

ROBUST AUTONOMOUS DETECTION OF THE FAULTY SENSORS OF A SENSOR ARRAY

Siddhartha Ghosh, Alex Freitas

Computing Laboratory, University of Kent
Canterbury, UK, CT2 7NF
sg55@kent.ac.uk

Ian Marshall

Centre for Environmental Informatics
Lancaster University
Lancaster, LA1 4YQ
I.W.Marshall@lancaster.ac.uk

ABSTRACT

We propose a technique for the autonomous detection of the faulty sensors of a sensor array that are aberrant relative to the rest. Our approach is based on probabilistically modeling the distribution of the differences between the sensor measurements as a mixture of gaussians and then classifying further instances of the sensor differences using a naive bayes classifier. We demonstrate the applicability of this technique to the diagnosis of the sensors/photosites of a CCD array, using sensor array data comprising of randomly selected images. Our technique performs well for different combinations of parameter settings at the detection of the faulty photosites of a CCD array.

1 Introduction

We propose a computational technique for autonomously detecting the faulty or aberrant sensors of a sensor array consisting of a number of homogeneous sensors, each of which measures the same physical quantity of interest. Our data-driven technique can autonomously identify a wide range of faulty sensors, as shown later, without the need for any defect model or apriori knowledge of the nature of the sensor faults. In order to do so, from the sensor data, we first compute pairwise differences between the measurements of every sensor and all the others - for n sensors there will be $n \times (n - 1)$ difference estimates. We probabilistically model the distribution of all such pairwise difference estimates as a gaussian mixture model (GMM) where every component of the GMM represents a cluster or grouping of the difference estimates; the choice of a GMM is motivated by the fact that a mixture model based formalism provides a natural framework for dealing with uncertain estimates [1]. We then classify further instances of the pairwise difference estimates using a naive bayes classifier and the MAP (maximum a posteriori) decision rule; finally, we aggregate the classifications of all the pairwise difference estimates of a sensor to determine whether or not the sensor is faulty. Since, the primary objective of our work is to accurately detect faulty sensors, our methodology is based on the integration of two well understood techniques as in [2]: the semi-parametric

mixture model based clustering and subsequent application of the naive bayes classifier, and the MAP decision rule - we have not reinvented the wheel. However, our use of the pairwise difference estimates instead of the raw sensor measurements - a transformation that evinces more information on the sensors' behaviour as compared to raw measurements, is a novel application. We explain this methodology in detail in Section 2. To demonstrate its applicability, we apply our technique to the detection of the faulty photosites of the CCD array of a monochrome digital camera - the CCD serves as a model of a sensor array, and in this paper we use the terms photosite and sensor interchangeably. We show that the technique yields good results at the detection of the the faulty sensors of the CCD array, in Section 3. Previously, the use of PCA to identify bad photosites from image data was undertaken by Alonso et al. [3]; their approach only worked with uniform images. Chapman et al. proposed a simple bayesian algorithm for the online identification of faults in their APS image sensor arrays, to which they applied different defect models [4]. The fundamental difference between their work and ours is that they used an apriori defect model rather than discovering faulty photosites autonomously from data, as we do. Their technique was susceptible to detection errors when there were clusters of faulty photosites. Koushanfar et al. proposed a mechanism for the online fault detection of sensor measurements using a model based cross-validation technique [5]; such techniques are computationally quite prohibitive for larger numbers of sensors than the ones we consider. Unlike some of the previous approaches, we do not make any assumptions of the nature of the noise sources afflicting the sensors or any use labeled training data; the results of our univariate analysis technique are easy to visualize and tangible.

2 The proposed methodology

2.1 Modelling the difference estimates

A single pairwise difference estimate between two sensors measuring the same same physical quantity highlights the similarity or in contrary, the dissimilarity between the mea-

measurements of the sensors. A pairwise difference estimate between two sensors that are ‘healthy’ or functioning normally, will be lower than one between a faulty sensor and a healthy one. Computing pairwise difference estimates may be particularly useful when the number of sensors in an array may not be very large as it increases the number of points for density estimation, as opposed to using individual sensor measurements. Consider a group of n sensors, measuring the same physical quantity of interest (eg. light, temperature etc.) - If every sensor makes p such measurements corresponding to the measurements at p distinct time instants, we can form an $an \times p$ matrix M . Every $1 \times p$ row in M , M_a (where $a = 1 \dots n$), will then represent a vector of p measurements of a single sensor. For every row vector M_a , we then compute the difference between M_a and all other M_b s (where $b = 1 \dots n$ and $a \neq b$). The Δ estimates between the two vectors are computed as follows.

$$\Delta_{ab} = \sum_{k=1}^p |M_{ak} - M_{bk}| \quad (1)$$

These Δ s are the difference estimates between two sensors computed from p different measurements, and indicate how close or far apart the measurements of the two sensors are. The distribution of all such pairwise difference estimates will have natural groupings or clusters since they have been computed from sensor measurements that are inherently noisy. Consequently, clusters of such estimates will not be well-defined or have distinct boundaries and we can probabilistically model the distribution of the Δ s as a GMM, with the number of mixture components C fixed as GMM’s are well-suited to modeling arbitrary multi-modal densities. The GMM formulation is $\sum_{k=1}^C \pi_k P_k(\Delta_{ab}|\theta_k)$, where π_k is the weighting factor of the k th component ($\pi_k \geq 0$ and $\sum_{k=1}^C \pi_k = 1$) and θ_k represents the parameter vector. $P_k(\Delta_{ab}|\theta_k)$ for our purpose is the equation of a univariate gaussian as in Eq 2.

$$P_k(\Delta_{ab}|\theta_k) = \frac{1}{\sigma_k \sqrt{2\pi}} e^{-(\Delta_{ab} - \mu_k)^2 / 2\sigma_k^2} \quad (2)$$

Every gaussian component of the mixture represents a partition or cluster in the distribution of the difference estimates and every difference estimate has a finite probability of belonging to every gaussian component. We use the standard EM algorithm with K-means initialization for estimating the GMM. The input to the EM algorithm is a shuffled form of the vector of the difference estimates Δ s, in which the inherent ordering of the pairwise difference estimates is not preserved; this avoids any positional bias when computing pairwise difference estimates.

2.2 Classification of the difference estimates

From the sensor measurements, we construct sensor data sets $S_1 \dots S_N$ where each set S_i is composed of $n \times (n - 1)$ pairwise difference estimates computed from p independent measurements according to Eq 1. Since estimating the

Table 1. Table showing the posterior probabilities and classification of sensor pair difference estimates

Sensor a	Sensor b	$P(C_1 \delta_{ab}, \theta_x)$	$P(C_2 \delta_{ab}, \theta_x)$	$P(C_3 \delta_{ab}, \theta_x)$	classify(δ_{ab})
1	2	0.1	0.7	0.2	2
1	3	0.15	0.8	0.05	2

GMM’s parameters using EM is computationally expensive, we use only one of the sets from all the constructed sets for estimating the GMM - we classify further instances of difference estimates computed from the remaining $N - 1$ sets into the components of the GMM using a naive bayes classifier. For the naive bayes classification, let Δ_{ab}^i be a difference estimate computed between two sensors a and b , according to Eq 1, for the set S_i ; every Δ_{ab}^i is estimated for each of the $N - 1$ data sets. Given the estimated GMM, every Δ_{ab}^i is independent of all the other Δ_{ab}^j s where $j = 1, \dots, N - 1, j \neq i$. For the $N - 1$ sets, we get a vector $\delta_{ab} = (\Delta_{ab}^1, \dots, \Delta_{ab}^{N-1})$. Subsequently, applying the naive bayes classification rule to the vector δ_{ab} yields the corresponding class posterior probabilities shown in Eq 3.

$$P(C_k|\delta_{ab}, \theta_k) = P(\theta_k) \prod_i^{N-1} P_k(\Delta_{ab}^i|\theta_k) \quad (3)$$

Each $P_k(\Delta_{ab}^i|\theta_k)$ is estimated according to Eq 2. The above process is repeated for all the sensor pairs $a = (1 \dots n)$ and $b = (1 \dots n)$, ($a \neq b$) in all the image sets $N - 1$, thereby, producing the corresponding $n \times (n - 1)$ Δ_{ab}^i s for each data set S_i , the δ_{abs} and the corresponding posterior probabilities. Table 1 is an excerpt from the table of the posterior class probabilities computed in the case where there are three classes C_1, C_2 and C_3 . Each row of the table contains the posterior probabilities of the difference estimates between two sensors in the three classes computed from the sensor data sets S_1 to S_{N-1} according to Eq 3. The last column of Table 1 is the final decision as to which of the three classes or categories the difference estimates between a sensor pair belongs to, shown as Eq 4, which is an application of the MAP decision rule that minimizes the probability of error or error rate under zero-one loss [6].

$$\text{classify}(\delta_{ab}) = \arg \max_k P(C_k|\delta_{ab}, \theta_k) \quad (4)$$

2.3 Sensor categorization

The posterior probability values of Eq 3 indicate how close or far apart the measurements of two sensors are. We amass the $(n - 1)$ final classifications of all the difference estimates (each of them computed according to Eq 4) of a sensor together. We then assign a sensor to the class C_k that is the most frequently occurring class amongst the $(n - 1)$ classifications of a single sensor’s $(n - 1)$ difference estimates, as per Eq 5.

$$\text{maxClass}(Sensor_a) = \text{maxOccurrence}_k(\text{classify}(\delta_{ab})) \quad (5)$$

where $b = 1, \dots, n, b \neq a$

A faulty sensor will always be assigned to the class C_k or the component of the GMM that has the smallest weighting factor - such a component is therefore a representative of the ‘outlier’ component of the GMM.

3 Experiments with a CCD array

3.1 Experimental setup

The CCD array of a monochrome USB digital camera serves as our sensor array; we took several images of different random scenes in a number of indoor rooms with the camera. In every image, the measured value obtained for a single photosite after digitization, is as an 8-bit integer, resulting in 256 possible different shades of grey ranging from black to white, which can be represented as a number in the [0-255] interval; the measurement of an individual photosite of the CCD in a single image therefore represents one sensor measurement. Bad sensors corresponding to the faulty photosites of a CCD array can be of several different types: Stuck photosite (a photosite that always reads high or maximum (typically 255) on all exposures), Hot photosite (a photosite that reads unusually high values at longer exposures), Dead photosite (a photosite that reads zero (black) at all exposures) and photosites that may be hypo or hyper sensitive and are particularly hard to detect. We applied our technique to the detection of different faulty photosites. Although the CCD has a resolution of 720×480 photosites, we restricted our analysis of the sensors of the CCD array to windows of $x \times y$ sensors (e.g. it might be a 10×20 window corresponding to 200 sensors) of the array. We took a number of different images under exposures of $1/25$, $1/250$ and $1/500$, constituting multiple image-banks of 60 to 140 images. Every image-bank was comprised of images taken under a particular exposure setting and every image in an image-bank represented a snapshot in time. We also formed a separate ‘combined’ image-bank consisting of images taken under all the three exposures. Following our methodology as described in Sec 2.2, for each image-bank, we formed N separate sets of images, each set comprising of the same number p of randomly selected images each, from which we compute the difference estimates. One of these sets was used for the GMM parameter estimation using EM and the remaining ones for the naive bayes classification (For e.g. Out of 60 randomly selected images, we might use $p = 12$ for GMM estimation and from the remaining images 48 images, we form 4 sets of $p = 12$ images each for naive bayes classification). We also found that the camera’s CCD did not have any dead or stuck photosites. Hence, for these two cases, we artificially introduced dead and stuck photosites in the images.

3.2 Model estimation and classification

Each of the n photosites has a measurement in a single image forming a $n \times 1$ vector for an image, that is then normal-

ized to zero mean and unit variance. As our technique relies on p measurements for every photosite, we take p images I_1, \dots, I_p and the measurements of all the n photosites for p images form a $n \times p$ matrix M - there are N such matrices for the N image data sets S_1, \dots, S_N as described in Sec 2.2. For a single row of the matrix M_i corresponding to p measurements of photosite i (where $i = 1, \dots, n$), we computed the difference estimates between every M_i and M_j ($j = 1 \dots n, i \neq j$). Using one of the N sets, we modelled the distribution of the difference estimates as a GMM, where we fixed the number of components or clusters C to be 3. The choice of the number of components was dictated by the need to arrive at a compromise between robustly detecting faulty photosites and estimating the number of components accurately (we applied BIC for this purpose as in [1]); we found that a choice of 3 components helped achieve this tradeoff. We discovered that the components of the gaussian mixture labelled as ‘low’ and ‘medium’ encompassed a large proportion of the difference estimates. These two components correspond to differences between the photosites’ measurements that are relatively low and the photosites that are responsible for these two categories of difference estimates have close associations with the other photosites (they are quite similar). However, the third gaussian component, labelled as ‘high’, corresponds to the component of the gaussian mixture that is the representative of the high difference estimates, that lie at the tails of the density - it accounts for a very small proportion of the total difference estimates and has a high variance - it serves as the ‘outlier’ component. Larger differences in measurements are produced by dissimilar photosites and consequently, we can say that such difference measurements are generated by photosites that are aberrant relative to the rest. We then applied the classification (Sec 2.2) and categorization (Sec 2.3) steps to additional sets of p images to identify the faulty photosites amongst the n photosites. We consider a photosite to be aberrant if majority of its difference estimates are classified in the ‘high’ category according to Eq 5.

4 Experimental Results and Discussion

4.1 Hot, dead and stuck photosites cases

In the images taken under long exposure (exposure greater than $1/50$), the hot photosites were prominent and appeared much brighter than the other photosites and the dead and stuck pixels were artificially introduced at certain location of the CCD. Since, the faulty photosites were either visible (hot) or known (dead and stuck), we evaluated the performance of our analysis on two metrics: the false positive error rate (FPR) and the false negative error rate (FNR). FPR was obtained by dividing the number of normal photosites wrongly classified in the ‘high’ category (false positives) as per Eq 5, by the total number of classifications of all the photosites. FNR was obtained by dividing the number of

Table 2. Performance at identification of hot photosites amongst 121 photosites for different numbers of images

Total No. images	p	FPP	FNP
8	4	0.0	0.0
10	5	0.0	0.0
12	6	0.0	0.0
24	8	0.0	0.0
48	12	0.0	0.0
60	12	0.0	0.0

Table 3. Performance for different exposure settings and different numbers of sensors

Exposures	Hot			Dead			Stuck		
	FPR	FNR	DD	FPR	FNR	DD	FPR	FNR	DD
1 By 25	0.0	0.0	1.19	0.0	0.0	1.19	0.0	0.0	1.19
1 By 500	NA	NA	NA	0.0	0.0	1.19	0.0	0.0	1.19
Random	0.0	0.0	1.19	0.0	0.0	1.19	0.0	0.0	0.0
No. Sensors									
273	0.0	0.0	0.73	0.0	0.0	1.1	0.0	0.0	1.1
75	0.0	0.0	2.67	0.0	0.0	2.67	0.0	0.0	2.67
12	0.0	2.0	16.7	2.0	12.0	16.7	0.0	0.0	16.7

faulty photosites not correctly classified in the ‘high’ category (false negatives) as per Eq 5, by the total number of classifications of all the photosites; the total number of classifications being the number of trials or runs times the number of photosites that are classified per run. At first, we assessed the performance of our technique at detecting the hot photosites in a set of total 60 images taken under 1/25 exposure for different numbers of images and different values of p . Table 2 shows this result of 20 runs of our technique applied to a window containing 121 photosites (the total classifications in this case is equal to $20 \times 121 = 2420$), that included the known hot photosites. We then evaluated our analysis for the detection of hot, dead and stuck photosites for images taken under different exposures and for different number of photosites - Table 3 shows the performance of our technique. As the hot photosites were not evident or conspicuous in the images taken under 1/500 exposure, we did not try to detect hot photosites in the 1/500 exposure image-bank. In Table 3, DD stands for the defect density: the fraction of faulty photosites in the total number of photosites and NA indicates ‘not applicable’. Our technique performed well under most conditions. For small numbers of photosites, the technique was prone to false negatives and for the cases when the faulty photosites were wrongly diagnosed to have a majority of their difference estimates in the ‘high’ category, we found that the majority of their difference estimates were in the ‘medium’ category. What this indicates is that, for smaller number of photosites, possibly, a choice of 2 mixture components might be a better one as the number of difference estimates are not high enough for the density to be approximated by 3 gaussian components.

4.2 Hyper/Hypo sensitive photosite case

We also tried to determine if the analysis can detect photosites that are slightly faulty or ‘suspect’; they may be more or less sensitive than the rest, but, are not visibly faulty. Since, they cannot be seen upon visual inspection, it is im-

Table 4. Performance at identification of ‘suspect’ photosite amongst 230 photosites

FPR	FNR	DD
0.0123	0.0023	0.0043

perative to discover the ‘suspect’ photosites on the CCD. We tried to determine these ‘suspect’ photosites from the images using a test for outlier based on the normal distribution of the photosite values from an unbiased random averaged image, having very little image formation. We created such an averaged random by accruing a large number of images of plain surfaces (such as a wall), and averaging all these frames to get a single frame. From this frame, we selected n photosites, such that n was quite large ($> 10,000$), normalized the photosite values (subtracting the mean of n and dividing by the standard deviation of n) and then plotted the distribution of these photosites. We applied the Jarque-Bera normality test at a 1 percent significance level to ensure that the distribution of the photosites was a standard normal. We tried to look for photosites whose measurements were greater than 3σ from the mean in multiple such averaged frames and found that photosite at location $X : 393$, $Y : 291$ on the CCD was greater than the 3σ in three such averaged frames; we deemed it to be ‘suspect’. We then tried to determine if our integrated technique could detect this ‘suspect’ photosite and evaluated its performance at doing so. From 25 runs of applying our technique to a window of 230 photosites containing the above ‘suspect’ photosite, in 96 randomly selected images from the ‘combined’ image-bank, we found that out of the 25 runs, 7 runs failed to identify the ‘suspect’ photosite as being aberrant (having majority of its difference estimates in the ‘high’ category) and 18 runs correctly did so, although the overall results were quite good as shown in Table 4. Unlike the test based on normal distribution, that is run on plain images, the image information in a set of randomly selected images does affect the diagnosis of the ‘suspect’ photosite. However, the former needed the measurements of 11,000 photosites to have a normal distribution before the ‘suspect’ photosites can be detected at the 3σ level, whereas our analysis could detect one such ‘suspect’ photosite amongst 230 other photosites.

5 References

- [1] C Fraley and A. E. Raftery, “How many clusters? which clustering method? answers via model-based cluster analysis,” *The Computer Journal*, vol. 41, no. 8, pp. 578–588, 1998.
- [2] T. Hastie and R. Tibshirani, “Discriminant analysis by Gaussian mixtures,” *Journal of the Royal Statistical Society series B*, vol. 58, pp. 158–176, 1996.
- [3] J. M. Lopez-Alonso and J. Alda, “Bad pixel identification by means of principal components analysis,” *Optical Engineering*, vol. 41, pp. 2152–2157, Sept. 2002.
- [4] J. Dudas, C. Jung, L. Wu, G.H. Chapman, I. Koren, and Z. Koren, “On-line mapping of in-field defects in image sensor arrays,” in *DFT ’06: Proceedings of the 21st IEEE International Symposium on Defect and Fault-Tolerance in VLSI Systems*, Washington, DC, USA, 2006, pp. 439–447, IEEE Computer Society.
- [5] F. Koushanfar, M. Potkonjak, and A. Sangiovanni-Vincentelli, “Online fault detection of sensor measurements,” in *Proceedings of IEEE Sensors 2003*, October 2003.
- [6] P. Domingos and M. Pazzani, “On the optimality of the simple bayesian classifier under zero-one loss,” *Mach. Learn.*, vol. 29, no. 2-3, pp. 103–130, 1997.