Scalable anomaly detection in large homogeneous populations

Henrik Ohlsson, Tianshi Chen, Sina Khoshfetratpakazad, Lennart Ljung and S. Shankar Sastry

Linköping University Post Print

N.B.: When citing this work, cite the original article.

Original Publication:

Henrik Ohlsson, Tianshi Chen, Sina Khoshfetratpakazad, Lennart Ljung and S. Shankar Sastry, Scalable anomaly detection in large homogeneous populations, 2014, Automatica, (50), 5, 1459-1465.

<http://dx.doi.org/10.1016/j.automatica.2014.03.008>

Copyright: International Federation of Automatic Control (IFAC)

<http://www.ifac-control.org/>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-108173>

Brief paper

Scalable anomaly detection in large homogeneous populations[☆]

Henrik Ohlsson^{a,b}, Tianshi Chen^{a,1}, Sina Khoshfetrat Pakazad^a, Lennart Ljung^a,
S. Shankar Sastry^b

^a Department of Electrical Engineering, Linköping University, SE-581 83 Linköping, Sweden

^b Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, CA, USA

ARTICLE INFO

Article history:

Received 3 October 2012

Received in revised form

21 December 2013

Accepted 19 February 2014

Available online 28 March 2014

Keywords:

Anomaly detection

Outlier detection

Multi-hypothesis testing

Distributed optimization

System identification

ABSTRACT

Anomaly detection in large populations is a challenging but highly relevant problem. It is essentially a multi-hypothesis problem, with a hypothesis for every division of the systems into normal and anomalous systems. The number of hypothesis grows rapidly with the number of systems and approximate solutions become a necessity for any problem of practical interest. In this paper we take an optimization approach to this multi-hypothesis problem. It is first shown to be equivalent to a non-convex combinatorial optimization problem and then is relaxed to a convex optimization problem that can be solved distributively on the systems and that stays computationally tractable as the number of systems increase. An interesting property of the proposed method is that it can under certain conditions be shown to give exactly the same result as the combinatorial multi-hypothesis problem and the relaxation is hence tight.

© 2014 Elsevier Ltd. All rights reserved.

1. Introduction

In this paper we study the following problem: we are given N systems and we suspect that $k \ll N$ of them behave differently from the majority. We do not know beforehand what the normal behavior is, and we do not know which k systems behave differently. This problem is known as an *anomaly detection problem* and has been discussed e.g., in Chandola, Banerjee, and Kumar (2009); Chu, Gorinevsky, and Boyd (2011); Gorinevsky, Matthews, and Martin (2012). It clearly has links to *change detection* (e.g., Basseville & Nikiforov, 1993; Gustafsson, 2001 and Patton, Frank, & Clark, 1989) but is different because the detection of anomalies is done by comparing systems rather than looking for changes over time.

The anomaly detection problem typically becomes very computationally demanding, and it is therefore of interest to study *distributed solutions*. A distributed solution is also motivated by that

many anomaly detection problems are spatially distributed and lack a central computational unit.

Example 1 (Aircraft Anomaly Detection). In this example we consider the problem of detecting abnormally behaving airplanes in a large homogeneous fleet of aircrafts. Homogeneous here means that the normal aircrafts have similar dynamics. This is a very relevant problem (Chu et al., 2011; Gorinevsky et al., 2012) and of highest interest for safety in aeronautics. In fact, airplanes are constantly gathering data and being monitored for this exact reason. In particular, so called *flight operations quality assurance* (FOQA) data are collected by several airlines and used to improve their fleet's safety.

As showed in Chu et al. (2011), faults in the angle-of-attack channel can be detected by studying the relation between the angle of attack, the dynamic pressure, mass variation, the stabilizer deflection angle, and the elevator deflection. The number of airplanes in a fleet might be of the order of hundreds and data from a couple of thousand flights might be available (200 airplanes and data from 5000 flights were used in Chu et al., 2011). Say that our goal is to find the 3 airplanes among 200 airplanes that are the most likely to be anomalous to narrow the airplanes that need manual inspection. Then, we would have to evaluate roughly 1.3×10^6 hypothesis (the number of unordered selections of 3 out of 200 airplanes). For each hypothesis, the likelihood for the observed data would then be maximized with respect to the unknown parameters and the most likely hypothesis accepted. This is clearly a very computationally challenging problem.

[☆] The material in this paper was partially presented at the 16th IFAC Symposium on System Identification (SYSID 2012), July 11–13, 2012, Brussels, Belgium. This paper was recommended for publication in revised form by Associate Editor Michele Basseville under the direction of Editor Torsten Söderström.

E-mail addresses: ohlsson@isy.liu.se (H. Ohlsson), tschen@isy.liu.se (T. Chen), sina.kh.pa@isy.liu.se (S.K. Pakazad), ljung@isy.liu.se (L. Ljung), sastry@eecs.berkeley.edu (S. Shankar Sastry).

¹ Tel.: +46 0 13 28 22 26; fax: +46 0 13 28 26 22.

Example 1 considers anomaly detection in a large homogeneous population and is the type of problem we are interested in solving in this paper. The problem has previously been approached using *model based anomaly detection* methods, see e.g., Chandola et al. (2009); Chu et al. (2011) and Gorinevsky et al. (2012). This class of anomaly detection methods is suitable to detect anomalies in systems, as opposed to non-model based methods that are more suitable for finding anomalies in data. Model based anomaly detection methods work under the assumption that the dynamics of normal systems is the same, or equivalently, that the population of systems is homogeneous. The normal dynamics is modeled from system observations and most papers assume that an abnormal-free training data set is available for the estimation, see for instance (Abraham & Box, 1979; Abraham & Chuang, 1989; Fox, 1972). Some papers have been presented to relax this assumption. In e.g., Rousseeuw and Leroy (1987), the use of a regression technique robust to anomalies was suggested.

The detection of anomalous systems is in model based anomaly detection done by comparing system observations and model predictions and often done by a statistical test, see e.g., Desforges, Jacob, and Cooper (1998) and Eskin (2000). However, in non-model based anomaly detection, classification based (Duda, Hart, & Stork, 2000; Tan, Steinbach, & Kumar, 2005), clustering based (Jain & Dubes, 1988), nearest neighbor based (Tan et al., 2005, Chapter 2), information theoretic (Arning, Agrawal, & Raghavan, 1996) and spectral methods (Parra, Deco, & Miesbach, 1996) are also common. See Chandola et al. (2009) for a detailed review of anomaly detection methods. Most interesting and similar to the proposed method is the more recent approach taken in Chu et al. (2011) and Gorinevsky et al. (2012). They simultaneously estimate the regression model for the normal dynamics and perform anomaly detection. The method of Chu et al. (2011) is discussed further in the numerical section. There has also been some work on distributed anomaly detection, e.g., Chatzigiannakis, Papavassiliou, Grammatikou, and Maglaris (2006), Chu et al. (2011) and Zimmermann and Mohay (2006).

The main contribution of the paper is a novel distributed, scalable and model based method for anomaly detection in large homogeneous populations. The method is distributed in the sense that the computations can be distributed over the systems in the population or a cluster of computers. It is scalable since the size of the optimization problem solved on each system is independent of the number of systems in the population. This is made possible by a novel formulation of the multi-hypothesis problem as a sparse problem. The method also shows superior performance and is easier to tune than previously proposed model based anomaly detection methods. Lastly, the method does not need a training data set and a regression model of the normal dynamics is estimated at the same time as abnormal systems are detected. This is particularly valuable since often neither a training data set nor a regression model for the normal dynamics is available.

The remainder of the paper is organized as follows. Section 2 states the problem and shows the relation between anomaly detection and multi-hypothesis testing. Section 3 reformulates the multi-hypothesis problem as a sparse optimization problem and Section 4 gives a convex formulation. The convex problem is solved in a distributed manner on the systems and this is discussed in Section 5. We return to Example 1 and compare to the method of Chu et al. (2011) in Section 6. Finally, we conclude the paper in Section 7.

2. Problem statement and formulation

Assume that the population of interest consists of N systems. Think for example of the N airplanes studied in Example 1. Further assume that there is a linear unknown relation describing the

relation between measurable quantities of interest (angle of attack, the dynamic pressure, mass variation, the stabilizer deflection angle, and the elevator deflection in Example 1):

$$y_i(t) = \varphi_i^\top(t)\theta_{i,0} + e_i(t), \quad i = 1, \dots, N, \quad (1)$$

where t is the time index, i indexing systems, $y_i(t) \in \mathbb{R}$ and $\varphi_i(t) \in \mathbb{R}^m$ are the measurement and regressor vector at time t , respectively, $\theta_{i,0}$ is the *unknown* model parameter, and $e_i(t) \in \mathbb{R}$ is the measurement noise. For the i th system, $i = 1, \dots, N$, let $\{(y_i(t), \varphi_i(t))\}_{t=1}^\Omega$ denote the collected data set and Ω the number of observations collected on each system. We assume that $e_i(t)$ is white Gaussian distributed with mean zero and some *unknown* variance σ^2 and moreover, independent of $e_j(t)$ for all $i \neq j$. However, log-concave distributed noise could be handled with minor changes.

We will in the following say that the population behaves *normally* and that none of the systems are abnormal if $\theta_{1,0} = \dots = \theta_{N,0} = \theta_0$. Conversely, if any system has a model parameter deviating from the *nominal* parameter value θ_0 , we will consider that system as *abnormal*.

To solve the problem we could argue like this: suppose we have a hypothesis about which k systems are the anomalies. Then we could estimate the nominal parameters θ_0 by least squares from the rest, and estimate individual θ_i for the k anomalies. Since we do not know which systems are the anomalies, we have to do this for all possible hypotheses: choosing k systems from a set of N leads to a total of

$$c(N, k) = N! / ((N - k)!k!) \quad (2)$$

possible hypotheses. To decide which is the most likely hypothesis, we would evaluate the total misfit for all the systems, and choose that combination that gives the smallest total misfit. If we let γ_j be the set of assumed abnormal systems associated with the j th hypothesis $j = 1, \dots, c(N, k)$, this would be equivalent to solving the non-convex optimization problem

$$\begin{aligned} & \text{minimize} \sum_{j=1, \dots, c(N, k)} \min_{s \in \gamma_j} \sum_{t=1, \dots, \Omega} \|y_s(t) - \varphi_s^\top(t)\theta_{j,s}\|^2 \\ & + \min_{\theta_{j,0}} \sum_{s \notin \gamma_j, t=1, \dots, \Omega} \|y_s(t) - \varphi_s^\top(t)\theta_{j,0}\|^2. \end{aligned} \quad (3)$$

Since we assume that all systems have the same noise variance σ^2 , this is a formal hypothesis test. If the systems may have different noise levels we would have to estimate these and include proper weighting in (3).

The difficulty is how to solve (3) when the number of systems N is large. As seen in Example 1, even for rather small examples ($k = 3, N = 200$), the number of hypothesis $c(N, k)$ becomes large and solving problem (3) becomes computationally intractable.

3. Sparse optimization formulation

A key observation to be able to solve the anomaly detection problem in a computationally efficient manner is the reformulation of the multi-hypothesis problem (3) as a sparse optimization problem. To do this, first notice that the multi-hypothesis test (3) will find the k systems whose data are most likely to not have been generated from the same model as the remaining $N - k$ systems. Let us say that j^* was the selected hypothesis and denote the parameter of the i th system by θ_i , $i = 1, \dots, N$. Then $\theta_{i_1} \neq \theta_{i_2}$ for all $i_1, i_2 \in \gamma_{j^*}$ and $\theta_{i_1} = \theta_{i_2}$ for all $i_1, i_2 \in \{1, \dots, N\} \setminus \gamma_{j^*}$. Note that $N - k$ systems will have identical parameters. An equivalent way of solving the multi-hypothesis problem is therefore to maximize the likelihood under the constraint that $N - k$ systems are identical.

This can be formulated as

$$\begin{aligned} & \underset{\theta_1, \dots, \theta_N, \theta}{\text{minimize}} \sum_{i=1}^N \sum_{t=1}^{\Omega} \|y_i(t) - \varphi_i^T(t)\theta_i\|^2 \\ & \text{s.t.} \left\| \begin{bmatrix} \|\theta_1 - \theta\|_p & \|\theta_2 - \theta\|_p & \dots & \|\theta_N - \theta\|_p \end{bmatrix} \right\|_0 = k, \end{aligned} \quad (4)$$

where θ denotes the nominal parameter, $\|\cdot\|_0$ is the zero-norm (pseudo norm) which counts the number of non-zero elements of its argument. In this way, the k systems most likely to be abnormal could now be identified as the ones for which the estimated $\|\theta_i - \theta\|_p \neq 0$. Note that this is exactly the same problem as (3) and the same hypothesis will be selected.

Since abnormal model parameters and the nominal model parameter θ_0 are estimated from the given data sets $\{(y_i(t), \varphi_i(t))\}_{t=1}^{\Omega}, i = 1, \dots, N$, that are subject to the measurement noise $e_i(t), i = 1, \dots, N$, they are random variables. Moreover, it is well-known from Ljung (1999, p. 282) that if the given data sets $\{(y_i(t), \varphi_i(t))\}_{t=1}^{\Omega}, i = 1, \dots, N$, are *informative enough* (Ljung, 1999, Definition 8.1), then for each $i = 1, \dots, N$, the estimate of θ_i converges (as $\Omega \rightarrow \infty$) in distribution to the normal distribution with mean $\theta_{i,0}$ and covariance M_i/Ω where M_i is a constant matrix which depends on $\theta_{i,0}$ and the data batch $\{(y_i(t), \varphi_i(t))\}_{t=1}^{\Omega}$. This implies that as $\Omega \rightarrow \infty$, (4) will solve the anomaly detection problem exactly.

In the case where Ω is finite, our capability of correctly detecting the anomaly working systems will be dependent on the scale of the given Ω and σ , even with informative enough data sets $\{(y_i(t), \varphi_i(t))\}_{t=1}^{\Omega}, i = 1, \dots, N$. This is due to the fact that the variance of the estimate of $\theta_i, i = 1, \dots, N$, decreases as Ω and $1/\sigma$ increase. As a result, larger Ω and smaller σ allow us to detect smaller model parameter deviations and hence increase our ability to detect abnormal behaviors.

Remark 3.1. It should also be noted that if there is no overlap between observed features, anomalies cannot be detected even in the noise free case. That is, if we let Q be an $m \times m$ -matrix with all zeros except for element (q, q) , which equals 1, then if

$$\varphi_i^T(t_1)Q\varphi_j(t_2) = 0, \quad j = 1, \dots, i-1, i+1, \dots, N, \quad \forall t_1, t_2,$$

a deviation in element q in θ_i is not detectable. It can be shown that this corresponds to that the data is not informative enough.

4. Convex relaxation

It follows from basic optimization theory, see for instance (Boyd & Vandenberghe, 2004), that there exists a $\lambda > 0$, also referred to as the regularization parameter, such that

$$\begin{aligned} & \underset{\theta_1, \dots, \theta_N, \theta}{\text{minimize}} \sum_{i=1}^N \sum_{t=1}^{\Omega} \|y_i(t) - \varphi_i^T(t)\theta_i\|^2 \\ & + \lambda \left\| \begin{bmatrix} \|\theta_1 - \theta\|_p & \|\theta_2 - \theta\|_p & \dots & \|\theta_N - \theta\|_p \end{bmatrix} \right\|_0, \end{aligned} \quad (5)$$

gives exactly the same estimate for $\theta_1, \dots, \theta_N, \theta$, as (4). However, both (4) and (5) are non-convex and combinatorial, hence unsolvable in practice.

What makes (5) non-convex is the second term. It has recently become popular to approximate the zero-norm by its convex envelope. That is, to replace the zero-norm by the one-norm. This is in line with the reasoning behind Lasso (Tibsharani, 1996) and compressive sensing (Candès, Romberg, & Tao, 2006; Donoho, 2006). Relaxing the zero-norm by replacing it with the one-norm leads to the following convex optimization problem

$$\underset{\theta, \theta_1, \dots, \theta_N}{\text{minimize}} \sum_{i=1}^N \sum_{t=1}^{\Omega} \|y_i(t) - \varphi_i^T(t)\theta_i\|^2 + \lambda \sum_{i=1}^N \|\theta - \theta_i\|_p. \quad (6)$$

In theory, under some conditions on $\varphi_i(t)$, k and the noise, there exists a $\lambda^* > 0$ such that the criterion (6) will work essentially as well as (4) and detect the anomalous systems exactly. This is possible because (6) can be put into the form of group Lasso (Yuan & Lin, 2006) and the theory from compressive sensing (Candès et al., 2006; Donoho, 2006) can therefore be applied to establish when the relaxation is tight. This is reassuring and motivates the proposed approach in front of e.g., Chu et al. (2011) since no such guarantees can be given for the method presented in Chu et al. (2011).

In practice, λ should be chosen carefully because it decides k , the number of anomalous systems picked out, which correspond to k non-zeros among $\|\hat{\theta} - \hat{\theta}_i\|_p, i = 1, \dots, N$. Here $\hat{\theta}, \hat{\theta}_i, i = 1, \dots, N$ denote the optimal solution of (6). For $\lambda = 0$, all $\hat{\theta}_i, i = 1, \dots, N$, are in general different (for finite Ω) and all N systems can be regarded as anomaly. It can also be shown that there exists a $\lambda^{\max} > 0$ (it has closed form solution when $p = 2$) such that all $\hat{\theta}_i, i = 1, \dots, N$, equal the nominal estimate $\hat{\theta}$ and thus there are no anomalous systems picked out, if and only if $\lambda \geq \lambda^{\max}$. As λ increases from 0 to $\lambda^{\max}$, k decreases piecewise from N to 0. To tune λ , we consider:

- if k is known, λ can be tuned by trial and error such that solving (6) gives exactly k anomalous systems. Note that making use of $\lambda^{\max}$ can save the tuning efforts.
- if k is unknown and no prior knowledge other than the given data is available, the tuning of λ , or equivalently the tuning of k , becomes a model structure selection problem with different model complexity in terms of the number of anomalous systems. This latter problem can then be readily tackled by classical model structure selection techniques, such as Akaike's information criterion (AIC), Bayesian information criterion (BIC) and cross validation. If necessary, a model validation process, e.g., Ljung (1999, p. 509), can be employed to further validate whether the chosen λ , or equivalently the determined k is suitable or not.

In what follows, we assume a suitable λ has been found and focus on how to solve (6) in a distributed way.

Remark 4.1. The solutions from solving the problem in either (4) or (5) are indifferent to the choice of p . However, this is not the case for the problem in (6). In general, $p = 1$ is a good choice if one is interested in detecting anomalies in individual elements of the parameter. On the other hand, $p > 1$ is a better choice if one is interested in detecting anomalies in the parameter as a whole.

5. An ADMM based distributed algorithm

Although solving (6) in a centralized way provides a tractable solution to the anomaly detection problem, it can be *prohibitively expensive* to implement on a single system (computer) for large populations (large N). Indeed, in the centralized case, an optimization problem with $(N + 1)m$ optimization variables and all the data would have to be solved. As the number of systems N and the number of data Ω increase, this will lead to increasing computational complexity in both time and storage. Also note that many collections of systems (populations) are naturally distributed and lack a central computational unit.

In this section, we further investigate how to solve (6) in a distributed way based on ADMM (Alternating Direction Method of Multipliers). An advantage of our distributed algorithm is that each system (computer) only requires to solve a sequence of optimization problems with only $2m$ (independent of N) optimization variables. It requests very low storage space and in particular, it does

not access the data collected on the other systems. Another advantage is that it can guarantee the convergence to the centralized solution. In practice, it can converge to modest accuracy within a few tens of iterations that is sufficient for our use, see [Boyd, Parikh, Chu, Peleato, and Eckstein \(2011\)](#).

We will follow the procedure given in [Bertsekas and Tsitsiklis \(1997, Section 3.4.4\)](#) to derive the ADMM algorithm. First define

$$x = [\theta_1^T \ \theta_2^T \ \cdots \ \theta_N^T \ \theta^T]^T \in \mathbb{R}^{(N+1)m}. \quad (7)$$

The optimization problem (6) can then be rewritten as

$$\underset{x}{\text{minimize}} \ G(x). \quad (8)$$

Here, $G(x)$ is a convex function defined as

$$G(x) = \sum_{i=1}^N \|Y_i - \Phi_i \theta_i\|^2 + \lambda \|\theta - \theta_i\|_p \quad (9)$$

where for $i = 1, \dots, N$,

$$Y_i = [y_i(1) \ y_i(2) \ \cdots \ y_i(\Omega)]^T, \quad (10)$$

$$\Phi_i = [\varphi_i(1) \ \varphi_i(2) \ \cdots \ \varphi_i(\Omega)]^T.$$

As noted in [Bertsekas and Tsitsiklis \(1997, p. 255\)](#), the starting point of deriving an ADMM algorithm for the optimization problem (8) is to put (8) in the following form:

$$\underset{x}{\text{minimize}} \ G_1(x) + G_2(Ax) \quad (11)$$

where $G_1 : \mathbb{R}^{(N+1)m} \rightarrow \mathbb{R}$, $G_2 : \mathbb{R}^q \rightarrow \mathbb{R}$ are two convex functions and A is a suitably defined $q \times (N+1)m$ matrix. The identification of $G_1(\cdot)$, $G_2(\cdot)$ and A from (8) is crucial for a successful design of an ADMM algorithm. Here, the main concern is two fold: first, we have to guarantee that the two induced alternate optimization problems are separable with respect to each system; second, we have to guarantee $A^T A$ is nonsingular so that the derived ADMM algorithm is guaranteed to converge to the optimal solution of (8). We will get back to the convergence of proposed algorithm later in the section.

Having the above concern in mind, we identify

$$G_1(x) = 0, \quad (12)$$

$$G_2(z) = \sum_{i=1}^N \|Y_i - \Phi_i \alpha_i\|^2 + \lambda \|\beta_i - \alpha_i\|_p, \quad (13)$$

where

$$z = [\alpha_1^T \ \alpha_2^T \ \cdots \ \alpha_N^T \ \beta_1^T \ \beta_2^T \ \cdots \ \beta_N^T]^T \in \mathbb{R}^{2Nm}. \quad (14)$$

From (12) to (14), we have $G(x) = G_1(x) + G_2(Ax)$, with

$$A = \begin{bmatrix} I_{Nm} & 0 & \cdots & 0 \\ 0 & I_m & \cdots & I_m \end{bmatrix}^T \in \mathbb{R}^{2Nm \times (N+1)m} \quad (15)$$

where I_{Nm} and I_m are used to denote Nm and m dimensional identity matrix, respectively. Now the optimization problem (11) is equivalent to the following one:

$$\underset{x,z}{\text{minimize}} \ G_1(x) + G_2(z) \quad (16)$$

subj.to $Ax = z$.

Following [Bertsekas and Tsitsiklis \(1997, p. 255\)](#), we assign a Lagrange multiplier vector $v \in \mathbb{R}^{2Nm}$ to the equality constraint $Ax = z$ and further partition v as

$$v = [u_1^T \ u_2^T \ \cdots \ u_N^T \ w_1^T \ w_2^T \ \cdots \ w_N^T]^T \in \mathbb{R}^{2Nm} \quad (17)$$

where for $i = 1, \dots, N$, $u_i \in \mathbb{R}^m$, $w_i \in \mathbb{R}^m$. Moreover, we consider the augmented Lagrangian function

$$L_\rho(\theta, x, v) = G_1(x) + G_2(z) + v^T(Ax - z) + (\rho/2)\|Ax - z\|_2^2. \quad (18)$$

Then according to [Bertsekas and Tsitsiklis \(1997, \(4.79\)–\(4.81\)\)](#), ADMM can be used to approximate the solution of (16) as follows:

$$x^{(k+1)} = \arg \min_x \{G_1(x) + (v^{(k)})^T Ax + (\rho/2)\|Ax - z^{(k)}\|_2^2\}, \quad (19a)$$

$$z^{(k+1)} = \arg \min_z \{G_2(z) - (v^{(k)})^T z + (\rho/2)\|Ax^{(k+1)} - z\|_2^2\}, \quad (19b)$$

$$v^{(k+1)} = v^{(k)} + \rho(Ax^{(k+1)} - z^{(k+1)}), \quad (19c)$$

where ρ is any positive number and the initial vectors $v^{(0)}$ and $z^{(0)}$ are arbitrary. Taking into account (7)–(17), (19) can be put into the following specific form:

$$\theta_i^{(k+1)} = \alpha_i^{(k)} - u_i^{(k)} / \rho, \quad (20a)$$

$$\theta^{(k+1)} = (1/N) \sum_{i=1}^N \beta_i^{(k)} - w_i^{(k)} / \rho, \quad (20b)$$

$$\{\alpha_i^{(k+1)}, \beta_i^{(k+1)}\} = \arg \min_{\alpha_i, \beta_i} \|Y_i - \Phi_i \alpha_i\|^2 + \lambda \|\beta_i - \alpha_i\|_p - (u_i^{(k)})^T \alpha_i - (w_i^{(k)})^T \beta_i + (\rho/2)\|\theta_i^{(k+1)} - \alpha_i\|_2^2 + (\rho/2)\|\theta^{(k+1)} - \beta_i\|_2^2, \quad (20c)$$

$$u_i^{(k+1)} = u_i^{(k)} + \rho(\theta_i^{(k+1)} - \alpha_i^{(k+1)}), \quad (20d)$$

$$w_i^{(k+1)} = w_i^{(k)} + \rho(\theta^{(k+1)} - \beta_i^{(k+1)}), \quad (20e)$$

where $i = 1, \dots, N$. It is worth to note that computations in (20) are separable with respect to each system. Therefore, we yield the following ADMM based distributed algorithm to the optimization problem (6):

Algorithm 1. On the i th system, $i = 1, \dots, N$, do the following:

- (1) Initialization: set the values of $\alpha_i^{(0)}$, $\beta_i^{(0)}$, $u_i^{(0)}$, $w_i^{(0)}$, ρ and $k = 0$.
- (2) $\theta_i^{(k+1)} = \alpha_i^{(k)} - u_i^{(k)} / \rho$
- (3) Broadcast $\beta_i^{(k)}$, $w_i^{(k)}$ to the other systems, $j = 1, \dots, i-1, i+1, \dots, N$.
- (4) $\theta^{(k+1)} = (1/N) \sum_{i=1}^N \beta_i^{(k)} - w_i^{(k)} / \rho$
- (5) $\{\alpha_i^{(k+1)}, \beta_i^{(k+1)}\} = \arg \min_{\alpha_i, \beta_i} \{\|Y_i - \Phi_i \alpha_i\|^2 + \lambda \|\beta_i - \alpha_i\|_p - (u_i^{(k)})^T \alpha_i - (w_i^{(k)})^T \beta_i + (\rho/2)\|\theta_i^{(k+1)} - \alpha_i\|_2^2 + (\rho/2)\|\theta^{(k+1)} - \beta_i\|_2^2\}$
- (6) $u_i^{(k+1)} = u_i^{(k)} + \rho(\theta_i^{(k+1)} - \alpha_i^{(k+1)})$
- (7) $w_i^{(k+1)} = w_i^{(k)} + \rho(\theta^{(k+1)} - \beta_i^{(k+1)})$
- (8) Set $k = k + 1$ and return to step 2.

5.1. Convergence of Algorithm 1

It is interesting and important to investigate if [Algorithm 1](#) would converge to the optimal solution of the optimization problem (6). The answer is affirmative. We have the following theorem to guarantee the convergence of [Algorithm 1](#), which is a straightforward application of [Bertsekas and Tsitsiklis \(1997, Chapter 3, Proposition 4.2\)](#) to the optimization problem (11).

Theorem 1. Consider (11). The sequences $\{x^{(k)}\}$, $\{z^{(k)}\}$ and $\{v^{(k)}\}$ generated by the ADMM based distributed [Algorithm 1](#) for any $\rho > 0$, initial vectors $v^{(0)}$ and $z^{(0)}$, converge. Moreover, the sequence $\{x^{(k)}\}$ converges to an optimal solution of the original problem (11). The

sequence $\{Ax^{(k)} - z^{(k)}\}$ converges to zero, and $\{v^{(k)}\}$ converges to an optimal solution of the dual problem of (11).

Proof. According to Bertsekas and Tsitsiklis (1997, Chapter 3, Proposition 4.2), we only need to check whether the optimal solution set of problem (11) is nonempty and $A^T A$ is invertible. These two assumptions are clearly satisfied. So the conclusion follows.

Remark 5.1. In practice, Algorithm 1 should be equipped with certain stopping criteria so that the iteration is stopped when a solution with satisfying accuracy is obtained. Another issue with Algorithm 1 is that its convergence may be slow. Faster convergence can be achieved by updating the penalty parameter ρ at each iteration. Here, the stopping criterion in Boyd et al. (2011, Section 3.3) and the updating rule of ρ in Boyd et al. (2011, Section 3.4.1) are used in Section 6.

6. Numerical illustration

In this section, we return to the aircraft anomaly detection problem discussed in Example 1 and illustrate how Algorithm 1 can distributedly detect abnormal systems. In particular, we consider a fleet of 200 aircrafts and have access to data from 500 flights. These data are assumed to be related through the linear relation (1) where $m = 4$, $N = 200$, $\Omega = 500$, and for the i th aircraft at the t th flight, $y_i(t)$ represents the angle-of-attack and $\varphi_i(t)$ consists of mass variations, dynamic pressure, the stabilizer deflection angle times dynamic pressure and the elevation deflection angle times dynamic pressure, and $e_i(t)$ is assumed to be Gaussian distributed with mean 0 and variance 0.83, i.e., $e_i(t) \sim \mathcal{N}(0, 0.83)$.

Similar to Chu et al. (2011), to test the robustness of the proposed approach, we allow some nominal variation and generate the nominal parameter $\theta_{i,0}$ for a normal aircraft as a random variable as $\theta_{i,0} \sim \mathcal{N}(\bar{\theta}, \Sigma_\theta)$ with

$$\bar{\theta} = \begin{bmatrix} 0.8 \\ -2.7 \\ -0.63 \\ 0.46 \end{bmatrix}, \quad \Sigma_\theta = \begin{bmatrix} 0.04 & 0.12 & -0.02 & 0.02 \\ 0.12 & 0.84 & -0.09 & 0.1 \\ -0.02 & -0.09 & 0.03 & 0 \\ 0.02 & 0.1 & 0 & 0.05 \end{bmatrix}.$$

To simulate the anomalous aircrafts in the fleet, we generate $\theta_{i,0}$ for an anomalous aircraft as a random variable as $\theta_{i,0} \sim \mathcal{N}(\tilde{\theta}, \Sigma_\theta)$ with $\tilde{\theta} = [3.5 \ -0.1 \ -3 \ 0.001]^T$. Moreover, aircrafts with tags 27, 161 and 183 were simulated as anomaly. The regressor $\varphi_i(t)$ for all aircrafts at all flights is generated as a random variable as $\varphi_i(t) \sim \mathcal{N}(\bar{\phi}, \Sigma_\phi)$ with

$$\bar{\phi} = \begin{bmatrix} 0.95 \\ -1.22 \\ -2.79 \\ 7.11 \end{bmatrix}, \quad \Sigma_\phi = \begin{bmatrix} 0.25 & -0.02 & 0.12 & -0.04 \\ -0.02 & 0.45 & 0.03 & -0.52 \\ 0.12 & 0.03 & 1.05 & -1.26 \\ -0.04 & -0.52 & -1.26 & 3.89 \end{bmatrix}.$$

It should be noted that this simulation setup was previously motivated and discussed in Chu et al. (2011) and Gorinevsky et al. (2012).

Fig. 1 shows the centralized solution (top plot) and decentralized solution (bottom plot) of (6). The solutions are close to identical, which is expected given the convergence result of Section 5.1. The achieved performance of Algorithm 1 was obtained within 15 iterations and λ was set to 150 to pick out the three most likely airplanes to be abnormal. The aircrafts for which $\|\hat{\theta}_i - \hat{\theta}\| \neq 0$ were the 27th, 161st and 183rd aircrafts, which also were the true abnormal aircrafts.

Fig. 1. The difference between the estimates of the nominal model parameter $\hat{\theta}$ and the parameter $\hat{\theta}_i$ for the i th aircraft, $i = 1, \dots, N$. The top plot shows the solution of (6) computed by a centralized algorithm. The plot at the bottom illustrates the distributed solution computed by Algorithm 1. Both the centralized and decentralized solutions were computed using a regularization parameter, λ , set to 150.

Fig. 2. The difference between the estimates of the nominal model parameter $\hat{\theta}$ and the parameter $\hat{\theta}_i$ for the i th aircraft, $i = 1, \dots, N$. The plots show the solution obtained by the method presented in Chu et al. (2011) with regularization parameters 10, 100 and 400, respectively from the top to bottom. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Fig. 2 illustrates the performance of the method presented in Chu et al. (2011) with three different regularization parameters, namely, 10, 100 and 400, shown from the top to bottom respectively. As seen in the plots, $\|\hat{\theta}_i - \hat{\theta}\| \neq 0$ also for normally working aircrafts leading to a result which is difficult to interpret. In fact, $\|\hat{\theta}_i - \hat{\theta}\|$, $i = 1, \dots, N$, will in general always be greater than zero, independent of the choice of regularization parameter. This is a well known result (see for instance Hastie, Tibshirani, & Friedman, 2001, Section 3.4.3) and follows from the use of Tikhonov regularization in Chu et al. (2011). We use a sum-of-norms regularization (which essentially is a ℓ_1 -regularization on norms) and solving (6) therefore gives, as long as the regularization parameter is large enough, $\|\hat{\theta}_i - \hat{\theta}\| = 0$ for some i 's. The result of Chu et al. (2011) could of course be thresholded (the solid red line in Fig. 2 shows a suitable threshold). However, this adds an extra tuning parameter and makes the tuning more difficult than for the proposed method.

Remark 6.1. A centralized algorithm for solving (6) would for this example have to solve an optimization problem with 804

Fig. 3. The computational time for solving (6) in a centralized and distributed manner, for different number of aircrafts, N . The dashed line shows the computational time for solving the problem in a centralized manner. The solid line depicts the average computational time for each aircraft when the problem is solved in a distributed manner.

optimization variables. If Algorithm 1 is used to distribute the computations, an optimization problem with 8 optimization variables would have to be solved on each system (computer). This comparison illustrates that Algorithm 1 imposes a much cheaper computational cost per iteration on each computer. In addition, the number of optimization variables is invariant to the number of systems and data. Hence, Algorithm 1 provides a computationally tractable and scalable solution to anomaly detection in large homogeneous populations. For illustration, we have compared the computational time for solving (6) in a centralized and distributed manner for different N , see Fig. 3. The data used for this experiment were randomly generated in the same way as discussed in the beginning of this section. The only difference was that we considered different values for N . In particular, N was set to [60, 70, 80, 90, 100, 125, 150, 175, 200, 250, 300, 350, 400, 450, 500], respectively.

Remark 6.2. We used 500 flights instead of 5000 flights (5000 flights were used in Example 1 and Chu et al., 2011) to make the estimation problem more challenging. Notice that the anomaly detection problem still leads to a computationally challenging multi-hypothesis problem with roughly 1.3×10^6 hypothesis.

7. Conclusion

This paper has presented a novel distributed, scalable and model based approach to anomaly detection for large populations. The motivation for the presented approach is:

- it leads to a scalable approach to anomaly detection that can handle large populations,
- it provides a purely distributed method for detecting abnormally working systems in a collection of systems,
- the method does not require a training data set, and
- the algorithm is theoretically motivated by the results derived in the field of compressive sensing.

The algorithm is based on ideas from system identification and distributed optimization. Basically the anomaly detection problem is first formulated as a sparse optimization problem. This combinatorial multi-hypothesis problem cannot be solved for practically interesting sizes of data and a relaxation is therefore proposed. The convex relaxation can be written as a group Lasso problem and theory developed for Lasso and compressive sensing can therefore be used to derive theoretical bounds for when the relaxation is tight.

Acknowledgments

This work is partially supported by the Swedish Research Council in the Linnaeus center CADICS, the Swedish Department of Education within the ELLIT project and the European Research Council under the advanced grant LEARN, contract 267381. Ohlsson is also supported by a postdoctoral grant from the Sweden–America Foundation, donated by ASEA's Fellowship Fund, and by a postdoctoral grant from the Swedish Science Foundation.

References

- Abraham, B., & Box, G. E. P. (1979). Bayesian analysis of some outlier problems in time series. *Biometrika*, 66, 229–236.
- Abraham, B., & Chuang, A. (1989). Outlier detection and time series modeling. *Technometrics*, 31, 241–248.
- Arning, A., Agrawal, R., & Raghavan, P. (1996). A linear method for deviation detection in large databases. In *Proceedings of 2nd international conference of knowledge discovery and data mining* (pp. 164–169).
- Basseville, M., & Nikiforov, I. V. (1993). *Detection of abrupt changes—theory and application*. Englewood Cliffs, NJ: Prentice-Hall.
- Bertsekas, D. P., & Tsitsiklis, J. N. (1997). *Parallel and distributed computation: numerical methods*. Athena Scientific.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. In *Foundations and trends in machine learning*.
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge University Press.
- Candès, E. J., Romberg, J., & Tao, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52, 489–509.
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: a survey. *ACM Computing Surveys*, 41(15), 1–58.
- Chatzigiannakis, V., Papavassiliou, S., Grammatikou, M., & Maglaris, B. (2006). Hierarchical anomaly detection in distributed large-scale sensor networks. In *Proceedings of the 11th IEEE symposium on computers and communications, ISCC'06* (pp. 761–767).
- Chu, E., Gorinevsky, D., & Boyd, S. (2011). Scalable statistical monitoring of fleet data. In *Proceedings of the 18th IFAC world congress* (pp. 13227–13232). Milan, Italy.
- Desforges, M., Jacob, P., & Cooper, J. (1998). Applications of probability density estimation to the detection of abnormal conditions in engineering. *Proceedings of Institute of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 212, 687–703.
- Donoho, D. L. (2006). Compressed sensing. *IEEE Transactions on Information Theory*, 52(4), 1289–1306.
- Duda, R. O., Hart, P. E., & Stork, D. G. (2000). *Pattern classification* (2nd ed.). Wiley-Interscience.
- Eskin, E. (2000). Anomaly detection over noisy data using learned probability distributions. In *Proceedings of the seventeenth international conference on machine learning* (pp. 255–262). Morgan Kaufmann Publishers Inc.
- Fox, A. J. (1972). Outliers in time series. *Journal of the Royal Statistical Society. Series B (Methodological)*, 34, 350–363.
- Gorinevsky, D., Matthews, B., & Martin, R. (2012). Aircraft anomaly detection using performance models trained on fleet data. In *2012 Conference on intelligent data understanding, CIDU* (pp. 17–23).
- Gustafsson, F. (2001). *Adaptive filtering and change detection*. New York: Wiley.
- Hastie, T., Tibshirani, R., & Friedman, J. (2001). *Springer series in statistics, The elements of statistical learning: data mining, inference, and prediction* (10th ed.). New York, NY, USA: Springer New York Inc.
- Jain, A. K., & Dubes, R. C. (1988). *Algorithms for clustering data*. Prentice-Hall, Inc.
- Ljung, L. (1999). *System identification—theory for the user* (2nd ed.). Upper Saddle River, NJ: Prentice-Hall.
- Parra, L., Deco, G., & Miesbach, S. (1996). Statistical independence and novelty detection with information preserving nonlinear maps. *Neural Computing*, 8, 260–269.
- Patton, R., Frank, P., & Clark, R. (1989). *Fault diagnosis in dynamic systems—theory and application*. Prentice Hall.
- Rousseeuw, P. J., & Leroy, A. M. (1987). *Robust regression and outlier detection*. New York, NY, USA: John Wiley & Sons, Inc.
- Tan, P.-N., Steinbach, M., & Kumar, V. (2005). *Introduction to data mining*. Addison-Wesley.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1), 267–288.
- Yuan, M., & Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 68, 49–67.
- Zimmermann, J., & Mohay, G. (2006). Distributed intrusion detection in clusters based on non-interference. In *Proceedings of the 2006 Australasian workshops on grid computing and e-research—Vol. 54, ACSW Frontiers'06* (pp. 89–95). Darlinghurst, Australia.

Henrik Ohlsson was born in Sweden in 1981. He received the M.Sc. degree in Applied Physics and Electrical Engineering in October 2006 and his Ph.D. degree in Automatic Control in November 2010, all from Linköping University, Sweden. He has held visiting positions at the University of Cambridge and the University of Massachusetts. His research interests are mainly within the areas of system identification and machine learning. Henrik is currently a visiting scholar at University of California at Berkeley and an Assistant Professor at Linköping University.

Tianshi Chen was born in China in November 1978. He received his M.E. degree from Harbin Institute of Technology in 2005, and Ph.D. from Chinese University of Hong Kong in December 2008. From April 2009 to March 2011, he was a postdoctoral researcher at the Automatic Control group of Linköping University, Linköping, Sweden. His research interests are mainly within the areas of system identification, statistical signal processing, and nonlinear control theory and applications. He is currently an Assistant Professor at Linköping University.

Sina Khoshfetrat Pakazad was born in Iran in 1985. He received his M.Sc. degree in Systems, Control and Mechatronics in January 2009 from Chalmers University, Sweden, and his Licentiate degree in Automatic Control in December 2011 from Linköping University, Sweden. His research interests include distributed and scalable optimization in control, system identification and signal processing. He is currently a Ph.D. student at Linköping University.

Lennart Ljung received his Ph.D. in Automatic Control from Lund Institute of Technology in 1974. Since 1976 he is Professor of the chair of Automatic Control in Linköping, Sweden. He has held visiting positions at Stanford and MIT and has written several books on System Identification and Estimation. He is an IEEE Fellow, an IFAC Fellow and an IFAC Advisor. He is a member of the Royal Swedish Academy of Sciences (KVA), a member of the Royal Swedish Academy of Engineering Sciences (IVA), an Honorary Member of the Hungarian Academy of Engineering, an Honorary Professor of the Chinese

Academy of Mathematics and Systems Science, and a Foreign Associate of the US National Academy of Engineering (NAE). He has received honorary doctorates from the Baltic State Technical University in St Petersburg, from Uppsala University, Sweden, from the Technical University of Troyes, France, from the Catholic University of Leuven, Belgium and from Helsinki University of Technology, Finland. In 2002 he received the Quazza Medal from IFAC, and in 2003 he received the Hendrik W. Bode Lecture Prize from the IEEE Control Systems Society, and he was the 2007 recipient of the IEEE Control Systems Award.

S. Shankar Sastry received his Ph.D. degree in 1981 from the University of California, Berkeley. He was on the faculty of MIT as Assistant Professor from 1980 to 1982 and Harvard University as a chaired Gordon Mc Kay professor in 1994. He is currently the Dean of Engineering at University of California, Berkeley. His areas of personal research are embedded control especially for wireless systems, cybersecurity for embedded systems, critical infrastructure protection, autonomous software for unmanned systems (especially aerial vehicles), computer vision, nonlinear and adaptive control, control of hybrid and embedded systems,

and network embedded systems and software. He has supervised over 60 doctoral students and over 50 M.S. students to completion. His students now occupy leadership roles in several places and on the faculties of many major universities. He has coauthored over 450 technical papers and 9 books. Dr. Sastry served on the editorial board of numerous journals, and is currently an Associate Editor of the IEEE Proceedings.