Estimating Bacterial and Cellular Load in FCFM Imaging

Sohan Seth

University of Edinburgh, School of Informatics, Edinburgh, EH8 9AB, UK

October 1, 2018

THE UNIVERSITY *of* EDINBURGH

Proteus: lighting up the lung, detecting disease

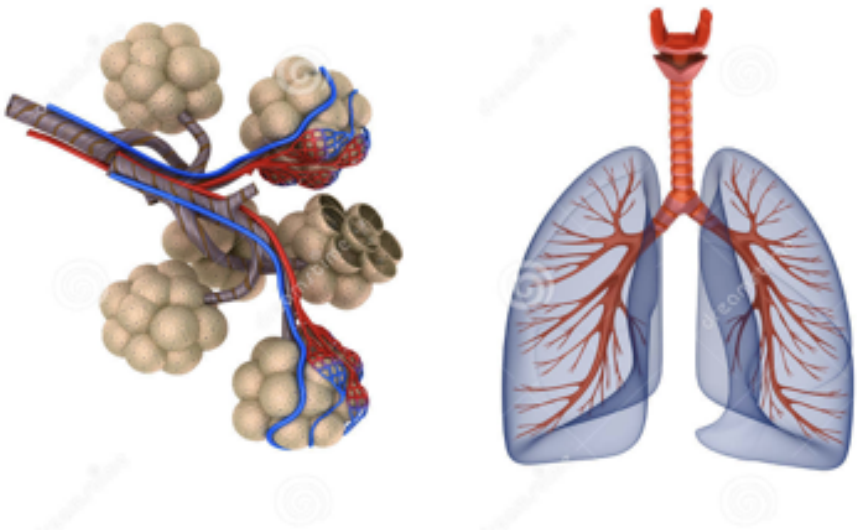

Clinical Need

"diagnosing bacterial infections relies on a slow process of detection followed by biopsy and lab-based culture growth – procedures that are prone to contamination and can result in late treatment."

Vision

"a fully integrated system that will provide the necessary rapid and accurate diagnosis of bacterial infection"

FCFM: How does it work?

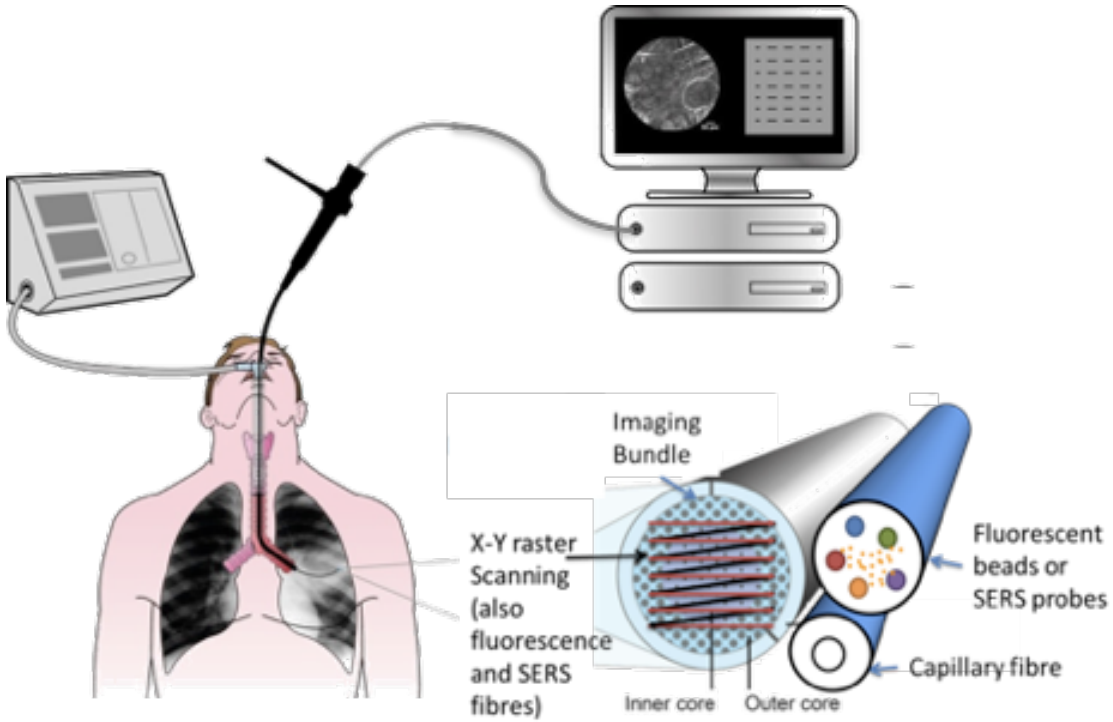

Fibered Confocal Fluorescence Microscopy

- Fiber optic imaging cable is inserted to the distal lung through a bronchoscope
- Imaging is performed by counting emitted photons through fibre optic
- Smartprobe (chemical compound) is delivered to make bacteria fluoresce

FCFM: What do we see?

Before smartprobe

After smartprobe

FCFM: What do we see?

Before smartprobe

After smartprobe

- Background composes of autofluorescent elastin (connective lung tissue)
- Autofluorescent cells appear as round objects
- Bacteria appear as ‘blinking dots’.
- Our objective is to detect and count the bacteria and cell in each frame.

Machine learning problem

Task: At each pixel, predict if there is an object (bacterium or cell)

Machine learning problem

Task: At each pixel, predict if there is an object (bacterium or cell)

Why is it interesting?

1. Noisy images: motion blur etc.
2. Limited annotations: time consuming etc.
3. Noisy annotations: manual error etc.
4. Imbalanced dataset: few objects per frame

Machine learning problem

Task: At each pixel, predict if there is an object (bacterium or cell)

Why is it interesting?

1. Noisy images: motion blur etc.
2. Limited annotations: time consuming etc.
3. Noisy annotations: manual error etc.
4. Imbalanced dataset: few objects per frame

Overview of approach:

1. Create a set of annotated images with the help of a clinician
2. Represent each pixel by an appropriate feature vector
3. Use the annotated images to train a classifier
4. Use cross-validation to find the best feature representation and classifier
5. Use the best classifier and feature representation to count objects in test videos

Machine learning problem

Task: At each pixel, predict if there is an object (bacterium or cell)

Why is it interesting?

1. Noisy images: motion blur etc.
2. Limited annotations: time consuming etc.
3. Noisy annotations: manual error etc.
4. Imbalanced dataset: few objects per frame

Overview of approach:

1. Create a set of annotated images with the help of a clinician
2. Represent each pixel by an appropriate feature vector
3. Use the annotated images to train a classifier
4. Use cross-validation to find the best feature representation and classifier
5. Use the best classifier and feature representation to count objects in test videos

Two same pipelines for cells and bacteria respectively.

1. Annotations

- We designed an interface for clinicians to annotate objects.
- To help the annotator capture the 'blinking effect', the present frame was annotated in the context of the past and future frames.
- We randomly chosen some frames to annotate.

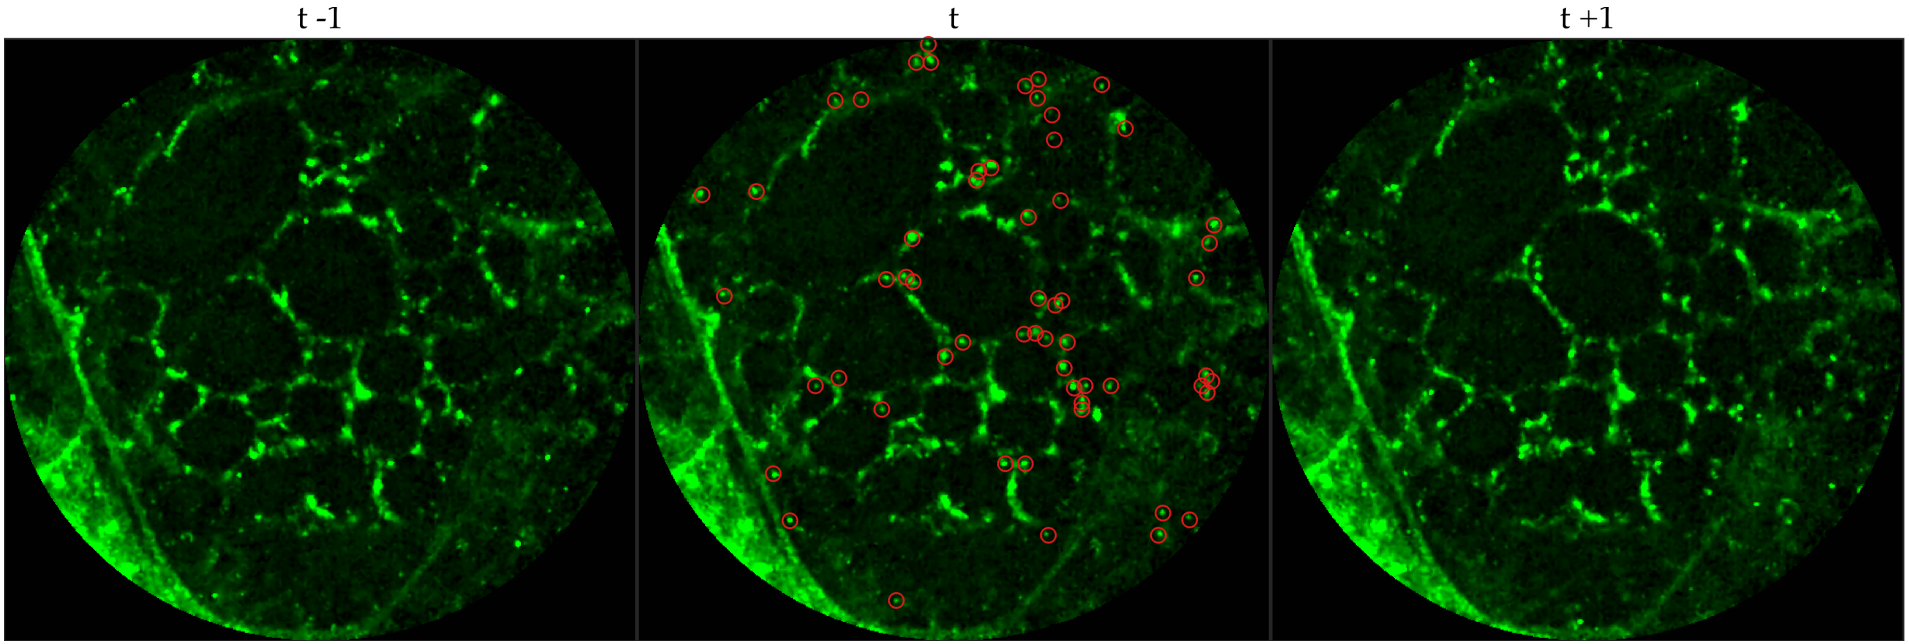

Figure: FCFM image frame with bacteria annotated by a clinician in circles.

2. Standardized multi-resolution spatio-temporal feature extraction

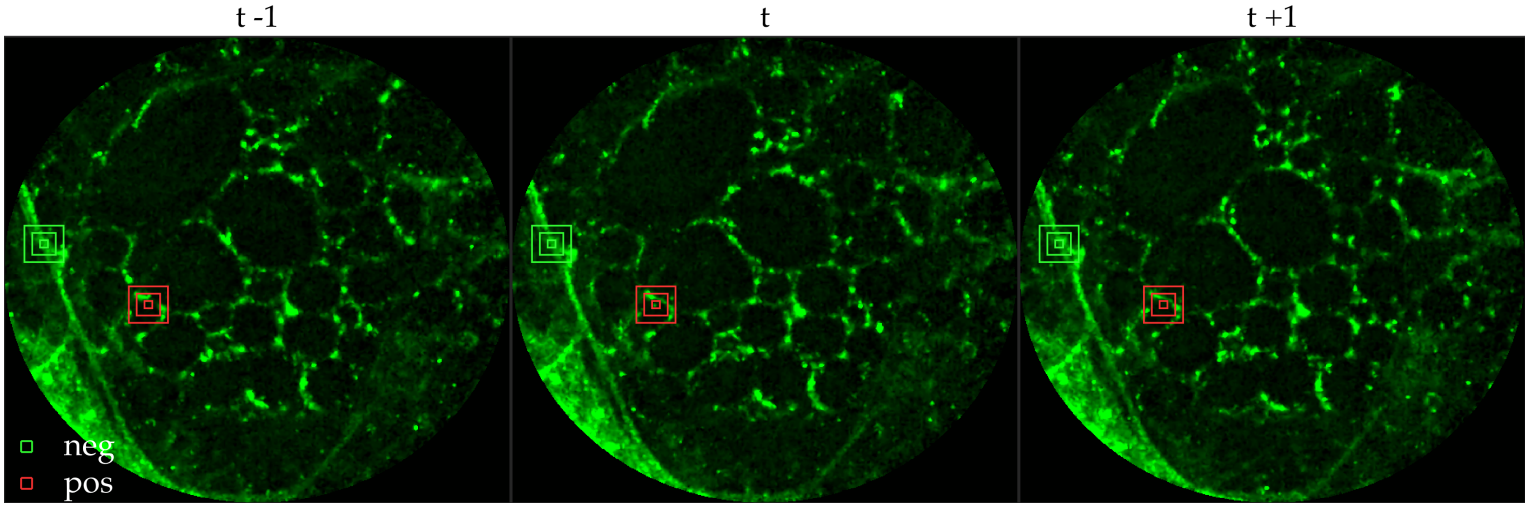

Image patches around annotated (red) and non-annotated (green) pixel

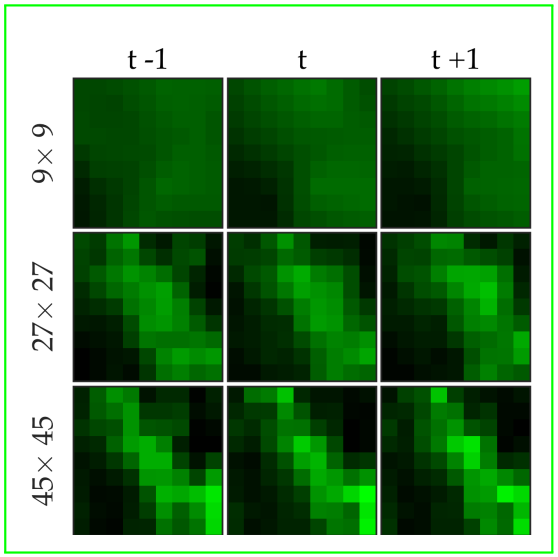

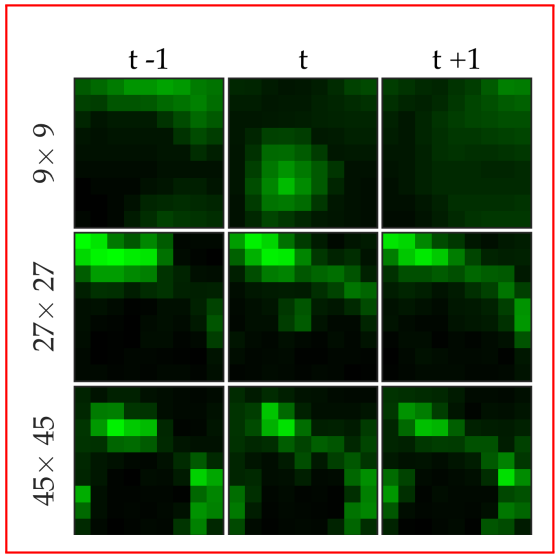

2. Standardized multi-resolution spatio-temporal feature extraction

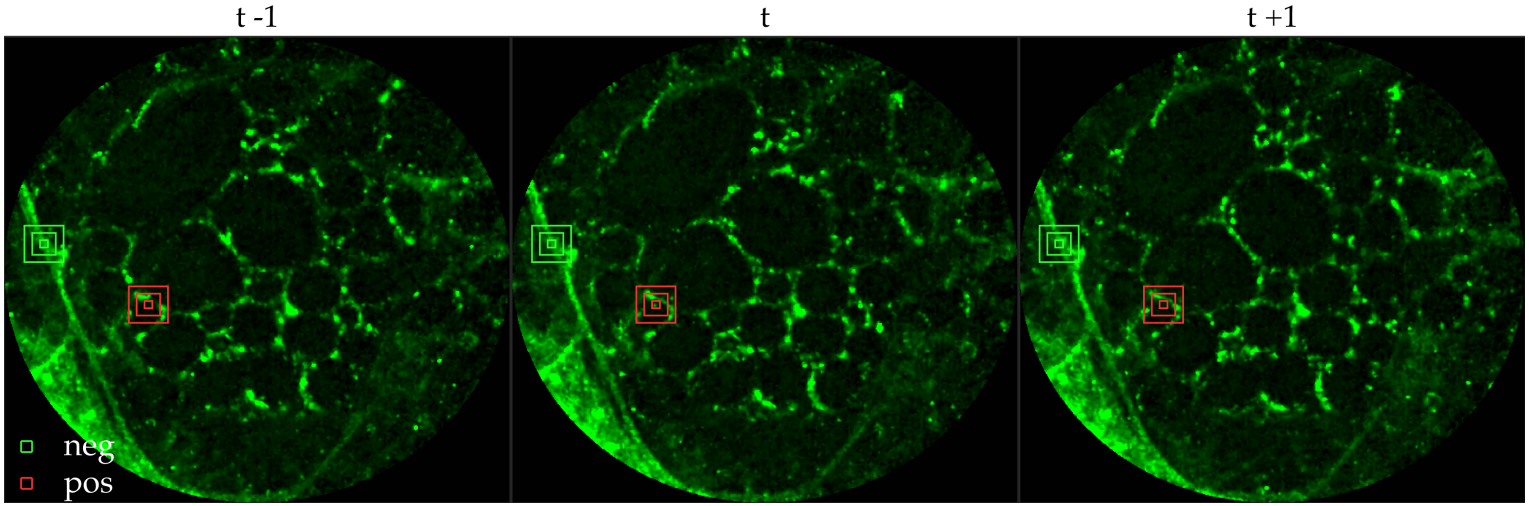

Image patches around annotated (red) and non-annotated (green) pixel

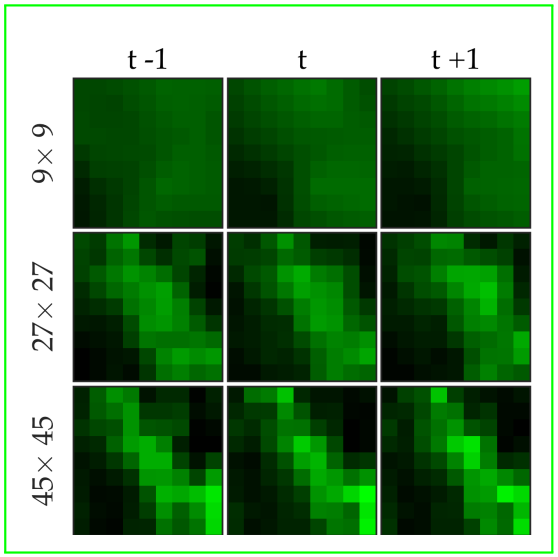

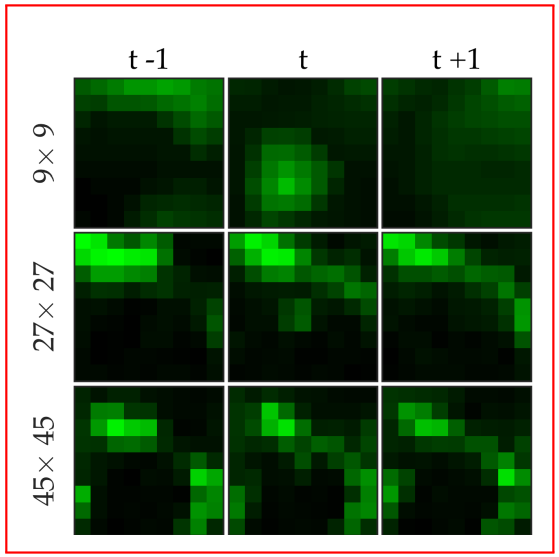

- Multiple resolutions provide context beside the object.
- Each spatio-temporal patch is standardized to give equal importance.

Task: Predict output $y \in \{0, 1\}$ from high dimensional feature vector $\mathbf{x} \in \mathbb{R}^d$

Classifier: (nonlinear) template matching with radial basis fuctions (RBF)

$$p(y = 1) = \sigma\left(\sum_{i=1}^{T} \alpha_i \kappa(\mathbf{x} - \mathbf{t}_i) + b\right)$$

- κ is measures similarity, e.g., if $\mathbf{x} = \mathbf{t}$ when $\kappa = 1$, and if $\mathbf{x} \neq \mathbf{t}$ when $\kappa = 0$, and
- $\mathbf{t}_i$s are feature **templates** learned using k-means.

Task: Predict output $y \in \{0, 1\}$ from high dimensional feature vector $\mathbf{x} \in \mathbb{R}^d$

Classifier: (nonlinear) template matching with radial basis fuctions (RBF)

$$p(y = 1) = \sigma\left(\sum_{i=1}^{T} \alpha_i \kappa(\mathbf{x} - \mathbf{t}_i) + b\right)$$

- κ is measures similarity, e.g., if $\mathbf{x} = \mathbf{t}$ when $\kappa = 1$, and if $\mathbf{x} \neq \mathbf{t}$ when $\kappa = 0$, and
- $\mathbf{t}_i$s are feature **templates** learned using k-means.

How to address class imbalance?

- k-means is done with all feature vector (about 1% with bacteria or cells)
- classifier is learned using balanced samples by subsampling 1% of the features without bacteria or cells.

3. Supervised learning

Task: Predict output $y \in \{0, 1\}$ from high dimensional feature vector $\mathbf{x} \in \mathbb{R}^d$

Classifier: (nonlinear) template matching with radial basis fuctions (RBF)

$$p(y = 1) = \sigma\left(\sum_{i=1}^{T} \alpha_i \kappa(\mathbf{x} - \mathbf{t}_i) + b\right)$$

- κ is measures similarity, e.g., if $\mathbf{x} = \mathbf{t}$ when $\kappa = 1$, and if $\mathbf{x} \neq \mathbf{t}$ when $\kappa = 0$, and
- $\mathbf{t}_i$s are feature **templates** learned using k-means.

How to address class imbalance?

- k-means is done with all feature vector (about 1% with bacteria or cells)
- classifier is learned using balanced samples by subsampling 1% of the features without bacteria or cells.

Consecutive pixels have high probabilities.

- The probability values at each pixel were thresholded
- Non-maximum suppression was applied to avoid multiple detections.

4. Cross validation using area under precision-recall curve

$+$ are ground truth, $\bullet$ are detections

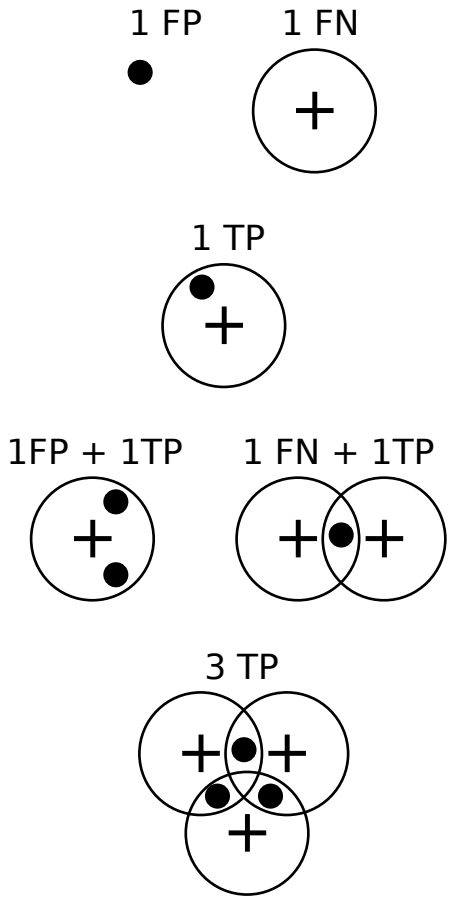

$$P = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}, \quad R = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}.$$

4. Cross validation using area under precision-recall curve

+ are ground truth, • are detections

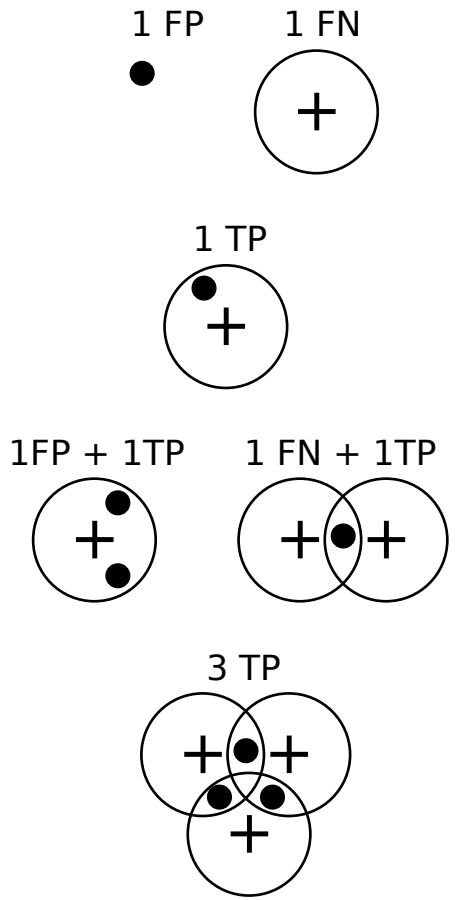

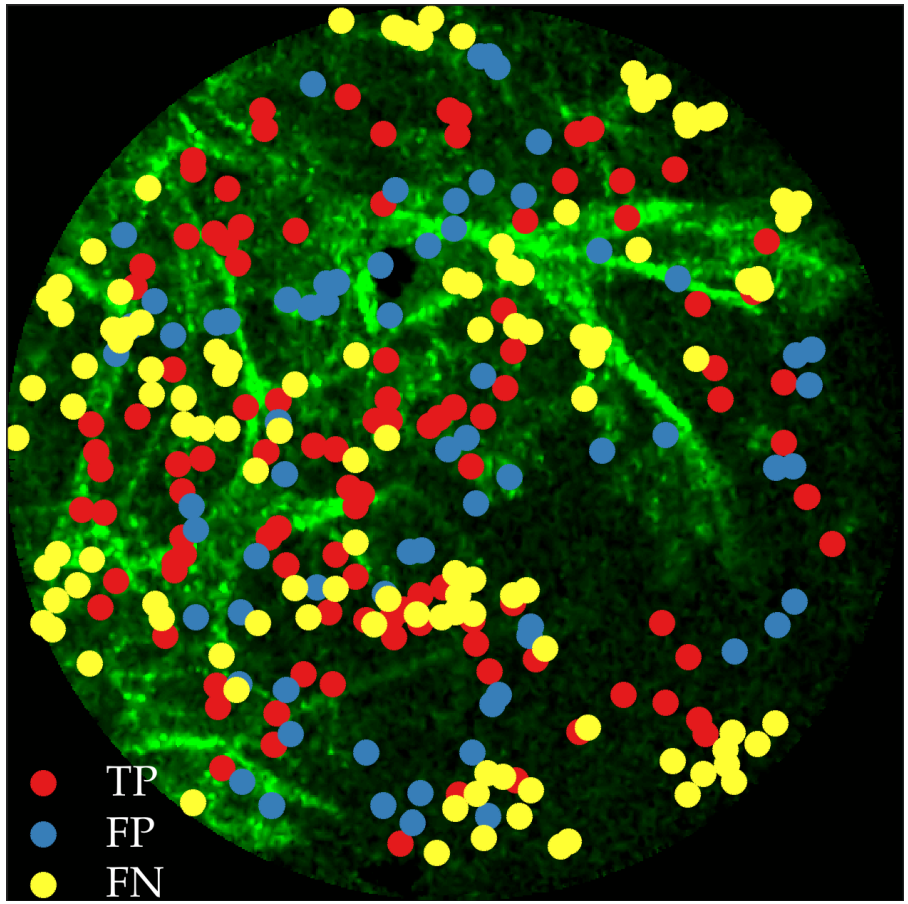

Figure: Ground truth vs. detections: Bacteria

$$P = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}, \quad R = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}.$$

4. Cross validation using area under precision-recall curve

+ are ground truth, • are detections

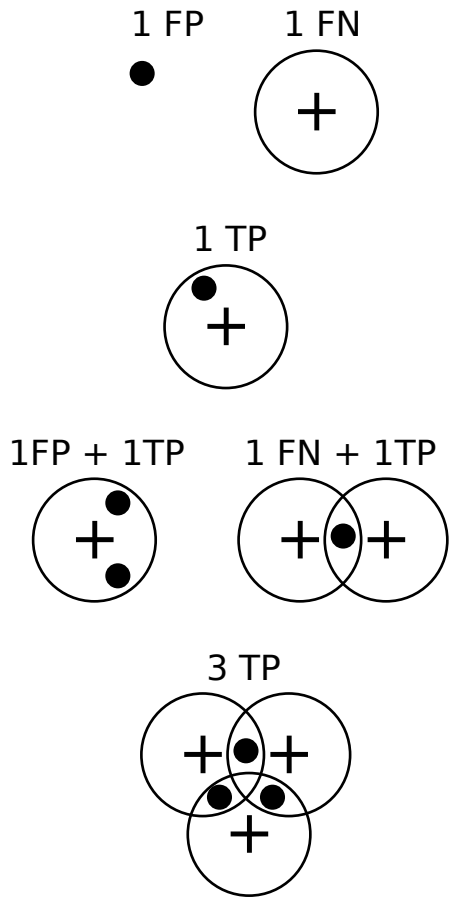

$$P = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad R = \frac{\text{TP}}{\text{TP} + \text{FN}}.$$

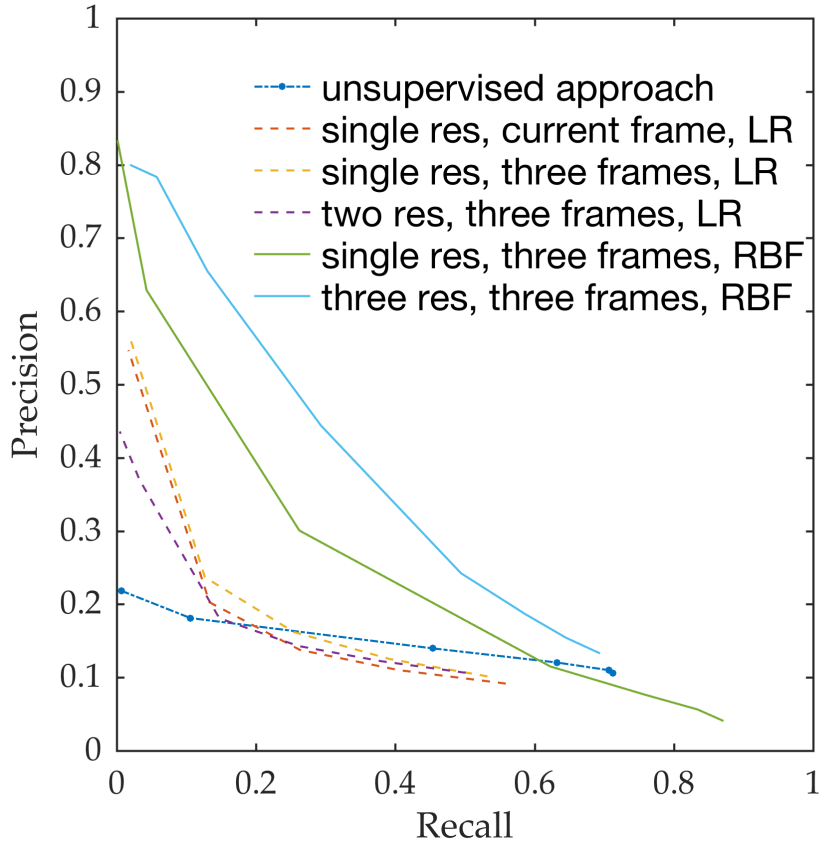

Figure: Precision recall curves for different methods

- Temporal information helps.
- Performs better than unsupervised approach.

4. Cross validation using area under precision-recall curve

+ are ground truth, • are detections

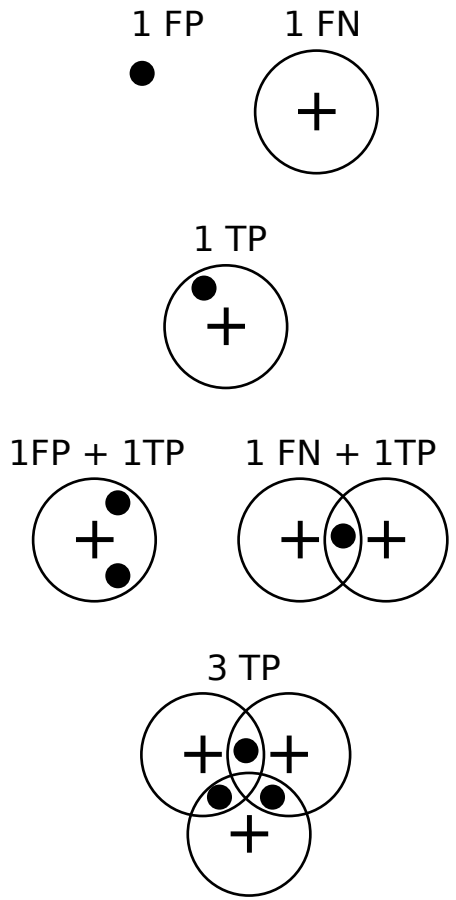

$$P = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FP}}, \quad R = \frac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}.$$

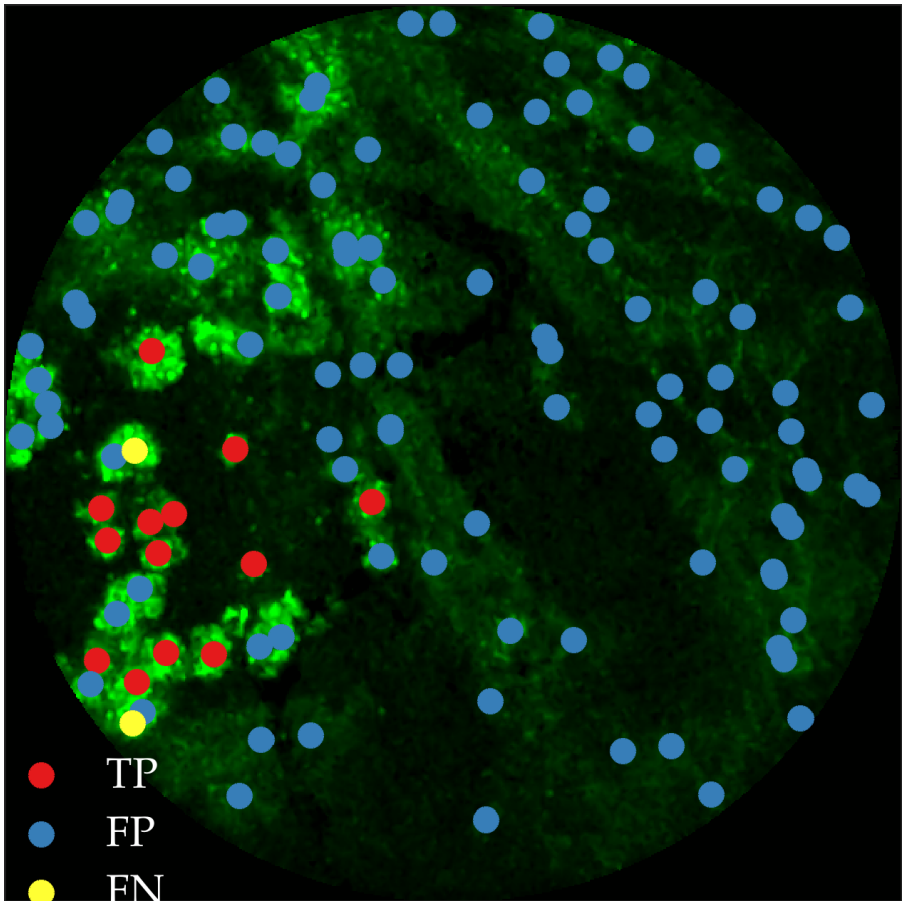

Figure: Ground truth vs. detections: Cells

4. Cross validation using area under precision-recall curve

+ are ground truth, • are detections

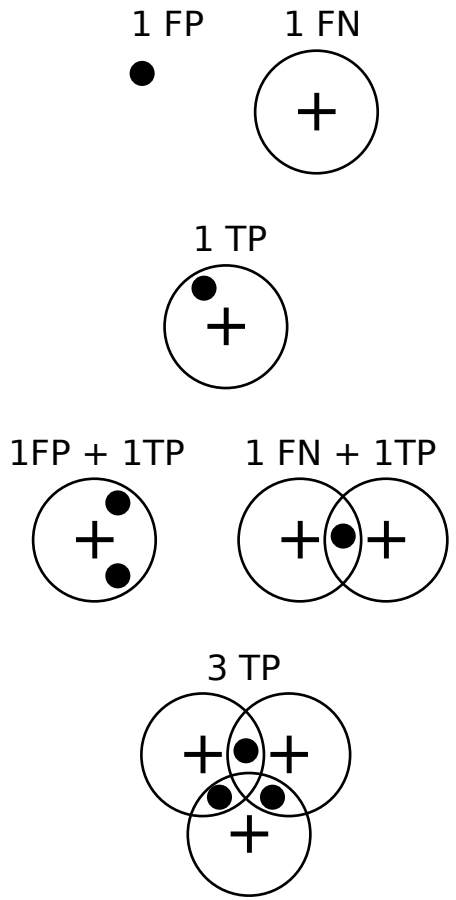

$$P = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad R = \frac{\text{TP}}{\text{TP} + \text{FN}}.$$

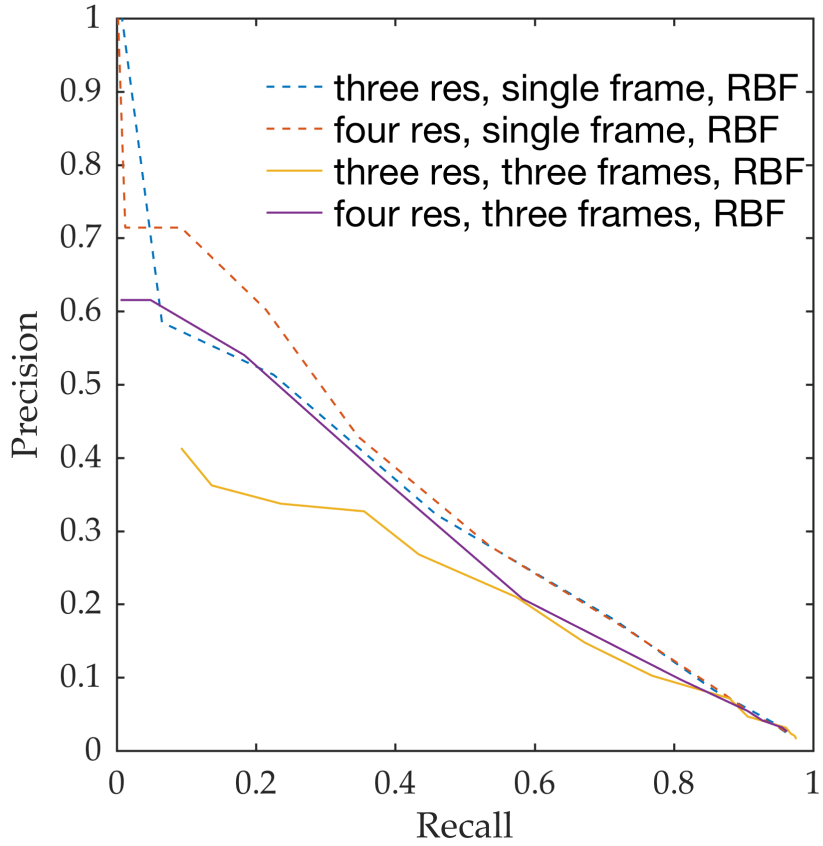

Figure: Precision recall curves for different methods

- Temporal information does not help.
- Larger spatial frame needed.

5. Bacterial load: FCFM videos

Before smartprobe

After smartprobe

5. Bacterial load: Case vs. control: pre- vs. post-substance: 1st cohort

Summary: 12 videos with ∼500 frames each, 3 videos (post-substance cases) with bacteria and 9 videos (pre-substance or control) without bacteria

	control	**control**	**control**	**case**	**case**	**case**
pre-substance	-	-	-	-	-	-
post-substance	-	-	-	+	+	+

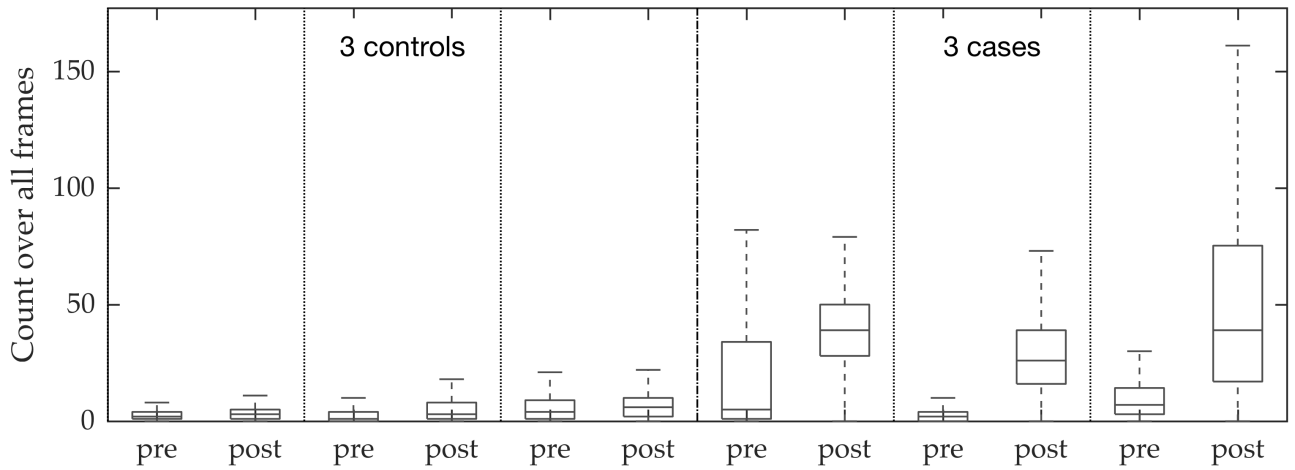

Figure: Change of bacterial load in 6 patients, pre- and post-substance

- We detect a statistically significant change in cases but not in controls.
- A fraction of frames were used for training.

5. Bacterial load: Case vs. control: pre- vs. post-substance: 2nd cohort

Summary: 10 videos with ∼500 frames each, 2 videos (post-substance cases) with bacteria and 8 videos (pre-substance or control) without bacteria

	control	**control**	**control**	**case**	**case**
pre-substance	-	-	-	-	-
post-substance	-	-	-	+	+

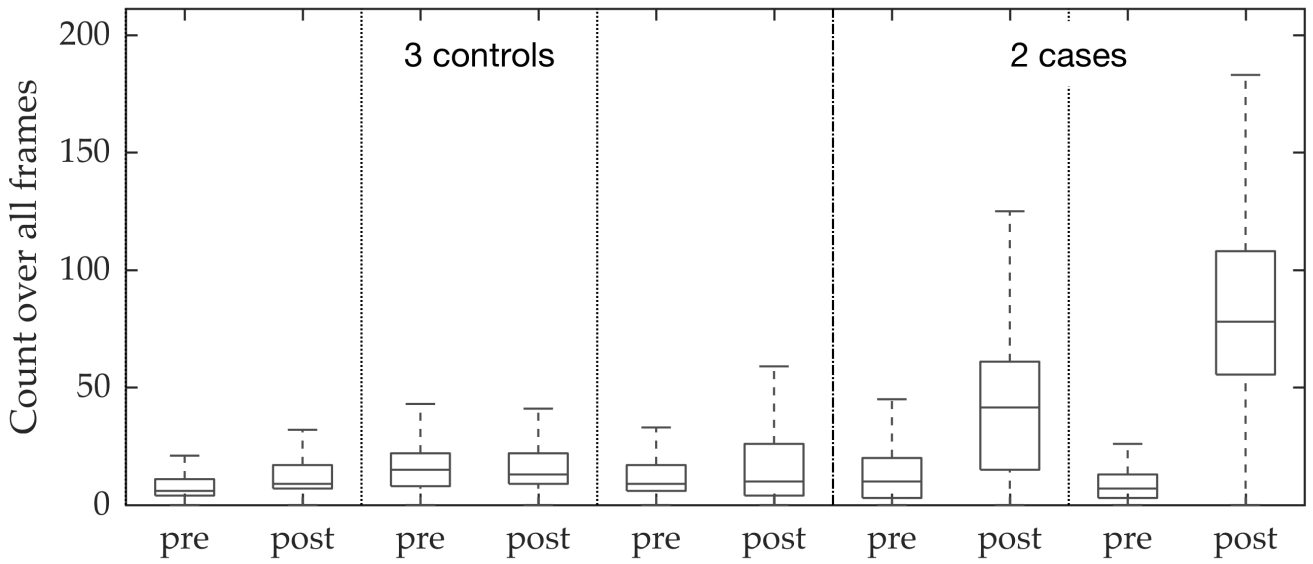

Figure: Change of bacterial load in 6 patients, pre- and post-substance

- We detect a statistically significant change in cases but not in controls.
- A fraction of frames from first cohort were used for training.

5. Cellular load: FCFM videos

Some cells

Very cellular

5. Cellular load: Comparison with clinical assessment

- 206 FCFM videos from 102 patients who have undergone bronchoscopy
- A fraction of frames were used for training
- A clinician annotated the videos according to their level of cellularity

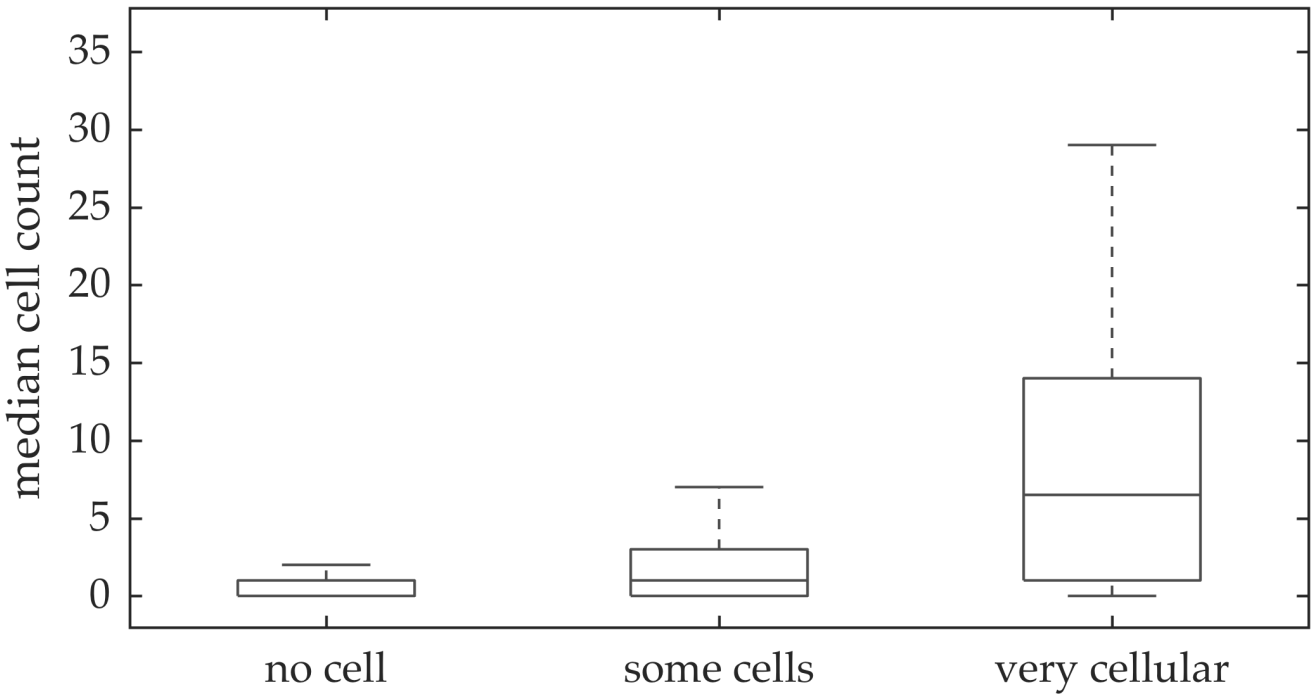

Figure: Comparison of median cell count against visual assessment of cellularity.

- The estimated cellular load agrees with the visual assessment of the clinician
- A fraction of frames were used for training.

Summary

- We address the task of estimating bacterial and cellular load in FCFM images.
- We create a database of annotated image frames.
- We observe significant fold change in the case videos before and after introducing smartprobe, which is not observed in the control group.
- We show that the estimated cellular load agrees with the clinician's assessment

What is next?

More data is becoming available

- 59 patients imaged in Edinburgh till date
- 19 in ICU: all mechanically ventilated
- Average duration of procedure: 8 minutes, 3-5 passes
- No significant adverse events

What is next?

More data is becoming available

- 59 patients imaged in Edinburgh till date
- 19 in ICU: all mechanically ventilated
- Average duration of procedure: 8 minutes, 3-5 passes
- No significant adverse events

- Distal lung is a “black hole”, we are developing approaches to understanding it
- Our system should be immediate, bedside, cheap, safe, accurate,and add to decision making, so it can become part of routine care.
- It should potentially be automated: i.e., self guided to different regions of the lung with clear sampling and external registration

Thank you!


Journal of
Imaging
Estimating Bacterial and Cellular Load in FCFM Imaging
Volume 4 - Issue 1 | January 2018
MDPI
mdpi.com/journal/jimaging
ISSN 2313-433X