Bayesian bacterial detection using irregularly sampled optical endomicroscopy images

Ahmed Karam Eldaly [a,b,c,*], Yoann Altmann [a,1], Ahsan Akram [b], Paul McCool [d], Antonios Perperidis [a,b], Kevin Dhaliwal [b], Stephen McLaughlin [a,2]

[a] Electrical and Electronic Engineering Department, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, United Kingdom
[b] MRC Centre for Inflammation Research, Queen's Medical Research Institute, The University of Edinburgh, Edinburgh, United Kingdom
[c] Biomedical Engineering Department, Faculty of Engineering, Helwan University, Cairo, Egypt
[d] Medical Devices Unit, NHS Greater Glasgow & Clyde, United Kingdom

ARTICLE INFO

ABSTRACT

Pneumonia is a major cause of morbidity and mortality of patients in intensive care. Rapid determination of the presence and gram status of the pathogenic bacteria in the distal lung may enable a more tailored treatment regime. Optical Endomicroscopy (OEM) is an emerging medical imaging platform with preclinical and clinical utility. Pulmonary OEM via multi-core fibre bundles has the potential to provide *in vivo, in situ,* fluorescent molecular signatures of the causes of infection and inflammation. This paper presents a Bayesian approach for bacterial detection in OEM images. The model considered assumes that the observed pixel fluorescence is a linear combination of the actual intensity value associated with tissues or background, corrupted by additive Gaussian noise and potentially by an additional sparse outlier term modelling anomalies (bacteria). The bacteria detection problem is formulated in a Bayesian framework and prior distributions are assigned to the unknown model parameters. A Markov chain Monte Carlo algorithm based on a partially collapsed Gibbs sampler is used to sample the posterior distribution of the unknown parameters. The proposed algorithm is first validated by simulations conducted using synthetic datasets for which good performance is obtained. Analysis is then conducted using two *ex vivo* lung datasets in which fluorescently labelled bacteria are present in the distal lung. A good correlation between bacteria counts identified by a trained clinician and those of the proposed method, which detects most of the manually annotated regions, is observed.

1. Introduction

Pneumonia is characterised by a host inflammatory response to a pathogenic infectious burden in the distal lung and is usually caused by bacterial infection (Akram et al., 2015). However, accurate quantitative diagnosis and monitoring of suspected pneumonia is a challenging task with currently available imaging/diagnostic tools (Thiberville et al., 2007). Structural imaging with X-ray or Computed Tomography (CT) scans are often non-specific. Consequently, current methodologies consist of culturing bronchoalveolar lavage fluid (BALF) retrieved from bronchoscopy, a process that often takes 48–72 h to yield results with low specificity and sensitivity (Fuchs et al., 2013). During this time, drugs and antimicrobials may be administered, some of which can be inappropriate to the underlying condition, potentially toxic, and may promote antimicrobial resistance (Akram et al., 2015). Thus, rapid and accurate diagnostic approaches to enable tailored treatment regimes are required.

Optical endomicroscopy (OEM) is an emerging imaging tool used both clinically and pre-clinically. Fibred Confocal Fluorescent Microscopy (FCFM), also referred to as probe-based confocal laser endomicroscopy (pCLE), is the most widely used platform and the only fiber-based endomicroscopic methodology approved for clinical use (Perperidis et al., 2018). The technology employs a laser

* Corresponding author at: Electrical and Electronic Engineering Department, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, United Kingdom.
E-mail addresses: ak577@hw.ac.uk (A.K. Eldaly), Y.Altmann@hw.ac.uk (Y. Altmann), Ahsan.Akram@ed.ac.uk (A. Akram), Paul.McCool@nhs.net (P. McCool), A.Perperidis@hw.ac.uk (A. Perperidis), Kev.Dhaliwal@ed.ac.uk (K. Dhaliwal), S.McLaughlin@ac.ac.uk (S. McLaughlin).

scanning proximal light source, linked to a flexible multi-core fibre. This fibre is passed through the working channel of endoscopes enabling microscopic imaging at its distal end. The lateral diameter of the fibre used in lung imaging applications is 1.4 mm. This miniaturisation enables the exploration of the distal pulmonary tract as well as the assessment of the respiratory bronchioles and alveolar gas exchanging units of the distal lung (Thiberville et al., 2007). OEM has been used clinically in the lung for the detection of lung cancer (Fuchs et al., 2013; Thiberville et al., 2009b) and has been used to image the distal lung (Thiberville et al., 2009c; 2009a) including the imaging of parenchymal lung diseases (Newton et al., 2012).

In pulmonary OEM, elastin and collagen generate auto-fluorescence when excited with a 488 nm laser source. While bacteria can potentially auto-fluoresce when exposed to light at the appropriate wavelength(s), the emitted fluorescent signal is usually too weak to be imaged effectively and to provide any diagnostic information. To address this limitation, bacteria can be specifically labelled with targeted fluorescently labelled ligands (SmartProbes) (Akram et al., 2015), which bind to the bacteria to generate fluorescence in response to a light excitation at pre-determined wavelengths (e.g., 488 nm laser excitation) (Akram et al., 2015). Stained bacteria then appear as bright dots in the images, whereas elastin and collagen are naturally auto-florescent and present a mesh-like structure in the distal lung, making the detection of bacteria a more challenging task.

Recently, there has been a strong research interest in the development and deployment of custom optical molecular SmartProbes to image *in vivo* targets in clinic. Several candidate SmartProbes have been designed to target bacteria (Akram et al., 2015) and fibrogenesis (Aslam et al., 2015). In this work, we assess different combinations of bacteria and SmartProbes which yield different fluorescence levels.

The commercially available OEM platform (Cellvizio, Mauna Kea Technologies, Paris-France) acquires images at 8–12 frames per second and clinical/preclinical OEM procedures often last minutes, generating thousands of frames, hence making their manual analysis a very labour-intensive process. There has therefore been an increased interest in the automated analysis of OEM images. OEM tailored approaches predominantly concentrate on (i) the reconstruction of fibre bundle images (Eldaly et al., 2018a; 2018b; Perperidis et al., 2017b; Le Goualher et al., 2004; Savoire et al., 2012; Parker et al., 2019), (ii) the retrieval and classification of frames sequences for a more efficient quantitative analysis (Ishijima et al., 2015; Perperidis et al., 2017a; Koujan et al., 2018; André et al., 2011; Désir et al., 2012), and (iii) the automated mosaicking of temporally adjacent frames to create an extended and continuous field of view (Vercauteren et al., 2006; 2005). Bacterial detection and quantification approaches developed for confocal microscopy (Lenseigne et al., 2007; Veropoulos et al., 1998) are not directly applicable to OEM data due to major differences in the image acquisition approaches employed by the two technologies. In particular, confocal microscopy employs a regular rectangular grid to acquire images of high lateral and axial optical resolutions which are 0.18μ m and 0.5μ m respectively (Fouquet et al., 2015). In contrast, OEM employs a sparse and irregular pattern (due to the organisation of the fibers in the fiber bundle) to acquire images of moderate lateral and axial optical resolutions of 0.5μ m and 3μ m respectively, within a constantly moving scene (breathing, heart beat and fibre motion) (Jabbour et al., 2012). The real datasets considered in this work are acquired by the Cellvizio, Mauna Kea Technologies (Paris, France), confocal endomicroscopy system (Ayache et al., 2006), which includes a set of preprocessing steps, which then makes the observation noise difficult to model analytically. Thus, we assume here that the noise is additive, white and Gaussian (McCool et al., 2016).

The object detection problem can be addressed in a variety of different annotation scenarios among which two widely used ones are the dot-annotation and the count-annotation (Arteta et al., 2014; von Borstel et al., 2016). While dot-annotation provides the location of objects that appear in the image, count-annotation only provides the number of objects in an image without explicitly revealing their locations. Dot-annotations are more informative given the small scale of the objects to be quantified. In this work, we assess the performance of our method by considering the number of detections (count-annotation) but also their estimated locations (dot-annotation).

In the literature, the object detection problem has been mainly addressed by either supervised approaches such as in (Arteta et al., 2016; Lenseigne et al., 2007; Veropoulos et al., 1998; Bonheur et al., 2019), or by unsupervised approaches such as in (Ding et al., 2011; Altmann et al., 2015; McCool et al., 2016; Lindeberg, 1998; He et al., 2010; Kimori et al., 2010; Bright and Steel, 1987; Eldaly et al., 2019). In supervised approaches, the dataset is usually divided into training and testing sets. During the training phase, a model is trained by pairing inputs with their expected outputs which are referred to as ground truth. This trained model is then used to estimate the output of the test dataset. While these algorithms are usually simple and fast, the trained model and its detection performance usually depends on the quality of the ground truth considered when training the model/classifier. Since bacteria annotation might be different from one clinician to another, considering such methods for cases where the ground truth can be highly subjective, can result in a biased trained model. Moreover, some of these models use convolutional neural networks (Arteta et al., 2016), which are problematic here due to lack of sufficient training images and bacteria annotations. On the other hand, several studies have considered unsupervised approaches to solve outlier detection problems. In these approaches, the objects to be detected are learned from the data by fitting them with suitable distributions without using explicitly-provided labels. Example of such approaches are hierarchical Bayesian models (Ding et al., 2011; Altmann et al., 2015; McCool et al., 2016) which offer a flexible and consistent methodology to deal with uncertainty in inference when limited amount of data or prior information is available. Moreover, other unknown parameters can be jointly estimated within the algorithm such as noise variance and regularization parameters. As such, they represent an attractive way to tackle ill-posed inverse problems. These methods rely on selecting an appropriate prior distribution for the unknown image and remaining unknown parameters. The full posterior distribution can then be derived using the Bayes' rule, and then exploited by optimization or simulation-based (Markov chain Monte Carlo) methods.

In this paper, motivated by the models considered in (Altmann et al., 2015; Ding et al., 2011; McCool et al., 2016), we develop an algorithm to detect bacteria in distal lung images, taken using an OEM system. We consider that each fibre core's spectrum is a linear combination of an actual intensity value corrupted by additive Gaussian noise, and possibly an additional term modelling outliers (which are considered to be candidate bacteria).

The main contributions of this work are threefold:

1. We develop an algorithm that can help to automatically detect labelled bacteria in optical endomicroscopy images. To the best of our knowledge, it is the first time this problem is addressed in a statistical framework by using a hierarchical Bayesian model. The statistical model does not rely on a particular spatial organisation of the fiber bundle and the algorithm is fully automatic in the sense that it allows the automated estimation of the model hyperparameters. Thus, it does not require the user to tune crucial parameters.
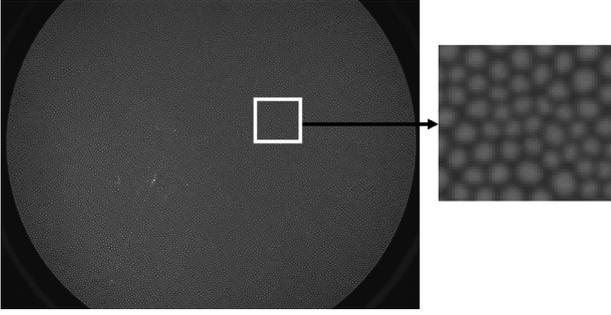
Fig. 1. A background image, and a zoomed part of the image.

2. We investigate different combinations of SmartProbes and bacterial concentrations including control cases for which no bacteria are present. Different SmartProbes that cause weak and stronger fluorescence are tested and the developed algorithm does not rely on strong assumptions about the SmartProbe used. We also show that the algorithm can differentiate controls from bacterial loads of different concentrations.

3. We compare the results of the proposed model with bacteria annotations performed by a trained clinician and four existing spot-detection algorithms, using both dot-annotation and count-annotation.

The remaining sections of the paper are organised as follows. Section 2 formulates the problem of bacteria detection in OEM images, followed by Section 3 which summarizes the likelihood and the prior distributions assigned to the unknown parameters of the model. The resulting joint posterior distribution and the partially collapsed Gibbs sampler used to sample that distribution are discussed in Section 4. Simulations conducted using synthetic and real datasets are presented in Sections 5 and 6 respectively. Limitations, conclusions and future work are finally reported in Sections 7 and 8 respectively.

2. Problem formulation

Fig. 1 depicts a background OEM image i.e., an image from a sample presenting constant intensity, and a zoomed-in region of this image. Bright and dark areas represent fiber cores and their cladding, respectively. This fibre bundle contains cores that are transmitting light and collecting fluorescence light simultaneously. Fibre cores contain information about the object being imaged, while the cladding (the space between the cores) does not. The underlying fluorescence profile is usually recovered by interpolating the fibre core intensities (Le Goualher et al., 2004; Eldaly et al., 2018a).

Fig. 2 shows a typical sequence of OEM frames which result from interpolated fibre core intensities, with labelled bacteria annotated by a trained clinician. We can observe that labelled bacteria appear as bright dots overlaid across a network of very bright elastin strands in these images.

Since only fibre cores hold information about the object being imaged, each image can be sparsely represented with only N measurements, associated with N fibers, where each fibre core is represented by a single intensity value. The size of OEM images in our case is 274×384 pixels and each of these images consists of $N = 12105$ fiber cores. This corresponds to approximately 8.69% of the original image pixels. The interpolated images contain the same amount of acquired information as the sets of core intensities, and these images are mostly used for visualization purposes. Moreover, processing only central core intensities, whose coordinates are known from factory pre-calibration, reduces the data vol-

ume to be processed and is expected to result in faster algorithms than when considering the whole interpolated images.

We consider a set of N observed sub-sampled intensities $\mathbf{y} = [y_1, \ldots, y_N]^T$. In a similar manner to (Altmann et al., 2015; Newstadt et al., 2014; Ding et al., 2011), each of these samples is assumed to result from a linear combination of an actual intensity value and additive Gaussian noise, potentially corrupted by additive outliers. The observation model can thus be expressed as

$$\mathbf{y} = \mathbf{x} + \mathbf{r} + \mathbf{w}, \tag{1}$$

with $\mathbf{x} = [x_1, \ldots, x_N]^T$, $\mathbf{r} = [r_1, \ldots, r_N]^T$ and $\mathbf{w} = [w_1, \ldots, w_N]^T$, where x_n is the actual intensity value of the nth sample, r_n represents a potential outlier (bacteria) and w_n represents the additive noise, which is assumed to be independently and identically distributed over the N fiber cores. The noise is assumed to be Gaussian distributed with covariance matrix $\sigma^2 \mathbf{I}_N$, denoted as $\mathbf{w} \sim \mathcal{N}(\mathbf{w}; \mathbf{0_N}, \sigma^2 \mathbf{I_N})$, where $\sim$ reads "is distributed according to".

The problem investigated in this paper is to estimate the outlier vector $\mathbf{r}$ in Eq. (1), but the intensity values $\mathbf{x}$ and the noise variance σ^2 are also unknown. Thus we estimate jointly $\mathbf{r}$, $\mathbf{x}$ and σ^2 from the observation vector $\mathbf{y}$. To solve this problem, we propose a hierarchical Bayesian model and a sampling method to estimate the unknown parameters.

3. Hierarchical Bayesian model

This section introduces the hierarchical Bayesian model proposed to estimate the unknown parameter $\mathbf{x}$, $\mathbf{r}$ and σ^2. This model is based on the likelihood function of the observations and on prior distributions assigned to the unknown parameters.

3.1. Likelihood

Eq. (1) implies that $\mathbf{y}|(\mathbf{x}, \mathbf{r}, \sigma^2) \sim \mathcal{N}(\mathbf{y}; \mathbf{x} + \mathbf{r}, \sigma^2 \mathbf{I}_N)$. Consequently, the likelihood function of $\mathbf{y}$ can be expressed as

$$f(\mathbf{y}|\mathbf{x}, \mathbf{r}, \sigma^2) = \left(\frac{1}{2\pi\sigma^2}\right)^{N/2} \exp\left(-\frac{\|\mathbf{y} - \mathbf{x} - \mathbf{r}\|_2^2}{2\sigma^2}\right). \tag{2}$$

3.2. Parameter prior distributions

3.2.1. Intensity field $\mathbf{x}$

In OEM images, the intensity values of the scene to be recovered are expected to be spatially correlated. A classical and convenient way to model spatially correlated intensities is to consider Markov random fields (MRF) to build a prior model for $\mathbf{x}$ (Rue and Held, 2005; Mardia, 1988; Eldaly et al., 2018a). MRFs assume that the distribution of a given intensity x_n, conditioned on the other intensity values of the image, reduces to its distribution conditioned on the values of its spatial neighbours, i.e., $f(x_n|\mathbf{x}_{\setminus x_n}) = f(x_n|\mathbf{x}_{\mathcal{V}_n})$, where $\mathcal{V}_n$ is the set of indices of the neighbours of x_n, $\mathbf{x}_{\setminus x_n}$ denotes the vector $\mathbf{x}$ whose element x_n has been removed, and $\mathbf{x}_{\mathcal{V}_n}$ is the subset of $\mathbf{x}$ composed of the elements whose indexes belong to $\mathcal{V}_n$. In this work, a Delaunay triangulation scheme (Preparata and Shamos, 2012) is used on the N samples to define the neighbourhood structure. Given the structure of the fiber bundle considered, each fibre core has between 2 and 7 neighbouring cores with mean distance between neighbours of 4.1 pixels. We specify $f(x_n|\mathbf{x}_{\mathcal{V}_n})$ as

$$f(x_n|\mathbf{x}_{\mathcal{V}_n}, \gamma^2) \propto \exp\left(-\frac{1}{\gamma^2}\sum_{n' \in \mathcal{V}_n}\frac{(x_n - x_{n'})^2}{d_{n,n'}}\right), \tag{3}$$

and where $\propto$ reads "is proportional to", $d_{n,n'}$ denotes the Euclidean distance between the spatial locations n and n', and the hyperparameter γ^2 controls the global spatial correlation between intensities. Eq. (3) promotes smooth intensity variations between neigh-
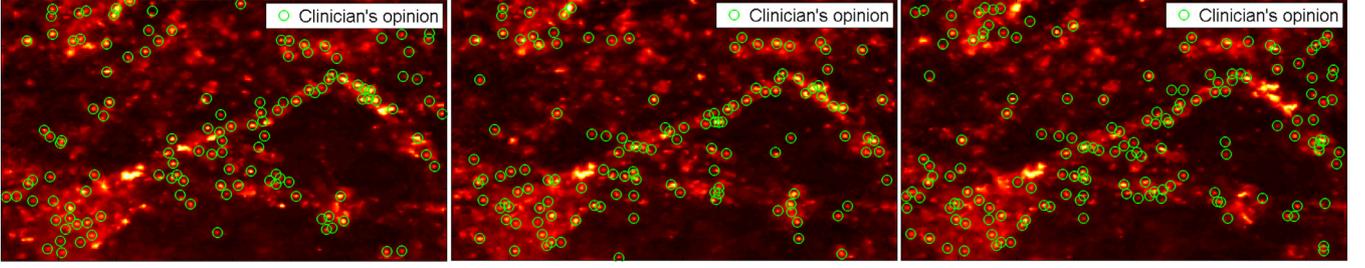
Fig. 2. A sequence of OEM frames with bacteria annotated by a trained clinician in circles.

bours while ensuring that the prior dependence between neighbours decreases as $d_{n,n'}$ increases. The resulting joint prior $f(\mathbf{x}|\gamma^2)$ can be expressed as

$$f(\mathbf{x}|\gamma^2) \propto \left(\gamma^2\right)^{-(N-1)/2} \exp\left(-\frac{\mathbf{x}^T \mathbf{\Delta} \mathbf{x}}{2\gamma^2}\right), \tag{4}$$

where

$$[\mathbf{\Delta}]_{i,j} = \begin{cases} \sum_{j \in \mathcal{V}_i} 1/d_{i,j} & \text{if } i = j, \\ 0 & \text{if } j \notin \mathcal{V}_i, \\ -1/d_{i,j} & \text{else.} \end{cases} \tag{5}$$

Note that each core has more than 2 neighbouring cores, and thus $\sum_{j \in \mathcal{V}_i} 1/d_{i,j} > 0, \forall i \in 1, \ldots, N$.

3.2.2. Noise variance σ^2

A Jeffreys' prior distribution (Jeffreys, 1946) is chosen for the noise variance σ^2, i.e.,

$$f(\sigma^2) \propto \sigma^{-2} \mathbf{1}_{\mathbb{R}^+}(\sigma^2), \tag{6}$$

where $\mathbf{1}_{\mathbb{R}^+}(\cdot)$ denotes the indicator function defined on $\mathbb{R}^+$, which reflects the lack of knowledge about this parameter. This non-informative prior can be easily replaced by conjugate inverse-Gamma prior to include knowledge available about the noise level.

3.2.3. Outliers $\mathbf{r}$

The concentration of bacteria in *in vivo* human lungs is expected to be such that only a small fraction of the fibers will be associated with bacteria detections. Hence the outliers are assumed to be sparse, i.e., for most of the spatial locations, the outliers are expected to be exactly equal to zero. To model the outlier sparsity, we factorise the outlier vector as

$$\mathbf{r} = \mathbf{z} \odot \mathbf{t}, \tag{7}$$

where $\mathbf{z} = [z_1, \ldots, z_N]^T \in \{0, 1\}^N$ is a binary label vector, $\mathbf{t} \in \mathbb{R}^N$ is the corresponding outlier amplitude vector, and $\odot$ denotes the Hadamard (term-wise) product. This decomposition allows one to decouple the location of the sparse components from their values. Precisely, $z_n = 1$ if a bacterium is present in the nth observed location with value equal to $r_n = t_n$, and $r_n = 0$ if there is no bacterium ($z_n = 0$). A conjugate Gaussian prior model is used for $\mathbf{t}$, i.e.,

$$f(\mathbf{t}|s^2) = \prod_{n=1}^N \mathcal{N}\left(t_n; \mu_t, s^2\right), \tag{8}$$

where μ_t and s^2 control the prior mean and variance of the outliers, respectively. In this work, we do not assume a particular spatial structure for bacteria positions. However, they have a priori the same probability of being present in any region of the scene. To model this prior belief, we assign each label z_n the following Bernoulli prior distribution

$$f(z_n|\omega) = \omega \delta(1 - z_n) + (1 - \omega)\delta(z_n), \quad z_n \in \{0, 1\}, \tag{9}$$

where $\delta(\cdot)$ denotes the Dirac delta function. Moreover, it is assumed that the probability of bacteria presence ω is also unknown and we include this parameter within the inference process.

3.2.4. Hyperparameters

To reflect the lack of prior knowledge about the outlier variance in Eq. (8) and regularisation parameter γ^2 in Eq. (4), the following weakly informative inverse-gamma priors are assigned to s^2 and γ^2

$$\begin{cases} s^2 \sim \mathcal{IG}(\eta, \nu), \\ \gamma^2 \sim \mathcal{IG}(\eta, \nu), \end{cases} \tag{10}$$

where (η, ν) are fixed to $(\eta, \nu) = (10^{-3}, 10^{-3})$. Note that we did not observe significant change in the results when changing these hyperparameters. Similarly, we assign μ the following conjugate truncated Gaussian prior

$$\mu_t|(\bar{\mu}, \xi^2) \sim \mathcal{N}_{\mathbb{R}^+}\left(\mu_t; \bar{\mu}, \xi^2\right), \tag{11}$$

where $(\bar{\mu}, \xi^2)$ are fixed and user-defined parameters (which depend on the dynamics of the image to be recovered). A truncated Gaussian prior on the positive set is considered as we expect the outlier mean to be positive. In this work, we fixed $(\bar{\mu}, \xi^2) = (0, 10^6)$. Finally, we assign the bacteria presence percentage ω a conjugate beta prior distribution

$$\omega \sim \mathcal{B}e(\omega; \alpha, \beta), \tag{12}$$

where (α, β) is fixed to $(\alpha, \beta) = (0.1, 1)$ as we expect the proportion of bacteria to be relatively small (the prior mean of ω is $\alpha/(\alpha + \beta)$).

3.3. Joint posterior distribution

Assuming the parameters $\mathbf{x}, \mathbf{z}, \mathbf{t}$ and σ^2 are *a priori* mutually independent, the joint posterior of the parameter vector $\mathbf{\Theta} = \{\mathbf{x}, \mathbf{z}, \mathbf{t}, \sigma^2\}$ and hyperparameters $\mathbf{\Phi} = \{\mu_t, s^2, \gamma^2, \omega\}$ can be expressed as

$$f(\mathbf{\Theta}, \mathbf{\Phi}|\mathbf{y}) \propto f(\mathbf{y}|\mathbf{\Theta})f(\mathbf{\Theta}|\mathbf{\Phi})f(\mathbf{\Phi}), \tag{13}$$

where

$$f(\mathbf{\Theta}|\mathbf{\Phi}) = f(\mathbf{x}|\gamma^2)f(\mathbf{z}|\omega)f(\mathbf{t}|\mu_t, s^2)f(\sigma^2), \tag{14}$$

$$f(\mathbf{\Phi}) = f(\gamma^2)f(\mu_t)f(s^2)f(\omega), \tag{15}$$

and $f(\mathbf{z}|\omega) = \prod_n f(z_n|\omega)$.

Fig. 3 depicts the directed acyclic graph (DAG) summarising the structure proposed Bayesian model.

The next paragraph presents a sampling strategy to estimate the unknown parameter vector $\mathbf{\Theta}$ and the hyperparameters $\mathbf{\Phi}$.

4. Bayesian inference

To overcome the challenging derivation of Bayesian estimators associated with $f(\mathbf{\Theta}, \mathbf{\Phi}|\mathbf{y})$, we propose to use an efficient Markov Chain Monte Carlo (MCMC) method to generate samples asymptotically distributed according to Eq. (13). In practice, strong correlations appear between $\mathbf{x}$ and $\mathbf{t}$, and between $\mathbf{z}$ and $\mathbf{t}$. Moreover, as $\mathbf{z}$
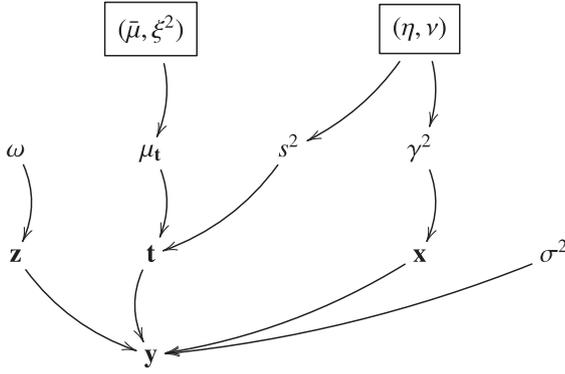
Fig. 3. Graphical model for the proposed hierarchical Bayesian model (fixed quantities appear in boxes).

is sparse, sampling $f(\mathbf{t}, \mu_t, s^2 | \mathbf{y}, \Theta_{\backslash \mathbf{t}}, \Phi_{\backslash(\mu_t, s^2)})$, where $\mathcal{H}_{\backslash \mathbf{u}}$ denotes the parameter vector $\mathcal{H}$ whose parameter $\mathbf{u}$ is omitted, using a traditional Gibbs sampler results in very slow convergence. Hence, we propose a partially collapsed Gibbs sampler (PCGS) which yields better mixing and convergence properties of the generated Markov chain. It has been shown in the literature that the PCGS is a better Bayesian inference tool for Bernoulli-Gaussian models than standard Gibbs samplers as it can explore the state space more efficiently (Kail et al., 2012; Lin et al., 2010; Kail et al., 2010; Ge et al., 2008; Robert and Casella, 2013). The PCGS used here samples groups of variables (e.g., $(\mathbf{x}, \mathbf{t})$) from their joint posterior distribution, in a similar fashion to block Gibbs samplers, which yields better mixing and convergence properties than sampling the variables (e.g., $\mathbf{x}$ and $\mathbf{t}$) sequentially from their conditional distributions. Sampling the joint distribution is achieved by first marginalising some variables which are then sampled from their full conditional distribution (Liu, 1994; Van Dyk and Park, 2008). Precisely, we propose to sample sequentially the elements of Θ and Φ using moves that are summarised in Algorithm 1.

Algorithm 1 Partially Collapsed Gibbs Sampling Algorithm For Bacteria Detection - Version I.

1: **Fixed input parameters**: Number of burn-in iterations N_{bi}, total number of iterations N_{MC}

2: **Initialization** ($k = 0$)
- Set $\mathbf{x}^{(0)}, \mathbf{z}^{(0)}, \mathbf{t}^{(0)}, \sigma^{2(0)}, \gamma^{2(0)}, \omega^{(0)}, \mu_t^{(0)}$

3: **Repeat** ($1 \leq k \leq N_{\mathbf{MC}}$)

(a) Sample $\left(s^{2(k)}, \mathbf{t}_0^{(k)}\right) | \left(\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}^{(k-1)}, \Phi_{\backslash s^2}^{(k-1)}\right)$

(b) Sample $\left(\mu_t^{(k)}, \mathbf{t}_0^{(k)}\right) | \left(\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}^{(k-1)}, \gamma^{2(k-1)}, s^{2(k)}, \omega^{(k-1)}\right)$

(c) Sample $\gamma^{2(k)} | \left(\mathbf{y}, \Theta^{(k-1)}, \mu_t^{(k)}, s^{2(k)}, \omega^{(k-1)}\right)$

(d) Sample $\omega^{(k)} | \left(\mathbf{y}, \Theta^{(k-1)}, \mu_t^{(k)}, s^{2(k)}, \gamma^{2(k)}\right)$

(e) Sample $\left(\mathbf{x}^{(k)}, \mathbf{t}^{(k)}\right) | \left(\mathbf{y}, \mathbf{z}^{(k-1)}, \sigma^{2(k-1)}, \Phi^{(k)}\right)$

(f) Sample $\left(\mathbf{z}^{(k)}, \mathbf{t}^{(k)}\right) | \left(\mathbf{y}, \mathbf{x}^{(k)}, \sigma^{2(k-1)}, \Phi^{(k)}\right)$

(g) Sample $\sigma^{2(k)} | \left(\mathbf{y}, \mathbf{x}^{(k)}, \mathbf{t}^{(k)}, \mathbf{z}^{(k)}, \Phi^{(k)}\right)$

4: **Set** $k = k + 1$.

In Algorithm 1, $\mathbf{t}_0$ denotes the elements of $\mathbf{t}$ whose corresponding labels in $\mathbf{z}$ are null. Similarly, $\mathbf{t}_1$ denotes the elements of $\mathbf{t}$ whose labels are equal to 1. We now detail each sampling step of Algorithm 1 as follow:

Sampling $(s^2, \mathbf{t}_0) | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2})$: As mentioned earlier, due to the sparsity of the outlier label vector $\mathbf{z}$, we propose to sample simultaneously $(s^2, \mathbf{t}_0) | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2})$, which improves the mixing and convergence properties of the Markov chains, rather than considering a Gibbs sampler to sample from $s^2 | (\mathbf{y}, \Theta, \Phi_{\backslash s^2})$ and $\mathbf{t}_0 | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi)$. This is done by using

$$f\left(s^2, \mathbf{t}_0 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2}\right) = f\left(s^2 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2}\right) f\left(\mathbf{t}_0 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi\right),$$

where

$$f\left(s^2 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2}\right) = \int f\left(s^2 | \mathbf{y}, \Theta, \Phi_{\backslash s^2}\right) d\mathbf{t}_0. \quad (16)$$

In other words, sampling from $(s^2, \mathbf{t}_0) | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2})$ can be achieved by sampling sequentially s^2 from $f\left(s^2 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2}\right)$, and then $\mathbf{t}_0$ from $f\left(\mathbf{t}_0 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi\right)$. In order to compute $f\left(s^2 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2}\right)$ in Eq. (16), we keep only the parameters that depend on s^2 in the joint posterior distribution in Eq. (13), this leads to

$$f\left(s^2 | \mathbf{y}, \Theta, \Phi_{\backslash s^2}\right) = \prod_{n=1}^N \mathcal{N}(t_n; \mu_t, s^2) \times \mathcal{IG}(\eta, \nu). \quad (17)$$

Due to the conjugacy of the distributions in Eq. (17), it is easy to show that sampling from $f\left(s^2 | \mathbf{y}, \Theta, \Phi_{\backslash s^2}\right)$ can be achieved by sampling from the following inverse-Gamma distribution (Gelman et al., 2013)

$$s^2 | (\mathbf{y}, \Theta, \Phi_{\backslash s^2}) \sim \mathcal{IG}\left(\eta + \frac{N}{2}, \nu + \frac{\sum_{n=1}^N (t_n - \mu_t)^2}{2}\right). \quad (18)$$

However, marginalizing $\mathbf{t}_0$ out the conditional distribution in Eq. (18) according to Eq. (16) gives

$$s^2 | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash s^2}) \sim \mathcal{IG}\left(\eta + \frac{N_1}{2}, \nu + \frac{\sum_{n \in I_1} (t_n - \mu_t)^2}{2}\right), \quad (19)$$

where $N_1 = \text{card}(\mathbf{t}_1)$ and $I_1 = \{n | z_n = 1\}$.

On the other hand, when $z_n = 0$, t_n does not appear in Eq. (1). Thus sampling $\mathbf{t}_0 | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi)$ reduces to sampling its elements independently from their Gaussian prior distribution defined in Eq. (8).

Sampling $(\mu_t, \mathbf{t}_0) | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t})$: In a similar fashion to the outlier variance, we consider a PCGS to sample simultaneously $(\mu_t, \mathbf{t}_0) | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t})$. This is done by using

$$f\left(\mu_t, \mathbf{t}_0 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t}\right) = f\left(\mu_t | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t}\right) f\left(\mathbf{t}_0 | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi\right),$$

where

$$f\left(\mu_t | \mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t}\right) = \int f\left(\mu_t | \mathbf{y}, \Theta, \Phi_{\backslash \mu_t}\right) d\mathbf{t}_0. \quad (20)$$

It is easy to show that sampling $\mu_t | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t})$ reduces to sampling from the following truncated Gaussian distribution

$$\mu_\mathbf{t} | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi_{\backslash \mu_t}) \sim \mathcal{N}_{\mathbb{R}^+}(\mu_t; \mathcal{M}, \mathcal{S}), \quad (21)$$

where

$$\mathcal{M} = \frac{\bar{\mu} s^2 + \xi^2 \sum_{n \in I_1} t_n}{\xi^2 N_1 + s^2}, \quad \mathcal{S} = \frac{s^2 \xi^2}{\xi^2 N_1 + s^2}. \quad (22)$$

Sampling from Eq. (21) can be achieved efficiently by using the method proposed in (Mazet et al., 2005) while sampling $\mathbf{t}_0 | (\mathbf{y}, \Theta_{\backslash \mathbf{t}_0}, \Phi)$ reduces to sampling from Gaussian distributions, as discussed above.

Sampling $\gamma^2 | (\mathbf{y}, \Theta, \Phi_{\backslash \gamma^2})$: By cancelling out the terms that do not depend on γ^2 from the posterior in Eq. (13), its conditional distribution reduces to the following inverse-gamma distribution

$$\gamma^2 | (\mathbf{y}, \Theta, \Phi_{\backslash \gamma^2}) \sim \mathcal{IG}\left(\eta + \frac{N-1}{2}, \nu + \frac{\mathbf{x}^T \Delta \mathbf{x}}{2}\right), \quad (23)$$

which is easy to sample from.

Sampling $\omega | (\mathbf{y}, \Theta, \Phi_{\backslash \omega})$: Due to the conjugacy of the hierarchical prior model $f(\mathbf{z} | \omega) f(\omega)$, the full conditional distribution of ω reduces to the following beta distribution

$$\omega | (\mathbf{y}, \Theta, \Phi_{\backslash \omega}) \sim \mathcal{B}e\left(\alpha + \sum_{n=1}^N z_n, \beta + N - \sum_{n=1}^N z_n\right). \quad (24)$$

Sampling $(\mathbf{x}, \mathbf{t})|(\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})}, \mathbf{\Phi})$: As mentioned earlier, we propose to sample simultaneously $(\mathbf{x}, \mathbf{t})|(\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})}, \mathbf{\Phi})$, rather than considering a Gibbs sampler to sample from $\mathbf{x}|(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{x}}, \mathbf{\Phi})$ and $\mathbf{t}|(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi})$. This is done by using

$$f\left(\mathbf{x}, \mathbf{t}|\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})}, \mathbf{\Phi}\right) = f\left(\mathbf{x}|\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})}, \mathbf{\Phi}\right) f\left(\mathbf{t}|\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi}\right),$$

where

$$f\left(\mathbf{x}|\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})}, \mathbf{\Phi}\right) = \int f\left(\mathbf{x}|\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{x}}, \mathbf{\Phi}\right) d\mathbf{t}. \quad (25)$$

It is easy to show that $f(\mathbf{x}|\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})})$ is the following multivariate Gaussian distribution

$$\mathbf{x}|(\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{x},\mathbf{t})}, \mathbf{\Phi}) \sim \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \mathbf{\Sigma}), \quad (26)$$

where

$$\begin{cases} \boldsymbol{\mu} = \mathbf{\Sigma}(s^2\mathbf{Z} + \sigma^2\mathbf{I})^{-1}(\mathbf{y} - \mu_\mathbf{t}\mathbf{z}), \\ \mathbf{\Sigma} = \left((s^2\mathbf{Z} + \sigma^2\mathbf{I})^{-1} + \gamma^{-2}\mathbf{\Delta}\right)^{-1}, \end{cases} \quad (27)$$

and $\mathbf{Z} = \text{diag}(\mathbf{z})$. The full conditional distribution of $\mathbf{t}$, i.e., $f(\mathbf{t}|(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi}))$ reduces to the following multivariate Gaussian distribution

$$\mathbf{t}|(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi}) \sim \mathcal{N}(\mathbf{t}; \boldsymbol{\mu}, \mathbf{\Sigma}), \quad (28)$$

where

$$\begin{cases} \boldsymbol{\mu} = \left(\frac{(\mathbf{y}-\mathbf{x})^T \odot \mathbf{z}}{\sigma^2} + \frac{\mu_t}{s^2}\right)\mathbf{\Sigma}, \\ \mathbf{\Sigma} = \left(s^{-2}\mathbf{I}_N + \sigma^{-2}\mathbf{Z}\right)^{-1}. \end{cases} \quad (29)$$

Note that $\mathbf{\Sigma}$ is a diagonal covariance matrix, which is easy to construct.

Sampling $(\mathbf{z}, \mathbf{t})|(\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{z},\mathbf{t})}, \mathbf{\Phi})$: Updating simultaneously $(\mathbf{z}, \mathbf{t})$ is achieved using

$$f\left(\mathbf{z}, \mathbf{t}|\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{z},\mathbf{t})}, \mathbf{\Phi}\right) = f\left(\mathbf{z}|\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{z},\mathbf{t})}, \mathbf{\Phi}\right) f\left(\mathbf{t}|\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi}\right).$$

It can be seen from Eq. (13) that

$$f(z_n = m|\mathbf{y}, \mathbf{\Theta}_{\backslash(z_{n'},t_{n'})}, \mathbf{\Phi}) \propto \bar{\omega}_n^{(m)} \quad \forall(n), \quad (30)$$

where $m \in \{0, 1\}$ and

$$\log(\bar{\omega}_n^{(m)}) = -\frac{1}{2}\log\left(2\pi(\sigma^2 + ms^2)\right)$$
$$- \frac{(y_n - x_n - m\mu_\mathbf{t})^2}{2(\sigma^2 + ms^2)} + \log(p(m|\omega)). \quad (31)$$

Consequently, the label z_n can be drawn from its conditional distribution (where t_n has been marginalised) by drawing randomly from $\{0, 1\}$ with probabilities given by

$$f(z_n = m|\mathbf{y}, \mathbf{\Theta}_{\backslash(z_{n'},t_{n'})}, \mathbf{\Phi}) = \frac{\bar{\omega}_n^{(m)}}{\bar{\omega}_n^{(0)} + \bar{\omega}_n^{(1)}}. \quad (32)$$

Moreover, the elements of $\mathbf{z}$ can be updated in a parallel manner using the fact that $f(\mathbf{z}|\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi}) = \prod_n f(z_n|\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi})$. Sampling from $f(\mathbf{t}|\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi})$ is then achieved using Eq. (28).

Sampling $\sigma^2|(\mathbf{y}, \mathbf{\Theta}_{\backslash\sigma^2}, \mathbf{\Phi})$: In a similar fashion to the regularisation parameter γ^2, the noise variance σ^2 can be sampled from the following inverse-gamma distribution

$$\sigma^2|(\mathbf{y}, \mathbf{\Theta}_{\backslash\sigma^2}, \mathbf{\Phi}) \sim \mathcal{IG}\left(\frac{N}{2}, \frac{\|\mathbf{y}-\mathbf{x}-\mathbf{r}\|_2^2}{2}\right). \quad (33)$$

Algorithm 2 is a detailed version of Algorithm 1 following the sampling steps explained above. Although Algorithm 2 seems more complex (more sequential steps) than Algorithm 1, several steps can be omitted since some generated variables are not actually used. For instance, the variables generated in steps (a2) and (b2) are not used during the following steps and thus can be omitted. Similarly, sampling $\mathbf{t}|(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}}, \mathbf{\Phi})$ in (e2) is omitted as this step is not required when sampling $\mathbf{z}|(\mathbf{y}, \mathbf{\Theta}_{\backslash(\mathbf{z},\mathbf{t})}, \mathbf{\Phi})$ in (f1). These simplifications result in the final algorithm shown in Algorithm 3.

**Algorithm 2** Partially Collapsed Gibbs Sampling Algorithm to Bacteria Detection for OEM - Version II.

1: **Fixed input parameters**: Number of burn-in iterations N_{bi}, total number of iterations N_{MC}
2: **Initialization** $(k = 0)$
  - Set $\mathbf{x}^{(0)}, \mathbf{z}^{(0)}, \mathbf{t}^{(0)}, \sigma^{2(0)}, \gamma^{2(0)}, \omega^{(0)}, \mu_t^{(0)}$
3: **Repeat** $(1 \le k \le N_{\mathbf{MC}})$

  (a1) Sample $s^{2(k)}|\left(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}_0}^{(k-1)}, \mathbf{\Phi}_{\backslash s^2}^{(k-1)}\right)$ from Eq. (19).

  (a2) Sample $\mathbf{t}_0^{(k)}|\left(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}_0}^{(k-1)}, s^{2(k)}, \mu_t^{(k-1)}, \gamma^{2(k-1)}, \omega^{(k-1)}\right)$ from Eq. (8).

  (b1) Sample $\mu_t^{(k)}|\left(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}_0}^{(k-1)}, \gamma^{2(k-1)}, s^{2(k)}, \omega^{(k-1)}\right)$ from Eq. (21).

  (b2) Sample $\mathbf{t}_0^{(k)}|\left(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}_0}^{(k-1)}, \mu_t^{(k)}, \gamma^{2(k-1)}, s^{2(k)}, \omega^{(k-1)}\right)$ from Eq. (8).

  (c) Sample $\gamma^{2(k)}|\left(\mathbf{y}, \mathbf{\Theta}^{(k-1)}, \mu_t^{(k)}, s^{2(k)}, \omega^{(k-1)}\right)$ from Eq. (23).

  (d) Sample $\omega^{(k)}|\left(\mathbf{y}, \mathbf{\Theta}^{(k-1)}, \mu_t^{(k)}, s^{2(k)}, \gamma^{2(k)}\right)$ from Eq. (24).

  (e1) Sample $\mathbf{x}^{(k)}|\left(\mathbf{y}, \mathbf{z}^{(k-1)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (26).

  (e2) Sample $\mathbf{t}^{(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \mathbf{z}^{(k-1)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (28).

  (f1) Sample $\mathbf{z}^{(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (32).

  (f2) Sample $\mathbf{t}^{(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \mathbf{z}^{(k)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (28).

  (g) Sample $\sigma^{2(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \mathbf{t}^{(k)}, \mathbf{z}^{(k)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (33).

4: **Set** $k = k + 1$.

Algorithm 3 Partially Collapsed Gibbs Sampling Algorithm to Bacteria Detection for OEM - Final Version.

1: **Fixed input parameters**: Number of burn-in iterations N_{bi}, total number of iterations N_{MC}
2: **Initialization** $(k = 0)$
  - Set $\mathbf{x}^{(0)}, \mathbf{z}^{(0)}, \mathbf{t}^{(0)}, \sigma^{2(0)}, \gamma^{2(0)}, \omega^{(0)}, \mu_t^{(0)}$
3: **Repeat** $(1 \le k \le N_{\mathbf{MC}})$

  (a1) Sample $s^{2(k)}|\left(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}_0}^{(k-1)}, \mathbf{\Phi}_{\backslash s^2}^{(k-1)}\right)$ from Eq. (19).

  (b1) Sample $\mu_t^{(k)}|\left(\mathbf{y}, \mathbf{\Theta}_{\backslash\mathbf{t}_0}^{(k-1)}, \gamma^{2(k-1)}, s^{2(k)}, \omega^{(k-1)}\right)$ from Eq. (21).

  (c) Sample $\gamma^{2(k)}|\left(\mathbf{y}, \mathbf{\Theta}^{(k-1)}, \mu_t^{(k)}, s^{2(k)}, \omega^{(k-1)}\right)$ from Eq. (23).

  (d) Sample $\omega^{(k)}|\left(\mathbf{y}, \mathbf{\Theta}^{(k-1)}, \mu_t^{(k)}, s^{2(k)}, \gamma^{2(k)}\right)$ from Eq. (24).

  (e1) Sample $\mathbf{x}^{(k)}|\left(\mathbf{y}, \mathbf{z}^{(k-1)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (26).

  (f1) Sample $\mathbf{z}^{(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (32).

  (f2) Sample $\mathbf{t}^{(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \mathbf{z}^{(k)}, \sigma^{2(k-1)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (28).

  (g) Sample $\sigma^{2(k)}|\left(\mathbf{y}, \mathbf{x}^{(k)}, \mathbf{t}^{(k)}, \mathbf{z}^{(k)}, \mathbf{\Phi}^{(k)}\right)$ from Eq. (33).

4: **Set** $k = k + 1$.
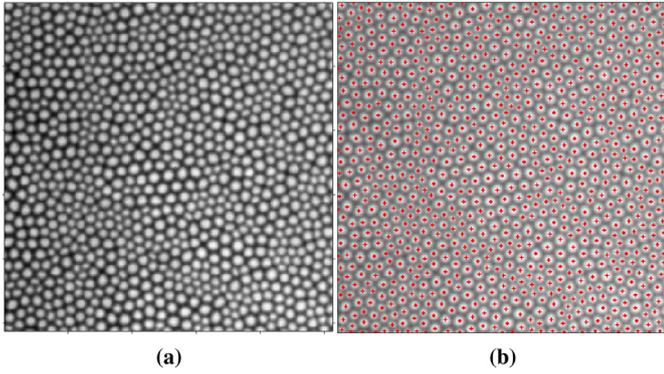
Fig. 4. (a) Example of background image of the optical endomicroscopy system (b) Image with detected fiber core centres superimposed (red crosses). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The algorithm is stopped after N_{MC} iterations, including N_{bi} burn-in iterations which correspond to the transient period of the sampler (determined visually from preliminary runs). The first N_{bi} samples are discarded and the remaining samples are used to approximate the following estimators. The label vector $\mathbf{z}$ is estimated using marginal maximum a posteriori (MAP) estimation. This estimator is then used to compute the minimum mean square error (MMSE) of $\mathbf{r}$ conditioned on $\mathbf{z} = \hat{\mathbf{z}}_{\mathrm{MAP}}$, i.e., $\hat{\mathbf{r}} = \left(\hat{\mathbf{r}}_{\mathrm{MMSE}} | \hat{\mathbf{z}}_{\mathrm{MAP}} \right) \odot \hat{\mathbf{z}}_{\mathrm{MAP}}$. Finally, the remaining parameters are estimated using the empirical averages of the generated samples (MMSE estimates), for instance, the MMSE of the actual intensity vector $\mathbf{x}$, denoted as $\hat{\mathbf{x}}$ is given by

$$\hat{\mathbf{x}} = \frac{1}{N_{\mathrm{MC}} - N_{\mathrm{bi}}} \sum_{t=N_{\mathrm{bi}+1}}^{N_{\mathrm{MC}}} \mathbf{x}^{(t)}. \tag{34}$$

5. Simulations using synthetic dataset

5.1. Data creation

In order to assess the performance of the proposed approach for outlier detection, two standard test images are used, namely the 'Elastin' (570×570 pixels) and the 'Coins' (492×600 pixels). Subsampled versions of these images are obtained by considering the sampling pattern of an actual OEM system (Krstajić et al., 2016) as shown in Fig. 4, which yield 4436 randomly subsampled pixels from the Elastin image (corresponding to 1.37%) and 4029 randomly subsampled pixels of the Coins image(corresponding to 1.36%). Fig. 5(a) and (e) show the original test images, (b) and (f) show natural neighbour interpolation (Sibson, 1981) of the randomly subsampled pixels of the images in (a) and (e) respectively, and (c) and (g) show examples of system output after applying the model in Eq. (1). In this case, the noise variance is $\sigma^2 = 10$, and 5% of outliers with mean ($\mu_{\mathbf{t}} = 128$), and variance ($s_{\mathbf{t}}^2 = 10$) are used. Note that the outliers and noise are added to the intensity field vector $\mathbf{x}$ before interpolation.

5.2. Evaluation criterion

To assess the outlier detection performance, considering the fibre cores that are neither known as outliers nor detected by the proposed algorithm as true negatives (TN) would result in a heavily unbalanced two-class problem. Hence, precision (also equivalent to positive predictive value) and recall (also equivalent to sensitivity or true positive rate) measures are considered, which are

computed as follows

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \qquad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}},$$

where TP, FN, and FP refer to the number of true positives, false negatives, and false positives respectively. Given the locations of cores where an outlier has been detected and the reference outlier locations, we consider

- any detection that matches reference outlier location as TP.
- any detection that does not match any of the reference outlier locations as FP.
- any reference outlier location that does not match with any of the algorithm detections as FN.

Since each bacterium corresponds to a set of connected cores, each group of connected detections is counted as a single detection. These connected detections arise due to the spread of some bacteria fluorescence light over neighbouring cores, which decreases as distance from the central illuminated core increases. Once these connected detections are identified, they are replaced by a single detection at the mean of their locations, which gives the total number of detected outliers in the test image.

5.3. Performance analysis

In order to assess the performance of the proposed method, weak and strong outliers and noise are tested and the precision and recall measures are computed accordingly. We evaluated the algorithm by selecting $\sigma^2 \in \{0, 1, 10\}$, $\mu_{\mathbf{t}} \in \{0, 128, 255\}$, $s_{\mathbf{t}}^2 \in \{10, 100, 10000\}$ and outlier percentage of 1%, 5% and 10%. We show the results for $\sigma^2 = 10$, and outlier percentage of 5%. Results for outlier percentages of 1% and 10% are not shown since the trends in behaviour are the same as with 5%. The precision and recall measures of the tested scenarios are represented by the empirical precision-recall graph shown in Fig. 6 for the standard Elastin image. We can observe a satisfactory detection accuracy for all of the tests except for the cases where outlier mean and/or variance is near to noise variance, in which the outliers are of low amplitude, and hence could not be detected. This should not be a problem as in most OEM images, including the one investigated in this work, outlier mean and variance are positive and much larger than that of the noise. Similar trends are observed for the standard Coins image and hence they are not presented here. Fig. 5(d) and (h) show examples of detections when $\sigma^2 = 10$, outlier presence probability of 5% and outlier mean and variance of $\mu_{\mathbf{t}} = 128$ and $s_{\mathbf{t}}^2 = 10$ respectively, for the Elastin and Coins image respectively.

5.4. Comparison with existing approaches

In this subsection, we compare the proposed approach with four existing spot-detection methods, namely the sparse coding approach proposed in (Eldaly et al., 2019), the Laplacian of Gaussian (LoG), its approximation; the difference of Gaussians (DoG) filters (Lindeberg, 1998; He et al., 2010), and the grey scale opening top-hat filter (GSOTH) (Kimori et al., 2010; Bright and Steel, 1987) which are all unsupervised detection methods. The sparse coding approach was used for bacterial detection to the same datasets processed in this work (Eldaly et al., 2019). This approach splits each image into a set of overlapping patches (each of size 27×27 with 50% overlap in this case) and assumes that observed intensities are linear combinations of the actual intensity values associated with background image structures, corrupted by additive Gaussian noise and potentially by a sparse outlier term modelling bacteria. The actual intensity term representing background structures is modelled as a linear combination of a few atoms drawn from a dictionary which is learned from bacteria-free data and
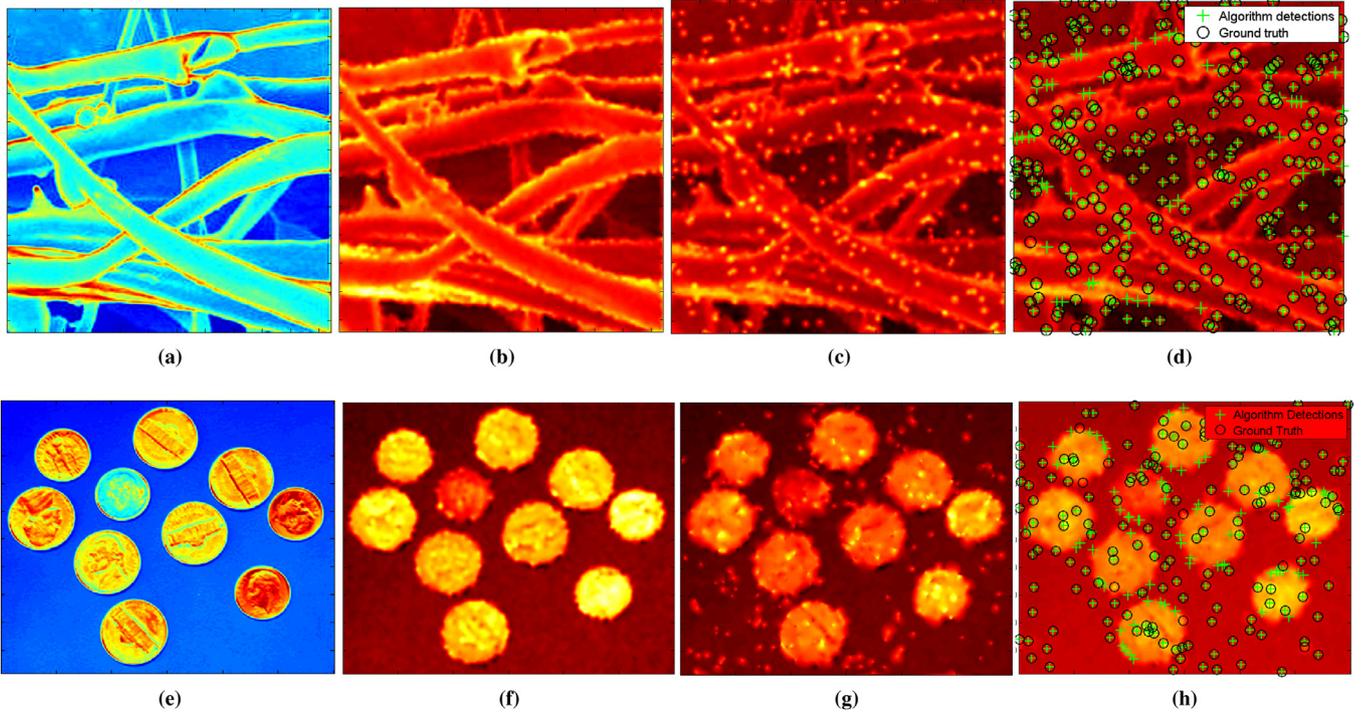
Fig. 5. Creation of the synthetic data: (a) and (e) Original Elastin and Coins images respectively. (b) and (f) Natural neighbour interpolation from roughly 1.36% samples of the images in (a) and (e). (c) and (g) Corruption by noise ($\sigma^2 = 10$) and 5% of outliers (μ, s_t^2) = (128, 10). (d) and (h) Results of outlier detection using MCMC, black circles represent true outlier locations whereas green pluses represent detections using the proposed approach. Note that the outliers and noise are added to the intensity field vector **x** before interpolation. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)
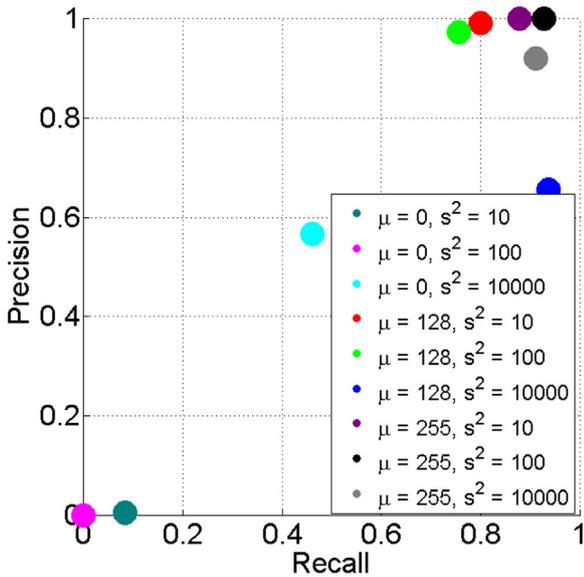
Fig. 6. Precision-recall plot for tested outlier means and variances for $\sigma^2 = 10$, and outlier proportion of 5%.

then fixed while analyzing new images. The bacteria detection task is formulated as a minimization problem and an alternating direction method of multipliers (ADMM) is used to estimate the unknown parameters. Different number of dictionary atoms including 50, 100 and 150, and different outlier regularization parameter values ranging between 0.01 and 15 are tested. Thus, we provide here the best results obtained. On the other hand, the LoG filter is implemented by employing a 5×5 kernel of standard deviation of

0.8 to each frame. Similarly, the DoG filter is implemented by considering the difference of two 5×5 Gaussian kernels of standard deviations of 0.5 and 0.8 respectively. The GSOTH is employed by first smoothing the input image by a Gaussian kernel to reduce the noise, then by computing the morphological opening of the input image by employing a 3×3 flat disc, which achieves the best detection results and then subtracts the result from the original image. Each of these four methods provides an outlier amplitude image that needs to be thresholded in order to determine bacterial locations. Different cut-off thresholds are tested and the same post processing steps described earlier (pixel grouping and computation of the barycenters) are also employed. The precision-recall curves are then constructed accordingly and the area under curves (AUC) are computed. The comparison with the proposed approach in this work is conducted in terms of precision-recall measures, as well as in terms of the computation time.

Bacteria detection performance

Due to the absence of AUC measures for the proposed approach, since it can provide instant detections that do not require to be thresholded to get outlier locations, we first provide a comparison between the four existing methods described above in terms of AUC measures. We then compare the proposed approach with the existing method that provides the highest AUC. Table 1 presents AUC measures of the four existing methods using different outlier means and variances. It is clear that the sparse-coding approach outperforms the rest of the three existing methods and that the LoG method provides the second best. Fig. 7 illustrates an empirical plot of the precision-recall measures of the proposed approach and the precision-recall curves of the sparse coding approach which provides the highest AUC measures. It is clear that the Bayesian approach outperforms the rest of the methods. Moreover, it is fully automatic as the user is not required to tune any

Table 1

Area under curve measures of the resulting precision-recall curves of the four existing methods. Bold (resp. underlined) represent best (resp. second best) results.

μ_t	s^2	Sparse coding	LoG	DoG	GSOTH
			AUC		
0	10	0.0011	0.0011	0.0018	0.0012
	100	0.0030	0.0027	0.0046	0.0025
	10,000	0.15	0.15	0.11	0.09
128	10	0.74	0.59	0.54	0.41
	100	0.71	0.53	0.47	0.40
	10,000	0.50	0.44	0.33	0.29
255	10	0.78	0.67	0.43	0.49
	100	0.78	0.64	0.49	0.43
	10,000	0.75	0.56	0.51	0.44
Average		**0.5**	<u>0.4</u>	0.324	0.288
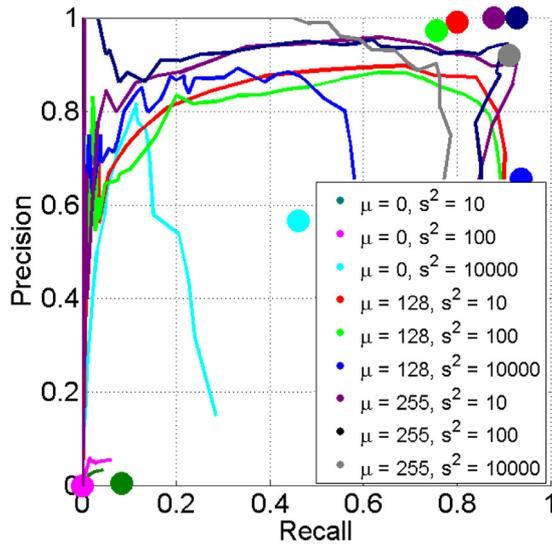
Fig. 7. Precision-recall curves of the sparse-coding approach (Eldaly et al., 2019) which provides best AUC measures among existing methods, and the empirical plot of the precision-recall measures of the proposed approach (shown as solid circles), for tested outlier means and variances for $\sigma^2 = 10$, and outlier proportion of 5% for the Elastin image.

crucial hyperparameters that would greatly affect the detection performance.

Computation time

Table 2 provides the average computation time of the five methods. For the sparse coding approach, the resulting number of test patches is 1521 yielding $\mathbf{Y} \in \mathbb{R}^{729 \times 1521}$, and the dictionary tested is $\mathbf{D} \in \mathbb{R}^{729 \times 150}$. The experiments were conducted on ACER core-i3-2.0 GHz processor laptop with 8 GB RAM. As the proposed approach can provide results from just a single run, the rest of the methods require thresholding the outlier amplitude image, which is difficult to convert into a precise computation time. However, the computation time of the sparse coding approach corresponds to the duration of five runs of different outlier regularization parameter values, which are used to select the best regularization parameter among the five values in terms of detection performance. Although the Bayesian approach proposed provides the highest computation time, it crucially brings the benefit of providing higher detection performance with respect to the other four methods, and being fully automatic, as there is no need to either set any regularization parameters nor threshold the resulting outlier amplitude image in order to identify outlier locations.

6. Simulations using real OEM images

6.1. Datasets

The proposed algorithm was assessed using two datasets of *ex vivo* ventilated whole ovine lungs instilled with bacteria. Dataset I contains seven videos assessing a combination of fluorescent dyes and bacterial types, including control segments. It contains (i) three videos of ovine lungs instilled with *Methicillin-sensitive Staphylococcus aureus* (*MSSA*) stained with a commercially available laboratory dye (PKH67, Sigma-Aldrich), a highly fluorescent cell membrane dye, (ii) two videos of ovine lungs instilled with bacteria (gram-positive *MSSA* and gram-negative *Pseudomonas* PA3284) stained in situ with an in-house bacterial detection SmartProbe (Akram et al., 2016), and (iii) two videos of ovine lungs without the presence of any bacteria. Videos 1–5 were instilled with a single concentration of bacteria, equivalent to Optical Density (OD595nm) of 2.

Dataset II contains four videos, each with an increasing bacterial concentration (OD595 nm 0.004, 0.04, 0.4, 4) all labelled with an in-house bacterial detection SmartProbe. Tables 3 and 4 summarise the details of Datasets I and II respectively.

The Cellvizio OEM imaging platform along with a 1.4mm, 12,105 core fibre bundle (Mauna Kea Technologies, Paris, France) was used to acquire all data in this study. Images sequences of size 274×384 pixels (306 µm $\times$ 429 µm) were captured at 12 frames per second. Furthermore, the fibre core locations were known from factory pre-calibration to sub-pixel accuracy in this image grid. Finally, prior to any frame selections, manual annotations and automatic processing, the uninformative frames were automatically removed (Perperidis et al., 2017a) from each video, ensuring only clinically relevant frames (~ 500 in each video) were considered for this work.

Random frames that are representative of each of the entire video sequences are chosen from each of the eight videos of Dataset I and the four videos of Dataset II by a trained clinician. Those comprises 133 image frames for Dataset I, and 58 frames for Dataset II as described in Tables 3 and 4. In each frame, a trained clinician marked the co-ordinates of phenomena that are thought to be bacteria with high confidence. Ambiguous points are ignored. Due to the size of the imaged bacteria (< 3 µm), a bright fluorescent 'dot' is identified as bacterium if one or two cores present higher intensities than the surrounding area. The location of the annotation is rounded to the closest core location. The fibre core locations are provided by the imaging system with sub-pixel accuracy, so the fibre core centre coordinates were also rounded to the nearest pixel.

The increasing bacteria concentration dataset is considered to compare the annotations performed by the clinician and the results of the algorithm.

6.2. Algorithm evaluation

The algorithm was run for each of the annotated frames, yielding sets of detected bacteria, i.e potential labelled bacteria. The same post processing steps described earlier (connected detections grouping and barycentre computation) are applied, which gives the total number of detected bacteria in each frame. Due to the subjectivity of the clinician's annotations, it cannot be considered as ground truth, therefore, in order to evaluate the performance of the proposed approach, we consider only count-annotation criterion by performing statistical comparison of bacterial counts performed by the trained clinician and the algorithm output.

Count-annotation effect

For Dataset I, the algorithm counts are compared with the clinician counts in each frame as shown in Fig. 8. We can observe an

Table 2
Average computation time (in seconds) for the proposed method and four existing spot-detection methods. In order to maintain a fair comparison between the five algorithms, the computation time of the sparse coding approach corresponds to the duration of five runs (used to select the best regularization parameter among the five values). Moreover, all of these methods apart from the Bayesian model requires manual thresholding of the outlier amplitude image in order to identify bacteria locations, but this can not be easily converted into precise computation time. Bold and underlined values are fastest and second fastest computation times respectively.

Hierarchical Bayesian model	Sparse-coding approach	LoG	DoG	GSOTH
10.9	5 x 1.98 = 9.9	0.6	**0.13**	<u>0.31</u>

Table 3
Description of dataset I.

Video	# of frames	Bacteria concent. (OD)	Fluoro-phore	Bacteria
1	26	2	PKH	*Staphylococcus aureus*
2	19			
3	13			
4	32		Smart- Probe	*Pseudomonas aeruginosa*
5	19			*Staphylococcus aureus*
6	12	NA	NA	Control
7	12			

Table 4
Description of dataset II.

Video	# of frames	Bacteria concen. (OD)	Fluoro- phore	Bacteria
1	14	0.004	Smart- Probe	*Pseudomonas aureus*
2	14	0.04		
3	15	0.4		
4	15	4		
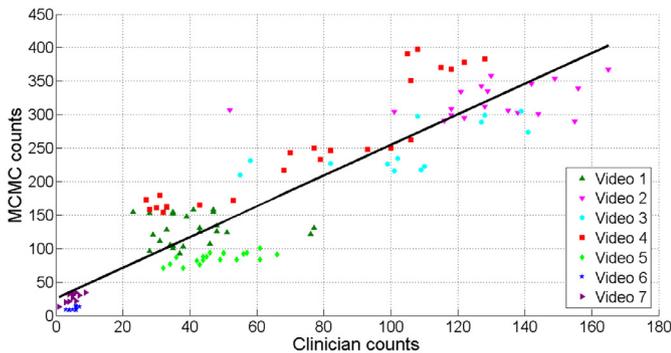
Fig. 8. Plot of clinician bacteria count versus algorithm bacteria count for dataset I. Dots correspond to frames, and colours correspond to videos.
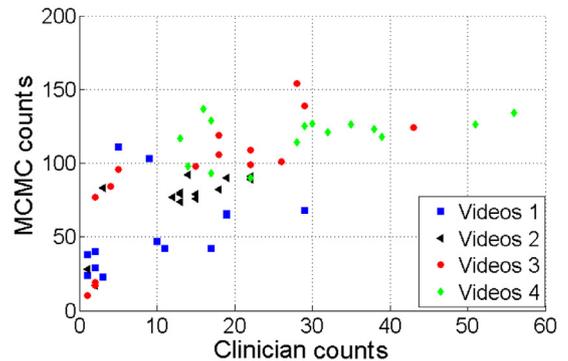
Fig. 9. Plot of clinician bacteria count versus algorithm bacteria count for dataset II. Dots correspond to frames, and colours correspond to videos.

almost linear relationship between the clinician counts and algorithm counts. Indeed the empirical correlation between the manually and automatically detected anomalies is 0.912.

Moreover, for videos 1, 2 and 3 in which a highly fluorescent SmartProbe is used, and videos 4 and 5 in which an in-house SmartProbe which produces weaker fluorescence signal is used, a similar trend is observed between the clinician and the algorithm. This also applies to the type of bacteria the samples are labelled with. Videos 6 and 7 which are controls, demonstrate minimal counts using both the clinician and the algorithm, which reflects the ability of the algorithm to differentiate bacterial loads from control. For Dataset II, Fig. 9 shows a plot of the clinician's bacteria counts versus those of the algorithm, we observe that each concentration occupies a broad count range in terms of both counts from the opinion of the clinician and algorithm output. Moreover, Fig. 10 depicts plots of the total bacteria counts performed by the clinician and that by the algorithm versus different bacterial concentrations. We can observe that both of the clinician and the algorithm counts

increase with the bacteria concentration which reflects the agreement between the approach considered and the clinician's annotations.

In the two datasets, we can observe that the algorithm counts are higher than that of the clinician, as we expect the algorithm to be able to identify dots that are hard to be seen by the naked eye. Moreover, the clinician did not annotate ambiguous dots in a number of cases. This, along with false positives is most likely the main reasons why the algorithm counts are higher than the clinician counts.

Fig. 11 shows examples of detections in two frames from Dataset I using the proposed approach. Fig. 11 (a) and (c) show the original frames and (b) and (d) show the algorithm detections and the clinician's annotations superimposed. We observe the accuracy of the algorithm as almost all of the cores annotated by the clinician are covered by the algorithm's detections.
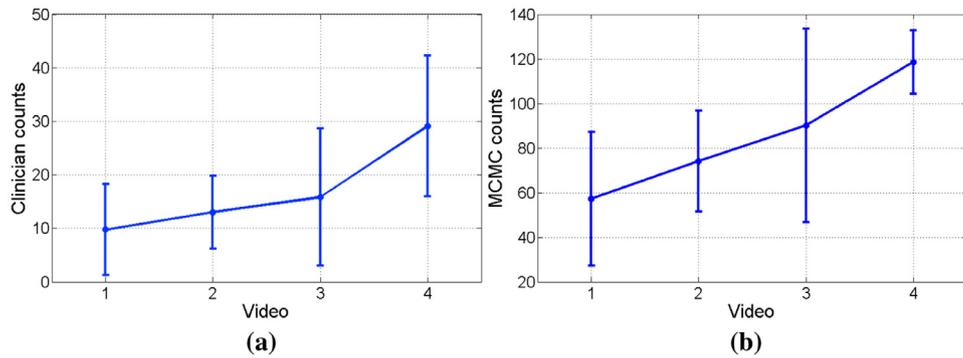
Fig. 10. Mean number of detections per selected frames in videos 1 to 4 of Dataset II and the corresponding standard deviation. (a) clinician's opinion, (b) proposed method.
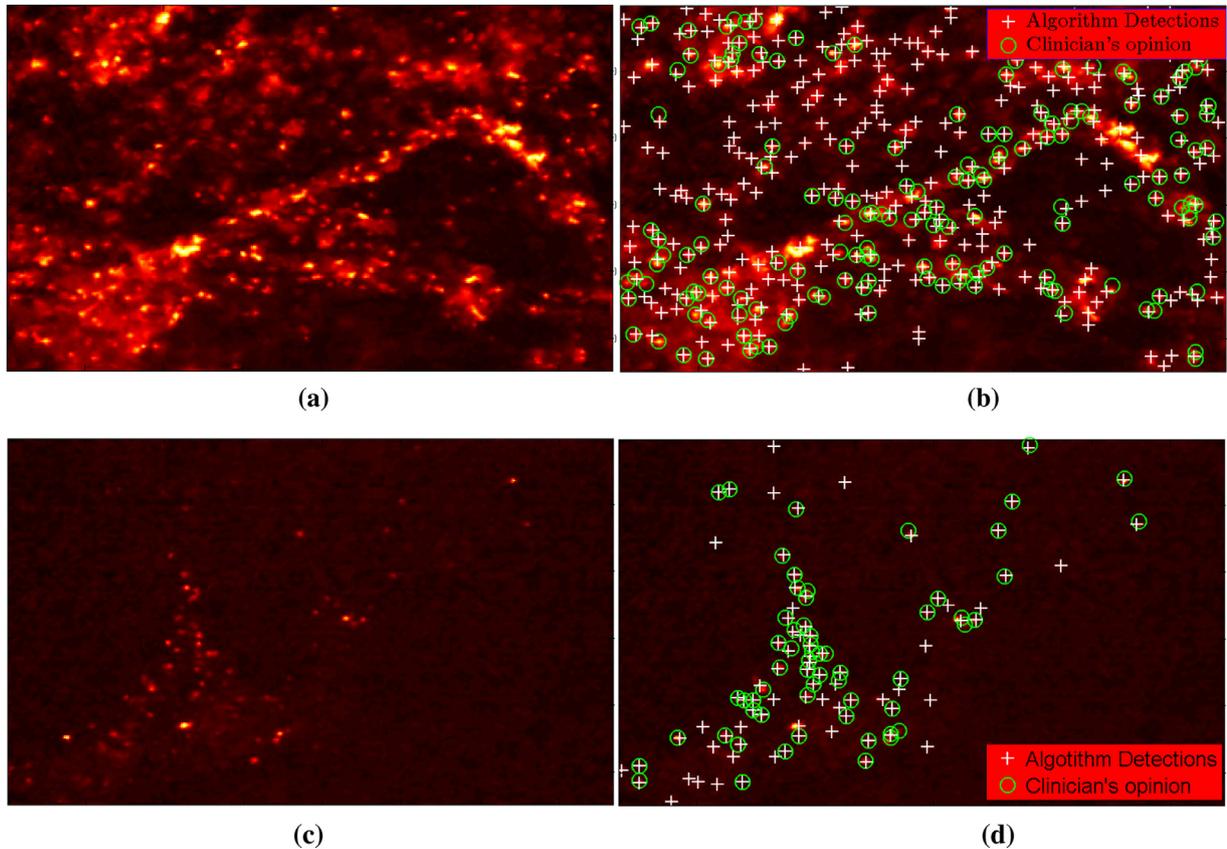

Fig. 11. Examples of detection in two frames of Dataset I (top: video 2, bottom: video 5). (a) and (c) original frames, (b) and (d) final detections.

6.3. Comparison with existing approaches

Since the sparse coding approach described earlier provides the best results in comparison with three existing methods (Eldaly et al., 2019), in this section, we compare its performance with the proposed approach in terms of count-annotation effect. Fig. 12 depicts the proposed algorithm's counts versus the clinician's counts in (a) and that of the sparse coding approach at cut-off thresholds of 0.07 and 0.05 in (b) and (c) respectively, all for Dataset I. The empirical correlation coefficients of the three cases are 0.91, 0.82 and 0.65 respectively. It can be seen that the proposed approach provides the highest correlation between algorithm's counts and clinician's counts. Moreover, the user would need to carefully tune the cut-off threshold of the sparse coding approach to obtain reasonable correlation. It can also be observed

that the proposed approach provides minimal counts for the control cases (videos 6 and 7). Although increasing the cut-off threshold will result in fewer counts, these counts are not necessarily true positives (this can be observed in Fig. 7, which indicates that having high precision might result in having small recalls).

In a similar fashion, Fig. 13 shows the mean number of detections per selected frames in videos 1–4 of Dataset II and the corresponding standard deviation for counts derived from the clinician's opinion in (a), proposed method in (b) and sparse coding approach with different cut-off thresholds in (c). The two methods are in agreement with the increasing concentration behaviour. Although the sparse coding approach provides smaller counts compared to the proposed approach, this comes to the cost of successful tuning of cut-off threshold to provide reasonable precision-recall measures.
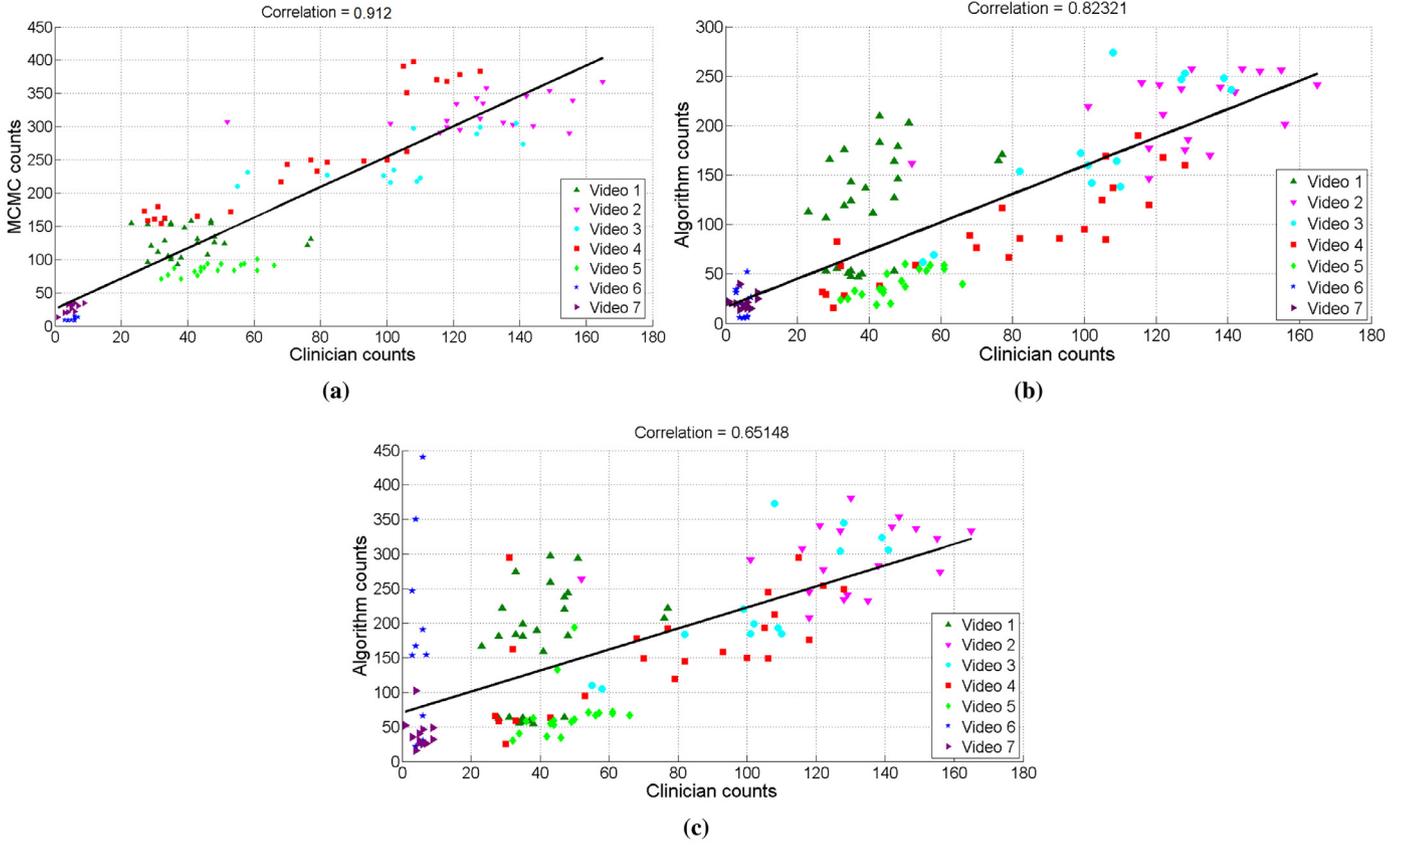
Fig. 12. Plot of clinician bacteria count versus algorithm bacteria count for dataset I. (a) proposed approach, (b) the sparse-coding approach presented in (Eldaly et al., 2019) with cut-off threshold of $\ell_d = 0.07$, (c) $\ell_d = 0.05$. Dots correspond to frames, and colours correspond to videos.
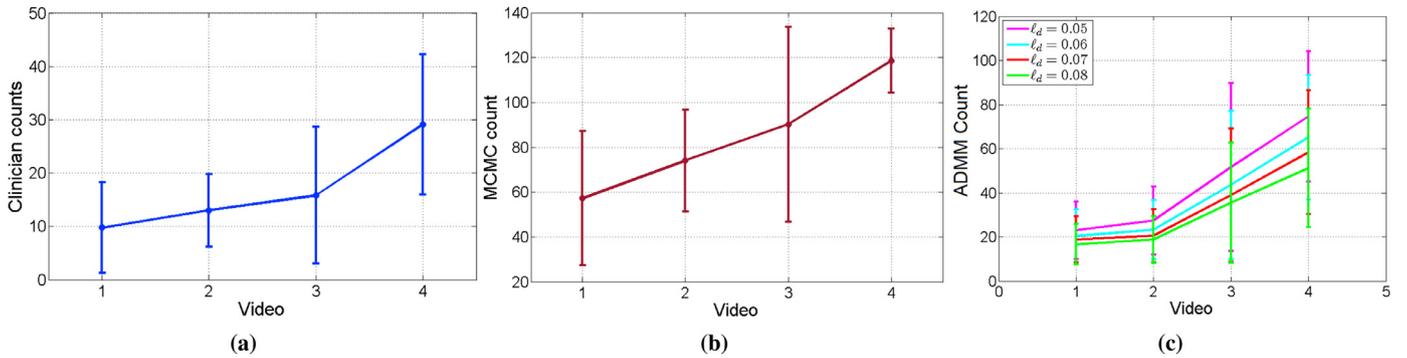

Fig. 13. Mean number of detections per selected frames in videos 1 to 4 of Dataset II and the corresponding standard deviation. (a) clinician's opinion, (b) proposed method, (c) sparse-coding approach (Eldaly et al., 2019) with different cut-off thresholds.

7. Limitations

Although the benefits of the proposed approach in detecting outliers are clear, a few limitations remain. The first is related to the noise model considered in this work. Due to the lack of background images from the system that the data was collected with, an exact noise model could not be inferred, hence a Gaussian noise model was assumed due to its simplicity. Therefore, validation of this noise model is still required. The second is related to the lack of ground truth in the real datasets investigated in this work due to the difficulty in annotating the bacteria. Although these datasets are annotated, the annotations cannot be considered as ground truth. While annotating, it is likely that the annotator, i.e. the clinician, will makes errors: they will either falsely annotate a bacterium when it is noise, or simply miss-annotate a bacterium due

to their overwhelming numbers in each frame, since our target objects are 'dots' with similar structure. These types of error are common in any annotation process. We note that in clinical context no ground truth will exist and it is a topic of further work to explore how the annotations from an automatic method would be used to assist a clinician.

8. Conclusions

In this paper, we illustrated the performance of a Bayesian algorithm for bacterial detection in OEM images of distal lung tissue. Prior distributions were chosen for the unknown model parameters and their corresponding hyperparameters. An MCMC method based on a partially collapsed Gibbs sampler was proposed to sample from the resulting posterior distribution in order to estimate the

unknown model parameters. The algorithm was validated using synthetic dataset for which it showed good detection performance. Subsequent analysis conducted using two ovine lung datasets, including a combination of fluorescent dyes, bacterial species and bacterial loads, demonstrated that the estimated bacterial count correlates with the bacterial counts performed by a clinician. The proposed algorithm is fully automatic in the sense that it does not require the user to set sensitive/crucial parameters as all model parameters are estimated within the MCMC algorithm. Comparisons with existing approaches showed superior performance of the proposed approach. Although the proposed approach is simulation-based for which a sophisticated MCMC method has been considered, it provided a competitive computation time compared to fast optimization methods, with the benefit of being fully automatic. Future work includes bacterial tracking by taking advantage of the temporal information in the datasets. Multiple concurrent imaging wavelengths can also be tested to extend to detection in multispectral datasets by using appropriate spectral unmixing algorithms.

Declaration of Competing Interest

This manuscript is an original work and therefore it has never been published previously and it is not under consideration for publication in another journal. We confirm that the manuscript has been read and approved by all named authors and that there are no other persons who satisfied the criteria for authorship but are not listed. We further confirm that the order of authors listed in the manuscript has been approved by all of us. We confirm that we have given due consideration to the protection of intellectual property associated with this work and that there are no impediments to publication, including the timing of publication, with respect to intellectual property. In so doing we confirm that we have followed the regulations of our institutions concerning intellectual property. We understand that the Corresponding Author is the sole contact for the Editorial process (including Editorial Manager and direct communications with the office). He is responsible for communicating with the other authors about progress, submissions of revisions and final approval of proofs. We confirm that we have provided a current, correct email address which is accessible by the Corresponding Author and which has been configured to accept email from (AK577@hw.ac.uk).

References

Akram, A.R., Avlonitis, N., Craven, T., Vendrell, M., McDonald, N., Scholefield, E., Fisher, A., Corris, P., Haslett, C., Bradley, M., Dhaliwal, K., 2016. Structural modifications of the antimicrobial peptide ubiquicidin for pulmonary imaging of bacteria in the alveolar space. Lancet 387, S17. doi:10.1016/S0140-6736(16)00404-9. Spring Meeting for Clinician Scientists in Training 2016

Akram, A.R., Avlonitis, N., Lilienkampf, A., Perez-Lopez, A.M., McDonald, N., Chankeshwara, S.V., Scholefield, E., Haslett, C., Bradley, M., Dhaliwal, K., 2015. A labelled-ubiquicidin antimicrobial peptide for immediate in situ optical detection of live bacteria in human alveolar lung tissue. Chem. Sci. 6 (12), 6971–6979. doi:10.1039/C5SC00960J.

Altmann, Y., McLaughlin, S., Hero, A., 2015. Robust linear spectral unmixing using anomaly detection. IEEE Trans. Comput. Image 1 (2), 74–85. doi:10.1109/TCI.2015.2455411.

André, B., Vercauteren, T., Buchner, A.M., Wallace, M.B., Ayache, N., 2011. A smart atlas for endomicroscopy using automated video retrieval. Journal of Medical Image Analysis 15 (4), 460–476.

Arteta, C., Lempitsky, V., Noble, J. A., Zisserman, A., 2014. Interactive object counting, 504–518.

Arteta, C., Lempitsky, V., Zisserman, A., 2016. Counting in the wild. In: European Conference on Computer Vision (ECCV). Springer, Amsterdam, The Netherlands, pp. 483–498.

Aslam, T., Miele, A., Chankeshwara, S.V., Megia-Fernandez, A., Michels, C., Akram, A.R., McDonald, N., Hirani, N., Haslett, C., Bradley, M., et al., 2015. Optical molecular imaging of lysyl oxidase activity–detection of active fibrogenesis in human lung tissue. Chem. Sci. 6 (8), 4946–4953.

Ayache, N., Vercauteren, T., Malandain, G., Oberrietter, F., Savoire, N., Perchant, A., 2006. Processing and mosaicing of fibered confocal images. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI): Workshop on Microscopic Image Analysis with Applications in Biology (MIAAB). Springer, Copenhagen, Denmark, pp. 1–5. Invited

Bonheur, S., Eldaly, A.K., Westerfeld, J., Wilson, D., Dhaliwal, K., McLaughlin, S., Perperidis, A., 2019. Segmentation and quantification of cellular load pulmonary endomicroscopic images using convolutional neural networks. In: IEEE International Symposium on Biomedical Imaging (ISBI), Venice, Italy, pp. 1–5.

von Borstel, M., Kandemir, M., Schmidt, P., Rao, M.K., Rajamani, K., Hamprecht, F.A., 2016. Gaussian process density counting from weak supervision. In: European Conference on Computer Vision (ECCV). Springer, Amsterdam, The Netherlands, pp. 365–380.

Bright, D.S., Steel, E.B., 1987. Two-dimensional top hat filter for extracting spots and spheres from digital images. J. Microsc. 146 (2), 191–200. doi:10.1111/j.1365-2818.1987.tb01340.x.

Désir, C., Petitjean, C., Heutte, L., Salaun, M., Thiberville, L., 2012. Classification of endomicroscopic images of the lung based on random subwindows and extra-trees. IEEE Trans. Biomed. Eng. 59 (9), 2677–2683.

Ding, X., He, L., Carin, L., 2011. Bayesian robust principal component analysis. IEEE Trans. Image Process. 20 (12), 3419–3430. doi:10.1109/TIP.2011.2156801.

Eldaly, A.K., Altmann, Y., Akram, A., Perperidis, A., McLaughlin, S., 2019. Patch-based sparse representation for bacterial detection. In: IEEE International Symposium on Biomedical Imaging (ISBI), Venice, Italy, pp. 1–5.

Eldaly, A.K., Altmann, Y., Perperidis, A., Krstajic, N., Choudhary, T.R., Dhaliwal, K., McLaughlin, S., 2018. Deconvolution and restoration of optical endomicroscopy images. IEEE Trans. Comput. Image 4 (2), 194–205. doi:10.1109/TCI.2018.2811939.

Eldaly, A.K., Altmann, Y., Perperidis, A., McLaughlin, S., 2018. Deconvolution of irregularly subsampled images. In: IEEE Statistical Signal Processing Workshop (SSP), Freiburg, Germany, pp. 303–307.

Fouquet, C., Gilles, J.-F., Heck, N., Santos, M., Schwartzmann, R., Cannaya, V., Morel, M.-P., Davidson, R.S., Trembleau, A., Bolte, S., 2015. Improving axial resolution in confocal microscopy with new high refractive index mounting media. PLoS One 10 (3), 1–17. doi:10.1371/journal.pone.0121096.

Fuchs, F., Zirlik, S., Hildner, K., Schubert, J., Vieth, M., Neurath, M., 2013. Confocal laser endomicroscopy for diagnosing lung cancer in vivo. Eur. Respir. J. 41 (6), 1401–1408. doi:10.1183/09031936.00062512.

Ge, D., Idier, J., Le Carpentier, E., 2008. A new mcmc algorithm for blind bernoulli–gaussian deconvolution. In: 2008 16th European Signal Processing Conference. IEEE, pp. 1–5.

Gelman, A., Stern, H.S., Carlin, J.B., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. Bayesian Data Analysis. Chapman and Hall/CRC.

He, F., Xiong, B., Sun, C., Xia, X., 2010. A Laplacian of Gaussian-based approach for spot detection in two-dimensional gel electrophoresis images. In: International Conference on Computer and Computing Technologies in Agriculture. Springer, Beijing, China, pp. 8–15.

Ishijima, A., Schwarz, R.A., Shin, D., Mondrik, S., Vigneswaran, N., Gillenwater, A.M., Anandasabapathy, S., Richards-Kortum, R., 2015. Automated frame selection process for high-resolution microendoscopy. J. Biomed. Opt. 20 (4). 046014–046014

Jabbour, J.M., Saldua, M.A., Bixler, J.N., Maitland, K.C., 2012. Confocal endomicroscopy: instrumentation and medical applications. Ann. Biomed. Eng. 40 (2), 378–397. doi:10.1007/s10439-011-0426-y.

Jeffreys, H., 1946. An invariant form for the prior probability in estimation problems. Proc. R. Soc. Lond. Ser.A Math. Phys. Sci.x 186 (1007), 453–461. doi:10.1098/rspa.1946.0056.

Kail, G., Tourneret, J.-Y., Hlawatsch, F., Dobigeon, N., 2010. A partially collapsed Gibbs sampler for parameters with local constraints. In: 2010 IEEE International Conference on Acoustics, Speech and Signal Processing. IEEE, pp. 3886–3889.

Kail, G., Tourneret, J.-Y., Hlawatsch, F., Dobigeon, N., 2012. Blind deconvolution of sparse pulse sequences under a minimum distance constraint: a partially collapsed Gibbs sampler method. IEEE Trans. Signal Process. 60 (6), 2727–2743.

Kimori, Y., Baba, N., Morone, N., 2010. Extended morphological processing: a practical method for automatic spot detection of biological markers from microscopic images. BMC Bioinform. 11 (1), 1–13.

Koujan, M., Ahsan, A., McCool, P., Westerfeld, D., Dhaliwal, K., McLaughlin, S., Perperidis, A., 2018. Multi-class classification of pulmonary endomicroscopic images. In: IEEE International Symposium on Biomedical Imaging (ISBI), Washington DC, USA, pp. 1–4.

Krstajić, N., Akram, A.R., Choudhary, T.R., McDonald, N., Tanner, M.G., Pedretti, E., Dalgarno, P.A., Scholefield, E., Girkin, J.M., Moore, A., et al., 2016. Two-color widefield fluorescence microendoscopy enables multiplexed molecular imaging in the alveolar space of human lung tissue. J. Biomed. Opt. 21 (4). 046009–046009

Le Goualher, G., Perchant, A., Genet, M., Cavé, C., Viellerobe, B., Berier, F., Abrat, B., Ayache, N., 2004. Towards optical biopsies with an integrated fibered confocal fluorescence microscope. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Springer, Saint-Malo, Brittany, France, pp. 761–768.

Lenseigne, B., Brodin, P., Jeon, H.K., Christophe, T., Genovesio, A., 2007. Support vector machines for automatic detection of tuberculosis bacteria in confocal microscopy images. In: IEEE International Symposium on Biomedical Imaging: From Nano to Macro (ISBI), Arlington, VA, USA, pp. 85–88. doi:10.1109/ISBI.2007.356794.

Lin, C., Mailhes, C., Tourneret, J.-Y., 2010. P-and t-wave delineation in ecg signals using a Bayesian approach and a partially collapsed Gibbs sampler. IEEE Trans. Biomed. Eng. 57 (12), 2840–2849.

Lindeberg, T., 1998. Feature detection with automatic scale selection. Int. J. Comput. Vis. 30 (2), 79–116. doi:10.1023/A:1008045108935DO.

Liu, J.S., 1994. The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem. J. Am. Stat. Assoc. 89 (427), 958–966.

Mardia, K., 1988. Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. J. Multivar. Anal. 24 (2), 265–284. doi:10.1016/0047-259X(88)90040-1.

Mazet, V., Brie, D., Idier, J., 2005. Simulation of positive normal variables using several proposal distributions. In: IEEE Statistical Signal Processing Workshop (SSP), Bordeaux, France, pp. 37–42.

McCool, P., Altmann, Y., Perperidis, A., McLaughlin, S., 2016. Robust Markov random field outlier detection and removal in subsampled images. In: IEEE Statistical Signal Processing Workshop (SSP), Palma de Mallorca, Spain, pp. 1–5. doi:10.1109/SSP.2016.7551766.

Newstadt, G.E., Hero, A.O., Simmons, J., 2014. Robust spectral unmixing for anomaly detection. In: IEEE Statistical Signal Processing Workshop (SSP), Gold Coast, Australia, pp. 109–112. doi:10.1109/SSP.2014.6884587.

Newton, R., Kemp, S., Yang, G.-Z., Elson, D., Darzi, A., Shah, P., 2012. Imaging parenchymal lung diseases with confocal endomicroscopy. Respir. Med. 106 (1), 127–137. doi:10.1016/j.rmed.2011.09.009.

Parker, H.E., Perperidis, A., Karam-Eldaly, A., Altmann, Y., Dhaliwal, K., Thomson, R.R., McLaughlin, S., Tanner, M.G., 2019. Characterising cross-coupling in coherent fibre bundles. In: European Conferences on Biomedical Optics, Munich, Germany, pp. 1–3.

Perperidis, A., Akram, A., Altmann, Y., McCool, P., Westerfeld, J., Wilson, D., Dhaliwal, K., McLaughlin, S., 2017. Automated detection of uninformative frames in pulmonary optical endomicroscopy. IEEE Trans. Biomed. Eng. 64 (1), 87–98.

Perperidis, A., Dhaliwal, K., McLaughlin, S., Vercauteren, T., 2018. Image computing for fibre-bundle endomicroscopy: areview. arXiv:1809.00604.

Perperidis, A., Parker, H.E., Karam-Eldaly, A., Altmann, Y., Dhaliwal, K., Thomson, R.R., Tanner, M.G., McLaughlin, S., 2017. Characterization and modelling of inter-core coupling in coherent fiber bundles. Opt. Express 25 (10), 11932–11953. doi:10.1364/OE.25.011932.

Preparata, F.P., Shamos, M.I., 2012. Computational Geometry: An Introduction. Springer Science & Business Media.

Robert, C., Casella, G., 2013. Monte Carlo Statistical Methods. Springer Science & Business Media.

Rue, H., Held, L., 2005. Gaussian Markov random fields: theory and applications. CRC press.

Savoire, N., Andre, B.G., Vercauteren, T., 2012. Online blind calibration of non-uniform photodetectors: application to endomicroscopy. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI), Nice, France: Springer, pp. 1–8.

Sibson, R., 1981. A brief description of natural neighbor interpolation. Interpret. Multivar. Data 21–36.

Thiberville, L., Moreno-Swirc, S., Vercauteren, T., Peltier, E., Cavé, C., Bourg Heckly, G., 2007. In vivo imaging of the bronchial wall microstructure using fibered confocal fluorescence microscopy. Am. J. Respir. Crit. Care Med. 175 (1), 22–31. doi:10.1164/rccm.200605-684OC.

Thiberville, L., Salaun, M., Lachkar, S., Dominique, S., Moreno-Swirc, S., Vever-Bizet, C., Bourg-Heckly, G., 2009. Human in vivo fluorescence microimaging of the alveolar ducts and sacs during bronchoscopy. Eur. Respir. J. 33 (5), 974–985. doi:10.1183/09031936.00083708.

Thiberville, L., Salaün, M., Lachkar, S., Dominique, S., Moreno-Swirc, S., Vever-Bizet, C., Bourg-Heckly, G., 2009. In vivo confocal fluorescence endomicroscopy of lung cancer. 13th World Conference on Lung Cancer, San Francisco, United States. 1–1

Thiberville, L., Salaun, M., Lachkar, S., Dominique, S., Moreno-Swirc, S., Vever-Bizet, C., Bourg-Heckly, G., 2009. Confocal fluorescence endomicroscopy of the human airways. Proc. Am. Thorac. Soc. 6 (5), 444–449.

Van Dyk, D.A., Park, T., 2008. Partially collapsed Gibbs samplers: theory and methods. J. Am. Stat. Assoc. 103 (482), 790–796. doi:10.1198/016214508000000409.

Vercauteren, T., Perchant, A., Malandain, G., Pennec, X., Ayache, N., 2006. Robust mosaicing with correction of motion distortions and tissue deformations for in vivo fibered microscopy. Med. Image Anal. 10 (5), 673–692.

Vercauteren, T., Perchant, A., Pennec, X., Ayache, N., 2005. Mosaicing of confocal microscopic in vivo soft tissue video sequences. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). Springer, California, USA, pp. 753–760.

Veropoulos, K., Campbell, C., Learmonth, G., Knight, B., Simpson, J., 1998. The automated identification of tubercle bacilli using image processing and neural computing techniques. In: International Conference on Artificial Neural Networks (ICANN). Springer, Skovde, Sweden, pp. 797–802. doi:10.1007/978-1-4471-1599-1_123.