
Robust Mixtures in the Presence of Measurement Errors

Jianyong Sun

School of Computer Science and School of Physics & Astronomy, University of Birmingham, Birmingham, UK

J.SUN@CS.BHAM.AC.UK

Ata Kabán

School of Computer Science, University of Birmingham, Birmingham, UK

A.KABAN@CS.BHAM.AC.UK

Somak Raychaudhury

School of Physics & Astronomy, University of Birmingham, Birmingham, UK

SOMAK@STAR.SR.BHAM.AC.UK

Abstract

We develop a mixture-based approach to robust density modeling and outlier detection for experimental multivariate data that includes measurement error information. Our model is designed to infer atypical measurements that are not due to errors, aiming to retrieve potentially interesting peculiar objects. Since exact inference is not possible in this model, we develop a tree-structured variational EM solution. This compares favorably against a fully factorial approximation scheme, approaching the accuracy of a Markov-Chain-EM, while maintaining computational simplicity. We demonstrate the benefits of including measurement errors in the model, in terms of improved outlier detection rates in varying measurement uncertainty conditions. We then use this approach for detecting peculiar quasars from an astrophysical survey, given photometric measurements with errors.

1. Introduction

The goal in robust unsupervised data modeling is to capture the structure of the typical observations while dealing with atypical or outlying observations in an automated manner. Outliers can occur for various reasons, such as unusually large measurement errors or the existence of peculiar objects in a data set. If atypical observations exist and are not properly dealt with, they lead to biases in the parameter estimates

and poor generalization of the structure inferred from the data. Therefore, a great deal of effort has been invested into modifying existing unsupervised methods to provide them with robustness properties. In statistics and statistical machine learning communities, the Student t -distribution was put forth and adopted as a robust building block, for clustering (Peel & McLachlan, 2000; Svensen & Bishop, 2005), visualization (Vellido et al., 2006) and robust projections (Archambeau et al., 2006). The t -distribution has heavy tails, hence it gives non-zero probability to observations that are far away from the bulk of the density.

Apart from the issue of robustness of the parameter estimates, the ability of detecting outliers is of special interest in certain scientific areas such as in Astrophysics (Djorgovski et al., 2001), where finding peculiar objects from large archives of multi-wavelength astronomical images provide a unique means of identifying candidates of possibly new types of objects that deserve more detailed follow-up study (e.g. using spectroscopy). However, a bottleneck already anticipated in Djorgovski et al. (2001) is ‘the likely overabundance of interesting objects found’ – the interpretation and understanding of which will necessitate costly detailed analysis. Indeed, not every atypical observation is truly interesting. One reason for this lies in measurement errors resulting from uncertainties in instrument calibration, physical limitations of devices and experimental conditions. These errors are typically carefully recorded in the case of scientific data and are available. Yet, most existing data analysis methods have no natural ways of taking these into consideration. In turn, neglecting the error information holds the risk of compromising the accuracy with which *genuine* outliers can be detected, since there is nothing to prevent us from confusing erroneous measurements with potentially interesting rare or peculiar ones.

Appearing in *Proceedings of the 24th International Conference on Machine Learning*, Corvallis, OR, 2007. Copyright 2007 by the author(s)/owner(s).

In classical statistics, models known as ‘errors in variables’ exist, such as the total least square approach for robust regression (Huffel & Lemmerling, 2002). Probabilistic approaches able to propagate uncertainty have also started to appear recently (Liu et al., 2006; Girard et al., 2003) for certain problems, and their benefits have been convincingly demonstrated. However we are aware of no work on including knowledge of observational errors specifically for unsupervised robust density modeling. Due to the importance of this issue in scientific data mining, this paper makes an attempt to fill in this gap.

2. Robust Mixtures for Data with Errors

Consider a data set in which each individual measurement is an estimate of the form $t_{in} \pm \sqrt{s_{in}}$, where $n = 1, \dots, N, i = 1, \dots, d$, N is the number of object instances and d is the number of features. It is conceptually justified to assume that the error associated with these individual measurements is normally distributed (e.g. Taylor (1996)). Organizing the square of errors into diagonal matrices \mathbf{S}_n , for each measured d -dimensional data point \mathbf{t}_n , the following heteroscedastic noise model can be written.

$$p(\mathbf{t}_n | \mathbf{w}_n) = \mathcal{N}(\mathbf{t}_n | \mathbf{w}_n, \mathbf{S}_n); \quad (1)$$

where $\mathcal{N}(\mathbf{t}_n | \mathbf{w}_n, \mathbf{S}_n)$ denotes the normal distribution with unknown mean \mathbf{w}_n and known diagonal covariance matrix \mathbf{S}_n .

The unknown mean values \mathbf{w}_n represent the clean, error-free version of the data. Since these cannot be directly observed, we will treat them as latent variables. The genuine outliers, in which we are interested, must be those of the density of \mathbf{w} rather than those of the density of \mathbf{t} . We will therefore model the hidden clean density as a robust mixture of Student t -distributions

(MoT)¹: $p(\mathbf{w}) = \sum_{k=1}^K \pi_k S_t(\mathbf{w} | \mu_k, \Sigma_k, \nu_k)$, where

$$S_t(\mathbf{w} | \mu_k, \Sigma_k, \nu_k) = \frac{\Gamma(\frac{\nu_k + d}{2}) |\Sigma_k|^{-1/2} \Gamma(\frac{\nu_k}{2})^{-1} (\nu_k \pi)^{-\frac{d}{2}}}{\left(1 + \frac{(\mathbf{w} - \mu_k)^T \Sigma_k^{-1} (\mathbf{w} - \mu_k)}{\nu_k}\right)^{\frac{\nu_k + d}{2}}}$$

By the use of t -densities, we make no assumptions on the distribution of outliers. Outliers are instances outside the high density ‘cluster’ regions.

As noted in Liu and Rubin (1995), with the introduction of an auxiliary hidden variable u , the t -distribution can be re-written as a convolution of a Gaussian with a Gamma placed on its precisions,

$$S_t(\mathbf{w} | \mu, \Sigma, \nu) = \int_0^\infty \mathcal{N}(\mathbf{w} | \mu, \frac{\Sigma}{u}) \mathcal{G}(u | \frac{\nu}{2}, \frac{\nu}{2}) du; \quad (2)$$

¹The instance indices n will be dropped for convenience, whenever their presence is obvious from the context.

where \mathcal{G} is the Gamma density, $\mathcal{G}(u | a, b) = b^a u^{a-1} \frac{\exp(-bu)}{\Gamma(a)}$. This re-writing has been exploited for developing an exact ML estimation algorithm for the MoT model (Peel & McLachlan, 2000).

In our model, the distribution of the observed data \mathbf{t} can be obtained by integration over \mathbf{w} . So we have:

$$p(\mathbf{t}) = \sum_k \pi_k \iint \mathcal{N}(\mathbf{t} | \mathbf{w}, \mathbf{S}) \mathcal{N}(\mathbf{w} | \mu_k, \frac{\Sigma_k}{u}) \mathcal{G}(u | \frac{\nu_k}{2}, \frac{\nu_k}{2}) dud\mathbf{w} \quad (3)$$

Thus, given a set of training data $\mathcal{Y} = (\mathbf{t}_1, \dots, \mathbf{t}_N)$, the complete probability model of the observed variable \mathbf{t} and the latent variables \mathbf{w}, u, z will have the following factorized form:

$$\mathcal{L}_C = \prod_n \prod_k [p(\mathbf{t}_n | \mathbf{w}_n) p(\mathbf{w}_n | u_n, z_n = k)]^{\delta(z_n = k)} \times \prod_n \prod_k [p(u_n | z_n = k) p(z_n = k | \pi)]^{\delta(z_n = k)} \quad (4)$$

where $\delta(\cdot)$ is the Kronecker delta. The plate diagram representation of this model is shown on the right-hand plot of Fig. 1, along with that of the MoT model.

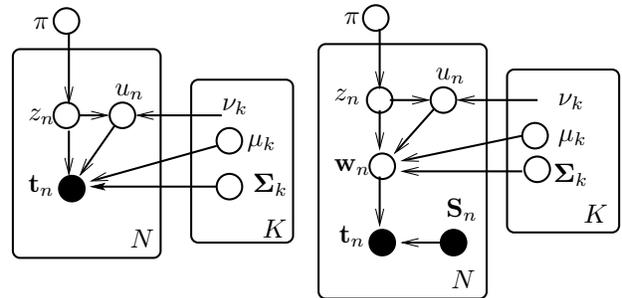


Figure 1. Plate diagrams of MoT (left), and the proposed model (right).

3. A Structured Variational EM Solution

Since the integration in Eq. (3) is not tractable, we develop a generalized EM (GEM) algorithm (see e.g. Hogg et al. (2005)), with approximate E-step. In general terms, for each data point \mathbf{t}_n , its log-likelihood can be written as follows, for any distribution q :

$$\begin{aligned} \log p(\mathbf{t}_n | \theta) &= \int q(h_n) \log \frac{p(h_n, \mathbf{t}_n | \theta)}{q(h_n)} \frac{q(h_n)}{p(h_n | \mathbf{t}_n, \theta)} dh_n \\ &\geq \int q(h_n) \log \frac{p(h_n, \mathbf{t}_n | \theta)}{q(h_n)} dh \equiv \mathcal{F}(\mathbf{t}_n | q, \theta) \end{aligned}$$

where q is the free-form variational family (or variational posterior), \mathcal{F} is called the variational free energy function, h_n is the set of latent variables associated with \mathbf{t}_n , and θ is the set of parameters of the model. In our case, $h_n = (z_n, \mathbf{w}_n, u_n)$ and $\theta =$

$(\{\mu_k\}, \{\Sigma_k\}, \{\nu_k\}, \pi)$. The log-likelihood of the given data set \mathcal{Y} is then lower bounded by the free energy:

$$\log p(\mathcal{Y}|\theta) = \sum_n \log p(\mathbf{t}_n|\theta) \geq \sum_n \mathcal{F}(\mathbf{t}_n|q(h_n), \theta) \quad (5)$$

In the E-step of the $(k+1)$ -th iteration of a GEM algorithm, we maximize \mathcal{F} w.r.t the variational distribution q while fixing the parameters in the k -th iteration, θ^k :

$$q^{k+1}(h_n) = \arg \max_q \mathcal{F}(\mathbf{t}_n|q, \theta^k). \quad (6)$$

In the M-step, we maximize Eq. (5) w.r.t the parameters θ to obtain the new parameter values θ^{k+1} :

$$\theta^{k+1} = \arg \max_{\theta} \sum_n \mathcal{F}(\mathbf{t}_n|q^{k+1}, \theta). \quad (7)$$

3.1. Tree-structured Variational Distribution

Some tractable form needs to be chosen for q . The most common choice is a fully factorial form (Jordan et al., 1999). In our case, this would be $q(\mathbf{w}, u, z) \equiv q(\mathbf{w})q(u)q(z)$. In the context of robust mixtures, fully factorial variational posterior distributions have been employed in (Svensen & Bishop, 2005), though with a slightly different model specification. Let us observe, however, that under our model definitions, it is feasible to keep some of the posterior dependencies by choosing the following tree-structured variational distribution:

$$q(\mathbf{w}, u, z = k) = q(z = k)q(\mathbf{w}|z = k)q(u|z = k)$$

Structured variational distributions have been used previously in the context of various other latent variable models (Geiger & Meek, 2005; Bishop & Winn, 2003) and have been found more accurate compared to the fully factorial choice. Yet, their use is still not as popular as it could be. In the following, we denote $q(z = k)$ by $q(k)$, $q(\mathbf{w}|z = k)$ by $q(\mathbf{w}|k)$ and $q(u|z = k)$ by $q(u|k)$. Also, expectations w.r.t. $q(\mathbf{w}|z = k)$ will be denoted by $\langle \cdot \rangle_{\mathbf{w}|k}$ and similarly, those w.r.t. $q(u|k)$ by $\langle \cdot \rangle_{u|k}$, and those w.r.t. the joint $q(\mathbf{w}, u|k) = q(\mathbf{w}|k)q(u|k)$ by $\langle \cdot \rangle_{\mathbf{w}, u|k}$.

3.2. Deriving the GEM Algorithm

The free energy function $\mathcal{F}(\mathbf{t}|q, \theta)$ can be evaluated as

$$\mathcal{F}(\mathbf{t}|q, \theta) = \sum_k q(k) [\langle \log p(\mathbf{t}, \mathbf{w}, u, k) \rangle_{\mathbf{w}, u|k}] + H(q)$$

where $H(q)$ is the entropy of the variational distribution: $H(q) = -\sum_k q(k) [\langle \log (q(u|k)q(\mathbf{w}|k)q(k)) \rangle_{\mathbf{w}, u|k}]$. Defining $A_{\mathbf{t}, k} = \langle \log p(\mathbf{t}, \mathbf{w}, u, k) \rangle_{\mathbf{w}, u|k} - \langle \log q(u|k) \rangle_{u|k} - \langle \log q(\mathbf{w}|k) \rangle_{\mathbf{w}|k}$ then we have:

$$\mathcal{F}(\mathbf{t}|q, \theta) = \sum_k q(k) [A_{\mathbf{t}, k} - \log q(k)]. \quad (8)$$

3.2.1. VARIATIONAL E-STEP

Now, in order to find the optimal functional form of the posterior distribution terms, we take functional derivatives of $\mathcal{F}(\mathbf{t}|q, \theta)$ w.r.t. the terms of q , i.e. $q(\mathbf{w}|k)$, $q(u|k)$ and $q(k)$ respectively, and equate these to the identically null function. We obtain the following:

$$q(\mathbf{w}|k) = \frac{\exp(\langle \log [p(\mathbf{t}|\mathbf{w})p(\mathbf{w}|u, k)] \rangle_{u|k})}{\int \exp(\langle \log [p(\mathbf{t}|\mathbf{w})p(\mathbf{w}|u, k)] \rangle_{u|k}) d\mathbf{w}} \quad (9)$$

$$q(u|k) = \frac{\exp(\langle \log [p(\mathbf{w}|u, k)p(u|k)] \rangle_{\mathbf{w}|k})}{\int \exp(\langle \log [p(\mathbf{w}|u, k)p(u|k)] \rangle_{\mathbf{w}|k}) du} \quad (10)$$

$$q(k) = \frac{\exp(A_{\mathbf{t}, k})}{\sum_{k'} \exp(A_{\mathbf{t}, k'})} \quad (11)$$

It can be seen that $q(\mathbf{w}|k)$ and $q(u|k)$ depend only on variables in their Markov blanket. However, the distribution $q(k)$ depends on all other variables in the graph. Conveniently, the quantities required for computing Eq. (11) will be available from the computations that are needed for evaluating the free energy function — which in turn is useful for monitoring the convergence of the GEM iterations.

Due to the conjugacy properties of the distributions we used, and after simplification, we now can obtain q analytically. Let us define:

$$\Sigma_{\mathbf{w}|k} = \mathbf{S} \left[\frac{\Sigma_k}{\langle u \rangle_{u|k}} + \mathbf{S} \right]^{-1} \frac{\Sigma_k}{\langle u \rangle_{u|k}} \quad (12)$$

$$\langle \mathbf{w} \rangle_k = \Sigma_{\mathbf{w}|k} (\langle u \rangle_{u|k} \Sigma_k^{-1} \mu_k + \mathbf{S}^{-1} \mathbf{t}) \quad (13)$$

$$a_k = \frac{\nu_k + d}{2}; \quad b_k = \frac{\nu_k + C_k}{2} \quad (14)$$

where

$$C_k = (\langle \mathbf{w} \rangle_k - \mu_k)^T \Sigma_k^{-1} (\langle \mathbf{w} \rangle_k - \mu_k) + \mathbf{Tr} (\Sigma_k^{-1} \Sigma_{\mathbf{w}|k}). \quad (15)$$

Then we have:

$$q(\mathbf{w}|k) = \mathcal{N}(\mathbf{w}|\langle \mathbf{w} \rangle_k, \Sigma_{\mathbf{w}|k}); \quad q(u|k) = \mathcal{G}(u|a_k, b_k). \quad (16)$$

3.2.2. THE VARIATIONAL LIKELIHOOD BOUND

$A_{\mathbf{t}, k}$ can be evaluated as follows:

$$\begin{aligned} A_{\mathbf{t}, k} &= \langle \log p(\mathbf{t}|\mathbf{w}) \rangle_{\mathbf{w}|k} + \langle \log p(u|k) \rangle_{u|k} + \log \pi_k + \\ &\quad \langle \log p(\mathbf{w}|u, k) \rangle_{\mathbf{w}, u|k} - \langle \log q(\mathbf{w}|k) \rangle_{\mathbf{w}|k} - \\ &\quad \langle \log q(u|k) \rangle_{u|k} \\ &= Q_1 + Q_2 + Q_3 + Q_4 + Q_5 + Q_6 \end{aligned} \quad (17)$$

where

$$Q_1 = -\frac{d}{2} \log(2\pi) - \frac{1}{2} \log |\mathbf{S}| - \frac{1}{2} \mathbf{Tr}(\Sigma_{\mathbf{w}|k} \mathbf{S}^{-1})$$

$$\begin{aligned}
 & -\frac{1}{2} \left[(\langle \mathbf{w} \rangle_k - \mathbf{t})^T \mathbf{S}^{-1} (\langle \mathbf{w} \rangle_k - \mathbf{t}) \right]; \\
 Q_2 &= \left(\frac{\nu_k}{2} - 1 \right) \langle \log u \rangle_{u|k} - \frac{\nu_k}{2} \frac{a_k}{b_k} \\
 & + \frac{\nu_k}{2} \log \left(\frac{\nu_k}{2} \right) - \log \Gamma \left(\frac{\nu_k}{2} \right); \\
 Q_3 &= -\frac{d}{2} \log(2\pi) - \frac{1}{2} \log |\boldsymbol{\Sigma}_k| + \frac{d}{2} \langle \log u \rangle_{u|k} \\
 & - \frac{1}{2} \frac{a_k}{b_k} \left[(\langle \mathbf{w} \rangle_k - \mu_k)^T \boldsymbol{\Sigma}_k^{-1} (\langle \mathbf{w} \rangle_k - \mu_k) \right] \\
 & - \frac{1}{2} \frac{a_k}{b_k} \text{Tr}(\boldsymbol{\Sigma}_{\mathbf{w}|k} \boldsymbol{\Sigma}_k^{-1}); \\
 Q_4 &= \log \pi_k; \\
 Q_5 &= \frac{d}{2} + \frac{d}{2} \log(2\pi) + \frac{1}{2} \log |\boldsymbol{\Sigma}_{\mathbf{w}|k}|; \\
 Q_6 &= -[(a_k - 1) \langle \log u \rangle_{u|k} + a_k \log b_k - a_k - \log \Gamma(a_k)];
 \end{aligned}$$

and where $\langle \log u \rangle_{u|k} = \psi(a_k) - \log b_k$ and $\psi(\cdot)$ is the di-gamma function.

In summary, given a data \mathcal{Y} , the log likelihood bound is computed cf. Eq. (8) as the following:

$$\mathcal{F} = \sum_n \sum_k q(z_n = k) [A_{\mathbf{t}_n, k} - \log q(z_n = k)] \quad (18)$$

where $A_{\mathbf{t}_n, k}$ is computed as in Eq. (17) for each data point \mathbf{t}_n . Eq. (18) is useful for monitoring the convergence.

3.2.3. M-STEP

The parameter re-estimates are obtained by solving the stationary equations of \mathcal{F} w.r.t μ_k , $\boldsymbol{\Sigma}_k$ and π_k , which yields:

$$\mu_k = \frac{\sum_{n=1}^N q(z_n = k) \langle u_n \rangle_{u_n|k} \langle \mathbf{w}_n \rangle_{\mathbf{w}_n|k}}{\sum_{n=1}^N q(z_n = k) \langle u_n \rangle_{u_n|k}} \quad (19)$$

$$\boldsymbol{\Sigma}_k = \frac{\sum_{n=1}^N q(z_n = k) \langle u_n \rangle_{u_n|k} \tilde{\boldsymbol{\Sigma}}_{n,k}}{\sum_{n=1}^N q(z_n = k)} \quad (20)$$

$$\pi_k = \frac{1}{N} \sum_{n=1}^N q(z_n = k) \quad (21)$$

where $\tilde{\boldsymbol{\Sigma}}_{n,k} = [(\mu_k - \langle \mathbf{w}_n \rangle_k)(\mu_k - \langle \mathbf{w}_n \rangle_k)^T + \boldsymbol{\Sigma}_{\mathbf{w}_n|k}]$. Finally, ν_k is re-estimated by solving the following non-linear equation.

$$\sum_n q(z_n = k) \left[\log \left(\frac{\nu_k}{2} \right) + 1 + \langle \log u_n \rangle_k - \frac{a_{nk}}{b_{nk}} - \psi \left(\frac{\nu_k}{2} \right) \right] = 0.$$

3.3. Scaling

Considering the time complexity of the algorithm, per iteration, computing the posterior mean and covariance ($\langle \mathbf{w} \rangle_k$ and $\boldsymbol{\Sigma}_{\mathbf{w}|k}$) for each data point \mathbf{t} takes $\mathcal{O}(d^3 K)$ operations. The computation of the parameters of $q(u|k)$, a_k and b_k take $\mathcal{O}(d^3 K)$, and the responsibility $q(k)$ needs $\mathcal{O}(d^3 K)$ time as well. In total, this is

$\mathcal{O}(d^3 KN)$. For comparison, the maximum likelihood estimation of MoT (Peel & McLachlan, 2000) takes $\mathcal{O}(d^3 K)$ to compute $p(u|k, \mathbf{t})$ and $\mathcal{O}(d^3 K)$ to compute $p(k|\mathbf{t})$, which totals a complexity of $\mathcal{O}(d^3 KN)$ — same as that of proposed algorithm. Moreover, using a full factorial approximation in our model also results in the same theoretical complexity per iteration. So the only extra burden of our proposed method is the computation of the posterior mean and covariance of the clean data \mathbf{w} . The most expensive operation appears to be the matrix inversion, however, it should be noted, this is only required when $\boldsymbol{\Sigma}_k$ are modeled as a full covariances, which is feasible in relatively low-dimensional problems ($d \ll N$). If this model was to be used on high dimensional data, then a diagonal form $\boldsymbol{\Sigma}_k$ would need to be taken — in which case the cubic operation is no longer required since the matrices to be inverted become diagonal.

3.4. Accommodating New Data Points

Since the model is fully generative, it can also be applied to new, previously unseen data from the same source. For a given test data set, we need to calculate the posterior distributions of \mathbf{w}_n and u_n associated with each test point \mathbf{t}_n . To calculate these, we fix the parameters μ_k , $\boldsymbol{\Sigma}_k$ and π_k , $1 \leq k \leq K$ obtained from the training set and perform the E-step iterations until convergence. This typically converges at least an order of magnitude faster than the full training procedure.

3.5. Determining the Number of Components

To determine the number of mixture components, we can employ the minimum message length (MML) principle (Figueiredo & Jain, 2002) by maximizing the following criterion:

$$\begin{aligned}
 \mathcal{L}(\theta, \mathcal{Y}) &= -\frac{\hat{n}}{2} \sum_{k: \pi_k > 0} \log \left(\frac{N\pi_k}{12} \right) - \frac{k_{nz}}{2} \log \left(\frac{N}{12} \right) \\
 & - \frac{k_{nz}(\hat{n} + 1)}{2} + \log p(\mathcal{Y}|\theta) \quad (22)
 \end{aligned}$$

where $p(\mathcal{Y}|\theta)$ is the data log-likelihood, \hat{n} is the dimensionality of the parameters, k_{nz} is the number of non-zero-probability components. The free parameters involved in the proposed algorithm are the means and the full covariance matrices of $\mathcal{N}(\mathbf{w}_n|u_n, z_n = k)$. Thus the dimensionality of the k -th parameter $\theta_k = (\mu_k, \boldsymbol{\Sigma}_k)$, is $d + d(d-1)/2$. We use our approximation to the data likelihood, given earlier in Eq. (5). Replacing this in (22), leads to maximizing:

$$\begin{aligned}
 \mathcal{L}(\theta, \mathcal{Y}) &\geq -\frac{\hat{n}}{2} \sum_{k: \pi_k > 0} \log \left(\frac{N\pi_k}{12} \right) - \frac{k_{nz}}{2} \log \left(\frac{N}{12} \right) \\
 & - \frac{k_{nz}(\hat{n} + 1)}{2} + \sum_n \mathcal{F}(\mathbf{t}_n|q, \theta) \equiv \tilde{\mathcal{L}}(\mathcal{Y}|q, \theta).
 \end{aligned}$$

This maximization is similar to the GEM presented in Section 3.1 and algorithmically the only difference is in computing the mixing proportions π_k in the M-steps, which is now:

$$\pi_k = \frac{\max \left\{ 0, \sum_{n=1}^N q(z_n = k) - \frac{\hat{n}}{2} \right\}}{\sum_{j=1}^K \max \left\{ 0, \sum_{n=1}^N q(z_n = j) - \frac{\hat{n}}{2} \right\}} \quad (23)$$

Of course, only the non-zero-probability components of the mixtures will contribute to $q(\mathbf{w}_n|z_n)$, $q(u_n|z_n)$ and $q(z_n)$.

4. The Outlier Detection Criteria

Since we modeled the clean, error-free data by a mixture of t -distributions, we would expect that the model can find outliers w.r.t the clean data, rather than the contaminated data. Following Peel and McLachlan (2000), the posterior expectation of u is interpretable as an outlieriness indicator. Using our posterior approximations described earlier (i.e. $q(u, k) = q(u|k)q(k)$ to approximate $p(u, k|\mathbf{t})$), then the variational expectation of u will be employed to infer outlieriness. This is:

$$e \equiv \sum_k q(k) \frac{\nu_k + d}{\nu_k + \mathbf{Tr}(\boldsymbol{\Sigma}_k^{-1} \boldsymbol{\Sigma}_{\mathbf{w}|k}) + \boldsymbol{\Delta}_{\mathbf{w}|k}^2} \quad (24)$$

where a_k and b_k are defined in Eq. (14); and $\boldsymbol{\Delta}_{\mathbf{w}|k}^2 = (\langle \mathbf{w} \rangle_{\mathbf{w}|k} - \mu_k)^T \boldsymbol{\Sigma}_k^{-1} (\langle \mathbf{w} \rangle_{\mathbf{w}|k} - \mu_k)$. Therefore, a data point is considered to be an outlier if its corresponding e value is sufficiently small.

In contrast, recall that for MoT, the outlier criterion value (Peel & McLachlan, 2000) is

$$e_{MoT} \equiv \sum_k p(k|\mathbf{t}) \frac{\nu_k + d}{\nu_k + (\mathbf{t} - \mu_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{t} - \mu_k)} \quad (25)$$

So we see that, instead of the Mahalanobis distance between the mean μ_k and the data \mathbf{t} , as in Eq. (25), we have the distance between the center μ_k and the expected value of the clean data \mathbf{w} in (24).

Further, it can easily be seen, for consistency, that in the limit of zero observation error, our outlieriness criterion reduces to that of MoT. Indeed, whenever $\mathbf{S} = \mathbf{0}$, Eqs. (12), (13) and (15) can be written as:

$$\begin{aligned} \boldsymbol{\Sigma}_{\mathbf{w}|k} &= \mathbf{0}; & \langle \mathbf{w} \rangle_{\mathbf{w}|k} &= \mathbf{t}; \\ C_k &= (\mathbf{t} - \mu_k)^T \boldsymbol{\Sigma}_k^{-1} (\mathbf{t} - \mu_k); \end{aligned}$$

Then replacing the above equations into Eqs. (10), (19), (20) and (21), we can recover the posteriors as:

$$q(u|k) = p(u|\mathbf{t}, k); q(k) = p(k|\mathbf{t});$$

and so, the update formulas of the MoT are recovered (Peel & McLachlan, 2000).

If the size of the measurement error \mathbf{S} (we can measure the size of \mathbf{S} by its trace) is small, we expect

the difference between e and e_{MoT} is relatively small too. However, as the size of the measurement error gets larger, the difference between the two outlieriness criteria becomes larger and consequently the ranking they produce will be different. In particular, we can gain more insights and see the effects of a misspecification of the error by rewriting the data likelihood (3) by integrating over \mathbf{w} :

$$p(\mathbf{t}_n) = \sum_k \pi_k \int \mathcal{N} \left(\mathbf{t}_n | \mu_k, \frac{\boldsymbol{\Sigma}_k}{u_n} + \mathbf{S}_n \right) \mathcal{G} \left(u_n | \frac{\nu_k}{2}, \frac{\nu_k}{2} \right) du_n \quad (26)$$

The posterior expectations $\langle u_n \rangle$ are data instance-specific, ensuring the robustness of the parameter estimates, even if the errors (diagonals of \mathbf{S}_n) are misspecified. However, this also implies that a data instance with an underestimated \mathbf{S}_n gets picked as a false ‘interesting’ outlier ($\langle u_n \rangle$ gets smaller). Clearly, if all errors are specified at zero, our model reduces to MoT and produces unwanted false detections.

5. Experiments and Results

To test the performance of the proposed algorithm, first we experimentally assess the accuracy of the structured factorization employed. Second, we perform a set of controlled experiments on synthetic and semi-synthetic data sets, in order to demonstrate the ability of detecting genuine outliers. Finally, we shall present a real application of our approach in astronomy, for finding peculiar (high redshift) objects from the SDSS quasar catalogue (York, 2000).

5.1. Synthetic Data & Illustrative Experiment

A synthetic data set is constructed comprising of error-free values sampled from a mixture of three well separated Gaussians and a uniform distribution simulates the presence of genuine outliers. Then we add Gaussian noise to all points, to simulate measurement errors, and apply our algorithm to the resulting dataset. The aim is to recover the genuine outliers (along with the density of non-outliers), despite the Gaussian noise added. The leftmost plot of Fig. 2 shows the error-free data, with the Gaussian covariances of the clusters of non-outliers superimposed. Different markers are used for points in different clusters and the outliers are marked with stars. The central plot shows the effect of simulating measurement errors. The marker sizes are proportional to the size of errors. Notice that due to the errors, some outliers appear closer to the main density regions while some of the non-outliers ‘jump’ away from the bulk of density. Thus the measurement errors make the problem of recovering genuine outliers much more challenging. The rightmost plot of Fig. 2 shows the result of our estimation procedure described earlier, superimposed over the data with errors.

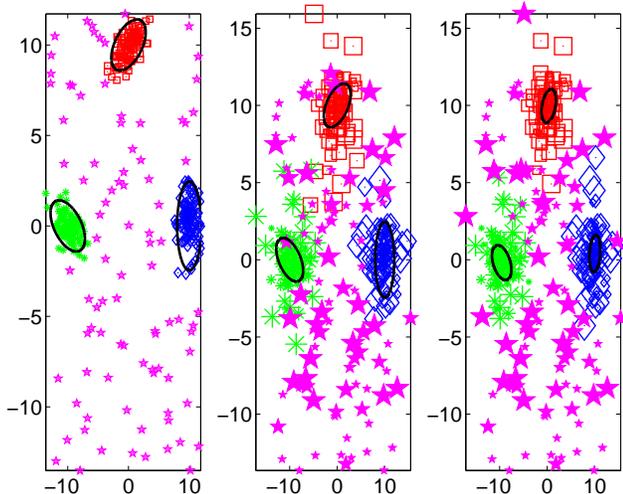


Figure 2. A synthetic data set with cluster structure and outliers. Left: Hidden error-free data with genuine outliers; Center: Data contaminated with measurement errors; Right: The estimated grouping and detected outliers.

5.2. Comparison of Alternative Approximate EM Methods

Now, we test the accuracy of the structured variational EM method developed here, against a fully factorial variational EM for the same model, and a Markov Chain EM (MCEM) realized through Gibbs sampling, the latter being considered to represent the ‘ground truth’. Fig. 3 shows the approximation of the log likelihood against iterations in a run on the synthetic data set shown earlier. For Gibbs sampling MCEM, $M = 10,000$ samples were used for computing the posterior estimates. The first 2000 samples were discarded as burn-in. All algorithms were started from the same initial parameter values. As expected, MCEM is supe-

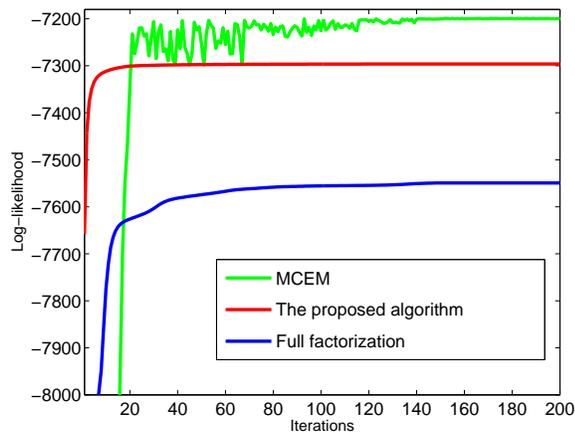


Figure 3. The optimization process of alternative approximate EM algorithms on the synthetic data set.

rior to variational methods, but at the price of a heavy computational demand and difficulties in determining its convergence. From the figures we also see the structured variational EM is closer to MCEM than the fully factorial variational EM. Therefore we use this method in the remainder of experiments reported².

5.3. Assessing the Accuracy of Detecting Genuine Outliers

To see how well we can detect genuine outliers, we start by carrying out a set of controlled experiments, varying the extent of measurement errors. We use our synthetic data sets and define five different measurement error levels: The diagonal elements of the error variance matrix \mathbf{S} will range between $[0,0.01]$, $[0,0.1]$, $[0,1]$, $[0,10]$ and $[0,100]$ respectively.

We perform receiver operating characteristics (ROC) analysis (Fawcett, 2004) to measure the performance. The area under the ROC curve (AUC) gives us the probability that a genuine outlier is detected. The MoT is employed as a baseline in our comparisons, in two instances. i) MoT applied to the clean data (which in real applications is not available) provides an idealized upper limit; ii) MoT applied to the data contaminated with observation errors provides a baseline against of which we measure our improvements. Fig. 4 summarizes the results obtained. For each of the 5 error conditions, the mean and standard deviation of the AUC values over repeated runs on 30 independent realizations of the data are shown: The upper plot shows the in-sample performance whereas the lower plot shows the out-of-sample performance, i.e. the ability to detect genuine outliers in previously unseen data from the same density model. The results are intuitive — we see a systematic and increasingly statistically significant improvement w.r.t. MoT/base, as the measurement uncertainty increases, both on in-sample data and on out-of-sample data.

In order to test our method further on data with a more realistic underlying density, while still being able to evaluate the benefits of using measurement error information in a controlled manner, we now apply our method on semi-synthetic data derived from real data, the lymphography data set (Blake & Merz, 1998). Originally, the data has four classes (148 data points in total and 18 dimensions), but two of them are quite small (2 and 4 data records), so we consider the two

²We also tested the variational methods, along with a MAP approach for u_n (using conjugate gradient optimization) for clustering accuracy, measured against the true labels. We found the structured variational EM superior and most stable in these tests as well.

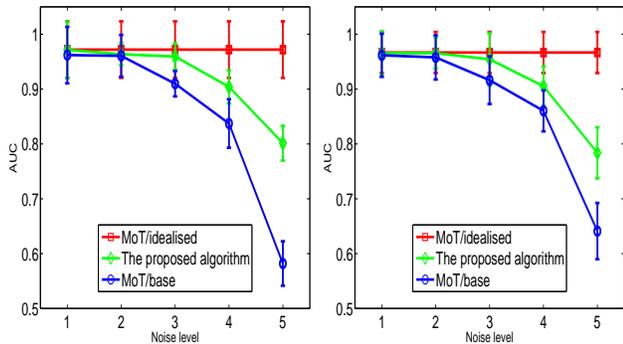


Figure 4. Comparison of MoT/idealized, MoT/base and the proposed algorithm on data sets with different levels of measurement error: in-sample (left) and out-of-sample (right).

small classes as outliers. We added heteroscedastic Gaussian noise with variances ranging between 0–0.1, to each observation, in order to simulate errors.

The algorithms were run on 10 independent realizations of the measurement errors, and the average ROC curves (Fawcett, 2004) and associated average AUC were then computed. The in-sample (93 data points) average AUC obtained by MoT/idealized is 0.9391, by MoT/base is 0.9005, whereas the proposed algorithm obtained 0.9555. The significance values of a t-test between MoT/idealized and the proposed algorithm was 0.39, while the value between the proposed algorithm and MoT/base was 9.6×10^{-5} . This suggests that the proposed algorithm performs comparably to MoT/idealized and significantly better than MoT/base in this experiment. Moreover, the out of sample (55 data points) performance is also of the same quality, illustrated in Fig. 5. We can conclude therefore, that knowledge of measurement errors is useful and can be exploited with the use of our approach to achieve a more accurate detection of genuine outliers.

5.4. Application to Detecting High-redshift Quasars from the SDSS Quasar Catalogue

In astrophysical measurements, there is no error-free situation (Taylor, 1996), but the measurement errors can be estimated for each feature and each object, from knowledge about observing conditions and instrumental limitations. Therefore, unlike above, error-free data is not available, which precludes the validation against an absolute ground truth. The data set analyzed here is extracted from a well-studied survey in astrophysics, the SDSS quasar catalog (York, 2000), which provides five magnitudes for a large number of quasars, representing their brightness measured with

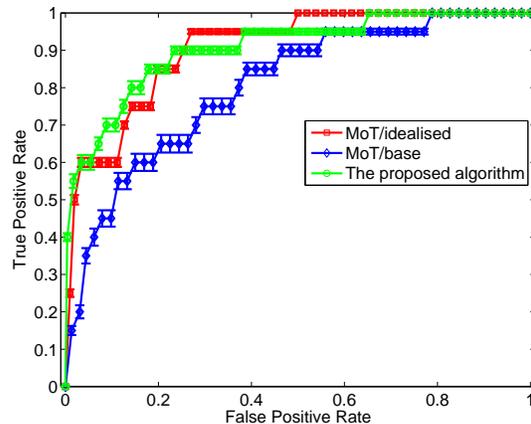


Figure 5. Out of sample average ROC curves of MoT/idealized, MoT/base and the proposed algorithm on the lymphography data set. The error bars represent one standard deviation.

five different optical filters u' g' r' i' and z' . From these, to avoid bias with brightness, we construct four features, each related to a color, by subtracting r' (which is the most reliably measured) from each of the others. In addition, spectroscopic redshift estimates are available for the subset we have extracted – these are not used within the algorithm, but are useful to in validating our results. The redshift is related to the distance of the object from the Earth, and very distant objects are rare. Given that with higher redshift, the entire spectral pattern is systematically shifted towards the redder end, there is physical reason for the rare quasars at very high redshift to be perceived as outliers in the overall density of quasars in the color space. This observation has been exploited in a number of previous studies for finding high redshift quasars in 2D projections of the multi-filter photometry (Fan, 2006). However, a comprehensive approach which both i) works in the multivariate feature space and ii) takes principled account of the measurement errors has not been available.

We apply our method to a sample of 10,000 quasars and compute the AUC values against a varying redshift threshold. The resulting relationship is shown in Fig. 6, for different choices of K . The optimal order determined by MML was $K = 2$, nevertheless, from the figure we see a remarkable robustness w.r.t. this choice. The y-coordinate of each point on these curves indicates the probability of detecting quasars of redshift greater than its x-coordinate. This plot shows clearly that our principled method in four-color space, using errors, can identify as outliers an overwhelming fraction of quasars already at a redshift of 2.5 (or higher), whereas the 2D projection methods, e.g.

Fan (2006), can manage to identify only those with $z > 3.5$, which are extremely rare, and obvious from naive projections. By being able to identify the latter category, when the SDSS galaxy catalogue is complete with four-color magnitudes, our method promises to retrieve an order of magnitude more interesting high-redshift quasars than existing methods would.

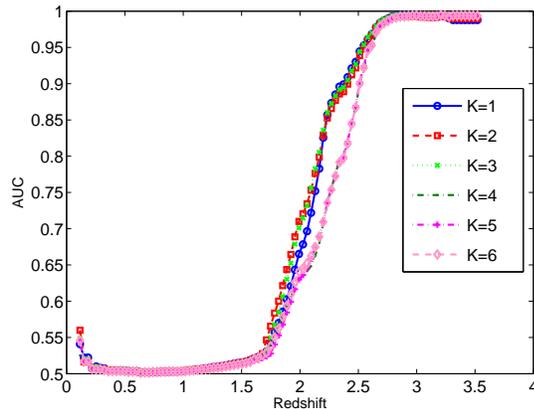


Figure 6. AUC versus possible redshift thresholds.

6. Conclusions

We propose a robust mixture model for multivariate data that includes error information, employing composite densities. We have derived a structured variational EM algorithm for inference and parameter estimation, which in the zero limit of the measurement errors reduces to maximum likelihood estimation of t -mixtures. Empirical results of a set of controlled experiments have shown a systematic and statistically significant improvement in terms of correct outlier detection rates in high measurement uncertainty conditions. Finally, a real application of our method to detecting peculiar, high-redshift quasars from the SDSS photometric quasar catalogue was demonstrated. Further work may concern extensions to robust projection models (Archambeau et al., 2006) and including an interactive visual element into the analysis of outliers for data with error information.

ACKNOWLEDGEMENTS

This research was supported by a grant from the Particle Physics and Astronomy Research Council, UK.

References

Archambeau, C., Delannay, N., & Verleysen, M. (2006). Robust probabilistic projections. *Proceedings of the 23rd International Conference on Machine Learning*.

- Bishop, C., & Winn, J. (2003). Structured variational distributions in VIBES. *Proc. AISTats*.
- Blake, C., & Merz, C. (1998). UCI repository of machine learning databases. www.ics.uci.edu/mllearn/MLRepository.html.
- Djorgovski, S., Mahabal, A., Brunner, R., Gal, R., & Castro, S. (2001). Searches for rare and new types of objects. *Virtual Observatories of the Future, ASP Conference Series, 225*.
- Fan, X. (2006). A survey of $z > 5.7$ quasars in the Sloan Digital Sky Survey. IV. discovery of seven additional quasars. *Astronomical Journal, 131*, 1203–1209.
- Fawcett, T. (2004). ROC graphs: Notes and practical considerations for researchers. *Machine Learning*.
- Figueiredo, M., & Jain, A. (2002). Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 24*, 381–396.
- Geiger, D., & Meek, C. (2005). Structured variational inference procedures and their realizations. In *Proc. AISTats*.
- Girard, A., Rasmussen, C., Candela, J. Q., & Murray-Smith, R. (2003). Multiple-step ahead prediction for non linear dynamic systems - a gaussian process treatment with propagation of the uncertainty. *Advances in Neural Information Processing Systems* (pp. 545–552).
- Hogg, R., McKean, J., & Craig, A. (2005). *Introduction to mathematical statistics*. Upper Saddle River, NJ: Pearson Prentice Hall.
- Huffel, S. V., & Lemmerling (Eds.). (2002). *Total least squares and errors-in-variables modeling*. Dordrecht: Kluwer Academic Publishers.
- Jordan, M., Ghahramani, Z., Jaakkola, T., & Saul, L. (1999). An introduction to variational methods for graphical models. *Machine Learning, 37*, 183–233.
- Liu, C., & Rubin, D. (1995). ML estimation of the t distribution using EM and its extensions: ECM and ECME. *Statistica Sinica, 5*, 19–39.
- Liu, X., Milo, M., Lawrence, N. D., & Rattray, M. (2006). Probe-level measurement error improves accuracy in detecting differential gene expression. *Bioinformatics, 22*, 2107–2113.
- Peel, D., & McLachlan, G. (2000). Robust mixture modelling using the t distribution. *Statistics and Computing, 10*, 339–348.
- Svensen, M., & Bishop, C. (2005). Robust bayesian mixture modelling. *Neurocomputing, 64*, 235–252.
- Taylor, J. (1996). *An introduction to error analysis*. University Science Books.
- Vellido, A., Lisboa, P., & Vicente, D. (2006). Handling outliers and missing data in brain tumor clinical assessment using t -GTM. *Computers in Biology and Medicine*.
- York, D. (2000). The sloan digital sky survey: Technical summary. *Astronomical Journal, 120*, 1579–1587.