

# Supervised Machine Learning Methods for Quantification of Pulmonary Fibrosis

Manuel Stritt, Roman Bär, Joël Freyss, Julia Marrie,  
Enrico Vezzali, Edgar Weber, and Anna Stalder

Actelion Pharmaceuticals Ltd.  
Hegenheimermattweg 91, CH-4123 Allschwil, Switzerland  
{manuel.stritt,anna.stalder}@actelion.com  
<http://www.actelion.com>

**Abstract.** Applied pattern recognition for microscopy images is a requirement for the automatization of image classification in a high content process. This paper describes methods and a system which classifies digital high resolution tissue slides. These methods cover the schedule line of global tissue regions and fine granular local tissue classification, down to the segmentation and classification of single cells. A system called Orbit, which implements these methods by making use of grid computing, is explained in detail and applied to a real world data set of lung tissue slides to quantify the degree of pulmonary fibrosis. A combined model which takes tissue classification, region discrimination, and mast cells into account results in a correlation coefficient of 0.81 relative to real fibrosis scores defined by a pathologist.

**Keywords:** machine learning, pattern recognition, tissue classification, cell quantification, grid computing, idiopathic pulmonary fibrosis

## 1 Introduction

Idiopathic pulmonary fibrosis (IPF) is a fatal life-threatening disease which involves the thickening and scarring of the lungs, typically from an unknown cause, which makes it hard for oxygen to enter the blood. Currently no specific treatments have yet been approved for this disease. It is known that fibrotic lungs have additional collagen in normal lung tissue and Hunt et al. reported an increase of mast cells [8]. In the context of biological studies it is necessary to determine the degree of fibrosis to quantify the effect of a disease model or drug. This can be done manually by pathologists looking through a microscope and analyzing the stained tissue structure. A digital tissue slide can be produced by staining the collagen in lung tissue and scanning an image with a digital microscope in a highly automated process. These digital tissue slides, captured at 20x magnification, are very high resolution images, typically around 100,000 x 60,000 pixels with three color channels. Since this procedure from the staining to image acquisition, is highly automated, up to 80 digital slides can be produced within a day. However, manual analysis of these slides slows down the whole

process. Therefore, we describe a system called Orbit that can store images centrally, distribute them among several clients, and perform image analysis on an in-house grid computing environment consisting of 300 nodes.

This paper is organized as follows: In Section 2, we describe the similarities and differences between related work and our system. In Section 3, the Orbit system and its underlying algorithms are described. The results of the IPF case study are presented in Section 4. We conclude with a discussion and thoughts on future research in Section 5.

## 2 Related Work

Automated image analysis as part of the content-based image retrieval (CBIR) process has been studied in several papers during the last decade [5]. In the last years, supervised learning in image analysis has become very important. In general, there are two main clusters: The first set of systems classifies image regions, for example, tissue classification. Susomboon et al. [17] studied pixel-based classification and applied it to CT images. The well-studied Haralick texture features [7] are used in many applications to generate features based on a set of spatial-dependence probability-distribution matrices. These features are very meaningful when the context size is large compared to the number of (gray) color shades. We are interested in local features and therefore use small regions around pixels. We use similar features as Haralick, but, instead of using the relative frequency distribution, we apply the statistics to the local region data directly (Section 3). Caicedo and Izquierdo studied several low-level features in tissue images and concluded that similarly distributed tissue features result in a boost of accuracy where a combination of fundamentally different image descriptors can lead to a performance decrease [3].

The second cluster identifies objects, for example cells, and classifies them [4, 9, 13, 10]. This method is also applied in the context of high content screening [15]. Snijder et al. report that the cell population context is important for the cell-to-cell variability in endocytosis and virus infection [16]. This very interesting outcome underlines the importance of extracting features on different levels. Applications which analyze and track cells in videos instead of single images also exist [1].

In contrast to these two clusters, our system has to struggle with both tissue and cell classification to determine both the degree of fibrosis, the quantification of a tissue class (Section 3.1), and the cell count of a particular cell type (Section 3.2) but each in specific tissue regions (Section 3.1). Orbit has been successfully applied to the identification of Cathepsin L as a potential sex-specific biomarker for renal damage [2]. The architecture and value to the drug discovery process of our in-house grid used by Orbit is discussed in [11].

### 3 Methods

The Orbit system processes images with three modules: The first module performs a pixel-based tissue classification which can be applied for fine granular local structure classification or, in a larger context, region discrimination (Section 3.1). The second module does an object segmentation which is usually used for cell segmentation. The combination of these modules allows a very convenient segmentation procedure without explicitly specifying a foreground/background-threshold (Section 3.2). A third module which can be used for cell classification classifies segmented objects (Section 3.3).

#### 3.1 Tissue Classification

Tissue classification is performed with supervised machine learning methods. In a first step, the user visually marks representative regions of tissue classes  $C$ . This can be, for instance, normal lung tissue, fibrotic lung tissue, and background. The marking is done by drawing shapes onto one or several images which works the same way as a brush tool in typical painting applications. For each class  $c \in C$ , several shapes  $s \in S_c$  exist which are possibly distributed among several images. The set of pixels  $P_{train}$  within the shapes labeled with the associated tissue class defines the training set, formal


$$P_{train} = \{p_t \in \bigcup_{c \in C} S_c\} \quad (1)$$

Where  $p \in s$  means pixel  $p$  is located in shape  $s$ . In the next step, a set of features is computed for each pixel. Therefore a window-size  $w$  is defined, usually between 3 and 20. All features are computed for a given (middle) pixel and the window-size which defines the local context around the pixel ( $(w/2) - 1$  pixel left, right, up and down). We use a combination of the following six features:

$$\begin{aligned} ident_p &= p & min_p &= \min p' \in W & max_p &= \max p' \in W & \mu_p &= \frac{1}{|W|} \sum_{p' \in W} p \\ \sigma_p &= \sqrt{\frac{1}{|W|-1} \sum_{p' \in W} (p' - \mu_p)^2} & edge_p &= \sqrt{\frac{1}{|W|-1} \sum_{p' \in W} (p' - p)^2} \end{aligned}$$

Where  $W$  is the set of pixels inside the window region.

These features are computed separately for each color channel. The images in the pulmonary fibrosis case study consist of three color channels (RGB) which lead to 18 numerical features per pixel. A support vector machine (SVM) is used to build the model based on the features of the training set  $P_{train}$  and to perform the pixel-based classification of the whole image. This procedure leads to the assignment of a tissue class to each pixel of the image:  $\forall_{p \in P} p \leftarrow c \in C$ . Figure 1 shows a lung tissue cut at 20x magnification and the computed classification displayed as an overlay.




**Fig. 1.** Extract of a lung tissue cut at 20x magnification (left). The collagen is stained blue. The right image displays the pixel-wise classification where the collagen tissue class is colored green.

**Region Discrimination** For pulmonary fibrosis, quantification it is not sufficient to quantify all the collagen in the lung because the majority of the collagen present is so-called structural collagen found around vessels. This kind of collagen is unrelated to pulmonary fibrosis-disease and is absolutely necessary for stabilizing the air and blood vessel borders. Unfortunately, this collagen is the same structure and color as the collagen generated by the fibrosis. Biological experiments have shown that it is not possible to discriminate between structural and non-structural collagen in the staining procedure.

Our first approach was to consider the vessels as objects and to apply a shape detection algorithm which finds 'holes' in the lungs and excludes these 'holes' and their border from further analysis. Unfortunately the results have been poor due to the high variability of vessel shape. Much better results were achieved by considering the vessel borders as a special tissue while looking at the lung at a 5x magnification. As mentioned, on a high resolution image, the different type of collagen could not be distinguished by structure and color. However, in a low resolution image combined with a big window-size ( $w = 25$ ), the supervised tissue classification procedure can detect the presence of collagen with a hard edge defined by the vessel border. In addition to the vessel borders, global tissue regions such as larger dense fibrotic regions, could be determined as well. Figure 2 shows a low resolution scan of a lung cut and the according discrimination regions classified by a SVM. The model has been trained on training regions of three other lung cuts.

This approach enables supervised region discrimination on low resolution images by applying tissue classification for high resolution images only to specific regions of the low resolution classification. This allows to mask the lung slide and to perform the collagen tissue quantification while excluding the structural collagen around vessels. An interesting fact is that tissue classification can be performed at multiple resolution levels which leads either to local features on high resolution images or region discrimination on low resolution scans. In addition to region masking as a pre-processing step for tissue classification, the region



**Fig. 2.** Low resolution lung tissue slide (left) and the region discrimination (right). Blue marks vessel borders, green normal lung tissue and yellow fibrotic lung regions.

discrimination information itself can be used as a feature in a prediction model because it adds information about the region ratio (Section 4.1).

### 3.2 Object Segmentation

Object segmentation in the case of the pulmonary fibrosis example means cell segmentation. In general, an object is a set of coherent pixels of the same tissue class  $c$ . Normally, object segmentation requires a binary image with foreground and background which is achieved by defining a threshold that distinguishes between these two classes. Here we generalize and use one specific tissue class  $c_f$  as foreground and the union of all other classes as background:  $c_b = \cup_{i \neq f} c_i$ . This allows us to use the supervised tissue classification procedure which produces a much more precise foreground definition compared to a simple threshold setting.

Shape detection is a three-step procedure: Applying the 'right hand' rule, refining the region, and executing the watershed algorithm to segment connected regions, for example overlapping cells. The 'right hand' rule scans the image pixel by pixel (from top left to bottom right) and checks if the current pixel is a foreground pixel,  $p_{x,y} \in c_f$ . If so, a new polygon shape object is created. Then neighbouring pixels are checked in clockwise order (top-right, right, ..., top) and, if they are of the same tissue class as the start pixel, the position is added to the polygon. All pixels added to the polygon are marked as *used* so that they will not be part of another polygon. The result of this algorithm is a set of polygons which define objects of the same specific tissue class, for example cells.

One of the problems we noticed is that when only applying the 'right hand' rule, objects often either have holes inside or are too narrow. To eliminate these drawbacks, we apply a region refinement step that performs a dilate step  $n_d$  times followed by  $n_e$  erode steps. A typical setting for  $n_d$  and  $n_e$  is three.

A frequent application of object segmentation is cell segmentation. Unfortunately, cells are often not well separated, but instead overlap each other. The so-called watershed algorithm [14] is often used to segment overlapping cells [12] and is the third step in our segmentation process. Basically this algorithm computes an inverse distance map from object to background. This map has low values in the middle of objects, the valleys, and high values at the border, the mountains. Now the valleys are successively flooded with water until the water reaches a height allowing two valleys to merge. A virtual dam is build there, which defines the segmentation of objects. Overlapping convex shapes are well separated with this procedure.

For cell segmentation, we obtained the best results by concatenating these three steps. However, applying the shape refinement step, the watershed algorithm or a combination of them is optional. Section 3.3 describes how object classification can be used to improve the accuracy of object segmentation for a specific cell type.

### 3.3 Object Classification

We previously described object segmentation. The result is a list of shapes, in particular polygons. For cells, these polygons describe circles more or less. However, there are several phenotypes expressed in different shape contours and arbitrary shape textures. Once more, we use supervised machine learning methods to classify these objects.

The extracted features are based on the shape and texture of the object shapes. The combination of shape descriptors and texture descriptors makes it possible to differentiate between several cell phenotypes and between correct and incorrect cell segmentations, for example residues. The texture descriptors are the same as for the tissue classification, which are described in Section 3.1. In addition we compute six shape descriptors for a shape  $s$  as follows:

$$bbArea_s = bb(s)_{width} * bb(s)_{height}$$

$$ratio_s = \max(bb(s)_{width}, bb(s)_{height}) / \min(bb(s)_{width}, bb(s)_{height})$$

$$area_s = \frac{1}{2} \sum_{i=0}^{|s|-1} s_{i_x} s_{i+1_y} - s_{i+1_x} s_{i_y}$$

$$perimeter_s = \sum_{i=1}^{|s|} \Delta(s_i, s_{i-1})$$

$$circularity_s = area_s / perimeter_s^2$$

$$roundness_s = \frac{1}{|s|} \sum_{i=0}^{|s|} |\Delta(s_i, centroid_s) - \bar{\Delta}_{(s, centroid_s)}|$$

Where  $s_i$  is the  $i$ -th point of the polygon  $s$ ,  $bb(s)$  is the bounding box of the shape  $s$ ,  $\Delta$  is the euclidean distance function, and  $\bar{\Delta}_{(s, centroid_s)}$  is the mean distance from each point in  $s$  to the centroid of  $s$ .

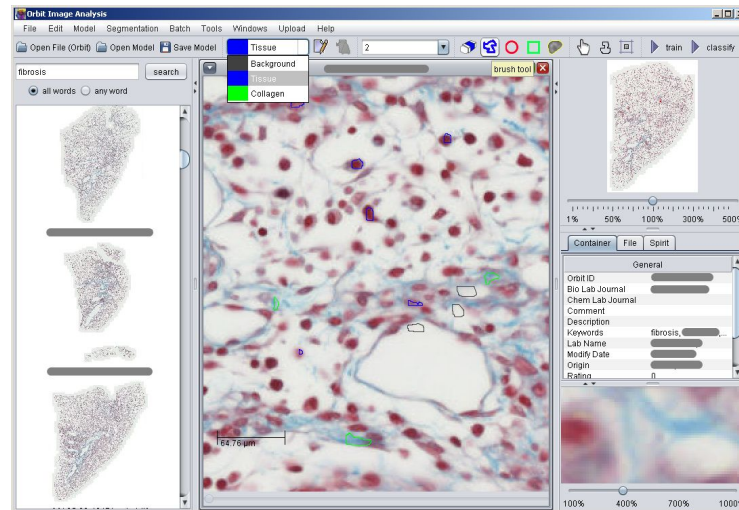
After object segmentation, the user has the possibility of defining a set of object classes. A training set is built by visually selecting a class label and clicking on a set of objects, e.g. cells, according to the selected class. In our tests, it was sufficient to mark about ten objects per class. A SVM gets the object features as input and is used to build a model based on the train set. Afterwards all other objects can be classified using this model. For cell classification the class set normally defines the cell phenotypes. However, we noticed that it is a good idea to add an additional 'residues' class. Some segmented objects that are incorrect, e.g. dust particles or objects that look wired due to segmentation errors, can be assigned to that particular class. The system is then able to classify error objects and one can exclude them from further analysis.

The object classification model has been built particularly with regard to the classification of arbitrary cell phenotypes for high content screening scenarios. In the fibrosis quantification case it is used to improve the counting of mast cells. Giemsa staining colors the mast cells a dark violet color. Object segmentation is used to segment the mast cells and count them. However, some other tissue parts have been segmented as 'mast cells' as well because the segmentation algorithm only takes the texture into account, and not the shape of a specific object. However, the object classification applied afterwards, is able to distinguish between real mast cells and residues based on the shape descriptors. This concatenated procedure is used for the mast cell count described in Section 4.1.

### 3.4 Implementation

Orbit is an in-house developed image analysis and management system. It consists of three main components: Image acquisition, image server, and image analysis. The main focus is on the ability to handle very large images, especially images which do not fit completely into the system memory. Further design goals are standardization and centralization. The entire system is written in Java, enabling platform independence. The in-house grid allows the parallel execution of grid tasks distributed on around 300 workstations. Figure 3 shows a screenshot of the Orbit Image Analysis module.

**Image Acquisition** Small client programs called agents are responsible for reading images files from a local computer or a machine. All images are converted into a standard sRGB color space in the case of color images or grey-color images. The tile-based TIFF container is used with JPEG as the compression format. Experiments have shown that with a JPEG compression factor of 0.85 there is no drawback for image analysis compared to uncompressed raw data, but it has the advantage that the images take only about one-ninth of the disk-space. All the meta data of the image as well as additional meta data from the machines, for example the microscope magnification, is read out and stored in the central meta data database of the image server. The image data itself is transferred to the image server via Java RMI in several blocks. The idea is that specific readers for arbitrary imaging systems, such as microscopes or high content image readers,



**Fig. 3.** The Orbit Image Analysis application showing a lung tissue (middle), search bar (left), overview image (right top), meta data (right middle) and a magnifier (right bottom). In the toolbar tissue classes can be selected and marked on the image by making use of drawing tools. With the train and classify button (right top) a statistical model can be trained and applied.

which handle specific image and meta data formats, exist. The output of the agents is always the same format regarding image container, compression, color-space, tile size and alignment, and meta data. For our pulmonary fibrosis case study, the tissue slides are scanned at 20x magnification with a digital microscope that can scan eight slides at a time. About 48 GB of compressed image data were produced for the 61 slides of the case study (Section 4.2).

**Image Server** The centralized image server plays an important role in the system: It holds all the image data on an iSCSI attached file server, handles search requests, and delivers image data to the clients. Large images are delivered in several tile pieces to the clients. The transfer of the image data to the clients is implemented as a servlet to maximize the stability, even with hundreds of (grid) clients accessing image data in parallel.

**Image Analysis** The image analysis module retrieves image data from the image server and applies the supervised pattern recognition methods described in Section 3. It can be executed as a stand-alone application or in grid mode. In the latter case, several image analysis clients act as grid clients and work in parallel on the grid. Java Advanced Imaging (JAI) is used for image manipulation operations and Weka [6] is used as a machine learning library. To handle very big images, a custom image implementation is used that fetches it tile by tile from the image server and discards tile data when it is no longer used. With this



approach, all mentioned image analysis algorithms scale very well even on high resolution images.

With grid computing, the computational power can easily be distributed. However, the bottleneck is the distribution of the image data from the centralized image server to all the parallel working grid clients. To avoid too much network traffic at any given point in time, the clients never load an entire image at once. Instead, a tile (or several tiles if needed) is loaded and processed. The next tile is loaded after the partial result of the last tile is ready. This procedure distributes the network traffic over the whole analysis process which smooths the network load enormous.

The images on the image server are compressed. However, if the image server reads a specific tile out of the image, it has to decompress the tile data. The server can now either deliver the uncompressed tile data directly to the clients or it can compress the data again. Both procedures have drawbacks: The first option causes a lot of network traffic and the latter requires a lot of computational power on the server side. We tried out both methods with 300 grid clients working in parallel and noticed that we run into network problems when sending the tile data uncompressed, whereas compressing the data on the server side worked well on a state of the art eight core server. Thus, image data is always compressed before sending it over the network.

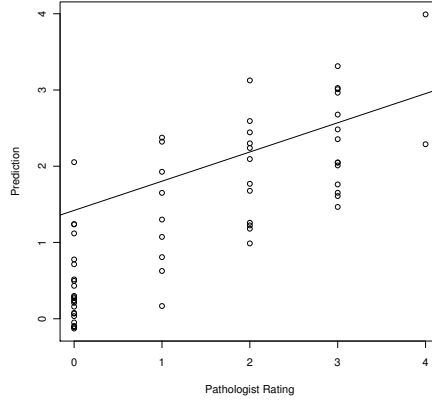
## 4 Results

For the pulmonary fibrosis case study, 61 lung tissue slides are stained with Trichrome staining coloring the collagen blue, and a further 61 slides from the same animals are stained with Giemsa staining to color the mast cells. The lung cuts stem from male Sprague Dawley rats of either a bleomycin treated group or a control group. The bleomycin is used to induce fibrosis and is dispensed from three days up to eight weeks. All slides are scanned with a digital microscope at 20x magnification. We make use of the pathology fibrosis rating to build a prediction model which is based on our predictions of region discrimination, tissue classification, and the cell count of a specific type of cells. This allows the prediction of a fibrosis degree score which we use to quantify the differences between the bleomycin group compared to the vehicle group over several weeks.

The analysis of one tissue slide takes about five hours on a standard 3GHz dual-core PC. This long runtime is due to the very large image dimensions and the fact that every single pixel is classified. To analyze all 61 tissue slides, the computation time is around 300 hours. By distributing the tasks in the grid some overhead is produced, for example, the image data has to be transferred over the network to each client. We experimented with different tile-per-job numbers and figured out that at 35 jobs per image, which is 2135 jobs in total, all images are analyzed within about three hours on the grid with around 300 clients. This means that an analysis on a single grid node takes about three times as long as the analysis on a standalone machine, but, the overall rate is increased by a factor of 100 which makes this type of application feasible.

#### 4.1 Pathologist Rating Model

A pathologist rated all slides manually with a discrete fibrosis stage from zero (no fibrosis) to four (heavy fibrosis). These values are considered the real fibrosis



**Fig. 4.** Pathologist rating vs. prediction. Correlation coefficient = 0.81.

scores. We use statistical models to predict the region discrimination, the tissue classification, and the number of mast cells as described in Section 3. Cell classification is used to improve the mast cell count by classifying each segmented mast cell as a real mast cell or a residue. With this method, some wrong segmented tissue with the same staining color as the mast cells could be detected as false positive and excluded from the mast cell count. All measured dimensions add to the degree of fibrosis, but it is unclear how much each component contributes. Therefore a linear model is built that makes use of all predictions as input, and outputs the pathologist scores. This function is called the combined model. Figure 4 shows the predictions of the combined model versus the real pathologist scores.

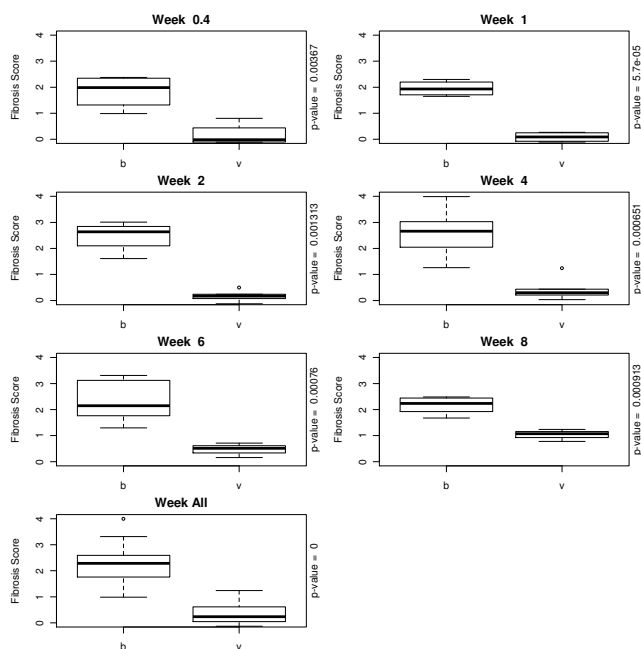
This combined model is evaluated with a 10-fold cross-validation. The evaluation shows a correlation coefficient of 0.79 and a mean error of 0.60. This is a good result due to the fact that the pathologist values are discrete values and the prediction model computes continuous values. Therefore, a mean error of 0.5 would be optimal. The correlation coefficient of the complete data is 0.81, which is slightly better than the cross-validation coefficient. Table 1 shows the evaluation for the single and combined models. It clearly shows that the combined model outperforms the single models and therefore all dimensions are needed to achieve a good fibrosis prediction model.

**Table 1.** 10-fold cross-validation results for predicting fibrosis scores.


Model	Correlation Coefficient	Mean Absolute Error
Tissue Classification	0.35	1.08
Region Discrimination	0.67	0.76
Mast Cell Count	0.37	1.05
<b>Combined Model</b>	<b>0.79</b>	<b>0.60</b>

## 4.2 Time Course Study

The combined model predicts fibrosis stages from zero to four which reflects the pathologist scores (Section 4.1). With these predictions, the treated versus vehicle tissues can be compared over time. Figure 5 shows the box plots for

**Fig. 5.** Treatment vs. vehicle for different treatment weeks.

each week and the according p-values from a one-sided Welch two sample t-test. It clearly shows a significant increase ( $p < 0.01$ ) of fibrosis in the treated compared to the vehicle group. The fibrosis evolves until week four and then decreases. For the vehicle group, the fibrosis increases with each week which can be explained with a correlation between age/weight and fibrosis. Figure 6 shows the mean values per week for each group over time which emphasizes the increase in fibrosis until week four which is followed by a decrease.



**Fig. 6.** Treatment vs. Vehicle mean values over time. The fibrosis evolves until week four and then decreases.

## 5 Conclusion and Future Work

In this paper, a system for automated image analysis using supervised machine learning methods has been presented. The system is in charge of image acquisition, image data distribution with a centralized image server, and image analysis which is performed on the in-house grid. One property of all image analysis algorithms is that they can be applied to the specific tiles of an image rather than the whole image. This allows the parallel execution of image analysis algorithms on grid clients that load image data from the centralized image server.

The main focus was on the quantification of idiopathic pulmonary fibrosis in lung cuts. We introduced the analysis modules tissue classification, region discrimination, object segmentation and object classification. Scored fibrosis stages from a pathologist have been taken as target values. We figured out that the combination of tissue classification, region discrimination, and mast cell count is a good predictor of bleomycin-induced pulmonary fibrosis. It is interesting to see that none of the single measurements is sufficient to predict a good score, but only the combination of all dimensions provides good results and outperforms the single models. The coherence of mast cell count and pulmonary fibrosis supports the outcome of Hunt et al. [8] which describes an increased number of mast cells in IPF.

The outcome of the trained model has been used to analyze a time course study where the bleomycin disease model was evaluated over eight weeks. The application of bleomycin showed an increase of fibrosis up to week four and a

decrease afterwards. The increase relative to the vehicle group is statistically significant for all weeks ( $p < 0.01$ ).

Cell classification has been applied to distinguish real mast cells from residues, which improved the mast cell count. However, this method can also be used for the classification of cell phenotypes in high content screening. This is another application of Orbit and part of our future work.

## 6 Acknowledgements

We thank Thomas Sander, Christoph Binkert and Oliver Peter for their trust and support. Their confidence made it possible to create a self-developed image analysis system. Furthermore, we thank Christian Rufener for inventing the in-house grid infrastructure and managing the servers, Tobias Fink for the database administration, and Modest von Korff for all his scientific input. We thank Isabelle Giraud for testing, educating the users and writing the documentation. We thank H el ene Pierre for proofreading and correcting this paper.

## References

1. Alworth, S., Watanabe, H., Lee, J.: Teachable, high-content analytics for live-cell, phase contrast movies. *Journal of Biomolecular Screening* (2010)
2. Bauer, Y., Hess, P., Qiu, C., Klenk, A., Renault, B., Wanner, D., Studer, R., Killer, N., Stalder, A., Stritt, M., Strasser, D., Farine, H., Kauser, K., Clozel, M., Fischli, W., Nayler, O.: Identification of cathepsin 1 as a potential sex-specific biomarker for renal damage. *Hypertension (In Press)* (April 2011)
3. Caicedo, J., Izquierdo, E.: Combining low-level features for improved classification and retrieval of histology images. *5th International Conference on Mass Data Analysis of Images and Signals in Medicine* (2010)
4. Carpenter, A., Jones, T., Lamprecht, M., Clarke, C., Kang, I., Friman, O., Guertin, D., Chang, J., Lindquist, R., Moffat, J., Golland, P., Sabatini, D.: CellProfiler: image analysis software for identifying and quantifying cell phenotypes. *Genome Biology* 7(10), R100+ (October 2006)
5. Datta, R., Joshi, D., Li, J., Wang, J.Z.: Image retrieval: Ideas, influences, and trends of the new age. *ACM Comput. Surv.* 40(2), 1–60 (April 2008)
6. Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I.H.: The weka data mining software: an update. *SIGKDD Explor. Newsl.* 11(1), 10–18 (2009)
7. Haralick, R.M., Shanmugam, K., Dinstein, I.: Textural Features for Image Classification. *IEEE Transactions on Systems, Man, and Cybernetics* 3(6), 610–621 (November 1973)
8. Hunt, L.W., Colby, T.V., Weiler, D.A., Sur, S., Butterfield, J.H.: Immunofluorescent staining for mast cells in idiopathic pulmonary fibrosis: quantification and evidence for extracellular release of mast cell tryptase. *Mayo Clin Proc.* 67(10), 941–948 (October 1992)
9. Jones, T.R., Kang, I.H., Wheeler, D.B., Lindquist, R.A., Papallo, A., Sabatini, D.M., Golland, P., Carpenter, A.E.: Cellprofiler analyst: data exploration and analysis software for complex image-based screens. *BMC Bioinformatics* 9 (2008)

10. Kapelner, A., Lee, P.P., Holmes, S.: An Interactive Statistical Image Segmentation and Visualization System. In: MEDIVIS '07: Proceedings of the International Conference on Medical Information Visualisation - BioMedical Visualisation. pp. 81–86. IEEE Computer Society, Washington, DC, USA (2007)
11. von Korff, M., Rufener, C., Stritt, M., Freyss, J., Bär, R., Sander, T.: Integration of distributed computing into the drug discovery process. *Expert Opinion on Drug Discovery* (In Press) (February 2011)
12. Malpica, N., de Solórzano, C.O., Vaquero, J.J., Santos, A., Vallcorba, I., García-Sagredo, J.M., Pozo, F.D.: Applying watershed algorithms to the segmentation of clustered nuclei. *Cytometry* 28(4), 289–297 (1997)
13. Rämö, P., Sacher, R., Snijder, B., Begemann, B., Pelkmans, L.: CellClassifier: supervised learning of cellular phenotypes. *Bioinformatics* 25(22), 3028–3030 (November 2009)
14. Roerdink, J.B.T.M., Meijster, A.: The watershed transform: definitions, algorithms and parallelization strategies. *Fundamenta Informaticae* 41, 187–228 (2001)
15. Shariff, A., Kangas, J., Coelho, L.P., Quinn, S., Murphy, R.F.: Automated image analysis for high-content screening and analysis. *Journal of Biomolecular Screening* 15(7), 726–34 (2010)
16. Snijder, B., Sacher, R., Ramo, P., Damm, E.M., Liberali, P., Pelkmans, L.: Population context determines cell-to-cell variability in endocytosis and virus infection. *Nature* 461(7263), 520–523 (August 2009)
17. Susomboon, R., Raicu, D.S., Furst, J.: Pixel-based texture classification of tissues in computed tomography. In: in CTI Research Symposium (2006)