A random forest-based method for selection of regions of interest in hyperspectral images of ex vivo human skin

Asgeir Bjorgan, Lise L. Randeberg

Asgeir Bjorgan, Lise L. Randeberg, "A random forest-based method for selection of regions of interest in hyperspectral images of ex vivo human skin," Proc. SPIE 10889, High-Speed Biomedical Imaging and Spectroscopy IV, 108891K (4 March 2019); doi: 10.1117/12.2506620

SPIE.

Event: SPIE BiOS, 2019, San Francisco, California, United States

A random forest-based method for selection of regions of interest in hyperspectral images of ex vivo human skin

Asgeir Bjorgan and Lise L. Randeberg

Department of Electronic Systems, NTNU Norwegian University of Science and Technology, Trondheim, Norway

ABSTRACT

Hyperspectral imaging is a useful tool for characterization of human tissue. However, the vast amount of data created makes it challenging and tedious to manually select spatial regions of interest for further processing. In this study, a random forest-based method was evaluated on basis of its ability to segment human skin regions from the background. The method was compared to the performance of two alternative methods, spectral angle mapper (SAM) and a K-means clustering-based method. The methods were tested on hyperspectral images of ex vivo and in vivo human skin in the wavelength range 400-1000 nm. The random forest approach was found to be robust and perform well regardless of image type. The method is simple to train, and requires minimal parameter tuning for good skin segmentation results.

Keywords: image segmentation, machine learning, tissue imaging, binary classification, spectral angle mapper, K-means clustering

1. INTRODUCTION

Hyperspectral imaging combines high spatial and spectral resolution in one modality, giving images with high spectral resolution in every pixel. This makes it a useful tool for characterization of human tissue.[1–5] Characterization can involve investigation of larger datasets that are acquired from many individual samples, or from the same sample over time. The large amount of data makes it challenging to identify systematic variations across images without additional analytical tools. Background clutter and non-relevant spectral signatures can influence the results of statistical algorithms. Segmentation is in general needed for proper visualization of the results. Visual selection of regions of interest is subjective, and automatic selection is therefore one important cornerstone of any processing chain.

Examples of hyperspectral data include reflectance images,[3,5] transmittance images[6,7] and autofluorescence images.[8–10] Each of these imaging techniques yield different types of spectral data, and a method for selection of regions of interest should be generic and robust for all kinds of images. Hyperspectral images have large dimensionality, with a large number of spectral bands that are highly correlated. Due to the correlation, using the spectral bands as features in classification algorithms can be challenging, typically requiring dimensionality reduction or feature engineering. It is desired to avoid individual tailoring and ad-hoc adjustments, and the required parameter tuning should be minimal.

Random forest classification is a robust technique which is able to obtain high prediction accuracy with little to no parameter tuning,[11,12] and which can be applied directly to data with high dimensionality and a relatively low number of samples.[13,14] This makes it attractive as a general technique which can be adapted for various types of images. An example of a simple segmentation technique would be the application of some kind of thresholding to the image bands. A decision tree[12] takes this further by considering a chain of thresholding operations that all are dependent on the results of the previous. Tree-based methods are promising due to their ability to model highly nonlinear relations, which could be adaptable for selection of regions of interest in complicated images. The random forest technique improves the test accuracy of decision tree-based methods by averaging the results from multiple decorrelated trees. The nature of the method makes it appropriate for a highly correlated predictor space such as hyperspectral data.[13,14]

Further author information: (Send correspondence to A.B.)
A.B.: E-mail: asgeir.bjorgan@ntnu.no

High-Speed Biomedical Imaging and Spectroscopy IV, edited by Kevin K. Tsia, Keisuke Goda, Proc. of SPIE Vol. 10889, 108891K · © 2019 SPIE · CCC code: 1605-7422/19/$18 · doi: 10.1117/12.2506620

The random forest technique can be used for both regression and classification tasks, and has earlier been used in hyperspectral imaging for classification in remote sensing[13,15] One study was also found on the application of detecting human skin in hyperspectral images.[16]

The primary aim of this study was to segment ex vivo human skin from the background, but the technique was also tested on images acquired from living humans. The technique was compared to SAM (spectral angle mapper) and a simplified K-means classifier. SAM classifies a spectrum based on the angle between the spectrum and a signature of interest. K-means classification clusters the data into K clusters based on centroid distances. The clusters can be used for classification after having training data labels assigned to them.

The techniques considered in this paper are all supervised classification methods, where a model $f(\boldsymbol{x})$ is built to predict a class label g from a spectrum $\boldsymbol{x}$. K-means clustering is an example of an unsupervised learning technique, where clusters, structures or relationships are learned from the data without having an associated response g. Associating such clusters with a response g makes the method supervised.

The theory behind the used classification methods is presented in section 2. The datasets used to evaluate the methods are presented in section 3, and the methods evaluation setup is presented in section 4. The results of the evaluation on the various datasets are presented in section 5.

2. METHODS

A general overview over the classification problem is first given, before the random forest, SAM and simplified K-means classifier methods are outlined.

2.1 Classification and target detection

The problem of segmenting human skin from background pixels is a binary classification problem, where every pixel $\boldsymbol{x}$ is assumed to belong to a class $g \in \{\text{background}, \text{skin}\}$. A pixel with K wavelengths can be considered a vector of predictors $\boldsymbol{x} = [x_1, \cdots, x_K]^T$ in a K-dimensional space. Classification effectively means that the vectors are labeled according their position in this space.[12] A classification method can be seen as a method to find region boundaries or decision surfaces for appropriate labeling,[12] or to model class probabilities and assign the class with the largest likelihood.[12]

The problem can also be considered to be similar to the target detection problem in remote sensing. Here, the task is to detect, enhance or calculate probabilities for a pixel to belong to a target of interest.[17] This problem corresponds to a binary classification problem if a hard decision is made.

The classification error can be defined as[12]

$$\text{Error} = \frac{1}{N} \sum_{\text{pixels}} (\hat{g}(i) \neq g(i)), \tag{1}$$

or the prediction accuracy as

$$\text{Accuracy} = 1 - \text{Error}. \tag{2}$$

The latter will be used to assess the performance of the algorithms investigated in this study. A prediction accuracy of 1 means that all pixels were assigned the correct class, while prediction accuracy of 0 means that none of them were. Prediction accuracy of 0.5 means that the classifier has the same performance as selecting the classes randomly.

Modified measures that weight background and target classification and misclassification differently exists, but are not considered in this study.

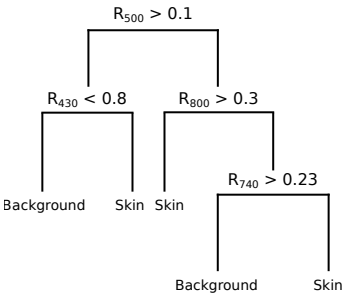

Figure 1. Example of decision tree for skin classification.

2.2 Random forest

A decision tree partitions the predictor space into rectangles, and assigns a class to each rectangle according to training data.[12] An example of a decision tree classifier for reflectance spectra is shown in Fig. 1. A class is assigned to a spectrum according to splits along reflectance values at specific wavelengths. This is effectively a series of thresholding operations on the band values of the image.

Small changes in the training data can lead to large changes in the structure of the tree due to a different series of splits being considered.[12] Several trees built based on bagged training data can be averaged to decrease this variance.[12] Such trees have identical probability distributions, and the expectation over the bagged estimate is the same as the expectation over a single tree.[12] The expected test error is therefore reduced only through reduction of the estimator variance, since the bias is constant.[12] Predictors that are strongly correlated with the response would typically always be selected during the first splits of the trees, leading to correlation between the trees.[18] This limits the variance reduction, and the potential test accuracy of the approach.[12]

The random forest technique[11] removes the correlation by always choosing m random predictors out of the total K predictors every time a new split in the tree is considered.[12] This decorrelates the trees, and improves the variance reduction and final prediction accuracy.[12]

In this study, 50 trees, also referred to as estimators, were used, unless otherwise specified. The parameter m was set to the square root of the total number of wavelengths available, following the typical recommendations for random forest classification.[12] The latter can be tuned with respect to test errors or cross-validation errors. This was not considered in this study in order to evaluate the method using only default parameters. The data were sum-normalized before application of the algorithm.

The final technique used for actual processing includes post-processing using binary hole filling, median filtering and selection of the largest contigous region. This is done in order to rectify for small misclassified regions. Post-processing was not applied to the results presented in this paper in order to properly evaluate misclassification in the various techniques. The Python method `sklearn.ensemble.RandomForestClassifier` was used as the random forest implementation. In practice, a random forest was trained and saved to file, and then applied to images in bulk in order to obtain masks for further application and processing.

2.3 SAM and simplified K-means clustering classification

The random forest algorithm was compared against two simpler methods, SAM and a K-means based method.

SAM assumes a hyperspectral pixel to be vector in a p-dimensional space. Considering two hyperspectral spectra $\boldsymbol{x}_1$ and $\boldsymbol{x}_2$ $(p \times 1)$, spectral angle can be defined as[19]

$$\theta = \arccos\left(\frac{\boldsymbol{x}_1^T \cdot \boldsymbol{x}_2}{\|\boldsymbol{x}_1\|\|\boldsymbol{x}_2\|}\right). \tag{3}$$

Taking θ to be a measure of spectral similarity, all spectra in an image can be compared against a spectrum $\boldsymbol{m}$ representing e.g. human skin. Thresholding θ yields a binary classification. The angle θ effectively represents a correlation measure between the two spectra that should be theoretically insensitive to illumination conditions. In this study, $\boldsymbol{m}$ was found from the training data by averaging over the pixels representing skin, and the

threshold was chosen by calculating the threshold yielding optimal separation between skin and background in the training data.

K-means classification is an unsupervised classification technique which assumes that K classes are present in the data, and finds K centroids within the predictor space.[12] Finding K such clusters within the data, assigning class labels to each of them and predicting class labels on new data according to the distance to the centroids can be used as a classification technique, and is considered a prototype method.[12] Selecting the optimal number K is not well-defined and has to be done through visual inspection of a within-class variance plot. Using $K = 2$ is not necessarily correct, since the background and the tissue of interest can consist of multiple clusters. In this study, a simplified variation of this technique was therefore used in order to get a general impression of the behavior of a similar method. The skin and background samples were averaged to yield spectra $\boldsymbol{m}_{\text{skin}}$ and $\boldsymbol{m}_{\text{background}}$. Class labels were then assigned according to the least Euclidean distance to each class centroid. This is not strictly the K-means classification technique proper, but vaguely related.

SAM and the simplified K-means classifier used in this paper are related, and differ only by the distance metric.

3. DATASETS

A number of datasets were used to train and test the methods. All data were acquired using a push-broom Hyspex VNIR-1600 camera (Norsk Elektro Optikk, Lillestrom, Norway). The images were acquired in the wavelength range 400-1000 nm, with a spectral resolution of 3.7 nm. The datasets are listed in table 1.

Table 1. List of datasets used in this study. Acquisition setup is described in the text.

Dataset	Subject	Training image	Test image
Arm images	Ventral side of the arm of a healthy, female volunteer (Caucasian, 39 years old)	Baseline	Image acquired after 5 minutes of pressure cuff occlusion
Skin model I	In vitro skin model with wound	Image acquired at day 1	Image acquired at day 22
Skin model II	In vitro skin model without wound	Image acquired at day 1	Image acquired at day 22

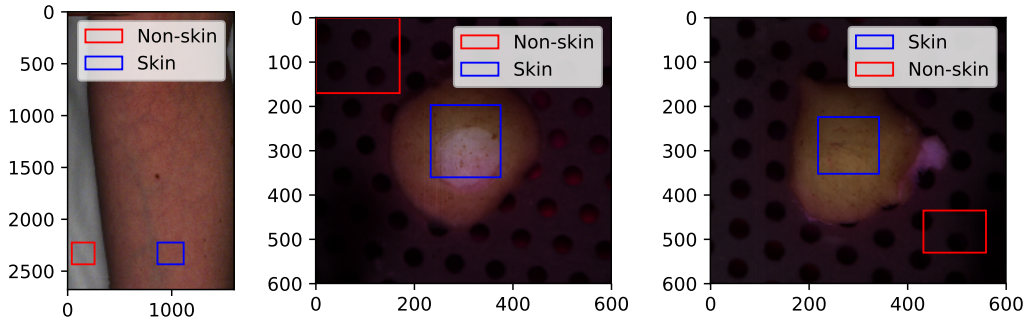

Figure 2. Training images labeled with the pixel subsets used for training: Arm, skin model I and skin model II.

The images were acquired with two linear light sources (Model 2900 Tungsten Halogen, Illumination Technologies, New York). Polarizers were mounted on the camera lens and the light sources (VLR-100 NIR, 450–1100 nm, Meadowlark Optics, Frederick, Colorado) in order to avoid specular reflection. The images were converted into reflectance and corrected for uneven illumination across the field of view using a Spectralon reflectance target (SRT-50-050 Reflectance Target, 12.7×12.7 cm, ACAL Bfi Nordic AB, Uppsala). The arm reflectance image has earlier been used in Bjorgan et al.[20] and Denstedt et al.[21] An earlier iteration of the in vitro wound model experiment has been presented in Randeberg et al.,[22] where the wound size was 5 mm. The data used in the current study had a wound size of 3 mm.[23]

4. EXPERIMENTAL EVALUATION SETUP

Various investigations were done into the accuracy and behavior of the methods.

Each method was trained on each of the training images listed in table 1 by training on the regions of interest shown in Fig. 2 and applying the methods on the corresponding test images. Test accuracies were calculated by comparing the segmentation results to a manual labeling of each pixel in the image. Training regions were selected in this way in order to emulate a typical application of the methods in a real data processing situation. Here, it is not desirable to manually label every pixel in the training image, but rather select e.g. rectangular bulk regions for each class.

In the next steps, each pixel in the training image of skin model I were manually labeled. A random subset without replacement of the training points were selected for each investigation, with equal number of data points in each class. First, the dependence on the size of the training data was found by training on a variable number of training points. Second, using a fixed number of training points, the random forest technique was tested as a function of the number of trees/estimators. Finally, to investigate the nature of the data used in this study, a PCA transform was fitted on the training data and applied on a random number of training and test data points for exploration of data clustering along the first two principal axes.

For all investigations, each method was trained on a specific training image, and tested on the corresponding test image acquired over the same object at a later day or after object manipulation, as outlined in table 1.

5. RESULTS AND DISCUSSION

The random forest method represents a tuning-free classification method which could be useful for selection of regions of interest in hyperspectral images of human tissue. The method is first compared against SAM and a simplified K-means classifier on the test datasets, and then investigated in detail in terms of its dependence on the training set and the number of estimators.

5.1 Comparison of the methods across images

Test accuracies for each of the methods are shown in Fig. 3. Corresponding visualizations of the classification results are shown in Fig. 4. The results show an overall higher test accuracy for the random forest method as compared to the other two methods in this study.

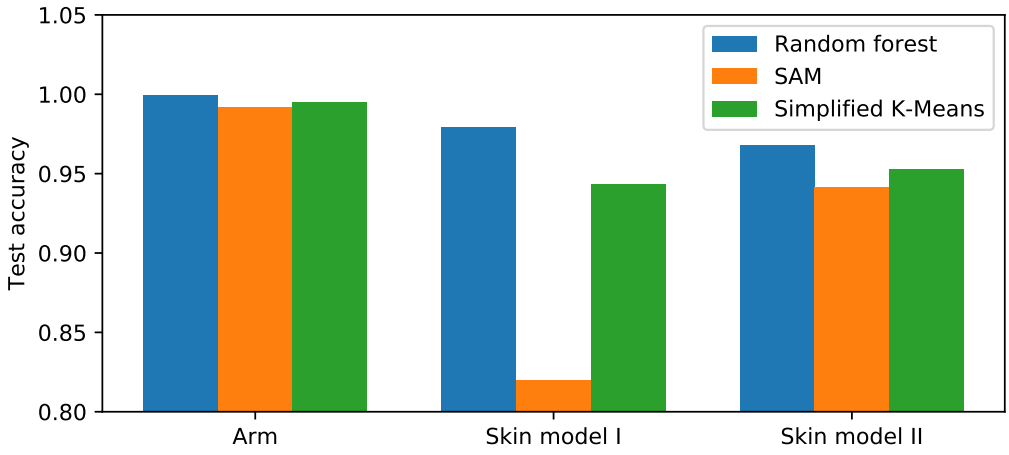

Figure 3. Test accuracy for models trained on the images shown in Fig. 2.

The accuracy is high for the arm image regardless of the classification method applied. This is due to the homogeneity of the sample. Similar behavior is seen for skin model II, which also is homogeneous over the region of interest, but has more complicated background.

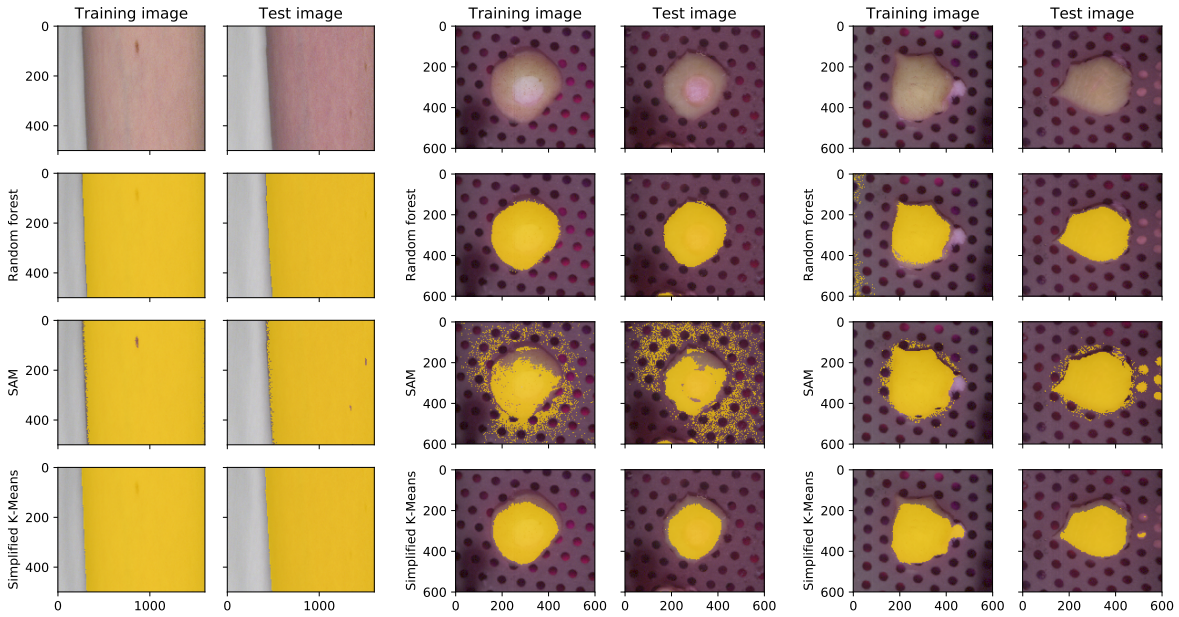

Figure 4. Segmentation results for the models trained on reflectance images: Arm (left), skin model I (center) and skin model II (right). Original images are shown on top, segmentation overlaid with original images in the following rows.

For the images of skin model I, random forest and the simplified K-means classifiers have test accuracies above 0.90, while SAM has a test accuracy below 0.85. This image has a more complicated surface than the image of skin model II due to the presence of a wound, causing the SAM spectrum to be less representative over the entire sample. The simplified K-means classifier is affected in the same way. SAM and the simplified K-means classifier are therefore susceptible to selection of training data points, while the random forest method adapts.

5.2 Application to a measurement series

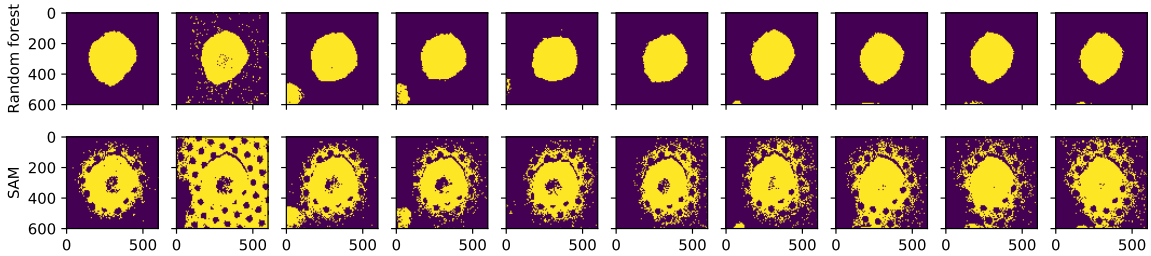

Figure 5. Application of random forest and SAM to a full reflectance image series of in vitro skin collected over several days. The methods were trained with 5000 datapoints from the first image.

The main application of this method is to segment larger series of images based on few training samples. Random forest and SAM were applied to a full measurement series in Fig. 5. SAM has higher misclassification of background, which is to be expected as the reflectance spectrum in the background is similar to the sample of interest. Background misclassification is suppressed for the random forest method. In both cases, application of proper spatial post-processing techniques would improve the results.

5.3 Dependence on the number of training points

In the initial tests, the methods were found to depend on the training samples. In order to investigate this, the test accuracy is plotted as a function of the size of the training set in Fig. 6.

The random forest method shows higher test accuracy with higher number of included points. This is to be expected for this method since random forest builds trees and generalizes based on training examples. The simplified K-means classifier and SAM, on the other hand, have no such dependence on training points. These methods work on mean spectra over the objects of interest. This leads to high test accuracy for simple situations, like the arm image, but taking the mean over more complex objects or backgrounds has low accuracy and changes with the samples included.

It would be appropriate to define multiple clusters within objects of interest and background in order to properly represent the complexity of the problem. For the simplified K-means classifier, this would be the same as using a conventional K-means classifier with some K defined. Such a method would have an increase in the test accuracy with the number of training points, but parameters like the number of clusters would have to be tuned. Random forest is already appropriate for representing such implicit multi-cluster situations without parameter tuning.

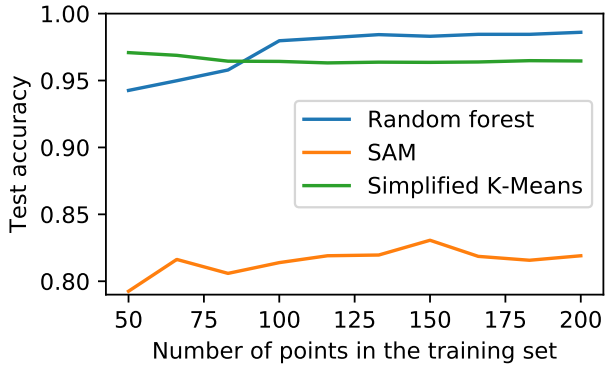

Figure 6. Test accuracy as a function of training set size for images of skin model I.

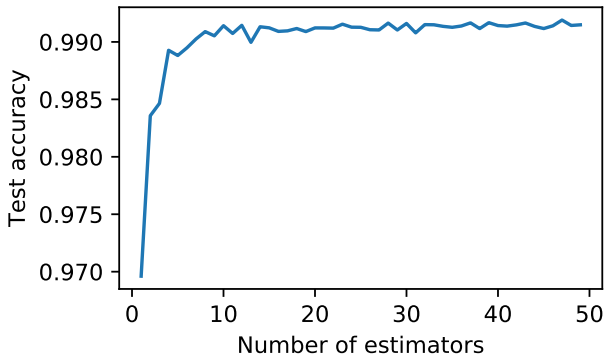

Figure 7. Test accuracy as a function of number of estimators in the random forest method for skin model I.

5.4 Feature importances and the dependence on the number of estimators

The test accuracy as a function of number of trees in the random forest method is plotted in Fig. 7. These results show that the trees converge towards a more or less constant test accuracy at around 10 trees.

Feature importances are plotted in Fig. 8, which is a measure for how many times a specific feature is used for splits in the trees. This implicitly shows which wavelenghts are important for the classification, and a method re-trained on these wavelengths would yield approximately the same test accuracy. The noisiness of the feature importances could indicate that the methods are slightly overfitted to the training data.

5.5 Investigation of the data in a lower dimensional space

Investigating the clustering behavior of the data could explain the behavior of the compared methods. A PCA transform was thus fitted to a centered and standardized reflectance image of skin model I in order to yield a rotated coordinate system along directions of the highest variance. Scores along the first two principal components are plotted in Fig. 9 for a random subset of each of the classes.

The classes consist of three clusters in the PCA space, with a gradual boundary in-between. The background and tissue would here be separable by a linear decision boundary. This explains why the simplified K-means classifier has a test accuracy comparable to the random forest classifier for the homogeneous reflectance data, as the clusters can be represented by their combined centroids. However, the elongation of the clusters also means that K-means clustering is not optimal for this kind of problem.

The data have a more or less homogeneous background and object of interest, and the simplicity of the problem is reflected with a linear decision boundary in the PCA space. More complex examples could be expected to have non-linear decision surfaces. The random forest classifier is well suitable for modeling non-linear decision

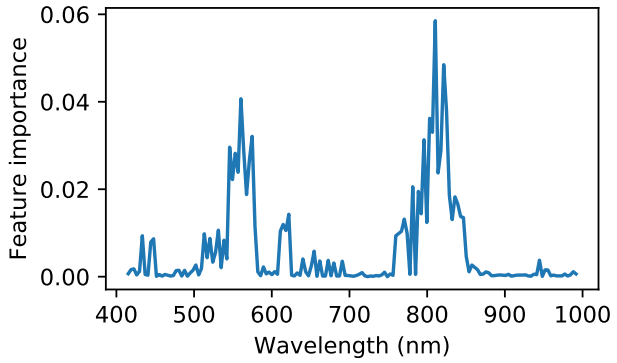

Figure 8. Feature importances in the random forest method for skin model I.

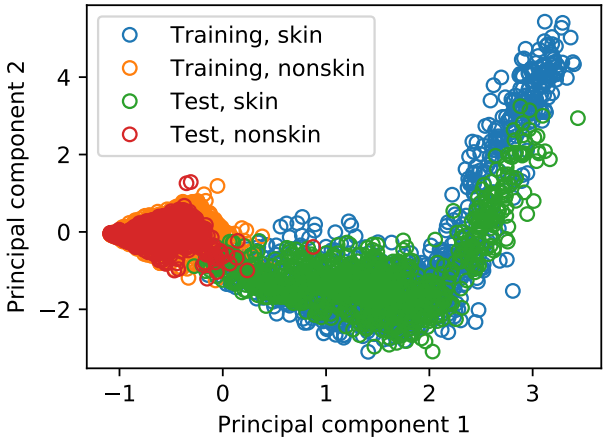

Figure 9. Data points corresponding to hyperspectral pixels along the first two components of a PCA space for reflectance images of skin model I.

surfaces, while methods like SAM or the simplified K-means classifier would have to be modified for implicit multi-class behavior.

Having more background clutter would be an example of a more complex image. This could typically be removed using pre- or post-processing, but such treatment would require manual tuning for each situation. Background clutter is not necessarily separable from the object of interest if they partially overlap. The random forest classifier automatically takes this into account. This avoids the need for manual tuning, and facilitates for automatic techniques.

5.6 General discussion

SAM and simplified K-means clustering are simple techniques that can be motivated from geometrical considerations of hyperspectral images, and which work well for simple scenes like homogeneous reflectance from a human arm. A flexible method like the random forest classifier seems to be appropriate for more complex situations wound images and complex backgrounds. The random forest technique is also appropriate for modeling of non-linear decision surfaces due to the underlying tree-based model. Having an adaptable technique with overall high classification accuracy across different types of images without parameter tuning is valuable.

However, the flexibility of the method can also make it less appropriate. The training data needs to be representative. The method can overtrain and not be able to find accurate decision surfaces. The PCA results showed that a linear decision boundary is appropriate for the examples shown in this study, meaning that restrictive methods would yield higher test accuracies. This study is not exhaustive, and linear classification techniques like LDA, QDA, logistic regression or SVM[12] could be appropriate for the examples shown. Gaussian clusters of data with equal or non-equal covariances would make methods like LDA and QDA optimal.[12] However, as seen in the reduced PCA space, each class is more likely to consist of multiple Gaussians, which makes the problem more challenging. Methods like LDA have challenges with large numbers of correlated features,[12] as is the case in hyperspectral imaging. Methods like SVM requires parameter tuning and selection of kernel functions. Other ensemble methods like boosting[12] might also be appropriate. Further work should involve testing more methods on both simpler and more complicated examples. SAM and the simplified K-means classifier are not necessarily optimal for the classification task at hand, and the random forest technique might have an unfair advantage.

Many classification techniques require preprocessing techniques like noise removal or dimensionality reduction. Dimensionality reduction would have to be fitted to the training data only, however, and could be less appropriate for the test data. Random forest is not dependent on preprocessing techniques, and is robust against the number

of features. The data were sum-normalized before application in order to make the spectra comparable and obtain some small improvements to the prediction accuracy, but no other preprocessing was applied.

Initial tests on fluorescence data (not shown here) showed that the random forest classifier had higher test accuracy than SAM and the simplified K-means classifier. The low photon count and highly correlated spectra of these data leads to a challenging classification problem, for which the random forest classifier was found to be well suitable. The classifier converged after 30-40 trees, whereas convergence was shown for 10 trees for the reflectance data used in this paper. The convergence rate thus depends on the complexity of the imaged samples. The required number of trees is expected to vary for reflectance data, and would depend on the required functional behavior of the decision surface and the behavior of the data points within the hyperspectral vector space.

Training and test images were chosen from the same data set. For both datasets, the test image was acquired over the same object as in the training data, but at a different time, and with changed spectral responses. It can be argued that these do not represent truly independent data since they are collected from the same object. They therefore do not properly measure the generalization ability of the method. However, selecting truly independent data would not necessarily yield a representative test accuracy, since the spectral responses could be too different across images. Statistical models generally do not necessarily extrapolate well outside the range of the training data.

The likely application of the method is to apply it within a specific measurement series over the same or similar objects, where it is convenient to manually label a small region in order to segment the rest of the images. Creating a truly generalizable method which can be trained once to be applied on any measurement series is outside of the scope of this study and the requirements of the method.

6. CONCLUSION

The random forest classifier has been found to be a robust technique for selection of regions of interest, and adapts well to different types of spectral data. There is no need for parameter tuning or consideration of preprocessing or modeling of implicit multi-cluster behavior in the data.

ACKNOWLEDGMENTS

Thanks to Lukasz Paluchowski for help in acquisition of the arm reflectance image. Thanks to Ingvild Haneberg and Matija Milanic for acquisition of in vitro skin model data. Thanks to Brita Pukstad for collaboration in the in vitro wound model experiment.

REFERENCES

[1] Lu, G. and Fei, B., "Medical hyperspectral imaging: a review," *J. Biomed. Opt.* **19**(1) (2014).

[2] Paluchowski, L. A., Nordgaard, H. B., Bjorgan, A., Hov, H., Berget, S. M., and Randeberg, L. L., "Can spectral–spatial image segmentation be used to discriminate experimental burn wounds?," *J. Biomed. Opt.* **21**(10) (2016).

[3] Denstedt, M., Pukstad, B. S., Paluchowski, L. A., Hernandez-Palacios, J. E., and Randeberg, L. L., "Hyperspectral imaging as a diagnostic tool for chronic skin ulcers," *Proc. SPIE* **8565** (2013).

[4] Yudovsky, D., Nouvong, A., and Pilon, L., "Hyperspectral imaging in diabetic foot wound care," *J. Diabetes Sci. Technol.* **4**(5), 1099–1113 (2010).

[5] Randeberg, L. L., Larsen, E. L. P., and Svaasand, L. O., "Characterization of vascular structures and skin bruises using hyperspectral imaging, image analysis and diffusion theory," *J. Biophotonics* **3**(1-2), 53–65 (2010).

[6] Milanic, M., Paluchowski, L. A., and Randeberg, L. L., "Hyperspectral imaging for detection of arthritis: feasibility and prospects," *J. Biomed. Opt.* **20**(9) (2015).

[7] Lu, R. and Ariana, D. P., "Detection of fruit fly infestation in pickling cucumbers using a hyperspectral reflectance/transmittance imaging system," *Postharvest Biol. Technol.* **81** (2013).

[8] Luthman, A. S., Dumitru, S., Quirós-Gonzalez, I., and Bohndiek, S. E., "Hyperspectral fluorescence imaging with multi wavelength LED excitation," *Proc. SPIE* **9711** (2016).
[9] Larsen, E. L. P., Randeberg, L. L., Olstad, E., Haugen, O. A., Aksnes, A., and Svaasand, L. O., "Hyperspectral imaging of atherosclerotic plaques in vitro," *J. Biomed. Opt.* **16**(2) (2011).
[10] Hernandez-Palacios, J., Randeberg, L. L., Baarstad, I., Loke, T., and Skauli, T., "Hyperspectral low-light camera for imaging of biological samples," *Work. Hypersp. Imag.* (2010).
[11] Breiman, L., "Random forests," *Machine Learning* **45** (2001).
[12] Hastie, T., Tibshirani, R., and Friedman, J., [*The elements of statistical learning*], Springer, second ed. (2008).
[13] Belgiu, M. and Dragut, L., "Random forest in remote sensing: A review of applications and future directions," *ISPRS J. Photogramm. Remote Sens.* (2016).
[14] Chan, J. C.-W., Bechers, P., Spanhove, T., and Borre, J. V., "An evaluation of ensemble classifiers for mapping natura 2000 heathland in belgium using spaceborne angular hyperspectral (chris/proba) imagery," *Int. J. Appl. Earth Obs.* (2012).
[15] Ham, J., Chen, Y., Crawford, M. M., and Gosh, J., "Investigation of the random forest framework for classification of hyperspectral data," *IEEE Trans. Geosci. Remote Sens.* (2005).
[16] Marques, I., Grana, M., Sanchez, S. M., Alkhatib, M. Q., and Velez-Reyes, M., "Person detection in hyperspectral images via skin segmentation using an active learning approach," *Proc. SPIE* **9472** (2015).
[17] Eismann, M. T., [*Hyperspectral remote sensing*], SPIE (2012).
[18] James, G., Witten, D., Hastie, T., and Tibshirani, R., [*An introduction to statistical learning*], Springer, first ed. (2013).
[19] Keshava, N. and Mustard, J., "Spectral unmixing," *IEEE Signal Proc. Mag.* **19**, 44 –57 (jan 2002).
[20] Bjorgan, A., Milanic, M., and Randeberg, L. L., "Estimation of skin optical parameters for real-time hyperspectral imaging applications," *J. Biomed. Opt.* **19**(6) (2014).
[21] Denstedt, M., Bjorgan, A., Milanic, M., and Randeberg, L. L., "Wavelet based feature extraction and visualization in hyperspectral tissue characterization," *Biomed. Opt. Express* **5**(12), 4260–4280 (2014).
[22] Randeberg, L. L., Hegstad, J. L., Paluchowski, L. A., and Pukstad, B. S., "Hyperspectral characterization of an in vitro wound model," *Proc. SPIE* **8926** (2014).
[23] Haneberg, I. J. A., *Hyperspectral imaging of in vitro wound models from human skin*, Master's thesis (2014).