Reconstructing piezoelectric responses over a lattice: adaptive sampling of low dimensional time series representations based on relative isolation and gradient size

Michael R. Lindstrom *[1], William J. Swartworth *[2], and Deanna Needell[3]

[1] University of California Los Angeles
mikel@math.ucla.edu
[2] University of California Los Angeles
wswartworth@math.ucla.edu
[3] University of California Los Angeles
deanna@math.ucla.edu

Abstract. We consider a $d-$dimensional lattice of points where, at each lattice point, a time series can be measured. We are interested in reconstructing the time series at all points of the lattice, to a desired error tolerance, by adaptively sampling only a subset of lattice points, each over a potentially short time interval. The method we develop is tailored to atomic force microscopy (AFM) data where time series are well-represented by constant functions at each lattice point. Through a convex weighting of a point's relative isolation and relative gradient size, we assign a sampling priority. Our method adaptively samples the time series and then reconstructs the time series over the entire lattice. Our adaptive sampling scheme performs significantly better than sampling points homogeneously. We find that for the data provided, we can capture piezoelectric relaxation dynamics and achieve a relative ℓ_2 reconstruction error of less than 10% with a 47% reduction in measurements, and less than 5% with a 25% reduction in measurements.

1 Introduction and Background

Atomic force microscopy (AFM) is an imaging paradigm used to study materials at the nanoscale [9]. The method is often costly in terms of the time it takes to measure each pixel, and thus efficient sampling and reconstruction approaches are desired. The example we consider here arises in the study of ferroelectric materials and involves the measurement of the piezoelectric response as a function of time, space and voltage. Piezoelectricity refers to an electric charge that collects in solid materials, such as the ferroelectric film of the compound lead titanate (PbTiO3), which motivates the challenge posed by the *2021 Smoky Mountains Computational Sciences and Engineering Conference Data Challenge* [12] that we address in this paper.

The challenge (problem 8) amounts to reconstructing the piezoelectric response of a ferroelectric film as a function of time, measured at points (x, y) on its surface at various voltages V by taking as few measurements as possible. Measurements are taken over some time duration (on the order of seconds) at a fixed (x, y, V) to obtain the piezoelectric response $p(t; x, y, V)$ representing the surface deformation [13] as a function of time. The method should quickly determine the next set of (x, y, V)-values at which to measure p. These problem data are based on ferroelectric polarization dynamics measured with dynamic piezoresponse force microscopy [7].

2 Proposed Solution

In solving this problem, we have developed Relative Isolation and Gradient Size Sampling (RIGS Sampling), a simple active learning method that can "fill in" the piezoelectric response from incomplete data — specifically, we work with a subset of points on the surface with corresponding voltages to reconstruct the piezoelectric response at unmeasured position-voltage pairings. While active learning typically focuses upon an algorithm selecting unlabelled points from which to learn in classification tasks [10][4], it can be used in contexts where measurements of complicated systems can only be applied sparsely such as in air quality measurements [1]. Here, RIGS Sampling learns the piezoelectric response by choosing positions and voltages at which to take measurements to improve its reconstruction accuracy.

At a high level, RIGS builds intermediate constructions, by using the simple and fast k-nearest neighbors regression [3]. It then uses these reconstructions to bias towards sampling points with large gradients. Importantly, we demonstrate that adaptive sampling substantially outperforms non-adaptive sampling.

* Equal contributions

We begin with some notation and assumptions. We assume the surface locations and voltages at which measurements can be taken belong to a lattice with $x-$values $0 \leq x_0 < x_1 < ... < x_{N_x-1} \leq L_x$, $y-$values $0 \leq y_0 < y_1 < ... < y_{N_y-1} \leq L_y$, and $V-$values $V_{\min} \leq V_0 < V_1 < ... < V_{N_V-1} \leq V_{\max}$ for some positive integers N_x, N_y, and N_V with lengths $L_x, L_y > 0$ and voltages $V_{\min} < V_{\max}$. Furthermore, the piezoelectric responses can be measured at time points $0 \leq t_0 < t_1 < ... < t_{N_t-1} \leq T_{\max}$ with $N_t > 0$ an integer and $0 < T_{\max}$. In the original dataset, x, y, V, t, and the piezoelectric response p had physical units but they are not directly relevant for our study. Here, we work with dimensionless x, y, and V: $x_0 = 0$ and $x_{N_x-1} = N_x - 1$ with regular spacings of 1 between x's and similarly for y. The voltages are made dimensionless by dividing physical voltages by 1 Volt. The times are measured in seconds and the piezoelectric responses are measured in picometers.

We define a sample point as a point of the lattice $(x, y, V) \in [0, L_x] \times [0, L_y] \times [V_{\min}, V_{\max}]$ and assume that at each sample point, measurements are taken at times $t_0, t_1, ..., t_{n_t-1}$ for some positive integer $n_t \leq N_t$. With a point (x, y, V) of the lattice, we denote $p(t; x, y, V)$ to be the piezoelectric response as a function of time t and $\hat{p}(x, y, V)$ to be a constant approximation of the piezoelectric response. Now for $i \in \{0, ..., N_x-1\}, j \in \{0, ..., N_y-1\}, k \in \{0, ..., N_V-1\}$ and $m \in \{0, ..., N_t\}\}$, we let $\mathcal{E}_{ijkm} = p(t_m; x_i, y_j, V_k)$ be a tensor of the full experimental measurements at all lattice points and all times and $\widehat{\mathcal{E}}_{ijkm}$ reconstructions at all lattice points and times. Much of our notation is summarized in Table 1.

Symbol	Value	Type	Meaning
N_x/N_y	60 each	given	Number of $x-$values/y-values
N_V	16	given	Number of voltage values
N_t	128	given	Number of time values that can be measured
$\mathcal{E}$	full dataset	given	tensor of response over lattice points and times
f	varies	hyperparameter	Fraction of lattice points sampled for reconstruction
n_t	varies	hyperparameter	Number of time values used in sampling
w	varies	hyperparameter	Weighting in $[0, 1]$ to give to gradient vs isolation
n_r	2000	fixed hyperparemeter	Number of steps between full reconstructions
n_n	4	fixed hyperparameter	Number of nearest neighbors for reconstructions
k	$k(u, s) = \exp(-\lvert u - s\rvert^2)$	fixed hyperparameter	Kernel for local density
$\hat{p}$	not set	computed from method	time-constant approximation to response
$\widehat{\mathcal{E}}$	not set	computed from method	tensor of response predictions over lattice points and times

Table 1: Parameters used in this study.

With these notations set up, we formulate our objective as follows: given a tolerance $0 < \tau$, we seek to find $n_t \leq N_t$ to generate a collection of sample points $\mathcal{S}$ with $|\mathcal{S}| < N_x N_y N_V$ with which we can accurately reconstruct the piezoelectric responses over the entire lattice and over the full time window $[0, T_{\max}]$ to within a relative ℓ_2-error of τ. We define both the relative ℓ_2-error r by

$$r = ||\widehat{\mathcal{E}} - \mathcal{E}||_F / ||\mathcal{E}||_F \tag{1}$$

where the F-subscript denotes the Frobenius norm, i.e., the square root of the sum of squared differences running over all 4 dimensions.

In deriving our methodology, we first seek to understand the innate dimension of the data. We begin by considering what the piezoelectric response looks like as a function of t. In sampling approximately 1% of $(x, y, V)-$values and performing a Principal Component Analysis [6], we find that over 98% of the total variance of the response time series $p(t; x, y, V)$'s can be expressed by the first principal component and that the first component is approximately constant as depicted in Figure 1. This suggests the time series can be well described by constant functions $p(t; x, y, V) = \text{const}$. Based on other experiments of piezoelectric responses (see e.g. [8]) where the response had a correlation length on

the order of ≈ 100 nm, we also expect that nearby points in space will have similar responses. This makes our later choice to fill in missing data through nearest neighbor regression seem plausible.

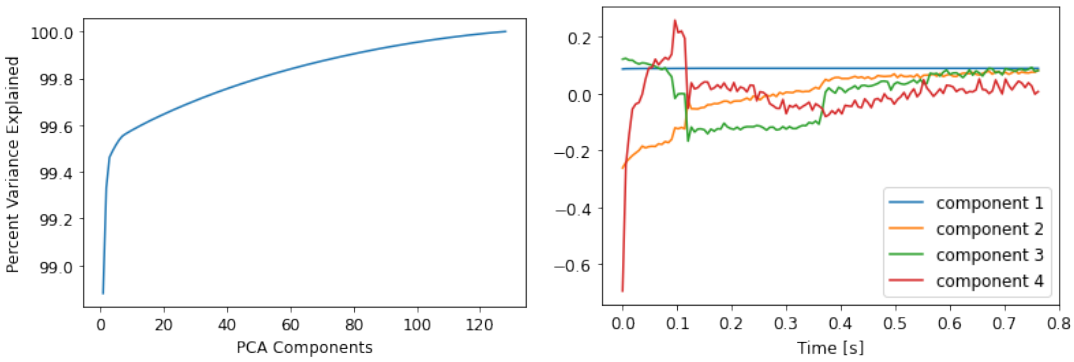

Fig. 1: Left: percentage of variance explained given number of principal components used in standardized samples of 1% of lattice points. Right: plots of first 4 principal components as a function of time.

At a high level, until some fixed fraction of lattice points have been sampled, RIGS Sampling does the following: (1) for each unsampled lattice point (x_i, y_j, V_k), compute a sampling priority based upon its distance to nearby sampled points and estimated gradients (with respect to x, y, and V) of the piezoelectric response at that point; (2) from the priorities assigned, identify lattice points (x, y, V) at which to to measure $p(t; x, y, V)$; then (3) take measurements up to a time t_{n_t} and approximate $p(t; x, y, V)$ by a constant $\hat{p}(x, y, V)$. Periodically, for use in computing the gradients, and at the end when all sampling has been done, a full reconstruction is done to generate $\widehat{\mathcal{E}}$ with a nearest-neighbor regression either by (a) using the time-constant approximates $\hat{p}$ (always done in intermediate steps and can be done as the final step) or (b) using the full time series at all sampled lattice points and reconstructing each fixed time slice separately (only as final step). We shall refer to (a) as a constant reconstruction and (b) as a full-time reconstruction.

3 Approach

`Python` code implementing our sampling scheme is available **at this Github repository** [11]. Algorithm 1 provides the full details of RIGS Sampling and makes use of the helper routines specified by Algorithms 2 and 3. In Algorithm 2, we wish to score how isolated a point is from other points that have been sampled. To do so, we ascribe a proxy for local density at each point by summing radially symmetric, decaying kernel functions centered at each sampled point and its reflections (to be explained later in this paragraph). Then unsampled points with the largest inverse density are deemed more isolated and should be sampled with higher priority. We find that with using the sampled points alone, points on the boundary will be scored as being more isolated even when they are not since they have fewer neighbors. To address this, we add reflection points in computing the density, i.e., to compute the density proxy, we use not only the sampled points but also their reflections $\mathcal{R}$ across each face of the cube. When gradient weight is not considered, this tends to give reasonably uniform samplings of points as can be seen in Figure 2. For simplicity, our kernel function is $k(u, s) = \mathrm{e}^{-|u-s|^2}$ but in the future we could consider choosing kernels based on Kernel Density Estimation [5]. In Algorithm 3, we score points based on how large $|\nabla \hat{p}|$ (computed in (x, y, V)-space) is relative to other points.

We note here that more sophisticated methods, such as those often used in image processing, could be used in place of a simple weighted average of nearest neighbors. However, we found such methods, like the Biharmonic Inpainting [2] approach, to be much slower while yielding comparable results. As another comment, it seems quite natural that instead of using constant approximations, we could instead compute the principal component representation of $p(t; x, y, V)$ at a sample of (x, y, V)-points and then use the first few principal components to reconstruct the piezoelectric response (with suitable adaptations such as using only n_t values to estimate the projections onto the components). This is especially appealing given that higher principal components have a nontrivial time-dependence; however, with regards to reconstruction errors, we did not find better performance in doing so and thus opt for the more simple approach. As a result, our method uses constant approximations, which is more or less equivalent to using only the first principal component.

Algorithm 1 Relative Isolation and Localized Gradient (RIGS)

|**Input**| the averaging window width n_t; an initially empty list of position-voltage points that have been sampled; a list of unsampled position-voltage points $\mathcal{U}$, initially including all lattice points; a gradient weight $0 \leq w \leq 1$; a sampling fraction threshold $0 < f \leq 1$; a reconstruction frequency n_r; a number of nearest neighbors for image reconstruction n_n; and whether or not to do a full time reconstruction, $full$.
$step = 0$
$m = (x_{round((N_x-1)/2)}, y_{round((N_y-1)/2)}, (V_{round((N_V-1)/2)})$, $\mathcal{S} = \{m\}$, $\mathcal{U} \leftarrow \mathcal{U} \setminus \{m\}$
Measure $p(t_0; m), p(t_1; m), ..., p(t_{n_t-1}; m)$ and set $\hat{p}(m) = \frac{1}{n_t} \sum_{j=0}^{n_t-1} p(t_j; u)$
while $|\mathcal{S}|/(|\mathcal{S}| + |\mathcal{U}|) < f$. **do**
  if $step$ is a multiple of n_r **then**
    Compute a reconstruction $\widehat{\mathcal{E}}$ over the entire lattice using the n_n nearest neighbors applied to the points in $\mathcal{S}$ with their constants $\{\hat{p}(s) | s \in \mathcal{S}\}$ using `Python`'s `KNeighborsRegressor` from the `sklearn` module.
    Compute the gradient score at each point according to Algorithm 3.
  end if
  Make a list of sampling priorities $\mathcal{P}$ corresponding to each unsampled point with scores given by $wg + (1-w)i$ where $g \in \mathcal{G}$ is the corresponding gradient score of the unsampled point and $i \in \mathcal{I}$ is the corresponding isolation score of the unsampled point — see Algorithms 2 and 3.
  At the $u \in \mathcal{U}$ with highest priority, measure $p(t_0; u), p(t_1; u), ..., p(t_{n_t-1}; u)$ and approximate the piezoelectric response at u by $\hat{p}(u) = \frac{1}{n_t} \sum_{j=0}^{n_t-1} p(t_j; u)$.
  $U \leftarrow \mathcal{U} \setminus \{m\}$, $\mathcal{S} \leftarrow \mathcal{S} \cup \{m\}$.
  Run Algorithm 2.
  $step \leftarrow step + 1$.
end while
if $full == true$ **and** $n_t == N_t$ **then**
  Compute the full-time reconstruction $\widehat{\mathcal{E}}$ over the entire lattice using the nearest neighbor algorithm as before in each time slice separately.
else
  Compute the constant reconstruction $\widehat{\mathcal{E}}$ over the entire lattice using the nearest neighbor algorithm as before using a constant-in-time approximation.
end if
|**Output**| $\widehat{\mathcal{E}}$

Algorithm 2 Density Algorithm: each unsampled point is ascribed an isolation score on $[0, 1]$. The set of reflecting functions $\mathcal{R}$ reflect sampled points across all 6-boundaries of the lattice.

|**Input**| a list of sampled lattice points $\mathcal{S}$, a list of unsampled lattice points $\mathcal{U}$ with corresponding isolation scores $\mathcal{I}$ all initially 0, a symmetric $k(u, s) = K(|u - s|)$ a decreasing function, and reflecting functions $\mathcal{R}$.
for all $u = (x, y, V)$ in $\mathcal{U}$ **do**
  $density = 0$
  for all $s = (x', y', V')$ in $\mathcal{S}$ **do**
    $density \leftarrow density + K(u, s) + \sum_{r \in \mathcal{R}} K(u, r(s))$.
  end for
  Set isolation score in $\mathcal{I}$ corresponding to u to $1/density$.
end for
$TotalIsolationSqd = 0$
for all i in $\mathcal{I}$ **do**
  $TotalIsolationSqd \leftarrow TotalIsolationSqd + i^2$
end for
for all i in $\mathcal{I}$ **do**
  $i \leftarrow i/\sqrt{TotalIsolationSqd}$
end for
|**Output**| $\mathcal{I}$

Algorithm 3 Gradient Algorithm: each point is ascribed a gradient score on $[0, 1]$.

|**Input**| a list $\mathcal{L}$ of all lattice points with corresponding gradient scores $\mathcal{G}$ all initially set to 0.
for all ℓ in $\mathcal{L}$ **do**
    Assign the corresponding gradient score of $\mathcal{G}$ to be the magnitude of the gradient at ℓ using `Python`'s `gradient` from the `numpy` module.
end for
$TotalGradSqd = 0$
for all g in $\mathcal{G}$ **do**
    $TotalGradSqd \leftarrow TotalGradSqd + g^2$
end for
for all g in $\mathcal{G}$ **do**
    $g \leftarrow g/\sqrt{TotalGradSqd}$
end for
|**Output**| $\mathcal{G}$

4 Results

Here we focus on investigating how the method performs at reconstructing the complete dataset $\mathcal{E}$ by sampling only a subset of lattice points or time series over a shorter duration. In sequence: we begin by observing how the gradient weight affects the lattice points being sampled; then, we study how the constant reconstructions fare under the choice of various hyperparameters; finally, we repeat the previous line of investigation for the full-time reconstruction.

We first see that as the relative weight ascribed to the gradient increases, the set of sampled points chosen changes from almost uniform to highly focused on areas with rapid changes in the piezoelectric response (see Figure 2). We speculate some of the circular sampling patterns with $w = 1$ stem from the gradients only being estimated every n_r iterations. Qualitatively, we also note that for a fixed time slice, a constant reconstruction accurately captures key features of the piezoelectric response. In Figure 3, we see that with only $n_t = 30$ time points used and by only sampling a fraction $f = 0.4$ of all lattice points, the constant approximation we obtain accurately depicts the response at time slice 64, which is beyond the range of times used in sampling. In particular we note that the reconstruction captures sharp edges from the original image quite well.

Among the hyperparameters we deem most important, we wish to study how the overall reconstruction error (Eq. 1) varies as we change the gradient weight w, number of time samples used n_t, and fraction of lattice points sampled f. We wish to identify which combinations of hyperparamameters reduce the reconstruction error to within various tolerances. In the case of using constant reconstructions, Figure 4 shows these results. With $w = 0.8$, $n_t = 50$ time points, and $f = 0.6$ lattice points sampled, the relative error is less than 15%, corresponding to measuring only 23% of all $N_x N_y N_V N_t$ possible lattice-time points. However, this fails to capture time-dynamics as the responses are constant in time for each fixed (x, y, V).

In Figure 5 and Figure 6 we consider the full-time reconstruction, which yields our highest quality results. Figure 5 shows the relative ℓ_2-loss for various fractions of points sampled and gradient weights. Unsurprisingly, the quality of the reconstruction strictly improves with the fraction of sampled points. More interestingly, we often observe improvement as the gradient weight w increases. This implies that our algorithm benefits from adaptivity. When $w = 0$ our algorithm samples non-adaptively, while adaptivity plays a progressively greater role as w increases. When only a small fraction of points are sampled, setting w to be too large damages the reconstruction accuracy. We suspect this is due to our algorithm leaving large low-gradient patches almost completely unsampled. For large fractions of sampled points f this effect is less significant as it is difficult to avoid sampling large patches. Notably we achieve less than 10 percent error by sampling only 53 percent of points, and achieve less than 5 percent error by sampling 75 percent of points. In addition Figure 6 shows that reconstructing time slices individually allows our reconstruction to capture the piezoelectric response as a function of time.

As a final remark, a `Python` implementation of RIGS Sampling for this dataset takes on average 2.7 ms per iteration with $n_r = 2000$ and $n_n = 4$, making it a viable choice in computing the optimal subsequent lattice sampling points.

5 Contributions

Our method is capable of reducing the number of experimental measurements necessary to describe the piezoelectric response measured at various positions and voltages. Notably, we have shown that in the example data, we can reconstruct the full dataset to within a relative ℓ_2-error of 10% with 47% fewer measurements than done to obtain the

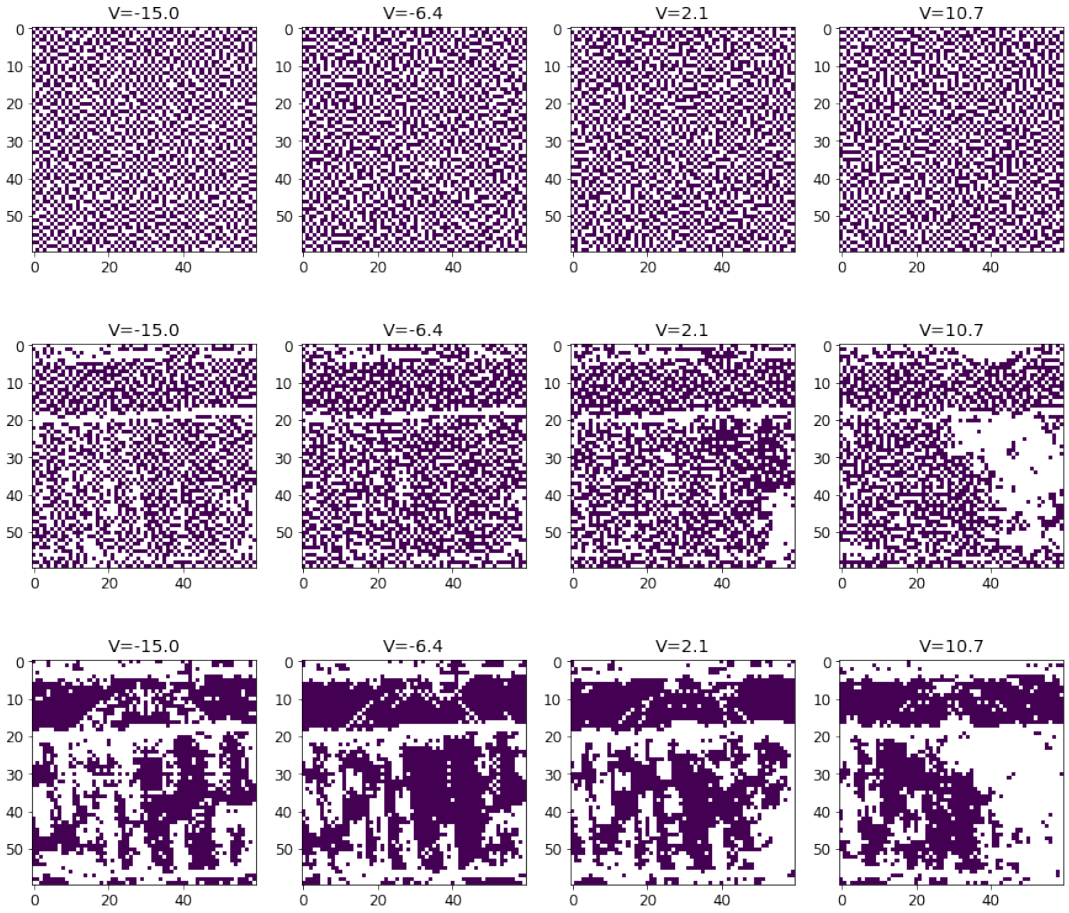

Fig. 2: The set of set of sampled points (shown in white) generated by sampling the first time series component ($n_t = 1$) from 50 percent of grid points, using gradient weights 0 (top row), 0.5 (middle row), 1.0 (bottom row). The dataset included 16 different voltages but for compactness only 4 are shown.

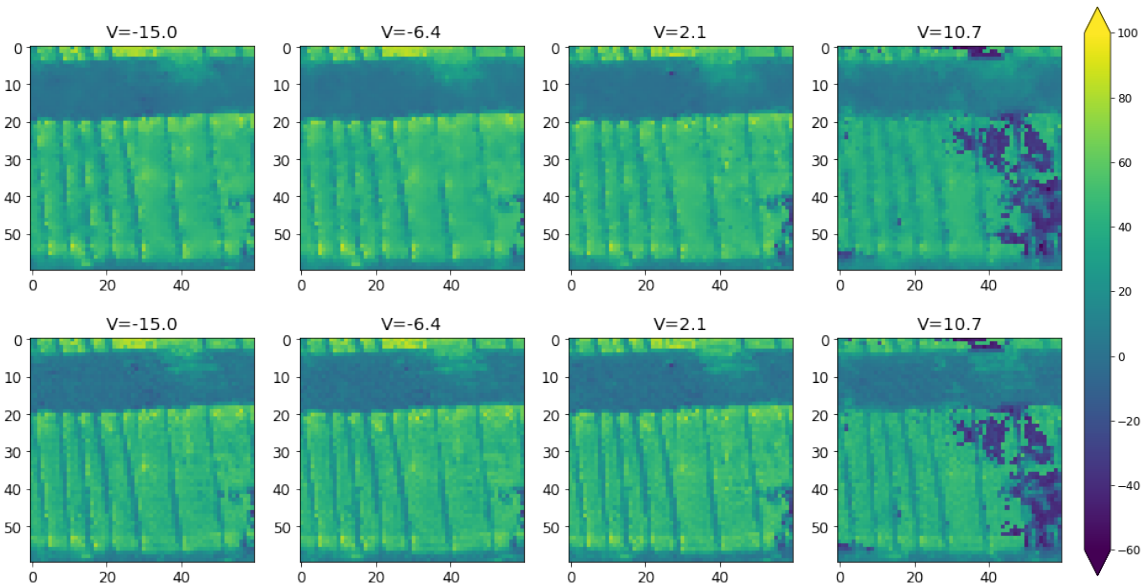

Fig. 3: Top row: The reconstructed piezoelectric response at time slice 64 using the first 30 time series points, sampling from 40 percent of the lattice, and with a gradient weight of 0.81. Bottom row: The true piezoelectric response at time slice 64.

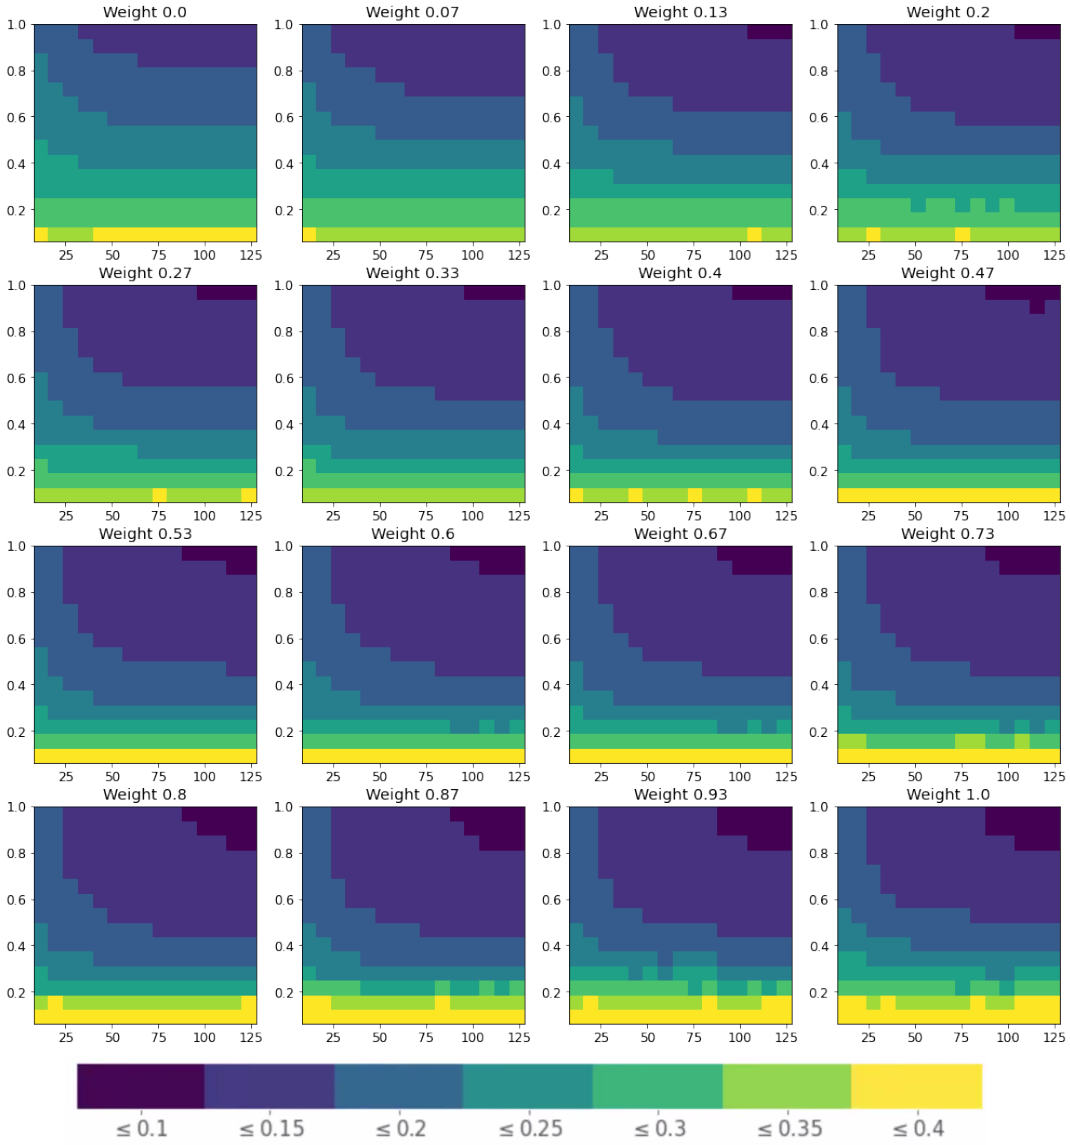

Fig. 4: Relative ℓ_2-error of the constant reconstructions as a function of the gradient weight w, the fraction of sampled points f (y-axes), and the number of time components n_t used (x-axes).

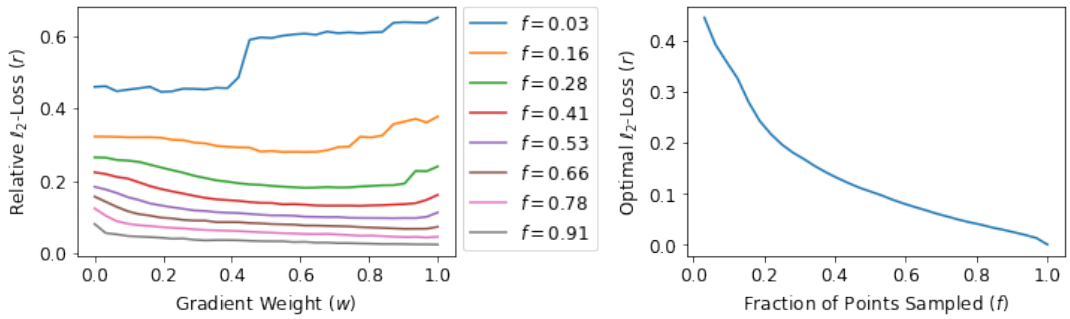

Fig. 5: Relative ℓ_2-error of full-time reconstructions as gradient weight w and sampling fraction f vary (left) along with the optimal reconstruction error attainable for a given sampling fraction f (right).

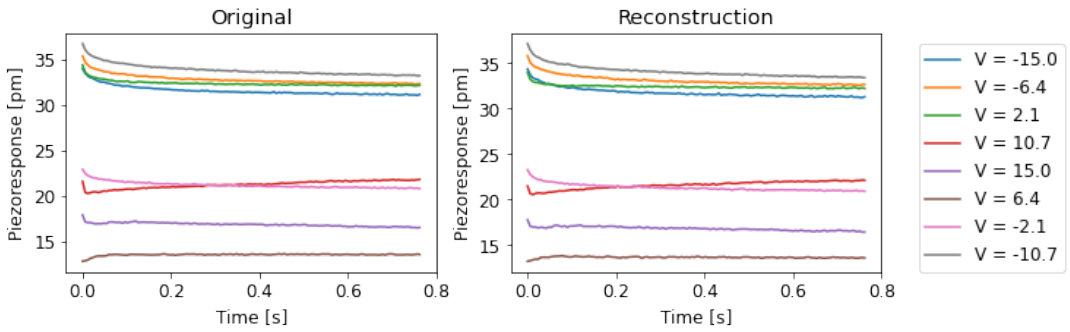

Fig. 6: Plots of the experimental and reconstructed piezoelectric responses vs time at 50 percent sampling, averaged over all (x, y)-lattice values at various voltages.

data. Additionally, we demonstrate that adaptive measurements can significantly reduce the measurement capacity required for piezoelectric microscopy, as compared to non-adaptive measurements.

Our work demonstrates that active learning based on gradient information is beneficial and that time series may have a low dimensional representation (constant in this case), which could be useful for other large multidimensional datasets.

Acknowledgments

WJS and DN acknowledge funding from NSF BIGDATA #1740325 and NSF DMS #2011140.

References

1. CHANG, H., Zhiwen, Y., Zhiyong, Y., Qi, A., Bin, G.: Air quality estimation based on active learning and kriging interpolation. Big Data Research **4**(6) (2018) 2018061
2. Damelin, S.B., Hoang, N.: On surface completion and image inpainting by biharmonic functions: Numerical aspects. International Journal of Mathematics and Mathematical Sciences **2018** (2018)
3. Fix, E., Hodges, J.L.: Discriminatory analysis. nonparametric discrimination: Consistency properties. International Statistical Review/Revue Internationale de Statistique **57**(3) (1989) 238–247
4. Hu, Y., Zhang, D., Jin, Z., Cai, D., He, X.: Active learning via neighborhood reconstruction. In: Proceedings of the Twenty-Third international joint conference on Artificial Intelligence, Citeseer (2013) 1415–1421
5. Hyndman, R.L., Zhang, X., King, M.L., et al.: Bandwidth selection for multivariate kernel density estimation using mcmc. In: Econometric Society 2004 Australasian Meetings. Number 120, Econometric Society (2004)
6. Jolliffe, I.T., Cadima, J.: Principal component analysis: a review and recent developments. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **374**(2065) (2016) 20150202
7. Kelley, K.P., Li, L., Ren, Y., Ehara, Y., Funakubo, H., Somnath, S., Jesse, S., Cao, Y., Kannan, R., Vasudevan, R.K., et al.: Tensor factorization for elucidating mechanisms of piezoresponse relaxation via dynamic piezoresponse force spectroscopy. npj Computational Materials **6**(1) (2020) 1–8
8. Kiselev, D., Bdikin, I., Selezneva, E., Bormanis, K., Sternberg, A., Kholkin, A.: Grain size effect and local disorder in polycrystalline relaxors via scanning probe microscopy. Journal of Physics D: Applied Physics **40**(22) (2007) 7109
9. Rugar, D., Hansma, P.: Atomic force microscopy. Physics today **43**(10) (1990) 23–30
10. Settles, B.: Active learning literature survey. (2009)
11. Swartworth, W., Lindstrom, M., Needell, D.: RIGS Code. https://github.com/wswartworth/active_imaging
12. Vasudevan, R., Kelley, K., Jesse, S., Kalinin, S., Ziatdinov, M.: High dimensional active learning for microscopy of nanoscale materials
13. Xu, T.B.: Energy harvesting using piezoelectric materials in aerospace structures. In: Structural Health Monitoring (SHM) in Aerospace Structures. Elsevier (2016) 175–212