CRYSTAL GROWTH & DESIGN

pubs.acs.org/crystal

Article

# Monitoring MBE Substrate Deoxidation via RHEED Image-Sequence Analysis by Deep Learning

Abdourahman Khaireh-Walieh, Alexandre Arnoult, Sébastien Plissard, and Peter R. Wiecha*

Cite This: https://doi.org/10.1021/acs.cgd.2c01132

Read Online

ACCESS | Metrics & More | Article Recommendations

**ABSTRACT:** Reflection high-energy electron diffraction (RHEED) is a powerful tool in molecular beam epitaxy (MBE), but RHEED images are often difficult to interpret, requiring experienced operators. We present an approach for automated surveillance of GaAs substrate deoxidation in MBE reactors using deep-learning-based RHEED image-sequence classification. Our approach consists of an nonsupervised autoencoder (AE) for feature extraction, combined with a supervised convolutional classifier network. We demonstrate that our lightweight network model can accurately identify the exact deoxidation moment. Furthermore, we show that the approach is very robust and allows accurate deoxidation detection for months without requiring retraining. The main advantage of the approach is that it can be applied to raw RHEED images without requiring further information such as the rotation angle, temperature, etc.

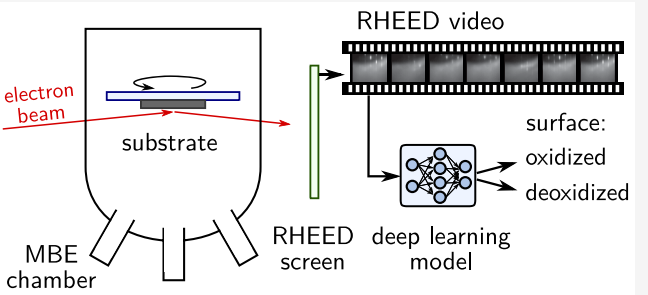




## 1. INTRODUCTION

Reflection high-energy electron diffraction (RHEED) is a widely used in situ control method in molecular beam epitaxy (MBE).[1−4] RHEED diffraction patterns provide information about the crystal surface with atomic resolution, and as the ultrahigh vacuum in typical growth chambers allows an easy integration of electron beam systems in MBEs, RHEED has become a standard in situ characterization instrument in MBE, enabling unprecedented accuracy in monitoring the crystal growth. RHEED is highly sensitive to several key MBE parameters such as the growth rate, the crystal structure, the lattice parameter and strain, etc.[5−9] However, RHEED images can be difficult to interpret, since the diffraction patterns produce information in the Fourier space. Furthermore, the actual recorded patterns are very sensitive to calibration, and often also dynamic variations in the patterns over several time scales contain valuable information, rendering their analysis even more challenging. Real-time exploitation of RHEED data is therefore often limited to easily accessible information such as the deposition rate. Sophisticated analysis is usually done a posteriori, on recorded RHEED images or videos. Due to the complexity of the task, RHEED interpretation usually requires experienced operators, possessing years of machine-specific training.

A common application of RHEED is the monitoring of the native oxide removal from commercial substrates prior to crystal growth. Surface oxidation of a few nanometers due to exposure to oxygen is unavoidable during transport of epitaxial substrates, which renders their surface noncrystalline. This oxide needs to be removed before any epitaxial material deposition, which is usually done by heating. In the case of gallium arsenide (GaAs), the substrate is slowly heated to around 610 °C, while stabilizing the crystal with a constant arsenic flux of around $1.2 \times 10^{-5}$ Torr, to avoid As evaporation.[10] Once the oxide is removed, in order to avoid damaging of the crystal, further temperature ramping needs to be stopped; usually the temperature is in fact decreased. To detect the deoxidation, the MBE operator supervises the RHEED image during temperature increase, and once the diffraction pattern of a crystalline surface starts to form, the operator manually ends the heating procedure. Not only is the constant presence of the operator required but also due to its manual character the deoxidation procedure is error-prone. Automatic detection of the deoxidation is challenging, first because RHEED patterns are often weak since the raw substrate surfaces are not atomically flat, and second because the RHEED image contrast is dependent on some parameters such as filament current or electron beam angle and hence is not exactly constant in each run. Finally, the substrate is usually lying on a rotating sample holder; hence, the RHEED pattern constantly changes.

Received: October 7, 2022
Revised: December 22, 2022

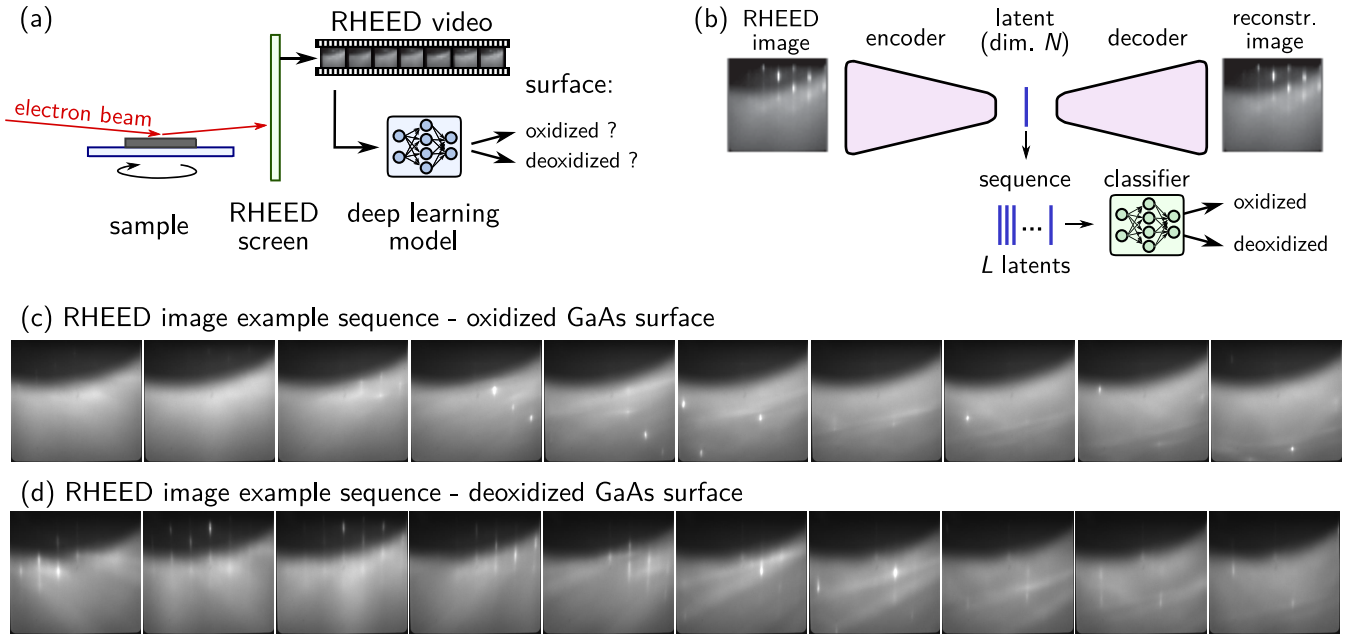


**Figure 1.** Deoxidation detection problem. (a) Short sequences of the RHEED video obtained from a rotating GaAs substrate are analyzed by a deep learning neural network (NN) model, to determine if the substrate deoxidation process has terminated. (b) The DL model consists of two stages. An autoencoder encodes every single RHEED image into a latent vector of dimension $N$ with the goal of reconstructing the original image. The latent vector thus contains all relevant features of the RHEED image. Sequences of $L$ latent vectors are then used for classification of deoxidized surfaces. The classifier network thus analyzes short RHEED videos covering a certain rotation angle. (c) Example sequence of 10 consecutive raw RHEED images obtained from an oxidized GaAs surface. (d) Example sequence of 10 consecutive raw RHEED images obtained after deoxidation from the same GaAs surface as shown in (c).

Methods from artificial intelligence, including deep learning, are being increasingly applied to nano and materials science.[11−15] Recently, first attempts have been reported to use statistical methods and machine learning for RHEED image interpretation.[16−20] Inspired by these pioneering works, we propose to use a deep-learning (DL) approach for classification of oxidized and deoxidized substrates via their RHEED patterns, to resolve the problems described above. As mentioned above, due to the sample rotation, the RHEED signal can confidently indicate deoxidation only during short moments, when the electron beam is aligned with a lattice direction of the crystal. In contrast to recent propositions to use DL with RHEED for surface reconstruction identification,[19,20] we therefore propose a model which analyzes *sequences* of RHEED patterns (i.e., videos), instead of single images. To this end, we propose a two-stage deep learning model. The first stage is an autoencoder, which compresses each full-resolution RHEED image into a low-dimensional latent vector. The second stage subsequently determines the oxidation state for a sequence of such latent vectors, hence for the compressed representation of a short RHEED video sequence. We provide a detailed analysis of the required latent and sequence lengths and demonstrate the accuracy of the model as well as its stability over a period of more than 6 months between training data sampling and testing. Finally, we provide online the data and codes as well as pretrained models to reproduce our results.[21]

## 2. RESULTS AND DISCUSSION

**2.1. Substrate Deoxidation Classification Problem.** On commercial substrates, a native oxide layer of a few nanometers encapsulates the crystal surface. The RHEED electron beam does not penetrate through this oxide and hence does not reach the crystalline lattice (in our case GaAs). The electrons are thus not diffracted but scattered. Independent of the sample rotation angle, the RHEED signal on the fluorescent screen is diffuse and no diffraction pattern occurs (cf. Figure 1c). On the other hand, without an oxide layer the RHEED electrons are diffracted by the atomic lattice of the now crystalline surface. However, the transition from oxidized to deoxidized is not instantaneous and during the deoxidation the classification is often difficult. Furthermore, due to the rotation of the sample, the diffraction pattern continuously changes and, especially during the deoxidation process, the pattern arises not similarly clearly for different rotation angles. Our operator classifies the surface as deoxidized when a clear diffraction pattern occurs repeatedly during at least one full rotation cycle of the substrate.

The general problem is schematically depicted in Figure 1a. Our goal is to precisely determine the moment of full oxide removal from a GaAs substrate by monitoring the RHEED pattern during the deoxidation process.

However, as mentioned above, the image dynamics due to the constant rotation of the sample is a challenge for an algorithmic evaluation. Furthermore, disordered bright spots can occur also from oxidized surfaces (see Figure 1c). Therefore, an algorithmic classification is not entirely trivial. By feeding short video sequences of several consecutive RHEED images to a classification neural network, we aim at determining the oxidation state of the substrate surface, in order to reduce the necessity of human supervision of the substrate cleaning process.

**2.2. Data Set.** To train a neural network on deoxidation reconnaissance, we generate a training data set by capturing RHEED videos before and after the oxide removal procedure. The images are collected in real time at 24 frames per second,


https://doi.org/10.1021/acs.cgd.2c01132
Cryst. Growth Des. XXXX, XXX, XXX−XXX

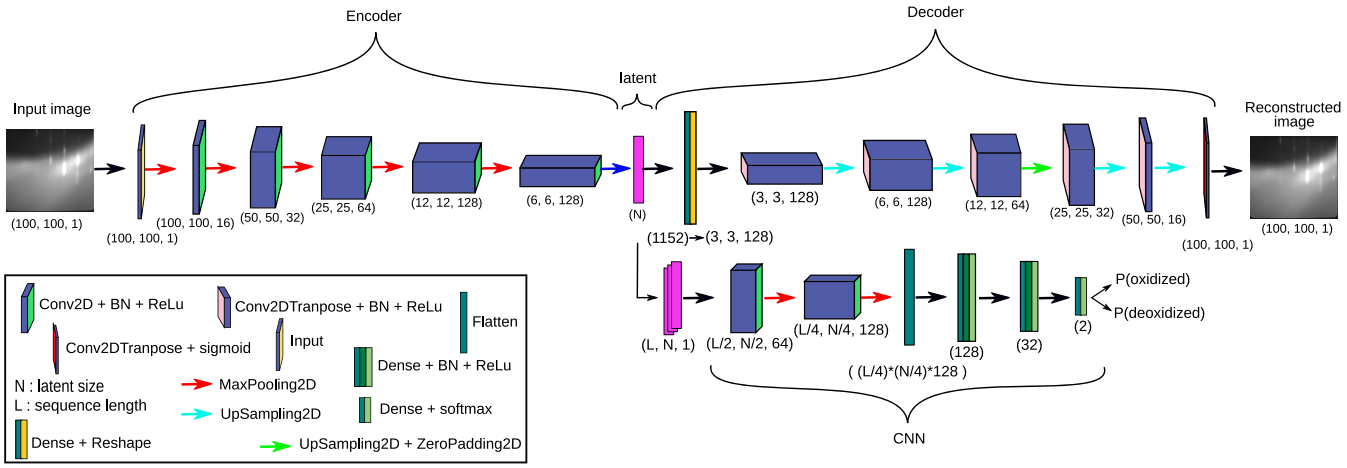


**Figure 2.** Detailed network architecture. In a CNN autoencoder (AE, top), each RHEED image is first compressed through an convolutional encoder into a 1D latent vector of length *N*. Training target of the AE is reconstruction of the original image from the latent vector (through the decoder stage, only used during training). The second stage of the model is a classifier CNN network (bottom right), taking as input a sequence of *L* latent vectors, corresponding to a series of RHEED images. These *L* latent vectors are stacked and passed into a CNN for classification into two classes: oxidized and deoxidized. All convolutions are followed by batch normalization (BN)[22] and ReLU activation. For descriptions of the different network layers we refer the interested reader to relevant literature.[23]

while the sample rotates with 12 rounds per minute; hence, we capture 120 images per full rotation. The RHEED video is thereby captured image by image, using a CMOS camera (Allied Vision Manta G319B) with 4 × 4 pixel binning, resulting in raw images of 416 × 444 pixels at 12 bit grayscale intensity resolution. Those images are simply converted to 8 bit format and scaled to 100 × 100 pixels. In total we collected videos containing a total of 7644 RHEED images from five substrate oxide removal procedures within a period of a few days. 3110 of these images correspond to deoxidized surfaces; the rest are images from GaAs surfaces which were covered by a native oxide layer. GaAs surface oxides decompose at temperatures of around 580−630° *C*.[24] On our commercial substrates we typically observe deoxidation at around 610−630 °C. Substrate temperatures at which deoxidized videos were taken are slightly lower, around 550−600 °C (videos were recorded during ramping of the temperature). Deoxidized images where taken directly after oxide removal, at around 610 °C. During and after deoxidation, the $As_4$ pressure for surface stabilization is held at $1.2 \times 10^{-5}$ Torr.

We use 20% of the data set for validation and the remaining 80% for training. In addition to the oxidized and deoxidized image sets, we also captured images during the full deoxidation procedure. These are not used during training and serve for testing of the algorithm. RHEED images during a further deoxidation were captured around 6 months after generation of the initial data set. These serve for an assessment of the long-term stability of the classification.

**2.3. RHEED Sequence Classifier Network Model.** Our deoxidation monitor deep-learning model is composed of two stages, as depicted schematically in Figure 1b. The first stage is a feature extractor network, compressing the large RHEED images into compact latent vectors. This is done separately image by image. The second stage is the actual classification network. Its inputs are sequences of latent vectors, corresponding to short RHEED videos. We implemented the models in Python using Keras with TensorFlow as the back end.[25,26] For preprocessing and data management we furthermore use the packages openCV, Scikit-image, Scikit-learn, and hdf5/h5py.[27−30]

*2.3.1. Image Feature Extraction.* As a feature extractor we use a deep convolutional autoencoder (AE) neural network, which has been reported to offer slightly superior compression quality compared to other dimensionality reduction methods such as principal component analysis (PCA), especially at high compression rates. We note, however, that AEs require in general more computational resources. Thus, if computation speed is crucial, PCA could be used instead, for situations where only a moderate reduction in encoding performance is expected.[31]

The model details of our AE are shown in the top row of Figure 2. A RHEED image goes through the encoder stage, being compressed into a latent vector of dimensionality *N*. For training, the latent vector is fed into a decoder stage, which is an exact mirror of the encoder, except for replacing convolution layers by transpose convolutions and applying zero padding if required, to maintain correct image dimension. Through nonsupervised training, the autoencoder learns to reconstruct the unlabeled input images from their learned latent vector representation.

Please note that we optimized the network for low parameter number, in order to have a computationally efficient model. To this end we do not double the number of channels once a depth of 128 filters is reached. We also compared the architecture with a ResNet,[32,33] replacing the single convolutions by residual convolutional blocks each of which employing a sequence of three convolutions. The performance is similar and offers no advantage in deoxidation classification. Please note that this applies to the specific problem discussed here; for other problems the slightly improved accuracy offered by a ResNet may very well be beneficial.

*2.3.2. Sequence Classification CNN.* The second stage of our model is the actual classification network. Because the MBE sample is rotating, the RHEED images are constantly varying. Especially during deoxidation, the surface is not atomically flat and signatures of oxide removal often occur in the RHEED images only when the electron beam is aligned with the crystal lattice of the substrate. For a high accuracy we therefore classify *sequences* of RHEED images: i.e., short videos. To do so, we compress the RHEED video image by image


https://doi.org/10.1021/acs.cgd.2c01132
Cryst. Growth Des. XXXX, XXX, XXX−XXX

using the trained AE. Using the results, we then create sequences of $L$ consecutive latent vectors. These sequences are arranged as 2D arrays of size $(L, N)$ and are fed into a 2D convolutional classifier network with two output classes, one for the oxidized and the second for deoxidized state of the surface. The details of the network are shown in the bottom right of Figure 2. It is trained in a supervised manner of the data set, with the goal of predicting the correct surface state from a sequence of compressed RHEED images.

**2.4. Results.** *2.4.1. Autoencoder Reconstruction Quality.* For the tuning of the autoencoder architecture we qualitatively tested AE layouts with varying numbers of layers and convolutional kernels by compressing and reconstructing random RHEED images that were not used for training. Once a layout was found that accurately reproduced the visual appearance of our RHEED images, we tested the reconstruction quality via peak signal-to-noise ratio (PSNR), where the original image is used as the signal and the difference between original and reconstructed image is used as noise. We obtained the best PSNR for $N = 70$ with a value of around 90; no further improvement was observed for larger latent dimensions.

*2.4.2. Classification Accuracy.* We test the classification accuracy of the full two-stage model (AE + classifier) with different values for the sequence length and fixed latent size $N = 50$, as well as for varying latent dimensions while fixing the sequence length $L = 15$. We find that the sequence length $L$ is indeed a crucial parameter. Single-image classification basically fails. Sequences of at least $L = 5$ images are required, to drop the error rates well below 1%, which can be seen in Figure 3a. Concerning the latent size, we require at least a compression dimensionality of $N = 10$ in order to get error rates well below 1%; increasing latent dimension further improves the accuracy only marginally. This is shown in Figure 3b.

Subsequently we test whether the network is capable of determining the exact moment of deoxidation on a set of images captured during the entire deoxidation procedure. First we fix the latent size to $N = 50$ and increase the sequence length successively from $L = 1$ to $L = 15$ (Figure 4, top row). In agreement with the former test (Figure 3a), we find that starting from sequence lengths of $L = 5$ the network works accurately and is essentially error-free. It detects the precise moment of deoxidation with an agreement of a few seconds compared to the estimation of the human operator. We then fix the sequence length to $L = 15$ and vary the latent dimension between $N = 1$ and $N = 70$ (Figure 4, bottom row). For latent vectors of dimension $N = 10$ or larger, we find again quasi-error-free classification and precise determination of the deoxidation moment.

We conclude that the smallest data size for accurate operation is a latent dimension of $N = 10$ and a sequence length of $L = 5$. At a video rate of 120 images per full rotation of the substrate holder, this sequence length ($L = 5$) corresponds to a rotation angle of 15° being concurrently processed by the classifier network.

*2.4.3. Temporal Stability of the Classification Accuracy.* An MBE instrument is a highly complex apparatus, housing many parts that have a potential impact on the long-term stability of the RHEED precision. To name a few, multiple pumps are required to maintain an ultrahigh vacuum, hot source chambers are distributed around the chamber, and mechanical stress is applied to the sample holder. It is not only constantly rotating but also repeatedly heated and cooled by temperature differences of many hundreds of degrees kelvin.

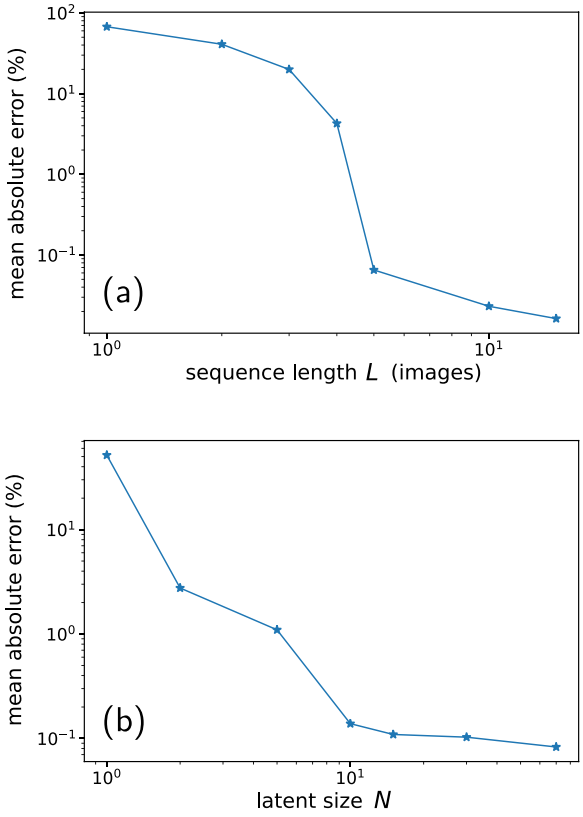


**Figure 3.** Prediction fidelity vs sequence length and latent size. Mean absolute error (MAE) on the test set for network models working with increasing sequence length and latent dimension. (a) MAE as a function of sequence length $L$. The latent size is fixed to $N = 50$. The actual classification error of sequences $L \geq 5$ is zero. (b) MAE as a function of latent space dimensionality $N$. The used sequence length is $L = 15$. The actual classification error with latent size $N \geq 10$ is virtually zero.

Furthermore, frequent (re)alignment of the RHEED source and camera also have an impact on the reproducibility of the diffraction patterns. In addition to all these mechanical perturbations, it is very hard to avoid deposition of the source material on the RHEED screen, slowly altering the diffraction images during an epitaxy campaign. In consequence, a constant RHEED image quality can usually not be guaranteed over a long term.

Deep learning being a data-based technique, the main advantage is generally a high robustness against noise and perturbations, provided that the training data are sufficiently rich.[34] We therefore expect our approach to deliver reliable classification results over significant time periods. In order to assess whether our approach can provide accurate deoxidation detection over durations of typical crystal growth campaigns, we recorded a deoxidation RHEED video around 6 months after the acquisition of the training data.

During the 6 months between generation of the training data and this test, around 40 epitaxies, equivalent to roughly 200 h of growth, have been performed on the machine. The most altered component in the system is in fact the RHEED screen, which suffers from metalization over time, implying a considerable deterioration of the RHEED image quality. In fact, our test on 6 months newer data was motivated by the observation of a degradation of the RHEED screen, especially


https://doi.org/10.1021/acs.cgd.2c01132
Cryst. Growth Des. XXXX, XXX, XXX−XXX

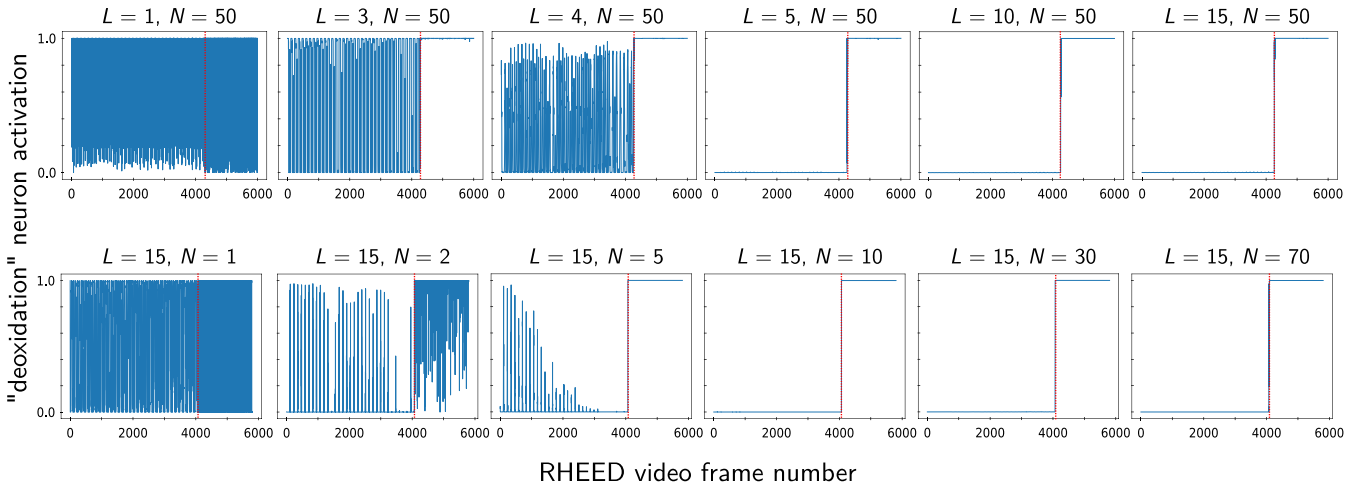


**Figure 4.** Detection accuracy of deoxidation moment. The impact of the sequence length as well as of latent vector dimension is tested on a video captured during a full deoxidation. The video consists of 31819 RHEED images, the last 6,000 of which are shown. Deoxidation occurs around 1800 frames before the sequence end (indicated by a red dotted line). The RHEED video is captured with 24 frames per second. Top row: with an increase in sequence length $L$, the latent dimension of the autoencoder is fixed to $N = 50$. Bottom row: with an increase in latent dimension $N$, the sequence length of the classifier is fixed to $L = 15$.

in the upper third of the diffraction images (see also Figure 1c,d).

A comparison of "original" and "six months newer" deoxidized GaAs RHEED images under (1,1,0) incidence is shown in Figure 5a,c. We find that the newer data is not only visually deteriorated as a result of screen metallization (c.f. upper regions of the images) but also the newer recordings show diffraction patterns which are shifted with respect to the original training data. The shift is probably due to repeated alignment corrections on the camera and the electron beam positions.

Figure 5b,d shows also the autoencoder reconstructions of the (1,1,0) diffraction patterns respectively from the training set and 6 months after. Interestingly, in the reconstruction of the 6 months newer image, the diffraction spots are shifted back to the positions at which they occurred at the time when the original data set was captured. It appears that the convolutional *encoder* stage correctly identifies the information about diffraction in the image. The *decoder*, on the other hand, has learned its spatial reconstruction only from the training set; therefore, it places the diffraction pattern according to the original RHEED alignment. In fact, the classifier still works correctly with shifted images. Since classification is done using latent vectors and not with the direct image information, we conclude that the relative diffraction spot/line positions are correctly interpreted by the encoder also from shifted images and mapped to the same latent variables as with the original training data.

In Figure 5e we show the classification results for a full video from a deoxidation run 6 months after network training. While we observe a slightly reduced classification certainty regarding the exact moment of deoxidation, also 6 months after initial training data recording, the nonretrained neural network still performs sufficiently well on the deoxidation detection. We want to note that after training the pretrained network for a few additional epochs on a small set of new images, the fidelity of the network reaches the same confidence as observed in the right-hand panels of Figure 4. In conclusion, convolutional neural networks offer a remarkable robustness for data characterization tasks that are susceptible to small perturbations and noise. CNNs are thus particularly interesting for long-term applications in real-world systems.

*2.4.4. Network Classification Speed.* Deep-learning frameworks such as TensorFlow/Keras are highly optimized for batch processing many samples in parallel. On prerecorded data, where our models can run parallel data processing, they are therefore extremely efficient, requiring only a few milliseconds of processing time per sequence. In an in situ operation, however, data needs to be processed "on the fly", hence image-by-image (or at least one video sequence at a time). For such sequential operation, TensorFlow is not optimized. Processing of a single sequence with 15 images by our autoencoder with subsequent classification takes in sum 220 ms. This is not enough for full real-time analysis of a 24 frames per second video stream as we receive from our RHEED camera. However, the time scales of MBE crystal growth dynamics are not faster than in the order of tens of seconds; thus, classification at full video speed is in fact not required. For the relevant time scales, our method is largely fast enough to practically perform real-time in situ control. Furthermore, the evaluation speed could probably be increased with some further optimization efforts. Finally, we want to note that our network models are compact enough to be used on conventional CPUs (central processing units). The full processing of a single 15-image sequence takes 220 ms on our NVIDIA 3070Ti GPU, against 260 ms on a 10th generation Intel i7 CPU. Hence, no additional hardware such as graphics processing units is necessary in an MBE control computer.

## 3. CONCLUSIONS

In conclusion, we presented a deep-learning model based on a 2D convolutional neural network to detect the surface oxidation state of GaAs substrates from raw RHEED image sequences, as typically available in molecular beam epitaxy. Our model consists of a first autoencoder neural network, which learns to compress individual RHEED images to a low-dimensional latent space. Sequences of consecutive, compressed RHEED images are then classified for their surface oxidation state through a second convolutional network. We


https://doi.org/10.1021/acs.cgd.2c01132
Cryst. Growth Des. XXXX, XXX, XXX−XXX

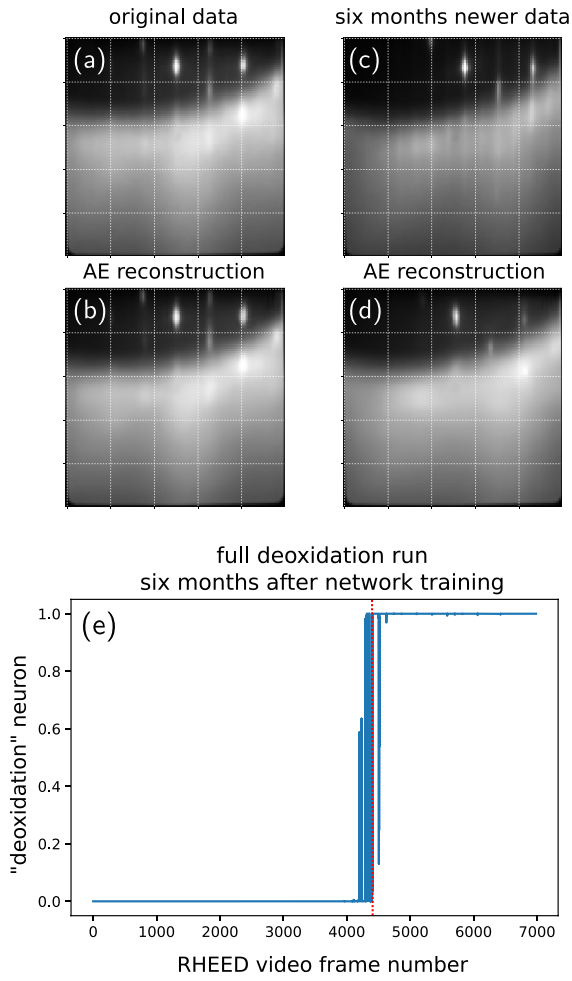


**Figure 5.** Test on 6 months newer data. (a) RHEED image of a deoxidized GaAs surface along (1,1,0) incidence, from training data, and (b) its reconstruction by the autoencoder with latent vector length $N = 50$. (c) RHEED image under the same incidence angle (along (1,1,0)), but 6 months later, after having performed more than 200 h of crystal growth. In particular an increased metalization of the upper third of the RHEED screen is clearly visible, along with a slight displacement of the diffraction pattern due to repeated alignment procedures. (d) Its reconstruction by the same autoencoder as used for (b), hence without retraining on new data. (e) Evaluation of the RHEED video from a full deoxidation run, recorded 6 months after the training data, without retraining the network model. Latent dimension $N = 50$, sequence length $L = 15$.

presented a systematic analysis of classification performance as a function of used compression ratio as well as RHEED video sequence length. We demonstrated that the model accurately identifies the exact surface deoxidation moment and that the performance is robust during at least 6 months of MBE operation without requiring retraining.

While our specific, trained network will of course work only with the MBE setup and RHEED screen used for our training data generation, a generalization to other growth chambers can simply be done by training the same models on the relevant data. Video recording is straightforward; we demonstrated that the approach works well with data from only five deoxidation recordings, and we proved it to function reliably during at least several months. In consequence, our approach is very appealing thanks to its simplicity and low computational cost. Without requiring additional hardware it can be easily set up in any RHEED-equipped MBE instrument.

## AUTHOR INFORMATION


### Corresponding Author

**Peter R. Wiecha** − *LAAS-CNRS, Université de Toulouse, CNRS, UPS, F-31400 Toulouse, France;* orcid.org/0000-0002-4571-0116; Email: pwiecha@laas.fr

### Authors

**Abdourahman Khaireh-Walieh** − *LAAS-CNRS, Université de Toulouse, CNRS, UPS, F-31400 Toulouse, France*

**Alexandre Arnoult** − *LAAS-CNRS, Université de Toulouse, CNRS, UPS, F-31400 Toulouse, France*

**Sébastien Plissard** − *LAAS-CNRS, Université de Toulouse, CNRS, UPS, F-31400 Toulouse, France*

Complete contact information is available at: https://pubs.acs.org/10.1021/acs.cgd.2c01132


### Notes

The authors declare no competing financial interest.

## ACKNOWLEDGMENTS

This work was supported by the Toulouse HPC CALMIP (grant p20010). This study benefited from the support of both the LAAS-CNRS micro and nanotechnologies platform, a member of the French RENATECH network, and the EPICENTRE common laboratory between Riber and CNRS.

## REFERENCES


(1) Ino, S. Some New Techniques in Reflection High Energy Electron Diffraction (RHEED) Application to Surface Structure Studies. *Jpn. J. Appl. Phys.* **1977**, *16* (6), 891.

(2) Horio, Y.; Hashimoto, Y.; Ichimiya, A. A new type of RHEED apparatus equipped with an energy filter. *Appl. Surf. Sci.* **1996**, *100-101*, 292−296.

(3) Braun, W.. Applied RHEED: Reflection High-Energy Electron Diffraction during Crystal Growth. *Springer Tracts in Modern Physics*. Springer Berlin Heidelberg: 1999.

(4) Ichimiya, A.; Cohen, P. I.; Cohen, P. I.. *Reflection High-Energy Electron Diffraction*; Cambridge University Press: 2004.

(5) NATO Advanced Study Institute on the Study of Surfaces, Interfaces by Electron Optical Techniques. In *Surface and Interface Characterization by Electron Optical Methods*; Valdre, U., Howie, A., Eds.; Plenum Press: Published in cooperation with the NATO Scientific Affairs Division: 1988.

(6) Jo, J.; Tchoe, Y.; Yi, G.-C.; Kim, M. Real-Time Characterization Using in situ RHEED Transmission Mode and TEM for Investigation of the Growth Behaviour of Nanomaterials. *Sci. Rep.* **2018**, *8* (1), 1694.

(7) Thelander, C.; Caroff, P.; Plissard, S.; Dick, K. A. Electrical properties of InAs1-xSbx and InSb nanowires grown by molecular beam epitaxy. *Appl. Phys. Lett.* **2012**, *100* (23), 232105.

(8) Daudin, B.; Feuillet, G.; Hübner, J.; Samson, Y.; Widmann, F.; Philippe, A.; Bru-Chevallier, C.; Guillot, G.; Bustarret, E.; Bentoumi, G.; Deneuville, A. How to grow cubic GaN with low hexagonal phase content on (001) SiC by molecular beam epitaxy. *J. Appl. Phys.* **1998**, *84* (4), 2295−2300.

(9) Ohtake, A.; Mano, T.; Sakuma, Y. Strain relaxation in InAs heteroepitaxy on lattice-mismatched substrates. *Sci. Rep.* **2020**, *10* (1), 4606.

(10) Bastiman, F.; Cullis, A. G. GaAs(001) planarization after conventional oxide removal utilising self-governed InAs QD site selection. *Appl. Surf. Sci.* **2010**, *256* (13), 4269−4271.

(11) Sacha, G. M.; Varona, P. Artificial intelligence in nanotechnology. *Nanotechnology* **2013**, *24* (45), 452002.

(12) Ren, J.-C.; Liu, D.; Wan, Y. Modeling and application of Czochralski silicon single crystal growth process using hybrid model of data-driven and mechanism-based methodologies. *Journal of Process Control* **2021**, *104*, 74−85.
(13) Schimmel, S.; Sun, W.; Dropka, N. Artificial Intelligence for Crystal Growth and Characterization. *Crystals* **2022**, *12* (9), 1232.
(14) Choudhary, K.; DeCost, B.; Chen, C.; Jain, A.; Tavazza, F.; Cohn, R.; Park, C. W.; Choudhary, A.; Agrawal, A.; Billinge, S. J. L.; Holm, E.; Ong, S. P.; Wolverton, C. Recent advances and applications of deep learning methods in materials science. *npj Computational Materials* **2022**, *8* (1), 1−26.
(15) Yann Battie, A; Valero, A. C.; Horwat, D.; Naciri, A. E. Rapid ellipsometric determination and mapping of alloy stoichiometry with a neural network. *Opt. Lett.* **2022**, *47* (8), 2117−2120.
(16) Brown, T.; Lee, K.; Dagnall, G.; Kromann, R.; Bicknell-Tassius, R.; Brown, A.; Dorsey, J.; May, G.. Modeling MBE RHEED signals using PCA and neural networks. In *Compound Semiconductors 1997. Proceedings of the IEEE Twenty-Fourth International Symposium on Compound Semiconductors*; IEEE: 1997; pp 33−36. DOI: 10.1109/ISCS.1998.711537.
(17) Vasudevan, R. K.; Tselev, A.; Baddorf, A. P.; Kalinin, S. V. Big-Data Reflection High Energy Electron Diffraction Analysis for Understanding Epitaxial Film Growth Processes. *ACS Nano* **2014**, *8* (10), 10899−10908.
(18) Provence, S. R.; Thapa, S.; Paudel, R.; Truttmann, T. K.; Prakash, A.; Jalan, B.; Comes, R. B. Machine Learning Analysis of Perovskite Oxides Grown by Molecular Beam Epitaxy. *Physical Review Materials* **2020**, *4* (8), 083807.
(19) Kwoen, J.; Arakawa, Y. Classification of Reflection High-Energy Electron Diffraction Pattern Using Machine Learning. *Cryst. Growth Des.* **2020**, *20* (8), 5289−5293.
(20) Kwoen, J.; Arakawa, Y. Multiclass classification of reflection high-energy electron diffraction patterns using deep learning. *J. Cryst. Growth* **2022**, *593*, 126780.
(21) Khaireh-Walieh, A.; Arnoult, A.; Plissard, S.; Wiecha, P. R. RHEED deoxidation dataset, codes and pre-trained models. *Figshare*, 2022
(22) Ioffe, S.; Szegedy, C. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. *arXiv:1502.03167 [cs]*, February 2015.
(23) Goodfellow, I.; Bengio, Y.; Courville, A. *Deep Learning*; MIT Press; 2016.
(24) SpringThorpe, A. J.; Ingrey, S. J.; Emmerstorfer, B.; Mandeville, P.; Moore, W. T. Measurement of GaAs surface oxide desorption temperatures. *Appl. Phys. Lett.* **1987**, *50* (2), 77−79.
(25) Chollet, F., et al. *Keras*, 2015.
(26) Abadi, M., et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. https://www.tensorflow.org/, 2015.
(27) Bradski, G. *Dr. Dobb's Journal of Software Tools*; The OpenCV library: 2000.
(28) van der Walt, S.; Schonberger, J. L.; Nunez-Iglesias, J.; Boulogne, F.; Warner, J. D.; Yager, N.; Gouillart, E.; Yu, T. Scikit-image: Image processing in Python. *PeerJ.* **2014**, *2*, No. e453.
(29) Pedregosa, F.; et al. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* **2011**, *12* (85), 2825−2830.
(30) Collette, A. *Python and HDF5*; O'Reilly: 2013.
(31) Fournier, Q.; Aloise, D. Empirical Comparison between Autoencoders and Traditional Dimensionality Reduction Methods. In *2019 IEEE Second International Conference on Artificial Intelligence and Knowledge Engineering (AIKE)*, pp 211−214, June 2019. DOI: 10.1109/AIKE.2019.00044.
(32) He, K.; Zhang, X.; Ren, S.; Sun, J. Deep Residual Learning for Image Recognition. *arXiv:1512.03385 [cs]*, December 2015.
(33) Szegedy, C.; Ioffe, S.; Vincent, V.; Alemi, A. Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*; pp 4278−4284, February 2016. arXiv:1602.07261.
(34) Kurum, U.; Wiecha, P. R.; French, R.; Muskens, O. L. Deep learning enabled real time speckle recognition and hyperspectral imaging using a multimode fiber array. *Opt. Express* **2019**, *27* (15), 20965−20979.


https://doi.org/10.1021/acs.cgd.2c01132
*Cryst. Growth Des.* XXXX, XXX, XXX−XXX