
Rapid Fitting of Band-Excitation Piezoresponse Force Microscopy Using Physics Constrained Unsupervised Neural Networks

Alibek T. Kaliyev
Lehigh University
Bethlehem, Pennsylvania
alk224@lehigh.edu

Ryan Forelli
Lehigh University
Bethlehem, Pennsylvania
rff224@lehigh.edu

Shuyu Qin
Lehigh University
Bethlehem, Pennsylvania
shq219@lehigh.edu

Yichen Guo
Lehigh University
Bethlehem, Pennsylvania
yig319@lehigh.edu

Pedro Sales
Massachusetts Institute of Technology (MIT)
Cambridge, Massachusetts
psales@mit.edu

Seda Ogrenci-Memik
Northwestern University
Evanston, Illinois
seda@northwestern.edu

Michael W. Mahoney
University of California
Berkeley, Berkeley, California
mmahoney@stat.berkeley.edu

Amir Gholami
University of California
Berkeley, Berkeley, California
amirgh@berkeley.edu

Rama K. Vasudevan
Oak Ridge National Laboratory
Oak Ridge, Tennessee
vasudevanrk@ornl.gov

Stephen Jesse
Oak Ridge National Laboratory
Oak Ridge, Tennessee
sjesse@ornl.gov

Nhan Tran
Fermi National Accelerator Laboratory
Batavia, Illinois
ntran@fnal.gov

Philip Harris
Massachusetts Institute of Technology (MIT)
Cambridge, Massachusetts
pcharris@mit.edu

Martin Takáč
Mohamed bin Zayed University
of Artificial Intelligence
Abu Dhabi, United Arab Emirates
martin.takac@mbzuai.ac.ae

Joshua C. Agar
Drexel University
Philadelphia, Pennsylvania
jca92@drexel.edu

This supplemental materials document contains additional details, results, and analyses that support the main paper “Rapid Fitting of Band-Excitation Piezoresponse Force Microscopy Using Physics Constrained Unsupervised Neural Networks.” The sections below expand on the neural network architecture and training, computational benchmarks, noise robustness tests, and other supporting content referenced in the main text.

1 S1. Experimental procedure

400 nm thick $PbZr_{0.2}Ti_{0.8}O_3/30nmBa_{0.5}Sr_{0.5}RuO_3/NdScO_3$ thin films were grown using pulsed laser deposition. The films have been the subject of prior studies[1][2][3]. These films exist in a typical tetragonal ferroelectric phase; however, $NdScO_3$ imposes a large tensile strain that drives strain-induced spinodal instability. The result is a hierarchical domain structure that is a mixture of $c/a/c/a$ and $a1/a2/a1/a2$ domains[4]. This results in a single film with large variations in piezoresponse and switching mechanisms. Within the primary out-of-plane polarized $c/a/c/a$ is a large vertical piezoresponse and classical ferroelectric switching mechanisms. Conversely, there is suppressed vertical piezoresponse in the $a1/a2/a1/a2$ domains. Since this material is at an energetic degeneracy between the $c/a/c/a$ and $a1/a2/a1/a2$, the application of bipolar-triangular switching waveforms results in a two-step, three-state ferroelastic switching process. In these films, the topography has a sawtooth-like structure with an amplitude of 4 nm and a periodicity of 900 nm because of the tetragonality and large difference in crystallographic orientation of domain variants. Furthermore, the highly asymmetric elastic modulus tensor increases elastic modulus variations by up to 23%. Further details regarding the structure, properties, and switching mechanisms of this material can be obtained in prior reports[1][2][3]. The mixed-phase $PbZr_{0.2}Ti_{0.8}O_3$ provides an ideal model system to stress test Band-excitation piezoresponse force microscopy (BE-PFM) due to the highly variable piezoresponse, switching mechanisms, and elastic modulus. Significant variances in local piezoresponse and tip-surface contact within the dataset shift the cantilever resonance frequency, necessitating the BE-PFM. The dataset used in this work, an original creation of the authors, has been the subject of prior reports, and publicly released under the open source creative commons attribution 4.0 License[1].

Briefly, in the original studies, BE-PFM Polarization Spectroscopy (BE-PS), in a square 60x60 grid was conducted using a conductive cantilever mounted tip (Figure 1a). To measure the piezoresponse a small-signal (1 V) BE waveform was applied to excite the cantilever at a band of frequencies (bandwidth = 200 kHz), near the cantilever resonance (Figure 1b-c). The resulting response was detected using a quadrant photodiode and processed using a data acquisition system sampling at (4 MS/s). An onboard FPGA was used to conduct real-time fast-Fourier transform, resulting in the real and imaginary components of the cantilever response (Figure 1d). This is commonly visualized as the magnitude and phase of the response (Figure 1e). BE is crucial for these measurements as it minimizes effects from changing tip-sample contact resonances that can alter the observed response, enabling consistent qualitative measurements. Two sequential bipolar-triangular-switching waveforms were applied to switch the sample while measuring the piezoresponse in the on and off-field state (Figure 1f). Following fitting the cantilever response to a simple harmonic oscillator (SHO) model, the piezoresponse amplitude, phase (A), phase (ϕ), cantilever resonance frequency (ω), and dissipation (q) can be determined. By computing the optimum rotation angle (θ) that maximizes the cosine function, the piezoelectric hysteresis loops at every pixel can be calculated using $P = A\cos(\phi + \theta)$ (Figure 1g). To extract further insight from the piezoelectric hysteresis loops, it is common to fit them to a 9-parameter empirical function[5]. This fitting function extracts parameters related to the shape of the loop. This function is merely empirical and does not have a solid physical basis.

2 S2. Least Squares Fitting method

Conventionally, the fits are computed using Least Square Fitting (LSQF) method $\min_x f(x) = 0.5 \sum_{i=0}^{m-1} \rho(f(x)^2)$ s.t. $lb \leq x \leq ub$ based on the Jacobian using the SciPy implementation[6]. These algorithms are subject to many pitfalls. They require initial priors in the form of ‘guesses.’ Designing robust guessing algorithms is challenging, particularly in real experimental data with low signal-to-noise (SnR). When the guesses are incorrect, algorithms can get stuck at a local rather than ideal global minimum.

3 S3. Computational benchmarks for piezoelectric hysteresis loops fitting

The equations used for this fitting are as follows:

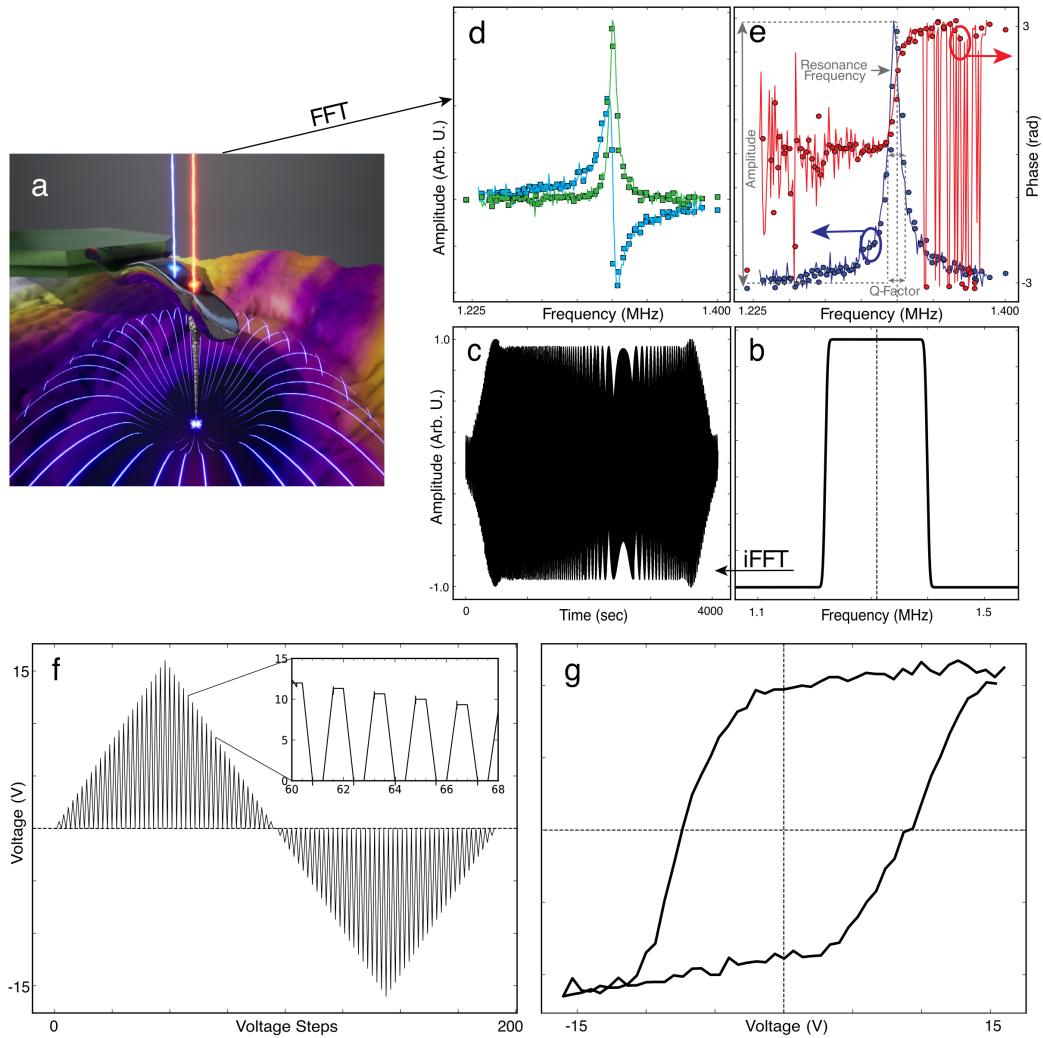


Figure 1: Schematic diagram of band-excitation piezoresponse force microscopy switching spectroscopy (BE-SS) a Artistic render of an AFM tip applying an electric field to the surface. b Band of excited frequencies excited. The dashed line shows the cantilever resonance frequency. c Band-excitation waveform used to excite the cantilever in time domain d Fast Fourier transform of a single-cantilever resonance during band-excitation piezoresponse force microscopy – shown as real and imaginary components. e Magnitude spectrum showing the amplitude and the phase of cantilever resonance. f Bipolar-triangular waveform used to switch the ferroelectric. The inset shows where the band-excitation waveform was applied in both the voltage-on and voltage-off states. g Example of a typical piezoelectric hysteresis loop obtained during BE-SS.

1. For g_1 , y_1 , and f_1 :

$$g_1 = \frac{b_1 - b_0}{2} \cdot \operatorname{erf}((V_1 - a_2)d + 1) + b_0 \quad (1)$$

$$y_1 = \frac{g_1 \cdot \operatorname{erf}\left(\frac{V_1 - a_2}{g_1}\right) + b_0}{b_0 + b_1} \quad (2)$$

$$f_1 = a_0 + a_1 y_1 + a_4 V_1 \quad (3)$$

2. For g_2 , y_2 , and f_2 :

$$g_2 = \frac{b_3 - b_2}{2} \cdot \operatorname{erf}((V_2 - a_3)d + 1) + b_2 \quad (4)$$

$$y_2 = \frac{g_2 \cdot \operatorname{erf}\left(\frac{V_2 - a_3}{g_2}\right) + b_2}{b_2 + b_3} \quad (5)$$

$$f_2 = a_0 + a_1 y_2 + a_4 V_2 \quad (6)$$

ADAM, a first-order method, didn't yield satisfactory results even after 20+ minutes. AdaHessian, a second-order method, performed slightly worse due to computational complexity.

A Trust Region (TR) CG optimizer outperformed ADAM and AdaHessian, achieving an order of magnitude lower MSE in just 300 seconds (0.00426 ± 0.00025 with a batch size of 256). After 650 epochs, similar to AdaHessian, it reached 0.003 MSE.

We visually compared DNN and LSQF reconstructions for best, median, and worst cases. TR-optimized DNN performed better in the worst-case scenario and showed equivalent performance in the best and median cases.

We used violin plots to compare NN and LSQF parameter distributions, demonstrating that NN consistently exhibited narrower, more Gaussian-like distributions with fewer outliers. The DNN's superior performance can be attributed to reduced sensitivity to noise, perturbing initial guesses and causing local minima. Trust Region CG outperformed ADAM threefold in predicting the correct loop shape. We plotted parameter maps using LSQF and TR-optimized neural networks, observing key domain structure features and their visual similarity.

Table 1: Piezoelectric hysteresis loops fitting results with the deep neural network predictions performed for different batch sizes on Tesla P100-PCIE GPU for 300 seconds each.

Optimizer	Batch size	# of epochs	Reconstruction Loss (MSE)	Params Loss (MSE)
Trust Region CG	64	8.4±0.5	0.00496±0.00042	0.5897±0.0789
	128	16.0±0.6	0.00452±0.00027	0.5235±0.0393
	256	27.6±1.6	0.00426±0.00025	0.4645±0.0268
	512	51.4±13.8	0.00396±0.00026	0.4875±0.0322
	1024	106.0±11.2	0.00389±0.00030	0.4366±0.0208
AdaHessian	64	292.0±1.0	0.00975±0.00132	2.55927±0.88428
	128	500.0±0.0	0.00803±0.00027	2.28119±0.63473
	256	647.0±2.4	0.01155±0.00361	3.61001±0.70472
	512	735.0±7.4	0.01340±0.00537	3.01150±1.38619
	1024	387.0±0.0	0.02671±0.01630	4.03416±1.54232
ADAM	64	560.0±125.6	0.00774±0.00042	2.41973±1.02710
	128	937.0±233.7	0.00835±0.00106	2.65518±0.32084
	256	1568.0±402.3	0.00830±0.00121	3.13807±0.71523
	512	2537.0±684.7	0.01134±0.00333	2.95574±1.05801
	1024	3214.0±937.2	0.01012±0.00166	3.24794±0.50383

4 S4. Computational benchmarks for Simple Harmonic Oscillator (SHO) fitting

We conducted comprehensive experiments to assess the computational performance of both LSQF and DNN models across a range of practical scenarios, each with varying hyperparameters. It's worth noting that designing benchmark studies in this context is a complex endeavor due to factors such as data variability, optimization techniques, and computing hardware constraints.

While both LSQF and DNN fitting methods for cantilever resonance responses exhibit remarkable parallelism, there are constraints on the degree of throughput enhancement achievable through

parallelization and practical hurdles in the implementation process. Nevertheless, it’s noteworthy that both LSQF and DNN fitting techniques can yield results quickly enough to inform subsequent experiments, typically taking only a few hours. Notably, DNNs, once trained, simplify computations during inference to highly parallelizable matrix multiplication operations, resulting in lightning-fast processing.

Our benchmarking commenced with evaluating LSQF’s SHO fits using the BGLib Python implementation, leveraging parallel processing capabilities. These benchmarks involved fitting the entire dataset comprising 1,382,400 spectra on a cluster of four 2.3 GHz CPU cores. The fitting process was completed in approximately 18 ± 1 minutes, equating to a rate of 1280 fits per second. It’s worth mentioning that single-fit latencies, crucial for real-time analysis, were measured at the millisecond per fit level.

Subsequently, we turned our attention to benchmarking the performance and speed of the DNN. All experiments were conducted on Google Colab, utilizing a single P100 PCIe GPU. Model training entailed a 70/30 train/test split, with batch sizes ranging from 64 to 1024, doubling at each interval, and all models were trained for five epochs.

Table 2: SHO fitting results with deep neural network predictions performed for different batch sizes using AdaHessian optimizer on Tesla P100-PCIE GPU. Results were computed by training and testing the inference of the model on 10 different random seeds (from 1 to 10) and by choosing 4 best performances.

Batch size	Training time (s)	Reconstruction Loss (MSE)	Params Loss (MSE)	Inference time (fit/s)
64	899.6±10.1	0.047±0.002	0.132±0.011	23,866.3±341.8
128	473.9±4.2	0.048±0.006	0.3±0.12	21,141.6±134.1
256	312.9±13.3	0.047±0.005	0.133±0.013	10,341.3±203.2
512	286.5±6.4	0.047±0.003	0.138±0.011	204,081.6±37,515.0
1024	314.3±2.4	0.051±0.005	0.139±0.140	476,689.7±113,497.5

Table 3: SHO fitting results with deep neural network predictions performed for different batch sizes using ADAM optimizer on Tesla P100-PCIE GPU. Results were computed by training and testing the inference of the model on 10 different random seeds (from 1 to 10) and by choosing 5 best performances.

Batch Size	Training Time (s)	Reconstruction Loss (MSE)	Params Loss (MSE)	Inference Time (fit/s)
64	309.6±1.4	0.0422±0.00021	0.1174±0.0056	30,389.5±453.6
128	170.6±1.1	0.0424±0.00012	0.1168±0.0029	57,951.5±583.9
256	109.1±0.8	0.0425±0.00020	0.1172±0.0036	105,380.7±1706.1
512	66.2±0.1	0.0429±0.00062	0.1169±0.0051	146,983.0±607.5
1024	58.2±0.2	0.0432±0.00009	0.1181±0.0061	174,417.8±519.2

We employed the Adam optimizer, a momentum-based first-order optimization method for model evaluation. The results from this evaluation are detailed in the *Table 3*.

Additionally, we examined the DNN’s training speed when utilizing AdaHessian, a second-order optimization method (*see Table 2*). Notably, training speed was significantly influenced by the batch size, with smaller batches resulting in substantially slower training. This is attributed to increased computational instructions, reduced parallelism, and increased minibatch updates. However, with a sufficiently large batch size (in this case, 256), the model converged to an acceptable Mean Squared Error (MSE) comparable to LSQF results. When assessing fit parameters, a more critical metric, we calculated the MSE loss between LSQF and DNN fits. Irrespective of the batch size, substantial differences were observed in the fit parameters between LSQF and DNN fits, primarily attributed to inadequacies in LSQF initialization, as indicated by our comprehensive analysis.

Furthermore, we evaluated inference latency as a function of batch size. Notably, as the batch size reached a sufficiently large value (in this case, >128), there were significant reductions in inference

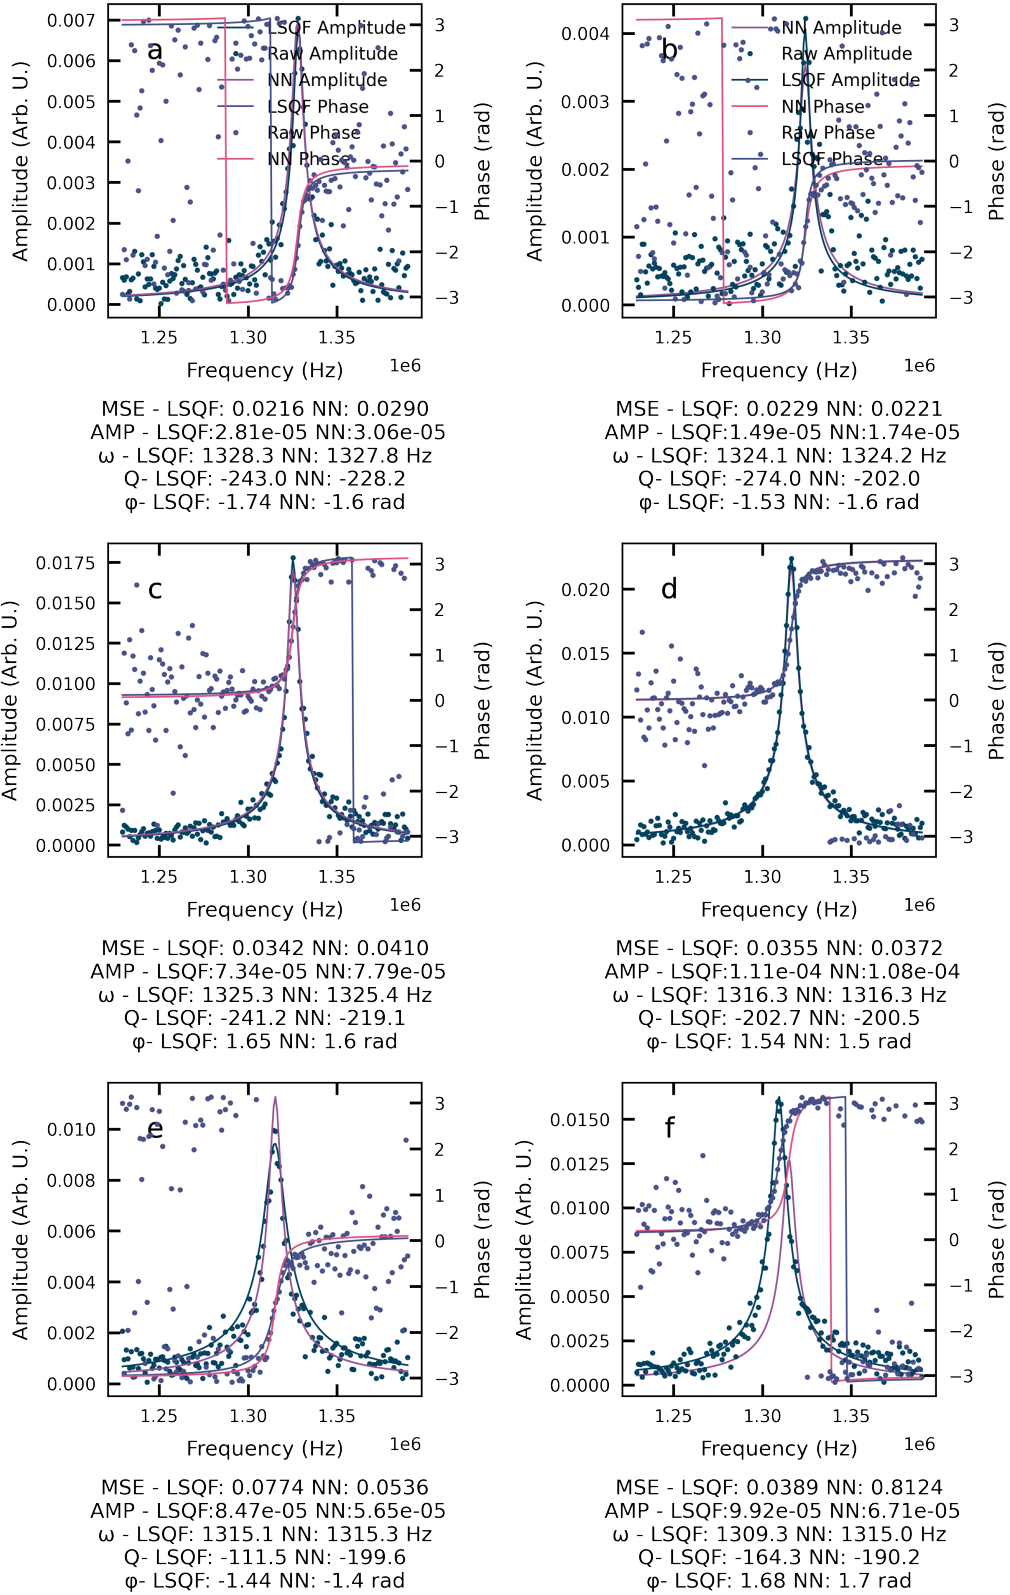


Figure 2: SHO fitting results of DNN in comparison with LSQF method's results. Best, median, and worst predictions of the LSQF method.

latency due to increased parallelism. For instance, with a batch size 1024, the model achieved a remarkable throughput of approximately 480,000 fits per second, completing the entire dataset in less than 3 seconds. This represents a speed improvement of more than two orders of magnitude compared to LSQF.

References

- [1] J. C. Agar, B. Naul, S. Pandya, S. van der Walt, J. Maher, Y. Ren, L.-Q. Chen, S. V. Kalinin, R. K. Vasudevan, Y. Cao, J. S. Bloom, and L. W. Martin. Revealing ferroelectric switching character using deep recurrent neural networks. *Nat. Commun.*, 10:4809, 2019.
- [2] J. C. Agar, Y. Cao, B. Naul, S. Pandya, S. van der Walt, A. I. Luo, J. T. Maher, N. Balke, S. Jesse, S. V. Kalinin, R. K. Vasudevan, and L. W. Martin. Machine detection of enhanced electromechanical energy conversion in $\text{pbzr}_{0.2}\text{ti}_{0.8}\text{o}_3$ thin films. *Advanced Materials*, 30:e1800701, 2018.
- [3] S. Qin, Y. Guo, A. T. Kaliyev, and J. C. Agar. Why it is unfortunate that “better faster and less biased linear machine learning models work so well in electromechanical switching of ferroelectric thin films (under review). *Adv. Mater.*, 2022.
- [4] J. C. Agar, A. R. Damodaran, S. Pandya, S. van der Walt, J. T. Maher, N. Balke, S. Saremi, Q. Li, J. Kim, M. R. McCarter, L. R. Dedon, T. Angsten, N. Balke, and S. V. Kalinin L. W. Martin S. Jesse, M. Asta. Three-state ferroelastic switching and large electromechanical responses in pbtio_3 thin films. *Adv. Mater.*, 29:1702069, 2017.
- [5] S. Jesse, H. N. Lee, and S. V. Kalinin. Quantitative mapping of switching behavior in piezoresponse force microscopy. *Rev. Sci. Instrum.*, 77:073702, 2006.
- [6] `scipy.optimize.leastsq` — `scipy v1.9.0` manual.