendritic patterns). Recent advances emphasize the need for operando imaging to capture these dynamics in real time [2]. However, automated analysis of such videos is non-trivial, since TMD domains have low contrast and fuzzy boundaries, and they often merge as they grow. Another major obstacle is the difficulty of monitoring crystal growth as it happens. Conventional experiments typically rely on post-growth characterization (e.g. ex situ microscopy), which provides only “before and after” snapshots and misses critical dynamic information about how nuclei form and domains expand [1][3]. Real-time, operando observation of the growth process is crucial for developing a true mechanistic understanding. Recent advances have begun to address this gap. For example, miniaturized CVD reactors with optical windows have enabled direct video monitoring of 2D crystal formation [4]. In one study, operando microscopy combined with machine learning was used to identify growth-governing factors and even predict conditions for large crystal synthesis [5][6].

To address this, we develop an AI-driven metrology pipeline that preprocesses image data, segments domain edges with sub-pixel accuracy, maps growth rates, and disentangles merged domains. This enables quantification of growth kinetics across entire flakes. Coupling these measurements with fluid-dynamics simulations of the reactor provides physical interpretation, including gas-flow and precursor-concentration fields that give rise to ring-like deposition fronts. The same comparison also establishes a quantitative feedback loop: deviations between simulated and measured ring periodicity or local edge velocity can be used to adjust sulfur/MoO₃ source positions, carrier-gas flow rate, or heating profiles in a future autonomous CVD controller. Together, these methods reveal how furnace conditions translate into local growth behavior, helping to optimize 2D TMD synthesis.




Fig. 1. Operando CVD platform and image-preprocessing workflow for dynamic TMD growth analysis. (a) Schematic of the CVD tube furnace setup for operando MoS₂ growth, illustrating the two-zone arrangement with MoO₃ and S precursors, the heated substrate, and the optical monitoring system. The camera views the substrate through a quartz window in real time while growth occurs. (b) Representative optical image frame captured during the CVD growth (top view of substrate). The triangular MoS₂ domains appear with distinct contrast against the SiO₂ background (lighter regions here indicate growing MoS₂ islands). The operando imaging captures the nucleation of many small domains and their expansion over time on the substrate. (c) Image-preprocessing workflow, including drift correction based on impurity-point registration and automatic localization and cropping of the target domain center using an improved Swin-Transformer detector with Shape-IoU. (d) Schematic illustration of the lateral growth boundary and the definition of the local epitaxial growth rate extracted from the time-resolved domain contours.

II. METHODS

A. Experimental Setup

Our CVD reactor is a standard horizontal tube furnace with a quartz tube (≈ 50 mm inner diameter) placed inside a resistive heating zone. Solid precursors are placed in alumina boats: molybdenum oxide (or Mo-containing substrate) at the downstream, and sulfur powder upstream at a fixed distance. Pure argon (Ar) is flowed through the tube at controlled sccm rates to carry sublimated species to the substrate. The furnace is heated to ~ 950 °C at the growth zone (± 50 °C) to volatilize S and Mo precursors, initiating MoS₂ nucleation on the Si/SiO₂ substrate. Mass flow controllers ensure constant gas flow (typically tens of sccm) and an inert environment. The reactor is operated at ambient pressure; other researchers have demonstrated both atmospheric and low-pressure CVD for TMDs, with different flow regimes.

Figure 1 illustrates the experimental apparatus and preprocessing steps. The quartz tube is imaged through a viewport window by a long-working-distance optical microscope. A CCD camera records a video of the growth area (bright-field illumination). Raw frames often show slight defocus and frame-to-frame misalignment due to thermal drift and stage jitter. The preprocessing module automatically adjusts focus (via digital filtering or deconvolution) and stabilizes the video. We employ a feature-free, optical-flow based registration (e.g. the RAFT deep network) to compute sub-pixel shifts between consecutive frames, effectively removing lateral drift and jitter. This ensures that the apparent growth of domains is due to crystal expansion only, not

camera motion. After stabilization, individual frames are extracted at regular time intervals (e.g. 1 frame/s). The pipeline then detects the domain region of interest (ROI) by simple intensity thresholding or edge localization in each frame. We crop around the ROI so that further analysis focuses on the active growth area (Figure 1). This preprocessing yields a time series of aligned, in-focus images containing the TMD flakes.

B. Data acquisition and preprocessing

The first module reads the raw microscopy video and outputs corrected frames. Sub-pixel registration is performed using the RAFT optical flow method [3][4], which produces near-perfect alignment of frames even without fiducial markers. Specifically, RAFT computes dense displacement fields and applies the median offset to each frame, eliminating stage drift and random jitter [6]. This step is crucial: without stabilization, domain edges would blur over time and growth-rate estimates would be inaccurate. After alignment, we optionally apply a defocus correction (e.g. deconvolution using known point-spread functions) if frames suffer mild focus drift. Next, frames are sampled at one-second intervals to balance temporal resolution with computation. For each frame, we automatically identify the domain ROI by detecting contiguous regions above a brightness threshold (the growing TMD flake) [7]. This ROI is then cropped and passed to the next module. Figure 1(d) schematically shows the result: a series of clean, aligned images focused on a growing triangular domain.




Fig. 2. Single-frame domain-boundary segmentation using Aegis-Net with MFAM and DFUAM. (a) Schematic illustration of the AI-driven operando metrology pipeline. The four modules (preprocessing, segmentation, kinetics mapping, and domain disentanglement) are shown as stages with example intermediate outputs. Raw video frames (left) are stabilized and filtered, then fed into the Aegis-Net segmentation network. The network outputs binary masks highlighting MoS₂ domains in each frame. These masks over time are compiled into a kinetics heatmap that encodes the local growth timing. Finally, overlapping domains are separated by analyzing their chronological boundaries. (b) Architecture of the Aegis-Net segmentation model used in this work. It consists of a U-Net-style encoder-decoder CNN equipped with multi-scale feature aggregation modules (MFAM) at multiple feature scales. MFAM integrates fine boundary texture with larger-scale domain-shape cues, enabling the network to detect both small early nuclei and large later domains with consistent fidelity. The decoder incorporates a dynamic feature up sampling attention module (DFUAM), the module labeled in the architecture diagram, to refine up sampled features and sharpen boundary details. This design allows reliable identification of TMD domain regions from noisy operando images, including overlapping or touching domains.

C. ML-based instance segmentation

Each ROI image is then segmented to find the precise domain boundary. Conventional methods like Otsu thresholding or Canny edge detection often fail here because the flake’s edges are partially blurred and the contrast is low [8]. Instead, we adopt the recent Aegis-Net model, which is tailored for microscopic image segmentation. In this work, Aegis-Net incorporates multi-scale feature aggregation modules (MFAM) to learn boundary and morphology cues across feature scales, allowing it to adapt to the varying sizes and shapes of domains [2][7]. The decoder uses a dynamic feature up sampling attention module (DFUAM) to restore high-resolution edge information during up sampling and to sharpen fuzzy or partially overlapping boundaries. In practice, we fine-tune Aegis-Net on a small set of labeled TMD images; the network then predicts a pixel-wise mask of the flake. The Swin-Transformer ROI detector and the MFAM–DFUAM segmentation backbone are also compatible with transfer learning across related TMD systems. For WSe₂ or WS₂, the detector can typically be reused after limited calibration of contrast and domain-size statistics, while Aegis-Net requires only a small number of additional annotated frames to fine-tune the final decoder and normalization layers rather than full retraining. Figure 2 shows the Aegis-Net architecture and the roles of MFAM and DFUAM. We find that Aegis-Net avoids the spurious gaps or extra blobs that plague simple thresholding; it yields crisp, connected boundaries around each triangular or hexagonal crystal. This precise edge localization is critical for downstream kinetics analysis.

D. Growth kinetics heatmaps

From the segmented frames, we extract the domain boundary in each image as a set of (x,y) coordinates. By comparing successive frames, we compute the local displacement of the edge over each time step. In a simple implementation, we project growth onto the normal direction of the contour, yielding a pixel-wise velocity (growth rate) map. By accumulating these velocity values, we construct a two-dimensional growth kinetics heatmap for each domain, as shown in Figure 3.




Fig. 3. Reconstruction of growth-kinetics heatmaps and multimodal validation of oscillatory growth behavior. (a) Frame-by-frame evolution of domain boundaries reconstructed from operando microscopy, together with finite-element simulation results showing good agreement with the experimentally extracted growth behavior. (b) Quantitative analysis of the

oscillatory lateral growth rate derived from the tracked domain contours. (c) Grain boundary position information obtained by merging domain boundaries and evolving frame by frame temporal sequence. (d) Photoluminescence (PL) image showing a ring-like structure that serves as independent evidence of oscillatory growth. (e) Scanning transmission electron microscopy (STEM) characterization of grain-boundary-related structural features. (f) Second-harmonic generation (SHG) imaging of grain boundaries, further validating the domain evolution and interface structures revealed by the AI-based analysis.

In this heatmap, each pixel’s color encodes the cumulative growth of that location (red = fast growth, blue = slow). For triangular or polygonal domains, this reveals patterns: e.g. edges that advance more rapidly produce bands of higher velocity. Such spatial rate maps provide insights into anisotropic growth (e.g. zigzag vs armchair edges) and diffusion effects. This quantitative kinetics mapping builds directly on prior studies of 2D material growth; for example, Dong et al. [3] review how attachment-limited vs diffusion-limited regimes control overall growth rates. Our heatmaps give a spatially resolved view of these kinetics, essentially measuring how local adatom attachment varies around the flake. In controlled experiments, we can correlate features in the heatmap (rings, fronts) with process parameters (temperature, gas flow).

E. Domain boundary disentanglement

As the deposition proceeds, multiple nucleation sites often grow and merge into polycrystalline films. When two triangular domains coalesce, their boundaries intersect and simple segmentation yields a single merged mask. To separately analyze each grain, we perform boundary disentanglement: we identify junction points where edges meet and split the contour graph into individual loops. Practically, we skeletonize the combined mask and apply a watershed-like split at concave corners to isolate grain boundaries. Each resulting loop is then treated as an independent domain. This allows computation of separate growth metrics for each crystal, even after coalescence. Figure 4 illustrates this: the left panel shows two merged domains with crossed edges, and the right panel shows the extracted individual boundaries. By tracking each grain separately, we can measure competitive growth (e.g. which grain dominates) and measure local strain boundaries.




Fig. 4. Grain-boundary extraction and domain disentanglement in coalesced crystal domains. (a) Comparison of grain-boundary features observed by operando microscopy, high-magnification microscopy, SHG imaging, and AI-based image segmentation. (b) Schematic illustration of domain coalescence and overlap-error evaluation for boundary separation, quantified

using the Hausdorff distance. (c) Statistical analysis of the average reconstruction error as a function of propagation distance, demonstrating the reliability of the proposed grain-boundary disentanglement method for merged domains.

F. Finite-element simulation of the CVD reactor

To interpret the observed growth patterns, we perform CFD simulations of the tube reactor. The simulation solves the compressible Navier–Stokes equations coupled with heat transfer and species transport. Specifically, we include continuity, momentum, and energy equations for the gas phase, along with convective–diffusive transport for the precursor vapors [9]. The solid precursors (sulfur, MoO_3) are represented via boundary conditions: for example, a source term (via a UDF) models the sublimation and reaction of sulfur at the solid surface. The flow is assumed laminar ($\text{Re} \ll 1000$) given the tube dimensions and low gas flow rates. The simulation mesh and pre-processing are done in Gambit, and the equations are solved in ANSYS Fluent [4]. We activate a Volume-of-Fluid (VOF) multiphase model to allow for any condensed-phase phenomena (e.g. liquid salt droplets), and use Fluent’s species transport for multi-component diffusion. User-Defined Functions (UDFs) are compiled into Fluent to include heterogeneous surface reactions: e.g., the $\text{MoO}_3 + \text{S} \rightarrow \text{MoS}_2$ deposition kinetics are implemented so that local deposition consumes vapor species at the substrate. Similar coupled CFD–DEM or CFD–CVD approaches have been used in particle-coating models [10][11], and here the Fluent UDF framework allows an analogous implementation in a fluid-phase reactor. [12]

Simulation results reproduce key features of the growth environment. Temperature and precursor concentration profiles match the expected gradients, and gas flow streamlines show how carrier gas carries vapor from the sulfur boat downstream to the substrate. Importantly, transient simulation captures oscillatory deposition fronts: as the heating ramp and flow change, rings of high deposition rate form and propagate from the central nucleation outward. These “growth rings” arise from periodic fluctuations in gas-phase precursor supply. The quantitative comparison is based on ring periodicity, defined as the average peak-to-peak interval in time-resolved deposition rate or the corresponding inter-ring spacing in the growth-time heatmap after spatial normalization. Thus, the reported agreement within 10% refers to ring periodicity, not absolute optical/PL intensity magnitude or exact pixel-level spatial position. In PL studies of TMD growth (e.g. WS₂ on sapphire), similar concentric halos were observed around salt particles [13][14], suggesting that our modeled mechanism is plausible, as shown in Figure 3(d). In summary, the CFD model corroborates that variations in flow and sublimation lead to spatial non-uniformity (e.g. rings) that the metrology pipeline quantitatively maps.

III. RESULTS

Applying the AI-driven pipeline to operando growth videos of MoS_2 revealed rich dynamic information that is inaccessible by traditional geometry endpoint characterization. Figure 1 shows the raw experimental setup and a sample video frame before/after preprocessing. The Aegis-Net segmentation as shown in Figure 2 accurately traces the

evolving triangular domain (white mask). From this, Figure 3 displays the computed growth kinetics heatmap: brighter bands at the triangular corners in PL map, which may indicate that edge advancement was fastest at those vertices. We observed that one vertex consistently lags (seen as a cooler color) due to local gas depletion effects. Furthermore, the Aegis-Net successfully detects nascent domains just a few pixels in size (tens of μm in reality) and continues to segment them as they grow.

A heatmap of growth dynamics was constructed based on this algorithm from frame-by-frame segmented data. The finite element simulation analysis in Figure 3(a) depicts a growth time map (heatmap) of a color ring that approximates the experimentally observed one. Each pixel’s color corresponds to the time at which that location became covered by MoS_2 . Several key observations can be made. First, nucleation is not spatially uniform: certain “hotspots” on the substrate nucleated earlier and produced larger domains, whereas other areas lagged. This non-uniformity may be caused by local variations in substrate condition or precursor concentration. Second, the heatmap reveals synchronized oscillatory growth behavior across the sample. Concentric bands of color implying periodic spurts of domain expansion can be seen (for instance, alternating rings of green and yellow in Figure 3(a)). This suggests that the growth proceeded in bursts – possibly due to cyclic depletion and replenishment of sulfur vapor in the furnace. To examine this, we extracted the area-vs-time curve for a representative MoS_2 domain (Figure 3(b)). The curve shows a stepwise increase in area: rapid growth phases followed by brief plateaus, in a repeating pattern. The period of these oscillations (~ 3.5 s) roughly matches the sulfur vapor replenishment cycle in our tube (as the S powder evaporates in pulses) [10]. Such oscillatory kinetics are difficult to detect without real-time monitoring and dense time sampling; here the AI pipeline makes it straightforward to quantify them. These results highlight how feedback between precursor supply and growth front dynamics can lead to periodic growth rates, an insight that could inform better control of precursor delivery in scaled-up processes.




Fig. 5. Substrate-dependent oscillatory growth kinetics on sapphire and silicon. Comparison of oscillatory lateral growth rates measured on different

substrates: (a) sapphire and (b) silicon. The shorter oscillation interval observed on sapphire is attributed to its higher interfacial binding-energy barrier and shorter surface diffusion length, indicating a pronounced substrate effect on the dynamic growth kinetics.

Comparing with simulations, Figure 3 summarizes representative phenomena: for instance, when the sulfur boat was moved closer to the substrate experimentally, both simulation and imaging show an increase in central growth rate and tighter ring spacing. The spatial profile of simulated deposition (heatmap from Fluent) matches the experiments within 10% in ring periodicity, calculated as the average relative error between corresponding adjacent-ring spacings after scale normalization. This metric does not claim a 10% absolute match in PL/optical intensity; rather, it validates the time-and-space periodicity of the oscillatory deposition fronts. We attribute this agreement to periodic changes in local precursor concentration due to intermittent mass transport bottlenecks [15]. In all cases, our operando measurements agree with the CFD prediction that the hottest growth regions occur where gas flow impinges on fresh edge sites. Notably, the MFAM-DFUAM-based Aegis-Net segmentation outperformed manual thresholding: classic binarization either missed faint edge segments or created spurious outlines in low-contrast regions. By contrast, Aegis-Net’s learned features provide continuous, smooth contours even under uneven illumination.

Initially, tens of small triangular MoS₂ nuclei appear randomly on the substrate surface. These nuclei subsequently expand laterally, with some coalescing into larger domains. By the end of growth, a significant fraction of the substrate is covered by monolayer islands, though gaps remain between them. After 40 min of growth, two domains merged; applying domain-boundary disentanglement produced the split boundaries shown in Figure 4. Quantitative analysis shows that the larger initial domain dominated growth, consistent with coalescence behavior known in polycrystalline films [1][2]. As growth nears completion, numerous MoS₂ domains begin to coalesce, forming a polycrystalline thin film. Identifying the final grain boundaries is crucial for decoupling the kinetic laws governing this coalescence process, particularly since grain boundaries significantly influence the material’s electrical properties. Our grain disentanglement module addresses this challenge by using optical-flow tracking to follow the relative positions of grain-boundary vertices across consecutive video frames. Figure 4(a) presents an optical micrograph of a substrate region after CVD growth, featuring an assembly of multiple triangular monolayer domains. Without additional analytical processing, it is difficult to accurately reconstruct the spatiotemporal evolution of the grain boundaries across consecutive frames. Figure 4(a) also displays the output of our disentanglement algorithm for this region: each distinct grain is pseudo-colored, thereby revealing where MoS₂ domains form single or double grain boundaries during coalescence. By tracking the contours of the domains forward across the video sequence and progressively extending outward to connect the vertices, the algorithm reconstructs the grain-boundary positions with sub-pixel precision. The reconstructed grain boundaries (yellow lines in Figure 4(a)) align with subtle contrast variations or linear defects visible in the optical image,

providing physical validation of the reconstruction accuracy. In Figure 3(c), we separately decouple and plot the epitaxial growth-rate evolution for each individual grain. Despite their eventual coalescence, these grains exhibit distinct growth kinetics: for instance, Grain 1 initially grew rapidly before slowing down, whereas Grain 2 nucleated later but ultimately attained a comparable final size. Information regarding the growth history of each individual grain is valuable for correlating specific grain characteristics, such as crystallographic orientation or defect density, with data obtained through ex-situ characterization techniques. In summary, our AI pipeline not only tracks the overall coverage area but also preserves domain-specific data within the polycrystalline environment, a capability essential for understanding coalescence and polycrystalline film formation [6].




Fig. 6. Representative growth phenomena revealed by AI-driven operando metrology. (a) Observation of multiple grain boundaries, including double and higher-order boundary configurations, identified from the reconstructed domain evolution. (b) Stacking behavior associated with multilayer growth. (c) Comparison of solid-state and liquid-mediated lateral growth rates, where liquid droplets are tracked using a Kalman-filter-based algorithm. (d) Reduced transverse growth rate observed in WSe₂, revealing material-dependent kinetic behavior captured by the proposed metrology framework.

Looking at the broader implications, the integration of AI with operando metrology opens new avenues for closed-loop control of 2D material synthesis. The measured ring periodicity, local edge velocity, and domain-coverage maps can serve as real-time feedback variables for an autonomous controller. For example, if the measured oscillation period is shorter than the target or the heatmap shows excessive central deposition, the controller could increase the source–substrate distance, reduce sulfur-zone temperature, or adjust Ar carrier-gas flow to moderate precursor delivery. Conversely, if edge growth stalls or coverage becomes non-uniform, the system could move the precursor boat closer, increase source temperature, or tune the flow profile to restore a more uniform vapor supply. The CFD model provides the forward map linking precursor placement and carrier-gas flow to concentration fields, while the AI metrology provides the experimental error signal needed for model-predictive or reinforcement-learning-based optimization.

Figure 5 further shows that oscillatory intervals differ on sapphire and silicon. This substrate dependence can be understood through differences in interfacial binding-energy barriers and surface diffusion lengths: stronger binding and

shorter diffusion length on sapphire confine adatom migration, producing shorter growth intervals and sharper local changes in edge velocity. These physical differences do not merely change the apparent optical contrast; they alter the temporal boundary displacement that the AI must resolve. The MFAM module helps maintain edge-detection precision by aggregating weak texture cues and domain-scale shape information, while DFUAM restores high-resolution boundary details during up sampling. This is particularly relevant for unidirectional epitaxy, where small substrate-induced differences in edge velocity and crystallographic alignment must be detected early enough to distinguish aligned domains from competing orientations.

The same pipeline can be adapted to 150-mm wafer-scale monitoring and industrial MOCVD reactors by combining multi-camera or line-scan optical acquisition with tiled ROI analysis. Since the Aegis-Net inference step is local to each ROI, wafer-scale operation can be implemented by stitching per-tile heatmaps into a full-wafer growth map and flagging abnormal regions such as stalled edges, dense secondary nucleation, or early grain-boundary formation. For materials beyond MoS₂, including WSe₂ and WS₂ as illustrated by Figure 6, the framework is expected to require limited transfer learning rather than complete redesign: the stabilization, ROI detection, boundary-disentanglement, and heatmap modules remain unchanged, while a small material-specific annotated set can recalibrate the segmentation decoder for different optical contrasts, edge fuzziness, or liquid-mediated growth modes.

IV. CONCLUSIONS

In this paper, we propose an AI-powered operando metrology framework to monitor and analyze the CVD operando growth of 2D MoS₂ in real time. By combining operando optical imaging with tailored deep-learning segmentation (Aegis-Net with MFAM and DFUAM) and custom data analysis modules, we obtained detailed insight into the dynamic processes of TMD crystal growth. The system automatically tracks the nucleation and expansion of dozens of domains, quantifies local growth rates across the substrate, and separates overlapping grains to preserve individual growth histories. Applied to MoS₂, the pipeline revealed non-uniform and oscillatory kinetics, which we traced to precursor supply dynamics and confirmed via reactor simulations. The results illustrate how AI-driven analysis can extract detailed kinetic parameters (growth rates, domain coalescence timing, ring periodicity, etc.) that were formerly inaccessible without laborious manual analysis or specialized equipment.

Our findings pave the way for intelligent control of 2D material fabrication. With real-time feedback, a CVD or MOCVD system could adjust parameters during growth, including precursor-source placement, source temperature, carrier-gas flow, and zone heating, to suppress oscillatory supply fluctuations and equalize lateral growth over the substrate. Such closed-loop control, informed by AI metrology and CFD modeling, could substantially improve film uniformity and quality, addressing a key challenge in scalable 2D manufacturing. Future work will focus on

integrating the pipeline with actuator controls for autonomous growth optimization, extending tiled imaging to wafer-scale substrates, and applying transfer learning to additional material systems such as WS₂ and WSe₂. Incorporating additional sensor modalities, such as spectroscopic signals, could provide a more comprehensive picture of the growth chemistry. We believe that the combination of operando data and AI analytics represents a powerful paradigm for studying and controlling nanoscale synthesis processes, moving the field toward the level of process insight and precision necessary for transitioning 2D materials from the laboratory to industry.

V. ACKNOWLEDGMENT

This work is supported by the Natural Science Foundation of China under Grant 62341408 and the Fundamental Research Funds for the Central Universities 2024300421.

REFERENCES

- [1] H. Wang et al., "In situ imaging of two-dimensional crystal growth using a heat-resistant optical microscope," *Nano Lett.*, vol. 24, no. 18, pp. 5498–5505, 2024, doi: 10.1021/acs.nanolett.4c00620.
- [2] X. Zhu, H. Wang, K. Wang, and L. Xie, "Progress on the in situ imaging of growth dynamics of two-dimensional materials," *Nanoscale*, vol. 15, pp. 11746–11758, 2023, doi: 10.1039/D3NR01475D.
- [3] J. Dong, D. Ding, C. Jin, Y. Liu, and F. Ding, "Edge reconstruction-dependent growth kinetics of MoS₂," *ACS Nano*, vol. 17, no. 1, pp. 127–136, 2023, doi: 10.1021/acs.nano.2c05397.
- [4] T. Kang et al., "Strategies for controlled growth of transition metal dichalcogenides by chemical vapor deposition for integrated electronics," *ACS Mater. Au*, vol. 2, no. 6, pp. 665–685, 2022, doi: 10.1021/acsmaterialsau.2c00029.
- [5] G. Xue et al., "Large-area epitaxial growth of transition metal dichalcogenides," *Chem. Rev.*, vol. 124, no. 17, pp. 9785–9865, 2024, doi: 10.1021/acs.chemrev.3c00851.
- [6] M. Xu et al., "Reconfiguring nucleation for CVD growth of twisted bilayer MoS₂ with a wide range of twist angles," *Nat. Commun.*, vol. 15, Art. no. 562, 2024, doi: 10.1038/s41467-023-44598-w.
- [7] L. Li et al., "Epitaxy of wafer-scale single-crystal MoS₂ monolayer via buffer layer control," *Nat. Commun.*, vol. 15, Art. no. 1825, 2024, doi: 10.1038/s41467-024-46170-6.
- [8] J. Chen et al., "TransUNet: Rethinking the U-Net architecture design for medical image segmentation through the lens of transformers," *Med. Image Anal.*, vol. 97, Art. no. 103280, 2024, doi: 10.1016/j.media.2024.103280.
- [9] P. Zheng et al., "Universal epitaxy of non-centrosymmetric two-dimensional single-crystal metal dichalcogenides," *Nat. Commun.*, vol. 14, Art. no. 592, 2023, doi: 10.1038/s41467-023-36286-6.
- [10] D. Moon et al., "Hypotaxy of wafer-scale single-crystal transition metal dichalcogenides," *Nature*, vol. 638, no. 8052, pp. 957–964, 2025, doi: 10.1038/s41586-024-08492-9.
- [11] H. Liu et al., "Boosting monolayer transition metal dichalcogenides growth by hydrogen-free ramping during chemical vapor deposition," *Nano Lett.*, vol. 24, no. 27, pp. 8277–8286, 2024, doi: 10.1021/acs.nanolett.4c01314.
- [12] M. C. Lemme, D. Akinwande, C. Huyghebaert, and C. Stampfer, "2D materials for future heterogeneous electronics," *Nat. Commun.*, vol. 13, Art. no. 1392, 2022, doi: 10.1038/s41467-022-29001-4.
- [13] K. P. O'Brien et al., "Process integration and future outlook of 2D transistors," *Nat. Commun.*, vol. 14, Art. no. 6400, 2023, doi: 10.1038/s41467-023-41779-5.
- [14] Y.-H. Lee et al., "Synthesis of large-area MoS₂ atomic layers with chemical vapor deposition," *Adv. Mater.*, vol. 24, no. 17, pp. 2320–2325, 2012, doi: 10.1002/adma.201104798.
- [15] J. Zhang et al., "Toward controlled synthesis of 2D crystals by CVD: Learning from the real-time crystal morphology evolutions," *Nano Lett.*, vol. 24, no. 8, pp. 2465–2472, 2024, doi: 10.1021/acs.nanolett.3c04016.