Supporting Information for

Machine Learning-Guided Discovery of Gas Evolving Electrode Bubble Inactivation

Jack R. Lake † [1], Simon Rufer † [1], Jim James [2], Nathan Pruyne [2], Aristana Scourtas [2,3], Marcus Schwarting[4], Aadit Ambadkar [2], Ian Foster [3,4], Ben Blaiszik [2,3], Kripa K. Varanasi [1]*

[1] Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA

[2] Globus, University of Chicago, 5801 S Ellis Ave, Chicago, IL 60637, USA

[3] Data Science and Learning Division, Argonne National Laboratory, 9700 S Cass Ave, Lemont, IL 60439, USA

[4] Department of Computer Science, University of Chicago, 5801 S Ellis Ave, Chicago, IL 60637, USA

† Authors contributed equally

AUTHOR INFORMATION

Corresponding Author

*E-mail: varanasi@mit.edu

This PDF file includes:
Figures S1 – S6
Table S1
Video captions SV1, SV2A, SV2B, SV3

Experimental Materials and Methods Detail

§1. Materials

§1.1. Electrochemical Testing Materials

For all electrochemical experiments, the following electrodes were used: counter electrodes were constructed from high purity platinum wire (99.997%) of 0.5 mm diameter which were purchased from Fisher Scientific, low-profile Silver Chloride (Ag/AgCl) reference electrodes which were purchased from Pine Research Instrumentation, and working Platinum electrodes were fabricated as described herein. For all tests, an acidic electrolyte was used, purchased as a standard solution of 0.5 M sulfuric acid from Fisher Scientific.

§1.2. Pt Working Electrode Fabrication Materials

Working Pt electrodes were fabricated using p-type <100> orientation silicon wafers purchased from WaferNet as the starting substrates for cleanroom processing. For Pt sputter deposition, a Pt target (99.99%) was used from Kurt Lesker. See section 4.2 for further details.

§2.   Experimental Setup and Imaging

§2.1. Electrochemical Cell

A custom fabricated Teflon electrochemical cell was used for all electrochemical testing, shown schematically in Figure 1A. The Ag/AgCl reference and Pt counter electrode are placed such that an unobstructed view of the Pt working electrode from the top down is possible. The electrochemical cell is placed beneath an upright microscope (Zeiss Axio Zoom) equipped with a camera for imaging bubble evolution. The electrochemical cell is open to the ambient air such the working Pt electrode can be directly imaged from top down through the electrolyte.

§2.2. Synchronous Imaging

In order to draw insights on the relationship between instantaneous bubble coverage and potential, temporal alignment of the recorded videos and chronopotentiometric data is required. A one millisecond duration square wave generated by the external digital output of the Potentiostat at the start of the experiment was used as a trigger. This signal was received by an Arduino, which then transmitted a trigger signal to the camera to begin recording. Given the one millisecond signal duration and the Arduino reading and transmission time of a few milliseconds, temporal alignment is expected to within approximately five milliseconds. All videos are recorded at 60 frames per second.

§3.   Electrochemical Measurements

All electrochemical experiments were carried out using a three-electrode setup and a potentiostat/galvanostat, purchased from Biologic (VSP-150), to accurately measure and control current and potential. The model electrochemical reaction used throughout this work was Pt-catalyzed OER in an acidic electrolyte of 0.5 M sulfuric acid. All experiments followed the same order of first cycling Pt working electrodes between -0.2 V and 1 V vs. Ag/AgCl for 16 cycles, followed by anodic cleaning of the Pt working electrode by chronoamperometry for 1 minute at an applied potential of 2.2 V vs. Ag/AgCl. Chronopotentiometry testing was then performed sequentially for each current density tested for the electrode. The applied current is calculated from the current density on a geometric basis using the area of the 3mm circular electrode. The electrolyte was not saturated with a specific gas prior to each experiment. The electrolyte was replaced between each subsequent chronopotentiometry experiment. None of the experiments herein are performed under iR-compensation.

§4. Pt Working Electrode Fabrication & Characterization

§4.1. Electrode Texturing

Silicon wafer substrates are used to fabricate both smooth and nanospot b50 and b200 electrodes. No additional texturing is required for the smooth electrodes, which proceed directly to the Pt and nitride layer deposition described below. To create the underlying nanotexture of the nanospot electrodes, a Reactive Ion Etching (RIE) etching process is used in conjunction with photolithography. Briefly, photoresist is patterned to expose the regions that will become the nanotextured spots at either 50 or 200 μm spacing. Next, RIE with Fluorine and Oxygen gasses is used to etch a nanotexture into the exposed regions. The photoresist is then removed to expose the silicon wafer patterned with nanospots.

§4.2. Patterning Pt Active Electrodes using Photolithography

A 400 nm layer of Pt was conformally deposited atop the entire textured silicon wafers using magnetron sputtering (AJA International) at a rate of 1.5 Angstroms/second under Ar atmosphere. Next, a plasma enhanced chemical vapor deposition (PECVD) process was used to deposit silicon nitride on the surface (Samco PD220). Photolithography was then used to selectively pattern the active and inactive portions of the electrode. By selectively etching back the nitride layer only in a 3mm diameter circle, the Pt surface is again exposed, forming the 3mm working electrode with all other areas of the wafer remaining inactivated. Briefly, photoresist is patterned to expose the regions that will become the active Pt working electrodes. Next, reactive ion etching (RIE) is used to etch through the inactivating nitride layer to expose the Pt layer beneath. Finally, the photoresist is removed and the Pt working electrodes have been successfully fabricated. This process was performed for both smooth and nanospot electrodes. A graphical workflow is shown for the photolithographic patterning of the Pt electrode active areas is shown in **Figure S1**.

§4.3. Fabricated Electrode Characterization

Scanning electron microscopy was performed to characterize the surface morphology following electrode fabrication (ZEISS Gemini SEM). **Figure S2** shows the active and inactive regions as well as the grid spacing of nanospots. **Figure S3** shows representative SEM images of the nanospot texture.

Atomic Force Microscopy was performed to assess the roughness and uniformity of the nanospot texture (Asylum Jupiter XR AFM). **Figure S4** shows a representative profile plot with an RMS of 150 nm.

§4.4. Pt Electrode Elemental Composition Confirmation

The scanning electron microscope used for electrode characterization (ZEISS Gemini SEM) was also equipped with an X-Ray detector to perform Energy-dissipative X-ray spectroscopy (EDS) for elemental composition information of the electrode's surface. Representative EDS confirmation of both the identity and spatial arrangement of the patterned electrodes is summarized in **Figure S5**. Using both a line scan and surface mapping, we were able to show that our samples have a clear and sharp delineation of an active Pt area and the inactive silicon nitride surrounding region.

Mass Transport and Ohmic Effects

§5. Justification of Neglecting Mass Transport and Ohmic Effects

In the 100mA/cm^2 experiment presented in Figure 4C and 4D, the overpotential fluctuates with a range of approximately 2mV. The overpotential variation may be due to a combination of bubble inactivation of the catalyst, mass transport effects, and ohmic effects. We reason that ohmic potential losses should increase as bubble coverage increases, due to decreased electrolyte area for ionic conduction. We reason that mass transport overpotentials should also correlate positively with bubble coverage: when bubble coverage is high, the available area for diffusion of reaction products away from the surface is decreased. Moreover, low bubble coverage coincides with bubble departure events, which introduce convective mixing of the electrolyte close to the electrode. By this reasoning, we expect mass transport overpotentials to be low when bubble coverage is low, and high when bubble coverage is high.

Because bubble inactivation losses, ohmic losses, and mass transport losses are all positively correlated with bubble projected area, their individual contributions to the total overpotential fluctuations should sum to 2mV. If mass transport and ohmic effects were significant, the contribution of bubble inactivation to the total overpotential fluctuations would be less than 2mV, pushing the coverage-derived data in Figure 4F to the left, deeper into the non-physical regime. The assumption that ohmic losses and mass transport are negligible is therefore justified as conservative.

Computational Methods Detail

§6. Generating Ground Truth Data

In order to successfully train supervised deep learning networks, it is necessary first to generate a sufficient quantity of high-quality ground truth data. To this end, we employed the Darwin V7 platform to aid in the manual annotation of bubbles in single still frames of a video. We selected still frames randomly from all recorded videos of catalytic bubble nucleation, with the specification that (i) still frames must contain bubbles and (ii) frames were chosen at a sufficient frame distance such that temporally-adjacent selected frames were sufficiently diverse. Presuming perfectly circular bubbles, we manually inscribed annotations on each image and recorded bubble radii and (x, y) center positions, all in units of pixels. While we assumed no maximum cutoff on bubble size, annotators loosely enforced an arbitrary minimum bubble size of a five-pixel diameter. This can be observed in the main body by the average size of a bubble at nucleation (or, more precisely, when the bubble was first observed to have nucleated). Figure **S6** shows one such still frame after manual annotation. Note that annotators labeled all bubbles even if they were partially clipped off the frame. In order to help minimize the subjectivity between annotators, we assigned separate reviewers to edit and adjust locations after an initial round of annotations.

In total we labelled all bubbles in 22 images on a grid-nucleated surface and in 17 images on a pristine platinum surface without nucleation spots. Even though this dataset of annotated images represent significant manual effort, it is hardly sufficient for data-hungry deep learning network architectures. We therefore leveraged a number of routine augmentation strategies including adding a Gaussian blur, cropping, rotation, and salt-and-pepper noise to increase the size of the dataset.

On the Minimum Detectable Bubble Size: The minimum bubble size detectable by the ML data collection system is constrained by three factors: i) camera resolution, ii) nanospot size, and iii) the procedures adopted for manual training set annotation. The resolution of the camera used in the experiments is 3.3 μm/pixel. A bubble must be several pixels wide to reasonably detect its presence, thus setting a lower-bound limit of ~10 μm. Moreoever, the nanospots used in the surface-engineered electrodes have a 10 μm diameter. In order to avoid detection of the nanospots as bubbles, any detected bubbles with diameters less than or equal to roughly 10 μm were discarded. This size/pixel limitation was also incorporated in the manual training set annotation: only bubbles >3-4 pixels (>~10 μm) were annotated.

§7.   Deep Learning Methodologies

While Hough transforms have been used for circle detection in the past[1], after preliminary testing and tuning we were able to qualitatively demonstrate that neural networks outperformed Hough. A tuned Hough algorithm achieved a mAP score of 0.729 on the smooth electrodes, while the MaskRCNN and FasterRCNN networks had mAP scores of 0.847 and 0.824, respectively. The family of networks we trained (including the networks FasterRCNN[2] and MaskRCNN[3]) all employ a similar process of down-selecting an image to a set of region proposal networks (RPNs), then refine the predictions within these RPNs to form a close bounding box around an object. These CNNs construct two sub-networks, a box classifier and a box regressor, both of which contain fully connected network (FCN) layers. In the case of MaskRCNN, there is an additional fully-connected sub-network used for masking. Running inference on an image through the RPN architecture yields a list of k bounding boxes that are each denoted by four Cartesian points in R^2.

The box classifier also returns a k sized vector, which is an *objectness* score for each box. Note that k is a hyperparameter denoting the maximum number of boxes per sliding window as well as the predefined box shapes (that is, the minimum and maximum size as well as aspect ratio range on a box) are also hyperparameters. Importantly, the RPN architecture shares the same weights for every sliding window, as a result of convolution. As such it is explicitly translation equivariant; that is, shifting the input image will result in a shift in the output region of interest. After the RPN sub-network has identified these k regions of interest (RoIs), the k regions are refined via a second FCN sub-network, which also performs pooling. The pooling process is somewhat different between architectures. FasterRCNN uses RoIPool, which is similar to a nearest neighbors downsampler. MaskRCNN uses RoIAlign, which acts as a bilinear interpolator and results in pixel-level translation equivariance. This pixel-level translation equivariance has been demonstrated to improve mask segmentation accuracy in most cases. In the case of MaskRCNN, the features from RoiAlign are then fed through a final FCN sub-network which up-samples the image back to the full resolution of the input window. From this up-sampling, the final fixed-size encoding is flattened and run through FCN layers to yield class scores and a final bounding box regression. Once bounding box locations are discerned, a final non-maximal suppression (NMS) procedure is applied. This procedure includes removing any concentric circles, removing overlapping circles based on an intersection-over-union (IoU), and additional thresholds based on confidences. Finally, we convert bounding boxes to circle predictions. Suppose one bounding box contains a bubble and has side lengths s_1 and s_2. We calculate the bubble center by computing the average of the four Cartesian coordinates of the box, and we calculate the bubble radius as $\dfrac{s_1 + s_2}{4}$.
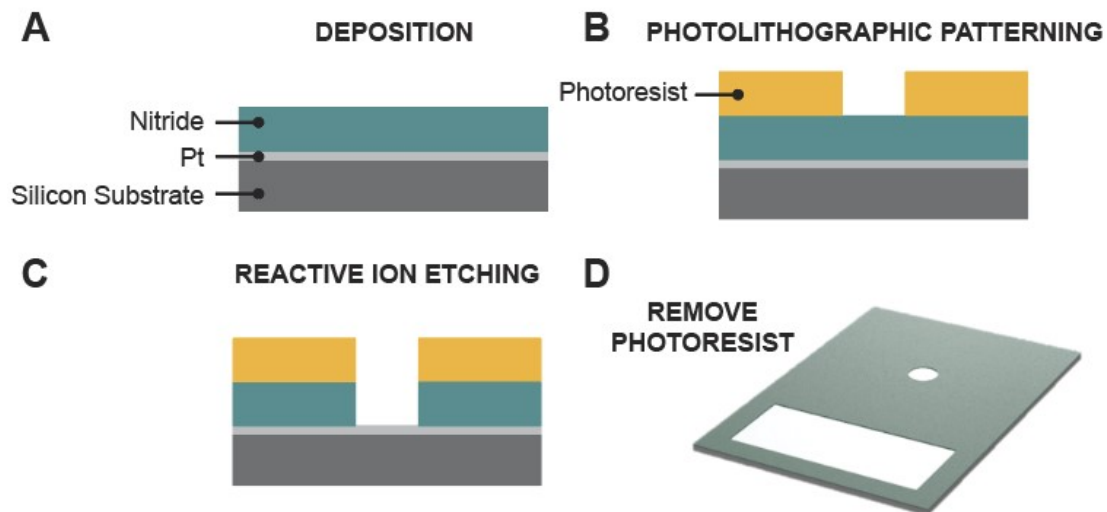
Figure S1. (A) First deposition of active Pt layer using magnetron sputtering, followed by inactive nitride layer via PECVD. (B) Photolithographic patterning of the regions to become the active area. (C) Reactive ion etching is used to etch through the inactivating nitride layer to expose the active Pt layer beneath. (D) After removing the photoresist, the final electrode sample with Pt patterned circular working electrode and exterior contact for connection to the potentiostat are shown.

Figure S2. Scanning electron micrograph highlighting the different electrode regions: the Pt active area with nanospots and the nitride inactivated surrounding area. Here, the b200 surface is shown with 200μm spacings of nanospots. Scale bar is 100μm.


Figure S3. Scanning electron micrograph of the nanotextures within a nanospot. Scale bar is 1μm.

Figure S4. Representative surface roughness illustration from AFM measurements of the nano texture.
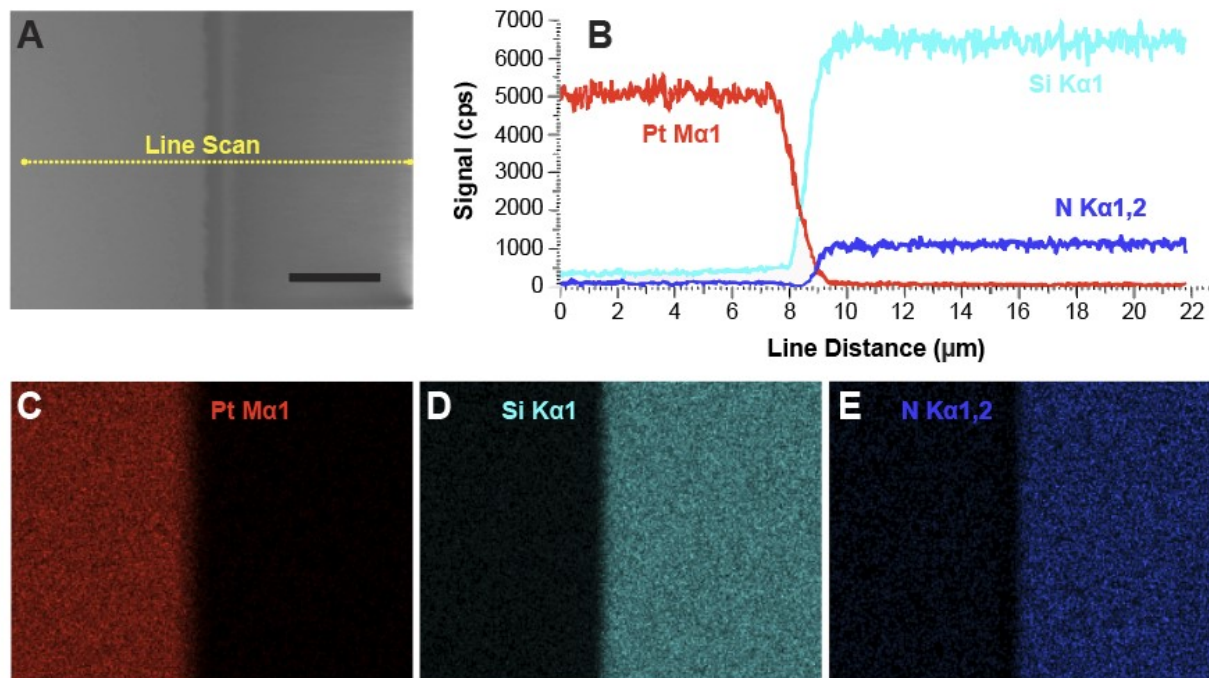
Figure S5. (A) Scanning electron image of the region analyzed using Energy-dispersive X-ray spectroscopy (EDS). Region on the left is the active Pt region and on the right is the inactive silicon nitride region. Annotation shown for location of line scan for spatial changes in composition of the EDS signal. Scale bar is 5 μm. (B) Line scan data for all elements measured during EDS line scan. (C-E) EDS mapping of the Pt, Si, and N characteristic EDS peaks, showing the Pt active region well-defined from the silicon nitride inactive region.

Figure S6: A still frame being manually annotated and reviewed within the Darwin V7 user interface.

Model	IoU Score	mAP Score
FasterRCNN V1	0.819	0.617
MaskRCNN V1	0.774	0.601
FasterRCNN V2	0.855	0.601
MaskRCNN V2	0.866	0.616

Table S1: IoU and mAP scores of various models tested in this work.

Supporting Information Video Captions

All Supporting Information Videos can be found here:
https://www.dropbox.com/scl/fo/s9m95a3znlgov70ab6765/h?dl=0&rlkey=wq8iuk49sfs46soekz9ica903

SI video SV1 Caption: Oxygen bubble evolution on the b200 nanospot electrode at 100mA/cm^2 showing coalescence-induced release of bubbles at consistent bubble diameters. Colored circles indicate where the ML algorithm detects a bubble. Electrode width is 3mm.

SI video SV2A Caption: Oxygen bubble evolution on the smooth Pt electrode at 25mA/cm^2. Colored circles indicate where the ML algorithm detects a bubble. Electrode width is 3mm.

SI video SV2B Caption: Oxygen bubble evolution on the smooth Pt electrode at 100mA/cm^2. Colored circles indicate where the ML algorithm detects a bubble. Electrode width is 3mm.

SI video SV3 Caption: Oxygen bubble evolution on the b50 nanospot electrode at 100mA/cm^2. The high number of bubbles and presence of bubble plumes obfuscates the surface, making detection difficult. Electrode width is 3mm.

REFERENCES

1. Illingworth, J. & Kittler, J. The Adaptive Hough Transform. *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-9**, 690–698 (1987).
2. Ren, S., He, K., Girshick, R. & Sun, J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *Adv. Neural Inf. Process. Syst.* **28**, (2015).
3. He, K., Gkioxari, G., Dollár, P. & Girshick, R. Mask R-CNN. *Proc. IEEE inter- Natl. Conf. Comput. vision,* 2961–2969 (2017).