

SEQUENTIAL ADAPTIVE DETECTION FOR IN-SITU TRANSMISSION ELECTRON MICROSCOPY (TEM)

Y. Cao^{*} S. Zhu^{*} Y. Xie^{**} J. Key[†] J. Kacher^{††} R. R. Unocic[‡] C. M. Rouleau^{‡‡}

^{*} School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, USA.

[†] School of Material Science and Engineering, Georgia Institute of Technology, Atlanta, GA, USA.

[‡] Oak Ridge National Laboratory, Oak Ridge, TN, USA.

ABSTRACT

We develop new efficient online algorithms for detecting transient sparse signals in TEM video sequences, by adopting the recently developed framework for sequential detection jointly with online convex optimization [1]. We cast the problem as detecting an unknown sparse mean shift of Gaussian observations, and develop adaptive CUSUM and adaptive SSRS procedures, which are based on likelihood ratio statistics with post-change mean vector being online maximum likelihood estimators with ℓ_1 . We demonstrate the meritorious performance of our algorithms for TEM imaging using real data.

Index Terms— Sequential detection, online algorithms, microscopy imaging

1. INTRODUCTION

TEM (Transmission Electron Microscopy) has long been a powerful tool for imaging material structure and characterizing material chemistry. However, the process to resolve structural features is laborious and time intensive, drastically limiting the characterization throughput. Recent advances in electron detector technology and computational capacity have facilitated the development of high-speed data collection with microsecond frame rate acquisition speeds. This advance in TEM technology has enabled new paradigms in data collection. Because of this, *in-situ* processing of the real-time collected data to detect emerging features become a highly desired property for the new TEM system. Currently, the data are captured real-time but analyzed off-line, limiting the experimentalist's ability to explore in detail regions of interest while at the microscope.

Sequential change-point detection that can be adaptive to data can revolutionize this process. The signal detection task in TEM has two characteristics. First, each observation is

a very high-dimensional vector so we need develop an algorithm that can handle a large amount of data sequentially. Second, the change is *sparse* in the sense that among all the parameters only a small proportion of them changes after the unknown change-point.

In this paper, we present a sequential adaptive change detection method for in-situ TEM signal detection. The method is developed by adapting the recent one-sample update based sequential detector in [1], by assuming Gaussian observations and the signal being a sparse mean shift to the Gaussian. Our method can precisely control false alarms and can be computed recursively, and thus automate the detection in real-time. We demonstrate meritorious performance of our methods for TEM imaging using real data.

Compared to the classic CUSUM procedure (see, e.g., [2]), which needs to pre-specify a post-change mean parameter and its performance can be impacted when there is parameter misspecification, our adaptive procedure is more robust since the mean is updated with sequential data. Compared to the classic generalized likelihood ratio (GLR) procedure when the plug-in estimators are exact maximal likelihood estimators (MLE), our method is much faster and memory efficient since our plug-in estimators are computed recursively with one-sample update (thus raw data needs not to be stored) using an online convex optimization algorithm.

2. SEQUENTIAL ADAPTIVE DETECTION

In this section, we present our two sequential adaptive detection algorithms, which are adapted from the one-sample update scheme in [1]. Assume a sequence of d -dimensional observations $X_1, X_2, \dots$ which are i.i.d. random variables from a multivariate normal distribution $\mathcal{N}(\theta, I_d)$ with *unknown* mean parameter $\theta \in \mathbb{R}^d$. We will estimate θ online to be adaptive.

Consider the sequential change-point detection problem that the underlying distribution of the data changes from a known state to an unknown state after at an unknown change-point ν . Without loss of generality, we assume that the pre-change mean is an all-zero vector. The post-change mean is

^{*}Corresponding author: yao.xie@isye.gatech.edu

[†]The work is sponsored by a Georgia Tech IMAT Seed Grant.

[‡]The TEM corrosion experiments were conducted at the Center for Nanophase Materials Sciences, which is a DOE Office of Science User Facility.




Fig. 1. Diagrams demonstrating the basic principles of bright field imaging in TEM. TEM can operate in two modes, illustrated in Left Panel: in the real space; Right Panel: in the diffraction space. The real space images can be computed from the diffraction space images. The incident beam of electrons passes through the sample and a lens/aperture system is used to form the image. Our algorithm can be applied to both type of data: the image sequence data and the diffraction pattern sequence data.

unknown and belong to a set $\mathcal{A}$ defined as $\mathcal{A} = \{\theta : \|\theta\|_0 \leq s\}$, where $\|\cdot\|_0$ is the number of non-zero entries of θ and s is a prescribed value to characterize the sparsity. Formally, we consider the following hypothesis test:

$$\begin{aligned} H_0 : X_1, X_2, \dots &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d), \\ H_1 : X_1, \dots, X_\nu &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, I_d), \\ X_{\nu+1}, X_{\nu+2}, \dots &\stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\theta, I_d), \quad \theta \in \mathcal{A}. \end{aligned} \quad (1)$$

The goal is to detect the change as quickly as possible after it occurs under the false alarm constraint. We will consider likelihood ratio based detection procedures which we call the adaptive CUSUM (ACM), and the adaptive SRRS (ASR) procedures, respectively.

Now we derive the detection statistics. For each putative change-point location k before the current time t , the post-change samples are $\{X_k, \dots, X_t\}$, and the post-change parameter is estimated as

$$\hat{\theta}_{k,i} = \hat{\theta}_{k,i}(X_k, \dots, X_i), \quad i \geq k. \quad (2)$$

Denote f_θ as the density function for $\mathcal{N}(\theta, I_d)$. The likelihood ratio at time t for a hypothetical change-point location k is given by (initialized with $\hat{\theta}_{k,k-1} = \theta_0$)

$$\Lambda_{k,t} = \prod_{i=k}^t \frac{f_{\hat{\theta}_{k,i-1}}(X_i)}{f_0(X_i)}, \quad (3)$$

where $\Lambda_{k,t}$ can be computed recursively since

$$\Lambda_{k,t} = \Lambda_{k,t-1} \cdot \frac{f_{\hat{\theta}_{k,t-1}}(X_t)}{f_0(X_t)}.$$

Since the change-point location ν is unknown, due to the maximum likelihood principle, we take the maximum of the

statistics over all possible values of k . We consider window-limited versions [3] to avoid infinite memory, by taking the maximum over $k \in [t-w, t]$, where w is a prescribed window size. This leads to the ACM procedure

$$T_{\text{ACM}}(b) = \inf \left\{ t \geq 1 : \max_{t-w \leq k \leq t} \log \Lambda_{k,t} > b \right\}, \quad (4)$$

where b is a pre-specified threshold.

The Shiryaev-Roberts (SR) procedure replaces the maximization over k in (4) with summation, which can be justified from a Bayesian prior assumption. By following the same strategy, we obtain the following ASR procedure [4]:

$$T_{\text{ASR}}(b) = \inf \left\{ t \geq 1 : \log \left(\sum_{k=t-w}^t \Lambda_{k,t} \right) > b \right\}, \quad (5)$$

where b is a pre-specified threshold. As shown in [1], the performance of the ACM and the ASR is very similar. However, the likelihood ratio in (3) can explode when d is very large. Thus, in practice we prefer to use the ACM procedure to avoid possible numerical issues.

The detection statistic relies on a sequence $\{\hat{\theta}_{k,t}\}$ of estimators constructed using online mirror descent (OMD). The main idea of OMD is that, at each time step, for any k , the estimator $\hat{\theta}_{k,t-1}$ is updated using the new sample X_t , by balancing the tendency to stay close to the previous estimate against the tendency to move in the direction of the greatest local decrease of the loss function. The advantages of OMD are (1) it allows a simple *one-sample update*: the update from $\hat{\theta}_{k,t-1}$ to $\hat{\theta}_{k,t}$ only uses the current sample X_t , and the update for the detection statistic has a simple recursive scheme. This is the main difference from the traditional GLR statistic [5] where each $\hat{\theta}_{k,t}$ is the exact MLE estimated using all the historical samples. (2) OMD is a generic algorithm for solving the online convex optimization (OCO) problem [6]. In [1], it is proven that even these approximate MLE schemes have very little statistical efficiency.

Here, we adapt the general ACM and ASR for exponential family distributions in [1] to the case when the signal is a sparse Gaussian mean shift, and set the constraint set for the unknown parameter to be $\Gamma = \{\theta : \|\theta\|_1 \leq s\}$ (refer to [7] for an automatic choice of s) as a convex relaxation of the non-convex set $\mathcal{A}$. Denote $\|\cdot\|_1$ and $\|\cdot\|_2$ as the ℓ_1 and ℓ_2 norms in the Euclidean space, respectively. The algorithms are summarized in Algorithm 1. The projection (step 7) onto ℓ_1 ball can be obtained via simple soft-thresholding [8].

3. RESULTS ON REAL-DATA

We test our methods on two TEM datasets: one consists of a sequence of real space images, and another one consists of a sequence of diffraction space images. The experimental setups are the exact same for the two datasets. The only difference is that we change the lens setting to collect diffraction

Algorithm 1 Online mirror-descent (OMD) for $\{\hat{\theta}_{k,t}\}$

Require: A sequence of data $X_k, \dots \in \mathbb{R}^d$; a closed and convex set $\Gamma \subset \mathbb{R}^d$ of the parameters; a decreasing sequence $\{\eta_t\}_{t \geq 1}$ of strictly positive step-sizes.

- 1: $\hat{\theta}_{k,k-1} = 0, \Lambda_{k,k-1} = 1$. {Initialization}
 - 2: **for all** $t = k, k+1, \dots$, **do**
 - 3: Acquire a new observation X_t
 - 4: Compute loss $\ell_t(\hat{\theta}_{k,t-1}) := \|\hat{\theta}_{k,t-1}\|_2^2/2 - \hat{\theta}_{k,t-1}^\top X_t$
 - 5: Compute $\Lambda_{k,t} = \Lambda_{k,t-1} \times f_{\hat{\theta}_{k,t-1}}(X_t)/f_0(X_t)$
 - 6: $\tilde{\theta}_{k,t} = \hat{\theta}_{k,t-1} - \eta_t(\hat{\theta}_{k,t-1} - X_t)$ {Dual update}
 - 7: $\hat{\theta}_{k,t} = \arg \min_{u \in \Gamma} \|u - \tilde{\theta}_{k,t}\|_2$ {Projected primal update}
 - 8: **end for**
 - 9: **return** $\{\hat{\theta}_{k,t}\}_{t \geq 1}$ and $\{\Lambda_{k,t}\}_{t \geq 1}$.
-

patterns instead of image-space images. We will develop different preprocess steps for these two datasets due to their different characteristics. After the preprocessing, we show that both become detecting a sparse signal in Gaussian noise.

Experiment set-up. The data is a sequence of metal corrosion images captured using bright-field transmission electron microscopy (TEM). The experiment setup is as follows. Iron thin films were sputtered at room temperature onto silicon nitride membranes compatible with an in-situ TEM liquid cell holder. 20 vol% acetic acid was introduced to the system to initiate corrosion. Imaging was performed using an FEI Titan at 300 kV with a Gatan OneView camera in either real space or diffraction space. The time-resolved diffraction patterns provide information on the formation of corrosion bi-products and the dissolution of crystalline material. For illustration purposes, we first select 23 gray images (2 images per second) in the bright-field TEM image sequence and downsize each image to 308-by-308 pixels. At some time point, corrosion initiates in the image sequence, which is emphasized by the red circle in Figure 2.




Fig. 2. A sequence of metal corrosion images captured using bright-field TEM. The time (index for the image in the sequence) is labeled; the corrosion initiates at time $t = 8$ (marked by the red circle) and develops over time. The corroded area has a higher intensity signal than the rest of the film.

We apply our ACM and ASR procedures, choosing $w = 200$ as the prescribed window size and setting the threshold b for detection procedures by simulation such that

the false alarm rate (the average-run-length, ARL, which expected number of observations between two false alarms) is about 10000.

3.1. Detection for real space images

Preprocessing. First, vectorize each image into a vector of dimension $308 \times 308 = 94864$. For each pixel, we take their value in the first 5 frames as the training data to compute the mean and standard deviation. Then we standardize samples for each pixel by subtracting the mean and dividing the standard deviation. We ignore the correlation between the pixels for this example and it turns out to be a good approximation.

After the pre-processing, we have a total of 23 such vectors: $X_1, \dots, X_{23} \in \mathbb{R}^d$, with $d = 94864$. The pre-change distribution is $\mathcal{N}(0, I_d)$ and the goal is to detect the unknown time ν at which the underlying distribution changes to $\mathcal{N}(\theta, I_d)$ for some unknown mean $\theta \neq 0$. We assume that on average each pixel has a unit shift after the standardization so we set $s = 10^5$ in Γ for our methods. We compare our algorithms with the standard multivariate CUSUM procedure [9] (the post-change mean parameter is set to be an all-one vector), and the GLR procedure [3].

Results. $T_{\text{ACM}}(b)$ and $T_{\text{ASR}}(b)$ with $\Gamma = \{\theta : \|\theta\|_1 \leq 10^5\}$ both stop at time $t = 8$, CUSUM procedure with an all-one post-change mean vector stops at time $t = 9$ and GLR procedure stops at time $t = 6$. Since we see from Figure 2 that the change happens at time $t = 8$, the detection delays of our methods are 0 (meaning it only takes one sample to detect the corrosion spot) while that of CUSUM procedure is 1 (meaning it takes two samples to detect). The GLR procedure raises a false alarm since it stops when there is no change. The possible reason for the GLR raising the false alarm is that the GLR is more easily affected by noise. This also shows that GLR procedure performs better in ideal case such as synthetic signals but may not perform well in practice.

3.2. Detection for diffraction space images.

Preprocessing. To detect weak signal (“sparse spot”) in diffraction image space, we need to effectively remove the background of the image since the change of interest will be tiny bright spot buried in between the bright rings. We develop a set of preprocessing steps that are tailored to the characteristics of the diffraction images. The most important part for the analysis is to remove the largest visible ring (neither the bright area near the center nor the dark area near the boundary of the image). This task is nontrivial for the following two reasons. First, there is a dark shadow of a irregularly shaped stick in the middle of the image so the center of the rings is hidden. Second, even if we find the center and remove the shadow we still need identify the the ring with the largest radius.

To overcome the first difficulty, we draw the histogram of the pixel values as shown in the left figure in Fig. 3. Sev-




Fig. 3. Left: a diffraction domain image; Middle: the histogram of the intensity; Right: The thresholded image if we only keep pixels of value in $(0, 0.2)$.

erall gaps between the brightness is observed in the histogram. For example, there is no point in the image with the brightness centered around 0.4, 0.6 and 0.8. In fact, these gaps are highly related to the rings in the image. The right plot in Fig. 3 shows the points with brightness between 0 and 0.1 and we surprisingly find the shadow. Therefore, we can throw away the points with that brightness in order to remove the stick in the original image. Then, Fig. 4 shows that we can separate the rings successfully by focusing on the points with separated ranges of the brightness.




Fig. 4. Background removal: we threshold a diffraction space image with different range of threshold values, and this yields rings at different radii. These concentric rings help to estimate their common center, and subsequently we subtract off the bright rights to remove these bright rings.

To overcome the second difficulties, we use the Hough transformation [10] to look for the centers. We run the Hough transformation on all the plots in Fig. 4 and then compute the center by averaging the 8 estimated centers. To find the largest visible rings, we apply the Canny edge detection algorithm [11] that identifies the boundary between the dark and bright area accurately. The final results after all the preprocessing procedures are shown in Fig. 5. Note that the change - a tiny bright spot can finally be revealed.

The bright spot is very weak and it is even hard to be observed by eyes. Fortunately, we know by domain knowledge that the bright spot usually appear near certain radius (but at an unknown angle). Therefore, we can “hunt” the bright spot around such radius. We focus our attention on the points with




Fig. 5. The diffraction image after preprocessing for background removal. The complete video is available at www.isye.gatech.edu/~yxie77/diffraction-video.mp4. The middle image is the 17th image of the 100 images. We zoom out the middle image to show the bright spot which represents the anomaly we would like to detect.

a specified radius r that is slightly larger than the radius of the ring. We then formulate the detection task as the detection of a sparse mean shift. Using polar coordinate transformation, for each preprocessed image we observe a 360-dimensional signal that represents the averaged pixel values in every angle for a fixed radius. The results are shown in Fig. 6. We can see clearly that one bright spot appears from about the 17th image for the angle equal to 171, and another bright spot appears from about the 49th image for the angle equal to 153. After the pre-processing, we have a total of 100 such vectors: $X_1, \dots, X_{100} \in \mathbb{R}^d$, with $d = 360$. The pre-change distribution is $\mathcal{N}(0, I_d)$ and the goal is to detect the unknown time ν at which the underlying distribution changes to $\mathcal{N}(\theta, I_d)$ for some unknown mean $\theta \neq 0$. In this example we just set $\Gamma = \mathbb{R}^d$.




Fig. 6. The extracted signals for the sequence of 100 images (selected angles).

Results. $T_{ACM}(b)$ and $T_{ASR}(b)$ with $\Gamma = \mathbb{R}^d$ both stop at time $t = 18$, CUSUM procedure with an all-one post-change mean vector stops at time $t = 24$ and GLR procedure stops at time $t = 4$. Domain knowledge tells us that the change happens at time $t = 17$. So the detection delays of our methods are 1 while that of CUSUM procedure is 7 (meaning it takes two samples to detect). The GLR procedure raises a false alarm because it is too sensitive to the noise.

4. REFERENCES

- [1] Yang Cao, Liyan Xie, Yao Xie, and Huan Xu, “Sequential change-point detection via online convex optimization,” *Entropy*, vol. 20, no. 2, pp. 108, 2018.
- [2] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*, CRC Press, 2014.
- [3] Alan Willsky and H Jones, “A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems,” *IEEE Transactions on Automatic control*, vol. 21, no. 1, pp. 108–112, 1976.
- [4] G. Lorden and M. Pollak, “Nonanticipating estimation applied to sequential analysis and changepoint detection,” *Annals of statistics*, pp. 1422–1454, 2005.
- [5] T.-Z. Lai, “Information bounds and quick detection of parameter changes in stochastic systems,” *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2917–2929, 1998.
- [6] E. Hazan, “Introduction to online convex optimization,” *Foundations and Trends in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.
- [7] Jordan Frecon, Nelly Pustelnik, Nicolas Dobigeon, Herwig Wendt, and Patrice Abry, “Bayesian-driven criterion to automatically select the regularization parameter in the ℓ_1 -potts model,” in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3839–3843.
- [8] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, “Efficient projections onto the ℓ_1 -ball for learning in high dimensions,” in *International Conference on Machine learning (ICML)*. ACM, 2008, pp. 272–279.
- [9] William H Woodall and Matoteng M Ncube, “Multivariate cusum quality-control procedures,” *Technometrics*, vol. 27, no. 3, pp. 285–292, 1985.
- [10] Hough Paul VC, “Method and means for recognizing complex patterns,” Dec. 18 1962, US Patent 3,069,654.
- [11] John Canny, “A computational approach to edge detection,” *IEEE Transactions on pattern analysis and machine intelligence*, , no. 6, pp. 679–698, 1986.