A MobileViT-based Detection of Anomaly in Temperature of Nuclear Power Plant Core

Rémi Cogranne

A MobileViT-based Detection of Anomaly in Temperature of Nuclear Power Plant Core

Rémi Cogranne[a]

[a]Troyes University of Technology, Troyes, France

ABSTRACT

This paper presents a simple model based on MobileViT-v2 for temperature monitoring within a nuclear power plant. Specifically, we propose utilizing MobileNet-v2 to detect a critical accident: a total and instantaneous blockage. We model the temperature effects of such an event and train a MobileViT-v2 model for detection. The trained classifier's results are then used in a sequential procedure to detect blockage as quickly and reliably as possible. We compare the performance of two sequential detection methods, namely sliding window and CUSUM, in terms of mean detection delay and probability of detection before a prescribed maximum detection delay. Experimental results, using actual temperature measurements from the Superphénix power station, demonstrate the effectiveness of the proposed detection method.

Keywords: Detection, Sequential methods, CUSUM, Empirical evaluation, Artificial Inteligence, Image processing

1. INTRODUCTION

Critical infrastructure monitoring is a top priority due to the severe consequences of potential incidents. Most countries, including those in the EU and US, have established offices and norms for critical national infrastructure protection, such as the EU Critical Infrastructure Protection (EUCIP) and US Critical Infrastructure Protection (CIP) programs. Among critical infrastructure, nuclear power plants are particularly sensitive due to the significant impact of any malfunction. This is especially true for new reactor types, such as the Sodium-cooled Fast Reactor (SFR), which is one of six models selected by the Generation IV International Forum (GIF) for future nuclear power plants.[1,2] Despite rigorous development processes, safety and reliability can still be challenged by extremely unlikely events, such as core cooling incidents. One such critical incident is the Total and Instantaneous Blockage (TIB), which, although extremely unlikely, can lead to the melting of several reactor rods if undetected. This paper focuses on the early detection of TIB. Change detection in complex systems, including thermal incidents, has been extensively studied in various fields, such as finance, econometrics, and manufacturing. In this paper, we address the general problem of abrupt change detection within a complex critical infrastructure in real time. Real-time monitoring of the nuclear reactor's core temperature is essential for detecting overheating and identifying various incidents, some of which are less obvious. In this operational context, the highest priority is to ensure the reliability of the detection system, requiring a detector with well-established statistical properties that guarantee a low false alarm rate and quick detection with minimal delay. In this context, we propose an effective use of autoencoders to model the temperature of nuclear cores, capturing complex temporal dependencies with precision. This approach enables the integration of statistical detection theory, allowing for accurate blockage detection. We believe that the combination of these two approaches offers a significant contribution to possible applications in a wide range of problems. The present paper is organized as follows. Section 2 provides a brief overview of the specific thermal anomaly event we aim to detect, introduces the main difficulties and the operational constraints. Section 3 states the problem of the quickest detection of coolant flow blockage, considering the very weak heat anomaly at the location of the temperature probes. The original method we had designed is presented and explained in detail in Section 4 which describes both the MobileNet-v2 based model for the detection of blockage and its use for sequential detection methods. Then Section 5 present and analyse the numerical results for both the detection of the blockage and the sequential detection methods. Finally, the Section 6 concludes the paper.

Further author information: Send correspondence to remi.cogranne@utt.fr.

2. DESCRIPTION OF THE CONTEXT

The Total and Instantaneous Blockage (TIB) event is a critical scenario that the proposed method aims to detect. This section provides an overview of the underlying challenges associated with TIB and introduces the notation used throughout the paper.

Figure 1 illustrates the structure of a Sodium-cooled Fast Reactor (SFR), highlighting the key components involved in the TIB event. The core consists of nuclear fuel rods surrounded by a coolant, which transfers heat to the alternator to generate electricity. Temperature measurements are taken from the output coolant, providing an indirect indication of rod temperatures. However, direct measurement of individual rods is impractical.

A TIB event occurs when the coolant flow is suddenly and completely blocked, causing the actual rod temperature to increase rapidly while the measured temperature remains relatively constant. As the coolant flow ceases, the measured temperatures no longer accurately reflect the actual rod temperatures. Although the temperatures increase due to thermal conductivity, the effect is subtle and gradual. If left undetected, this can lead to nuclear core melting and potentially catastrophic consequences. A notable example of TIB occurred on October 5, 1966, at the FERMI-1 reactor in Michigan, USA. According to the report,[3] during power ascension, zirconium plates at the bottom of the reactor vessel became loose, blocking sodium coolant flow to some fuel subassemblies. Two subassemblies began to melt, prompting a manual shutdown of the reactor.

Detecting TIB poses significant challenges, primarily due to the following difficulties:

- Temperature modelling: Accurately modelling the temperature of a nuclear power plant is extremely challenging, as it is influenced by numerous factors, including control rods, power adjustments, and external conditions.
- Weak signal detection: The impact of total blockage on temperatures produces a very weak initial signal, making detection even more challenging.
- Reliability requirements: The detection method must meet stringent reliability criteria, ensuring that the detection delay is below a critical threshold and the false alarm rate is extremely low to avoid unnecessary emergency shutdown procedures.
- Quickest possible detection: When the total blockage occurs, it is crucial to detect this event as quickly as possible. More precisely, a maximal detection delay is generally prescribed after which fatal consequences are inevitable, such as the partial melting of the core of the nuclear power plant.

To address these challenges, a highly accurate model of the temperature of all rods under normal operating conditions is required. Additionally, the model must be computationally efficient to enable real-time monitoring. Furthermore, some parameters, such as outside temperature and humidity, are not precisely known, adding complexity to the problem. Last but not least, we must detect a weak signal as quickly as possible. To address all these difficulties, the present paper proposes an original method based on the MobileViT-v2 deep learning model used in a sequential detection procedure.

2.1 A Brief Review on the Quickest Detection Problem

The traditional approach to sequential detection focuses on minimizing the average detection delay. However, in this case, we aim to maximize detection under a strict maximal delay constraint. This problem is discussed in more detail in the present section.

Real-time monitoring of nuclear rod temperatures is crucial for detecting blockages. This sequential detection setting involves analysing data one-by-one as they are received. Let $T_{i;x,y}$ denote the temperature of the rod at location $(x, y) \in \mathcal{X}$ at time $i \in [1, \dots, I]$. The temperature surface of all p rods at time i is denoted by $\mathbf{T}_i \in \mathbb{R}^p$. The goal is to model this complex surface accurately to distinguish regular changes from those induced by TIB. The sequential problem can be stated as follows:

$$\mathcal{H}_0 \qquad \mathbf{T}_i \sim \mathcal{P}_{\boldsymbol{\theta}_i} \ \forall i \leq I \tag{1}$$

$$\mathcal{H}_{1,\nu} \quad \mathbf{T}i \sim \begin{cases} \mathcal{P}_{\boldsymbol{\theta}_i} & \forall i < \nu \\ \mathcal{P}_{\boldsymbol{\theta}_i} + \mathbf{a}_{i-\nu} & \forall i \in \nu, \dots I, \end{cases} \tag{2}$$

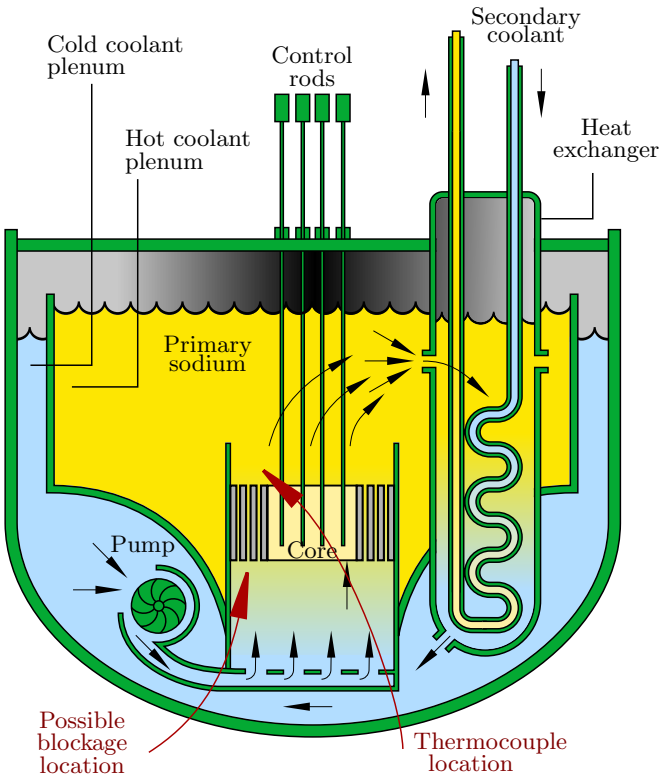

Figure 1: Illustration of the overall design of an SFR nuclear reactor, here using a pool architecture emphasizing the location of thermocouples with respect to the possible blockage. Source: based on an image in the public domain from www.Gen-4.org.

where $\mathcal{P}_{\boldsymbol{\theta}_i}$ is a known distribution parametrized by $\boldsymbol{\theta}_i$ describing the regular rod's temperature, and ν is the (unknown) time of change point. The quickest detection means that upon occurrence, hypothesis $\mathcal{H}_{1,\nu}$ must be detected with minimal delay $I-\nu$. The distribution parameter $\mathbf{a}_n$ represents the temperature drift after the TIB occurs and, of course, it changes over time. Note that the detection problem is complex because hypothesis $\mathcal{H}_{1,\nu}$ is composite, meaning that at each time i there are many possible possibilities for the change point, leading to different alternative hypotheses for the occurrence of the blockage. A sequential detection scheme is a stopping rule S on $\mathbf{T}_1, \mathbf{T}_2, \ldots, \mathbf{T}_i$ that indicates when to take a final decision. The stopping time t_S is the first time i at which the stopping rule S decides that the change point ν has been reached. False alarms occur if $t_S \leq \nu$, and the detection delay $t_S - \nu$ should be minimized. Two criteria are widely used for sequential detection: the Lorden criterion, which focuses on the worst possible detection delay

$$\sup_{\nu \in \mathbb{N}} \sup \mathbb{E}\left[t_S - \nu\right]. \tag{3}$$

and the Bayesian approach, which considers that the change point ν is drawn from a known statistical distribution $\boldsymbol{\pi} : \pi_k = \mathbb{P}[\nu = k], k \in \mathbb{N}$, and minimizes the average detection delay

$$\mathbb{E}_\pi\left[t_S - \nu\right]. \tag{4}$$

However, in our case, neither the worst detection delay nor the average detection delay are relevant criteria. The change point is a deterministic but unknown value, and the goal is to minimize the probability that the change point is detected with a prescribed strict maximal delay. This constraint is dictated by the operational context: if the TIB is not detected within 6 seconds, the core begins to melt, causing irreversible damage. Our goal is to find a stopping rule such that the worst probability of detecting a change point ν with maximal delay M is

maximized. The power function of the stopping rule S is:

$$\beta_S^{(M)} = \inf_{\nu \in \mathbb{N}} \mathbb{P}[t_S < \nu + M | t_S \geq \nu]. \tag{5}$$

Additionally, the largest probability of false alarm over a run length of L observations is:

$$\alpha_S^{(L)} = \sup_{i \in \mathbb{N}} \mathbb{P}[t_S \in i - L + 1, \ldots, i | \nu > i]. \tag{6}$$

To achieve this, we will use a sliding window likelihood ratio test (SW-LRT) which has been shown to be optimal for the worst possible power function (5), over all possible stopping time ν, under a prescribed mean time to false alarm (6), see details in[4–6] and the references therein. Before presenting the SW-LRT detection procedure in more detail, we need to explain how to deal with the nuisance parameters that normal temperatures are.

3. DETECTION OF TOTAL AND INSTANTANEOUS BLOCKAGE: STATEMENT OF THE PROBLEM

To understand better the contributions of this paper, this section outlines how we tackled the three primary difficulties mentioned in Section 2. Firstly, we discuss how to handle nuisance parameters in Section 3.1. Secondly, we present the model for the anomalous temperature caused by a blockage in Section 3.2.

3.1 Dealing With Nuisance Parameters

Handling nuisance parameters is a crucial issue in statistical detection, as discussed in.[7,8] A common approach is to assume linearity, which simplifies the problem. Specifically, we assume that under normal operating conditions, the temperature at time index i of all rods can be modelled as:

$$\mathbf{T}_i \sim \mathcal{N}(\boldsymbol{\mu}_i, \sigma^2 \mathbf{I}_p), \tag{7}$$

where $\boldsymbol{\mu}_i$ is the expected value of all rods' temperature at index time i, σ^2 is the variance of all rods, and $\mathbf{I}_p$ is the identity matrix of size $p \times p$ with p being the number of rods. The linear nuisance parameters model assumes that the average values of temperature can be represented with a small set of basis vectors:

$$\mathbf{T}_i \sim \mathcal{N}(\mathbf{H}\boldsymbol{\theta}_i, \sigma^2 \mathbf{I}_p), \tag{8}$$

where $\mathbf{H}$ is a known full-column rank matrix of size $p \times q$ and $\boldsymbol{\theta} \in \mathbb{R}^q$ is a vector representing the q-dimensional nuisance parameter. This model has been widely used in signal processing due to its simplicity and accuracy in several applications.

Note that when a blockage occurs, it creates an anomaly denoted $\mathbf{a}_{m,n}]$ in the temperature of the rods at location (m, n), see section 3.2 and, therefore, the statistical model of the core temperature becomes:

$$\mathbf{T}_i \sim \mathcal{N}(\mathbf{H}\boldsymbol{\theta}_i + \mathbf{a}_{m,n}, \sigma^2 \mathbf{I}_p), \tag{9}$$

Under this model, everything that falls within the column space spanned by $\mathbf{H}$, denoted $R(\mathbf{H})$, is considered as the nuisance parameter. Therefore, the rejection of the nuisance parameter can be carried out by simply projecting the vector of measured temperature $\mathbf{T}i$ onto the orthogonal complement $R(\mathbf{H})^\perp$ of the column space $R(\mathbf{H})$. This projection is defined with the orthonormal matrix $\mathbf{W} = (w_1, \ldots, wp - q)$, where w_i are the eigenvectors of the projection matrix $\mathbf{P}_\mathbf{H}^\perp = \mathbf{I}_p - \mathbf{H}(\mathbf{H}^\top \mathbf{H})^{-1}\mathbf{H}^\top$ corresponding to eigenvalues 1. The rejection of a linear nuisance parameter can be simply carried out as $\mathbf{W}^\top \mathbf{T}_i$ such that:

$$\mathbf{W}^\top \mathbf{T}i \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}p - q). \tag{10}$$

Again, we an anomaly is present, especially because of a blockage at rod location (m, n) the elimination of the nuisance parameters leads to

$$\mathbf{W}^\top \mathbf{T}i \sim \mathcal{N}(\mathbf{W}^\top \mathbf{a}_{m,n}, \sigma^2 \mathbf{I}p - q). \tag{11}$$

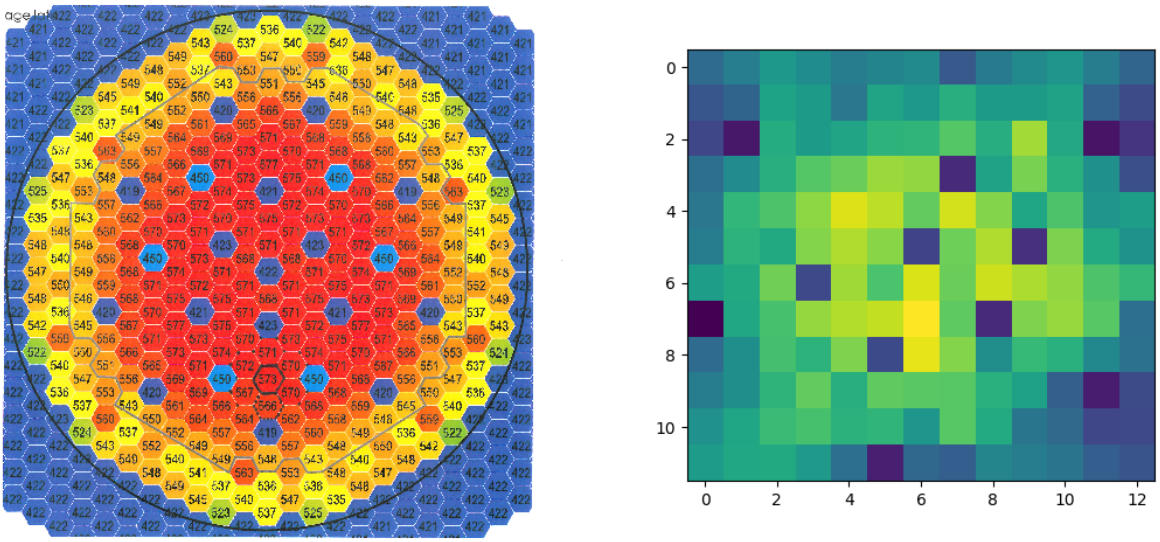

Figure 2: Illustration of the temperature of the rods of the nuclear power plant's core and its conversion into pixels: the temperature measurements are aligned on a grid. Note the temperature of the control rods with the lowest temperature in the middle of the core, and note that on the border of the image some pixels are duplicate to complete the grid. Note that the illustration on the left-hand side does not correspond to the images on the right-hand side: we had measurements over only 121 rods because of the small size of the experimental installation, while the right-hand side is the illustrates the simulated temperate for a larger, real core.

which clearly emphasizes that, on the one hand, the detection problem is reduced to the detection of a signal in a zero-mean noise (10)-(11) and, on the other hand, that the detection depends on the part of the anomaly, which can be distinguished from the non-anomalous background : $\mathbf{W}^\top \mathbf{a}_{m,n}$. It is also important to note that, in practice, the location of the blockage is not known; hence $\mathbf{a}_{m,n}$ is also unknown.

However, this approach is limited in cases of complex non-anomalous backgrounds, such as the temperatures of the rods, see, for instance,.[9] Recently, it was proposed to use a locally adaptive model that adjusts the linear model using the last measurements.[10–15] For the specific problem of total and instantaneous blockage detection, designing a linear model for representing the temperature measurements over 121 nuclear rods is not straightforward. To address this challenge, we proposed using the MobileViT-v2 deep learning model to implicitly learn the underlying nuisance parameters.
We have tried several alternatives for the deep learning model but, in the end, MobileViT is an extremely lightweight architecture yet far enough for our images of size 11×12 pixels. Its attention mechanism has also been experimentally shown very efficient to learn the reference temperature of the rods and eliminate this nuisance parameter.

Specifically, the temperature of the rods from the nuclear core is gathered and represented as the pixels' value of an image, see figure **??** for an illustration. The measures of the rods are aligned on a rectangular grid and presented to the MobileViT-v2 model as an image. We used this deep learning model for detecting the anomaly due to total and instantaneous blockage. To create a large dataset of temperature measurements containing the thermal footprints of a blockage, we accurately modelled the specific footprint of total and instantaneous blockage on temperature measurements, which is addressed in the next section.

3.2 Modelling Blockage Impact on Temperatures

To develop a comprehensive understanding of the impact of a blockage-induced anomaly (TIB) on temperatures, we leveraged numerical data collected at a Sodium Fast Reactor (SFR) station. Temperature measurements were obtained using K-type thermocouples, whose transfer function is well documented. The impulse response

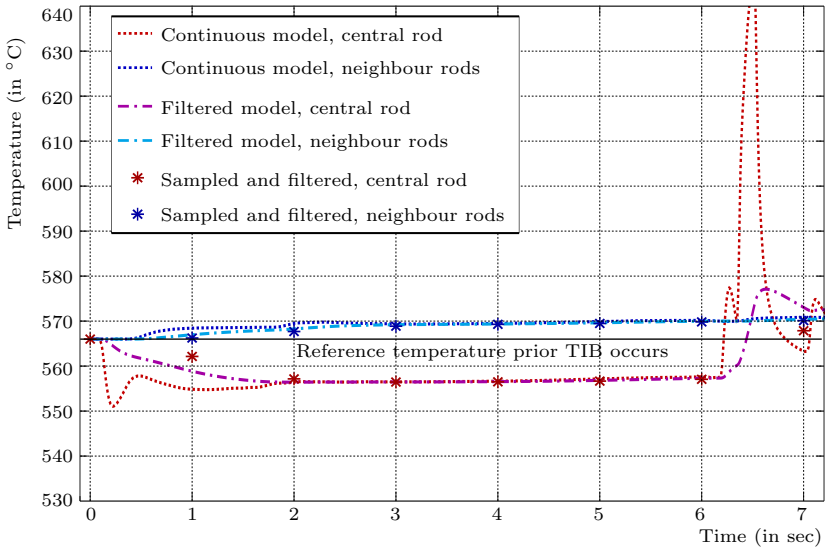

Figure 3: Illustration of the model of TIB contribution on both the central rod and its neighbours.

function of the thermocouples is given by:

$$h(t) = \frac{1}{\tau} \exp\left(\frac{-t}{\tau}\right), \tag{12}$$

where τ is the time constant of the thermocouple, approximately 0.5-1 seconds in our setting. We utilized a time constant of $\tau = 1$ second for our simulations. Given the impracticality of inducing a real TIB to measure its impact on rod temperatures, we relied on numerical simulations utilizing the ASTRID reactor model. These simulations accounted for relevant thermodynamic parameters, such as materials, assembly geometry, and sodium cooling flow. However, the simulated temperatures did not perfectly align with the measured temperatures due to discrepancies in sampling frequency and thermocouple response.
To mitigate this discrepancy, we applied the thermocouple response function, defined above, and subsampled the data by averaging temperatures over one second. This yielded a model for the anomalous temperature due to the TIB, denoted as $\mathbf{a}$. Figure 3 illustrates the temperature profile after the TIB occurs for both the rod at which the blockage happens and its neighbours.

Notably, the figure reveals that the temperature at neighbouring rods increases at a relatively constant rate of approximately 0.6°C per second, whereas the temperature of the central rod drops suddenly after TIB, followed by a rapid increase of about 150°C after approximately 7 seconds, attributed to local vaporization of sodium. The disparity between simulated and observable temperatures is evident, particularly for the central rod. We assumed a spatially invariant model of TIB impact, although the proposed methodology can be adapted to more accurate models by modifying the vector $\mathbf{a}$ according to the inspected rod position and temperature.

Note that in our simulation we did not use the exact profile of the abnormal temperature due to a blockage because with this setting we have a perfect detection, hence the difficulty to present, to analyse or even to compare the results. To avoid having a detection case which is "too simple", we multiplied the thermal footprint of a blockage for a factor of 1/6, which thus corresponds to an increase of about 0.1°C for the neighbouring rods and a drop in temperature of about 1.5°C on the rod subject to the blockage, see Figure 3.

4. PROPOSED METHOD FOR THE QUICKEST DETECTION OF A BLOCKAGE

We can now present the proposed method for the quickest detection of the total and instantaneous blockage. With the model of the impact of total blockage on nuclear core temperature, we created a dataset with additive anomalies. More specifically, for each temperature measurement, we added the anomaly at a random location due

to blockage after a random time between 1 and 6 seconds, because this corresponds to the maximal acceptable detection delay.

As already explained, we trained a MobileViT-v2 deep learning model that was trained over this dataset to distinguish images with an additive anomaly from original images. Since we have two set of measurements, each with a duration of one week, at a sampling frequency of one second, for a total of $604,800$ measurements each, we used one set of data for training and one set of testing so that we are certain that the training and testing set are independent from each other. The number of data is indeed suffisant for training the MobileViT-v2 model over very small images.

Regarding the implementation we used the timm[16] python package but, given the specificity of our image, we have noted that it is better not to use the pretrained weights. MobileViT-v2 model was trained using a decreasing learning rate following stochastic gradient descent with warm restarts (SGDR).[17–19] Given the small size of the image and the lightweight architecture of MobileVit-v2, the training over one epoch was rather very fast. However, given the peculiarity of the images we used, we used 155 epochs, which was far enough. The initial value of the learning rate was obtained using from the method initially proposed in[20] which estimates the optimal value using a one-cycle training with an increasing learning rate at each mini-batch, from a very low initial value to a final high value.
The learning rate scheduler is thus based on this initial guess and then slowing decreased over one cycle. The initial cycle length is set to 5; it is doubled for every cycle and the initial learning rate is divided by two after each cycle. We used five cycles for a total of 155 epochs, which is far enough for convergence.

4.1 Sequential Detection Scheme

The main originality of the present paper is that it is not aimed at detecting the blockage from a single measurement but rather to detect the occurrence of this event in less than a few seconds, in other words, over a set of successive temperature measures. In the literature this problem of sequential analysis is referred to as a change-point detection see, for instance,[21] and the reference therein for a complete introduction. To this end we propose two different approaches for the sequential detection of a blockage. First, it is proposed to use the "soft output" of the MobileViT-v2 model as an input for the well-known CUSUM (cumulative sum) control chart initially developed by Page.[22] Let us denote $x_1, \ldots, x_N$ the consecutive softmax output from the MobileViT-v2 model. The idea of the CUSUM is to compute a cumulative sum of the decision statistics x_n. After the occurrence of the change point, that is after the blockage occurs in our case, the MobileViT-v2 softmax output should be increased. Therefore the CUMSUM is computed as:

$$S_n = \max(0; S_n + x_n - c), \qquad S_0 = 0. \tag{13}$$

In the relation (13) one can note that the constant value c is subtracted from the MobileViT-v2 softmax output. This constant represents the minimum value above which it is considered that a blockage has been detected. In addition one can also note in the equation (13) that the CUSUM value S_n is reset to 0 when it is negative. This is a key concept of the CUSUM, which allows detecting the change point with a minimal delay because samples without blockage do not reduce the CUSUM value, which is assumed to remain at 0 before the event we want to detect occurs, in our case, the total and instantaneous blockage.

The alternative approach we used in the present paper is the sliding windows:

$$W_n = \sum_{n-\tau}^{n} x_n, \tag{14}$$

where τ is the size of the sliding window, which is defined in our case by $\tau = 6$ because, as explained in Section 3.2 the blockage must be detected in less than 6 seconds.

On the one hand, it has been shown that the CUSUM minimizes the average detection delay (3) for a fixed average false-alarm time.[23,24] On the other hand, it has been shown the sliding window method maximizes the probability of detecting a known change under a maximal prescribed delay (5). Both criteria are relevant in our case because, on the one hand, it is aimed at detecting the blockage in less than 6 seconds and hence to

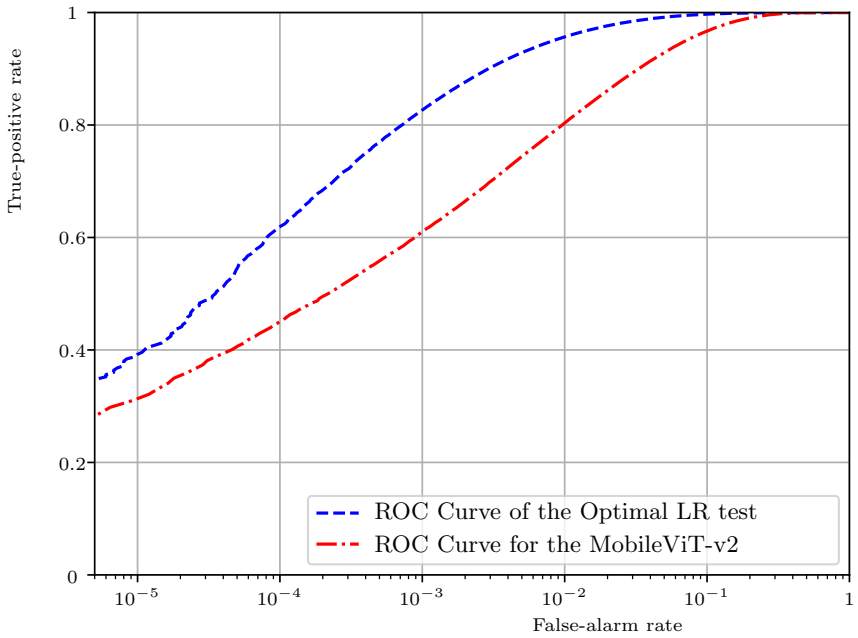

Figure 4: ROC Curves showing the detection accuracy obtained with the MobiletViT-v2 model and comparison with the empirical optimal detection power.

maximize the probability of detecting the blockage under this prescribed delay. On the other hand, minimizing the detection delay is also interesting, as, we shall see in the next Section 5 the CUSUM offers very good overall results.

5. NUMERICAL RESULTS AND ANALYSIS

First of all, we would like to show how well the MobileNet-v2 model performs for detecting the abnormal temperature due to the presence of a blockage we simulated in our dataset. To this end, Figure 4 shows the accuracy of the detection through a ROC (Receiver-Operating Characteristic) curve, which shows the true positive rate, some referred to as the recall or the sensitivity, as a function of the false-positive rate. Note that, for readability, the value of the x-axis is represented with a logarithmic scale in Figure 4.
Because we simulated the occurrence of a blockage, it is possible to compute the optimal most powerful test, which is simply a likelihood ratio test (LR test). In our case, when the anomaly due to blockage is perfectly known, the detection problem is reduced to a change of the mean in a Gaussian (8)-(9). Therefore, it is straightforward to establish that, provided that the non-anomalous back is also perfectly known (10)-(11), the LR test is given by the projection of the measurements $\mathbf{T}_i$ onto the thermal footprint of the blockage $\mathbf{a}_{m,n}$. For the sake of comparison, we computed empirically these decision statistics in order to provide, in the Figure 4 a comparison between the performance of the MobileViT-v2 detection and the optimal detection provided that the blockage and its location are perfectly known. It can be seen that the proposed detection method based on MobileViT-v2 performs extremely well, especially for low false-positive rates.

Of course, the most interesting results concern the sequential detection of total and instantaneous blockage using the MobileViT-v2 model. To this end Figure 5 first compares the performance of the CUSUM with the performance of the sliding window. In Figure 5 the average detection is presented, on the y-axis, as a function of the average mean time to false alarm. Note that experimentation is rather time-consuming because for each simulation we need to apply the CUSUM and the sliding window test over a long sequence of MobileViT-v2 output in order to measure the mean time to false alarm. For the sake of simplicity, we limit the analysis to 35000 samples. For each simulation we chose a random starting point in the sequence of $604,800$ measurements and we added Gaussian noise with standard deviation $\sigma = 0.25$ in order to generate $10,000$ repetitions to compute

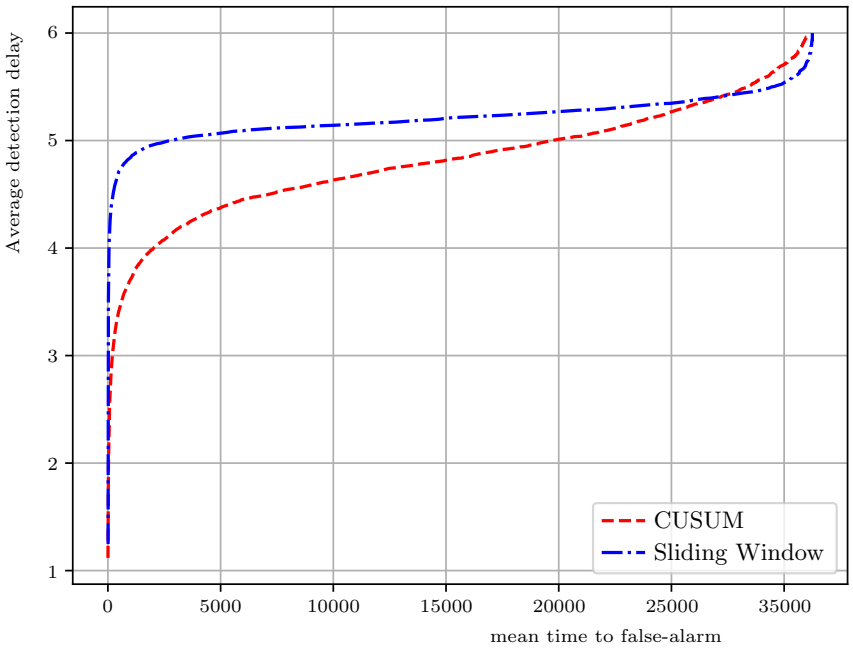

Figure 5: Average detection delay as a function of the mean time to a false alarm.

the mean time to false alarm as a function of the detection threshold.
We repeated the same procedure with an additive signal, at a random location, corresponding to the thermal footprint of a blockage. In the case we limit our analysis to 20 consecutive samples because, in practice, the blockage is always detected before such a long delay.

We can note in Figure 6 that the CUSUM sequential detection method always achieves a lower average detection delay than the sliding window method.

Last but not least, the Figure 6 presents the probability of detecting the blockage in 6 seconds or less as a function of the average time to false alarm. Surprisingly, one can note that the CUSUM sequential detection methods achieve a higher probability of blockage detection. This is especially surprising because the sliding window is designed to maximize the probability of detection under a maximal delay.[4–6,21] However, in the present case of blockage, the problem is difficult because the impact of blockage on the softmax output from the MobileViT-v2 is unknown; therefore, the sliding window is reduced to a simple sum (14) which is not the condition under which it is proved optimal.
All in all, even though we decrease significantly the impact of the blockage on the temperature in order to obtain meaningful results, it can be noted that the sequential detection methods allow obtaining very high detection performance, which is very encouraging.

6. CONCLUSION AND FUTURE WORKS

The present paper proposes a very original sequential method for the detection of a total and instantaneous blockage in the coolant of the nuclear power plant core. The detection method we propose is two-fold. First we trained a MobileViT-v2 model to detect the anomaly that is due to the blockage in the temperature of the core. This lightweight model leverages the attention mechanism to reach a very high detection power, which is close to the performance of the optimal LR test.
Then the main originality of the present paper is to use the softmax output from the MobileViT-v2 into a sequential statistical procedure. For the sequential detection, we used the CUSUM, which minimizes the average detection delay as well as the sliding window method, which maximizes the probability of detection under a prescribed detection delay provided that the signal to detect is perfectly known. Surprisingly, the CUSUM provides better performance in terms of average detection delay and probability of detection under the prescribed

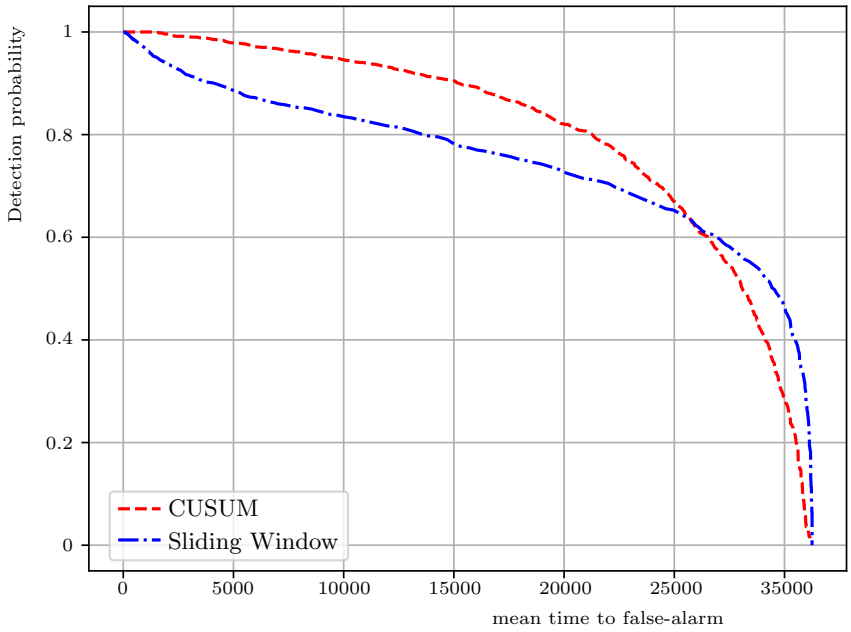

Figure 6: Probability of detection of blockage before the maximal delay of 6 seconds as a function of the mean time to false alarm.

delay of 6 seconds, which can be explained by the fact that the effect of the blockage on the softmax output from the MobileViT-v2 model is unknown.

REFERENCES

[1] Stanculescu, A., "GIF R&D outlook for generation iv nuclear energy systems: 2018 update," in [*Proceedings of the Generation IV International Forum, Paris, France*], 16–18 (2018).

[2] Pioro, I. L. and Rodriguez, G. H., "Generation IV international forum (GIF)," in [*Handbook of Generation IV Nuclear Reactors*], 111–132, Elsevier (2023).

[3] Bailey, E. N., "Independent confirmatory survey summary and results for the enrico fermi atomic power plant, unit 1 newport, michigan (final report no. 2)," tech. rep., Oak Ridge Inst. for Science and Education (ORISE), Oak Ridge, TN (United States) (2011).

[4] Guépié, B. K., *Sequential Detection of Transient Changes: application to water distribution network monitoring* in french, PhD thesis, Troyes University of Technology (UTT) (March 2013).

[5] Tout, K., Retraint, F., and Cogranne, R., "Wheels coating process monitoring in the presence of nuisance parameters using sequential change-point detection method," in [*2017 25th European Signal Processing Conference (EUSIPCO)*], 196–200 (2017).

[6] Tout, K., Retraint, F., and Cogranne, R., "Non-stationary process monitoring for change-point detection with known accuracy: Application to wheels coating inspection," *IEEE Access* **6**, 6709–6721 (2018).

[7] Frank, P. M., "Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy: A survey and some new results," *Automatica* **26**(3), 459–474 (1990).

[8] Fouladirad, M. and Nikiforov, I., "Optimal statistical fault detection with nuisance parameters," *Automatica* **41**(7), 1157–1171 (2005).

[9] Nikiforov, I., Harrou, F., Cogranne, R., Beauseroy, P., Grall, E., Guépié, B. K., Fillatre, L., and Jeannot, J.-P., "Sequential detection of a total instantaneous blockage occurred in a single subassembly of a sodium-cooled fast reactor," *Nuclear Engineering and Design* **366**, 110733 (2020).

[10] Tout, K., Cogranne, R., and Retraint, F., "Fully automatic detection of anomalies on wheels surface using an adaptive accurate model and hypothesis testing theory," in [*2016 24th European Signal Processing Conference (EUSIPCO)*], 508–512 (2016).

[11] Tout, K., Cogranne, R., and Retraint, F., "Statistical decision methods in the presence of linear nuisance parameters and despite imaging system heteroscedastic noise: Application to wheel surface inspection," *Signal Processing* **144**, 430–443 (2018).

[12] Nguyen, T., Cogranne, R., and Doyen, G., "An optimal statistical test for robust detection against interest flooding attacks in ccn," in [*2015 IFIP/IEEE International Symposium on Integrated Network Management (IM)*], 252–260 (2015).

[13] Nguyen, T., Mai, H.-L., Cogranne, R., Doyen, G., Mallouli, W., Nguyen, L., El Aoun, M., De Oca, E. M., and Festor, O., "Reliable Detection of Interest Flooding Attack in Real Deployment of Named Data Networking," *IEEE Transactions on Information Forensics and Security* **14**, 2470–2489 (Sept. 2019).

[14] Ghadban, N., Cogranne, R., and Doyen, G., "A decentralized approach for adaptive workload estimation in virtualized environments," in [*2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*], 1186–1194, IEEE (2017).

[15] Cogranne, R., Doyen, G., Ghadban, N., and Hammi, B., "Detecting botclouds at large scale: A decentralized and robust detection method for multi-tenant virtualized environments," *IEEE Transactions on Network and Service Management* **15**(1), 68–82 (2018).

[16] Wightman, R., "Pytorch image models." https://github.com/huggingface/pytorch-image-models (2019).

[17] Smith, L. N., "No more pesky learning rate guessing games," *CoRR, abs/1506.01186* **5**, 575 (2015).

[18] Loshchilov, I. and Hutter, F., "Sgdr: Stochastic gradient descent with warm restarts," *arXiv preprint arXiv:1608.03983* (2016).

[19] Smith, L. N., "Cyclical learning rates for training neural networks," in [*2017 IEEE winter conference on applications of computer vision (WACV)*], 464–472, IEEE (2017).

[20] Smith, L. N. and Topin, N., "Super-convergence: Very fast training of neural networks using large learning rates," in [*Artificial intelligence and machine learning for multi-domain operations applications*], **11006**, 369–386, SPIE (2019).

[21] Tartakovsky, A., Nikiforov, I., and Basseville, M., [*Sequential Analysis: Hypothesis Testing and Changepoint Detection*], Chapman & Hall/CRC Monographs on Statistics & Applied Probability, Taylor & Francis (2014).

[22] Page, E., "Continuous inspection schemes," *Biometrika* , 100–115 (1954).

[23] Lorden, G., "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics* , 1897–1908 (1971).

[24] Moustakides, G. V., "Optimal stopping times for detecting changes in distributions," *the Annals of Statistics* **14**(4), 1379–1387 (1986).