Application of smoothing splines for spectroscopic analysis in hyperspectral images

Asgeir Bjorgan, Lise L. Randeberg

Asgeir Bjorgan, Lise L. Randeberg, "Application of smoothing splines for spectroscopic analysis in hyperspectral images," Proc. SPIE 10873, Optical Biopsy XVII: Toward Real-Time Spectroscopic Imaging and Diagnosis, 108730O (4 March 2019); doi: 10.1117/12.2506618

SPIE.

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

Application of smoothing splines for spectroscopic analysis in hyperspectral images

Asgeir Bjorgan and Lise L. Randeberg

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

ABSTRACT

The spectral and spatial resolution of hyperspectral imaging is useful for investigation of tissue autofluorescence. The low-light, noisy conditions in fluorescence imaging usually necessitates noise removal for extraction of precise spectral signatures and peak shifts. However, noise removal techniques like low-pass filtering or the Maximum Noise Fraction transform might discard information or distort spectral features. In this study, smoothing splines is proposed as an alternative technique to avoid spectral distortion in analysis of hyperspectral fluorescence images in the wavelength range 400-1000 nm. Continuous tuning parameters and use of natural cubic splines makes the method advantageous for unbiased peak extraction. The method was tested on ex vivo images of atherosclerosis lesions and simulations. The method was used to estimate autofluorescence peak shifts, and found to perform well in comparison with MNF.

Keywords: fluorescence imaging, subresolution peak extraction, tissue optics, noise removal, maximum noise fraction transform

1. INTRODUCTION

Hyperspectral imaging is a valuable tool for investigation of tissue autofluorescence. It has the spectral resolution required to resolve multiple fluorophores, and can yield spatial maps over the locations of spectral signatures of interest. Applications of hyperspectral fluorescence imaging include imaging of atherosclerotic plaques,[1] food quality control[2,3] and various medical applications like characterization of cancer cells.[4] In the current study, smoothing splines is explored as a tool for removing noise and characterizing autofluorescence peaks in hyperspectral fluorescence images.

Imaging spectrometers require sufficient illumination. However, the only light available in imaged fluorescence scenes is the low-intensity autofluorescence light generated in the samples. Increasing the integration time of the camera or the power of the excitation light source is normally not feasible, as long exposure to strong radiation would bleach the fluorophores. Hyperspectral fluorescence imaging therefore leads to images with a low signal-to-noise ratio (SNR), where noise removal can be necessary as a pre-processing technique.

Automated characterization is important for processing larger datasets acquired as time series and across samples. An ultimate goal is to unmix the data with respect to the various fluorophores. This is challenging in tissue due to strong scattering and absorption at both excitation and emission wavelengths. Full unmixing requires full characterization of e.g. reflectance images acquired under a visible light source to estimate scattering and absorption properties, and the appropriate pairing of optical properties and apparent emission spectra. Manual decision making might be necessary.

Emission peak characterization in terms of peak position, peak height and peak width is an alternative technique for obtaining an overview over the data. This can be used as a part of the fluorophore characterization process, and for the initial visualization of the data. Comparability is an important requirement, where fluorophores across measurements should have the same peak behavior, and not be subject to e.g. pre-processing techniques that treat the data differently.

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

Optical Biopsy XVII: Toward Real-Time Spectroscopic Imaging and Diagnosis, edited by Robert R. Alfano, Stavros G. Demos, Angela B. Seddon, Proc. of SPIE Vol. 10873, 108730O · © 2019 SPIE · CCC code: 1605-7422/19/$18 · doi: 10.1117/12.2506618

Extracting peak positions at sub-resolution level could be useful for obtaining more accurate representations of peak shift distributions over larger regions. These are otherwise limited to counts over the two closest wavelengths due to the wavelength discretization. Sub-resolution peak shifts lead to subtle changes in the noise dynamics, which could be exploited by a suitable interpolating technique to estimate peak shifts. Thus, there are two goals in this study: To remove noise from the spectra in a comparable and objective way, and to estimate sub-resolution peak shifts, where the two problems are strongly inter-related.

MNF (Minimum noise fraction)[5] is a noise reduction technique often used for hyperspectral images, as it retains spatial and spectral resolution while yielding smooth spectra. MNF decomposes the image in terms of typically 8-10 less noisy, mixed components, with their own maxima and minima. Each pixel in the denoised hyperspectral image will be a linear combination of these. It is not guaranteed that the MNF decomposition is sensitive to small shifts in fluorescence emission peak position due to the way the images are decomposed. The MNF technique is subject to the available image statistics. Choosing the number of components has to be done based on visual inspection. MNF is not an interpolating technique, and the number of component cannot be objectively chosen by cross-validation or similar techniques, unless techniques further down the processing pipeline have properties that can be evaluated by cross-validation. In general, the subjectivity leads to images that are not necessarily comparable.

Smoothing splines[6,7] is proposed as an alternative technique for satisfying both denoising and peak detection aspects in fluorescence images. The method uses natural cubic splines to interpolate the data. Wavelength bands can be left out during fitting, to be evaluated by the fitted model at a later time. This makes the concept of independent test data well-defined. Subresolution peak positions are efficiently found using the analytical derivative. The spline coeffients are restricted according to a roughness penalty, in effect making the method a denoising method with the appropriate smoothing parameter. Smoothing splines is a linear method, and can thus be efficiently estimated on multiple spectra at the same time. Linearity also means that the leave-one-out cross-validation error is trivial to calculate without refitting the method. All spectra in the hyperspectral datacube can therefore efficiently be represented using spline coefficients, with crossvalidated tuning parameters that objectively can yield smoothed spectra that are comparable.

Smoothing splines, or variations of the technique, have been used in both hyperspectral imaging and in fluorescence spectroscopy. Marion et al.[8] uses smoothing splines as a basis for a method to recover reflectance values within major gas absorption bands in remotely sensed images. Berman[9] uses thin-plate smoothing splines to smooth the MNF transformation bands of remote sensing data. Lee et al.[10] compares various smoothing methods on autofluorescence data, including LOESS, COBS and robust smoothing splines, which uses a scale function on the mean squared error part of equation (1). Various methods are given for approximating the cross-validation error.[8–10]

The theory for the smoothing splines technique and considerations for application to hyperspectral images is outlined in section 2. The technique is then applied to images of atherosclerosis images and simulations, and compared to MNF.

2. SMOOTHING SPLINES

Given a data set with inputs $\boldsymbol{x} = [x_1, x_2, \ldots, x_K]^T$ and outputs $\boldsymbol{y} = [y_1, y_2, \ldots, y_K]^T$, the smoothing splines method estimates the function $f(x)$ which minimizes the penalized residual sum of squares (PRSS)[6]

$$\mathrm{PRSS}(f, \alpha) = \sum_{i=1}^{K} \{y_i - f(x_i)\}^2 + \alpha \int_{-\infty}^{\infty} \left\{ \frac{d^2}{dt^2} f(t) \right\}^2 dt. \tag{1}$$

The first term deals with the closeness of the function $f(x)$ to the outputs y_i, and the second term penalizes the curvature of the function.[6] Setting $\alpha = 0$ interpolates the outputs exactly (maximum curvature), while $\alpha = \infty$ yields a straight line (no curvature allowed). In practice, α controls the smoothness of the resulting function, and regularizes between under- and overfitting. An alternative, more convenient formulation is to express the PRSS as[7] $p \sum_{i=1}^{N} \{y_i - f(x_i)\}^2 + (1-p) \int_{-\infty}^{\infty} \left\{ \frac{d^2}{dt^2} f(t) \right\}^2 dt$, with $\alpha = (1-p)/p$. Here, $p = 0$ yields a straight line and $p = 1$ fits all datapoints exactly.

The parameter α is a tuning parameter which has to be estimated using independent validation data or through cross-validation. Minimization of $\mathrm{PRSS}(f, \alpha)$ would always set $\alpha = 0$, effectively overfitting the model on the data.

It can be shown that the solution $\hat{f}(x)$ is a natural cubic spline with K knots,[6] meaning that the function between the first and second datapoint and the function between datapoint $K-1$ and datapoint K are linear functions. The other line segments are cubic polynomials with continuous first and second derivatives. The linear segments reduce erratic behavior at the edges of the dataset.[6] The solution can be written as

$$\hat{f}(x) = \sum_{j=1}^{K} N_j(x)\hat{\theta}_j, \tag{2}$$

where $N_j(x)$ are the basis functions for the natural cubic splines. The coefficients $\hat{\theta}_j$ are given as[6]

$$\hat{\boldsymbol{\theta}} = [\hat{\theta}_1, \ldots, \hat{\theta}_K]^T = \underbrace{(\mathbf{N}^T\mathbf{N} + \alpha\boldsymbol{\Omega}_N)^{-1}\mathbf{N}^T}_{\stackrel{\text{def}}{=}\mathbf{H}_\alpha}\boldsymbol{y} = \mathbf{H}_\alpha\boldsymbol{y}. \tag{3}$$

The matrices $\mathbf{N}$ and $\boldsymbol{\Omega}_N$ consist of the elements $\{\mathbf{N}\}_{ij} = N_j(x_i)$ and $\{\boldsymbol{\Omega}_N\}_{jk} = \int \frac{d^2}{dt^2}N_j(t)\frac{d^2}{dt^2}N_k(t)dt$, respectively.[6] These are dependent only on x_i.

Evaluating $\hat{f}(x)$ at the inputs of the training set x_i yields

$$\hat{\boldsymbol{f}} = [\hat{f}(x_1), \hat{f}(x_2), \ldots, \hat{f}(x_K)]^T \tag{4}$$

$$= \sum_{j=1}^{K}[N_j(x_1), \ldots, N_j(x_K)]^T\hat{\theta}_j = \mathbf{N}\hat{\boldsymbol{\theta}} \tag{5}$$

$$= \underbrace{\mathbf{N}(\mathbf{N}^T\mathbf{N} + \alpha\boldsymbol{\Omega}_N)^{-1}\mathbf{N}^T}_{\stackrel{\text{def}}{=}\mathbf{S}_\alpha}\boldsymbol{y} = \mathbf{S}_\alpha\boldsymbol{y}. \tag{6}$$

Here, the matrix $\mathbf{S}_\alpha$ depends only on the training inputs x_i and α, and the fit $\hat{\boldsymbol{f}}$ is therefore linear in $\boldsymbol{y}$. The coefficients $\hat{\boldsymbol{\theta}}$ are also similarly linear in $\boldsymbol{y}$. Linearity has two main advantages which makes the method ideal for large datasets like hyperspectral images.

First, for a single α, it is computationally efficient to apply the method to multiple spectra. Given a hyperspectral image in matrix form, $\mathbf{B} = [\boldsymbol{y_1}, \ldots, \boldsymbol{y_N}]^T$ (N spectra $\times$ K wavelengths), the corresponding x_i are the same across all spectra. The matrices $\mathbf{N}$ and $\boldsymbol{\Omega}$ can therefore be calculated once for a specific α, and spline coefficients or function evaluations are obtained by a matrix multiplication:

$$\mathbf{B}_\alpha = \mathbf{S}_\alpha\mathbf{B} \tag{7}$$

$$\hat{\boldsymbol{\theta}}_\alpha = \mathbf{H}_\alpha\mathbf{B}. \tag{8}$$

An efficient BLAS (Basic Linear Algebra Subprograms)[11] implementation then ensures computational efficiency.

Second, linearity enables the calculation of the leave-one-out cross-validation error without refitting. Leave-one-out cross-validation estimates the test error of a given method by training on $K-1$ datapoints and testing on the remaining datapoint, and averaging over all combinations. This requires the model to be fitted K times for each parameter choice. For models that are linear in $\boldsymbol{y}$, as expressed in (6), the leave-one-out cross-validation error can be calculated directly from $\mathbf{S}_\alpha$:[6]

$$\mathrm{Err}_{\mathrm{CV}}(\hat{f}_\alpha) = \frac{1}{K}\sum_{i=1}^{K}\left(\frac{y_i - \hat{f}_\alpha(x_i)}{1 - S_\alpha(i,i)}\right)^2, \tag{9}$$

where $S_\alpha(i,i)$ are the diagonal elements of the matrix $\mathbf{S}_\alpha$. Thus, the smoothing splines for a given parameter choice α has to be fitted only once, and the full cross-validation error for this parameter choice is available using the above formula without refitting. For a given spectrum $\boldsymbol{y}$, the parameter α can objectively be found by minimizing (9). This yields the lowest cross-validation error, and provides an optimal choice between closeness and smoothness in the resulting spline.

2.1 Estimating optimal tuning parameters for all pixels

Choosing any α between 0 and ∞ yields a gradual transition from high-variance, noisy fits to smoother fits - and an optimal parameter in-between representing a denoised spectrum. This coincides with the parameter chosen through cross-validation.

The tuning parameter α is directly dependent on the signal to noise ratio in a given spectrum. An image might have the same noise variance in all pixels and bands, but the signal level varies. A constrained fit with high α would be required for low signal levels, where the noise dominates, while a less constrained fit and lower α is appropriate for high signal levels, where the relative noise contribution is low. An image containing various kinds of signals and signal amplitudes requires a different α for every single pixel.

Consider a typical bracketing routine for minimizing $\mathrm{Err}_{\mathrm{CV},\alpha}$ for a given spectrum, assuming a well-behaved error curve with a single, global minimum. All spectra would require estimation of the cross-validation error at a common set of lower and upper α. Multiple spectra would likely end up within the same parameter ranges, and go through a sequence of bracketing at the same α_i. All spectra share the same x_i, and a single $\mathbf{S}_\alpha$ can be calculated for a specific α_i. A divide-and-conquer approach may therefore be used.

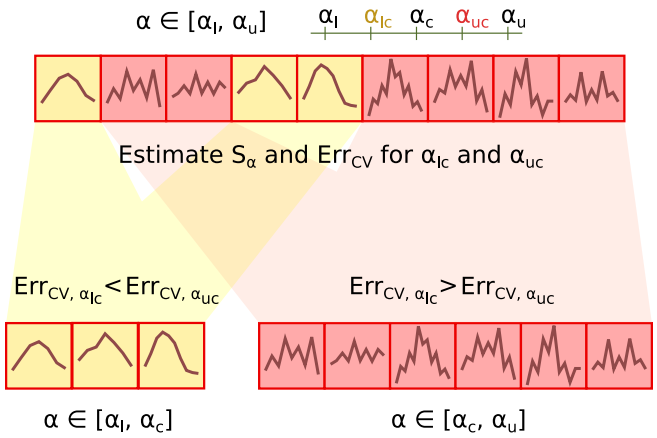

Figure 1. Split based on cross-validation error during divide-and-conquer strategy for estimating optimal α.

Assume lower and upper brackets α_l and α_u for the α providing the minimum cross-validation error for a range of pixels $\mathbf{B}$:

1. Calculate $\alpha_c = 0.5(\alpha_u + \alpha_l)$, $\alpha_{lc} = 0.5(\alpha_l + \alpha_c)$ and $\alpha_{uc} = 0.5(\alpha_u + \alpha_c)$.
2. Estimate $\mathbf{S}_\alpha$ for α_{lc} and α_{uc}.
3. Calculate the cross-validation error for all pixels for α_{uc} and α_{lc} using $\mathbf{S}_\alpha$.
4. Divide the pixel set in two subsets: Assign pixels to $[\alpha_l, \alpha_c]$ for the pixels which has the lowest cross-validation error for α_{lc}, and assign to $[\alpha_c, \alpha_u]$ for the pixels which has the lowest cross-validation error for α_{uc}.
5. Repeat the process from 1. on each of the pixel sets until convergence ($|\alpha_l - \alpha_u| < \alpha_c \cdot t$).

One iteration of the algorithm is illustrated in Fig. 1.

Thus, each pixel can have its optimal α estimated independently of each other to an arbitrary numerical precision, with full reuse of the same computations across pixels. This approach assumes that the cross-validation curve has a single minimum. This has been seen to be valid for most of the spectra encountered in this study, though there are cases where this is not valid.

Coefficients can either be kept during the cross-validation routine, or they can be obtained by reapplying the method on smoothed data using a smoothing parameter α close to 0, as illustrated in Fig. 2. The latter approach can also be applied on spectra denoised using e.g. MNF to obtain spline coefficients for peak characterization.

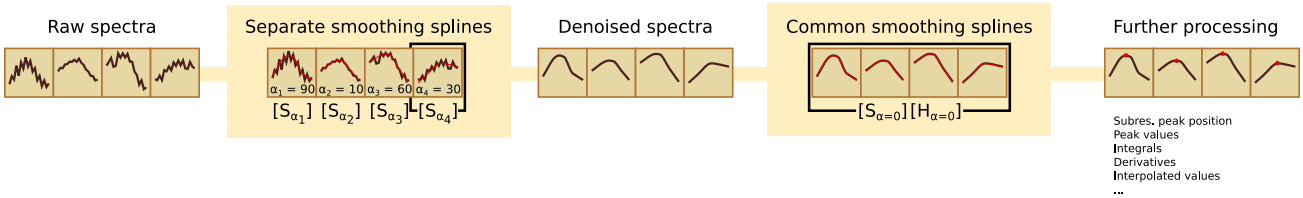

Figure 2. Flowchart for possible application of the method: Apply smoothing splines separately on each pixel and then apply a common smoothing spline matrix with low α to estimate spline coefficients for further analysis.

3. THE MINIMUM NOISE FRACTION TRANSFORM

The MNF transform $\mathbf{T}$ of an image $\mathbf{B}$ with band means $\bar{\mathbf{B}}$ can be expressed using a linear transformation matrix $\mathbf{A}$ as[5]

$$\mathbf{T} = \mathbf{A}(\mathbf{B} - \bar{\mathbf{B}}), \tag{10}$$

where the subraction of the band means $\bar{\mathbf{B}}$ are taken row-wise. The constituent vectors $\boldsymbol{a}$ of the transformation matrix $\mathbf{A} = [\boldsymbol{a_1}, \ldots, \boldsymbol{a_K}]^T$ are found by solving the eigenvalue problem[5]

$$\boldsymbol{\Sigma}_{\mathcal{N}}\boldsymbol{a} = \lambda\boldsymbol{\Sigma}\boldsymbol{a}, \tag{11}$$

where $\boldsymbol{\Sigma}_{\mathcal{N}}$ and $\boldsymbol{\Sigma}$ are the noise and image covariance matrices, respectively.

The bands in $\mathbf{T}$ are sorted in terms of decreasing SNR by sorting the eigenvalues.[5] Most of the signal is compressed in the first bands, slowly degrading into pure noise with increasing band index. Discarding the last $K - r$ bands can be used as a compression technique, and processing algorithms can be made to work on the r transformed bands instead of the full image. Using the first r components of $\mathbf{A}^{-1}$ on these r components would yield noise-reduced spectra within the original image space. If only the denoising properties of the technique are desired, this process can be expressed as

$$\mathbf{B}^* = \bar{\mathbf{B}} + \mathbf{A}^{-1}\mathbf{R}\mathbf{A}(\mathbf{B} - \bar{\mathbf{B}}) = \bar{\mathbf{B}} + \mathbf{D}(\mathbf{B} - \bar{\mathbf{B}}), \tag{12}$$

where $\mathbf{R}$ is the identity matrix with the last $K - r$ diagonal entries set to zero. All steps in the method can be implemented as matrix operations, which can make use of efficient linear algebra implementations.[12]

4. EXPERIMENTAL SETUP

The smoothing splines method was tested for its peak extraction capabilities on a fluorescence dataset. The noise removal capabilities were compared against MNF on the same dataset. Here, the ground truth is not known. Simulations were therefore also generated in order to evaluate the methods on denoising of spectra where the underlying models are known. For the simulated dataset, quantification was done using the mean squared error between the original, noise-free spectra and the denoised spectra.

4.1 Measurements

The fluorescence example dataset consists of measurements over atherosclerotic lesions from human aortic samples, previously published in Larsen et al.[13] and Randeberg et al.[14] The 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 samples were illuminated using a frequency tripled Nd:YAG-laser (Quanta Ray Lab-series L-190, Spectra Physics, Mountain View, California), providing UV-excitation at 355 nm. For more details on the experimental setup, see Larsen et al.[13]

4.2 Simulations

To emulate fluorescence spectra similar to the spectra encountered in the measurements, data were generated using the function

$$R(x) = \exp(mx \exp(-x/s + 1)/s), \tag{13}$$

with x ranging from 0 to 100 in discrete steps. Three baseline spectra were chosen using parameters $s = [10, 20, 60]$ and $m = [4.0, 3.8, 3.6]$, shown in Fig. 3.

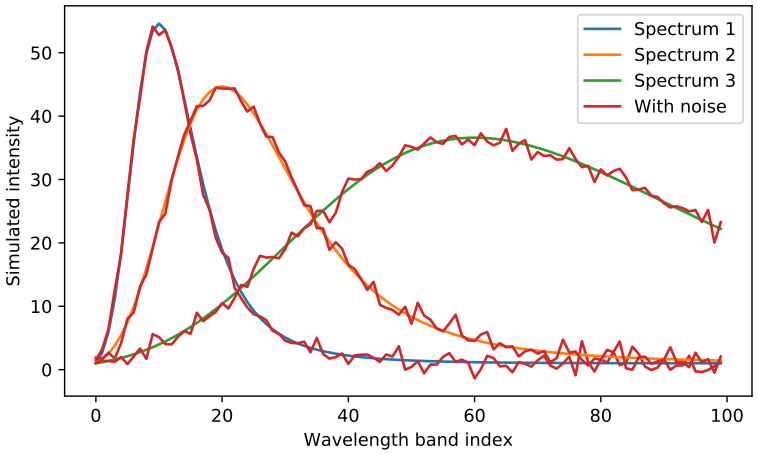

Figure 3. Three basic spectra used in the simulations.

An image with 3 rows and 500 columns was created. Each of the basic baseline spectra were assigned to a row. Each row was divided into 100 patches with 5 pixels in each patch, where parameters m and s were randomly generated from a Gaussian distribution with means corresponding to the baseline parameters and standard deviation 1, and used to yield variations of the baseline spectra. Constant variance ($\sigma = 1$) Gaussian noise was added to the spectra. The goal was to create an image with a large enough number of pixels to be able to estimate reliable image statistics, but let no specific spectrum be homogeneously overrepresented.

5. RESULTS AND DISCUSSION

The smoothing splines method is explored as a technique for noise removal and subresolution peak characterization in hyperspectral images, and compared against MNF. The method is first applied to autofluorescence measurements, and then to simulations of spectra with additive noise.

5.1 Peak position characterization in fluorescence data

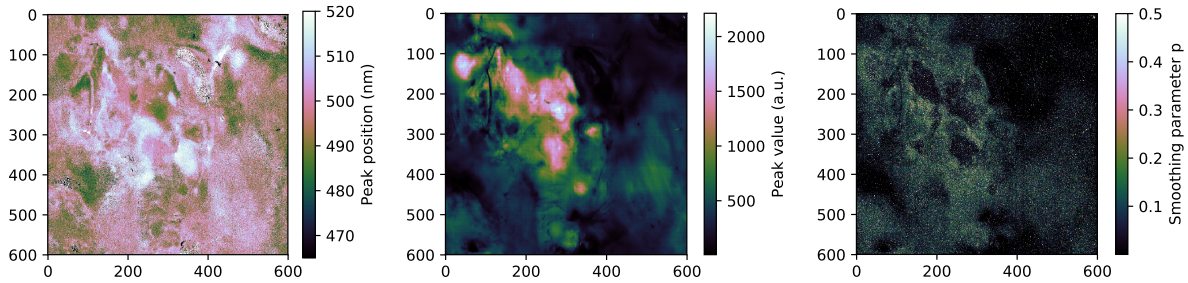

Figure 4. Peak positions (left), peak values (center) and smoothing parameters $p = \frac{1}{\alpha+1}$ as estimated using the cross-validation approach (right).

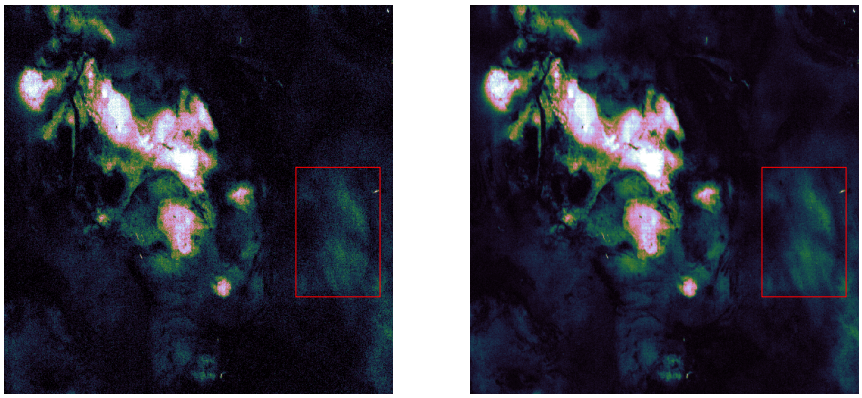

Figure 5. Comparison of the blue band before and after application of smoothing splines.

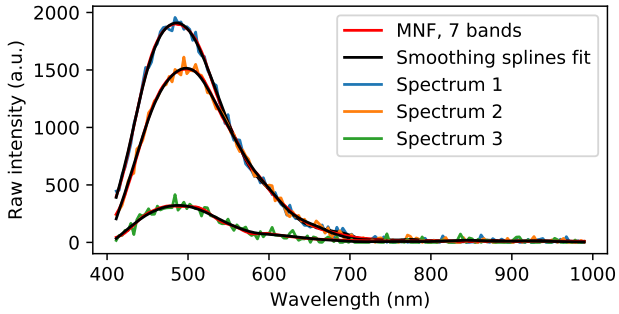

Figure 6. Raw spectra compared against smoothing splines and MNF denoised spectra.

Peak positions, peak values and the corresponding smoothing parameters for the measurements are shown in Fig. 4.

Band images before and after application of the technique for one of the noisier bands in the blue region are shown in Fig. 5. Spatial noise is reduced in these bands after application of the technique, due to the constraining effect of the splines. See for example the region marked in red, where graininess in the raw image is replaced by a smoother texture with more detail.

Raw fluorescence spectra were selected from three pixels for further comparison (positions $(x, y) = (279, 232)$, (272, 336), (598, 398) for spectrum 1, 2 and 3, respectively), and are compared against the smoothing splines fit in Fig. 6. The splines yield smooth spectra that well follow the trends in the spectrum.

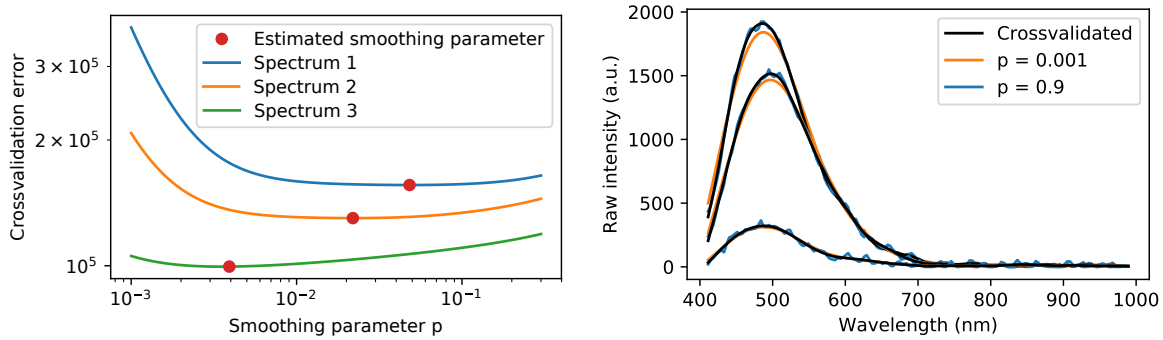

Figure 7. Cross-validation curves (left) and corresponding spectra (right) at the cross-validated smoothing parameters and parameters above and below this parameter.

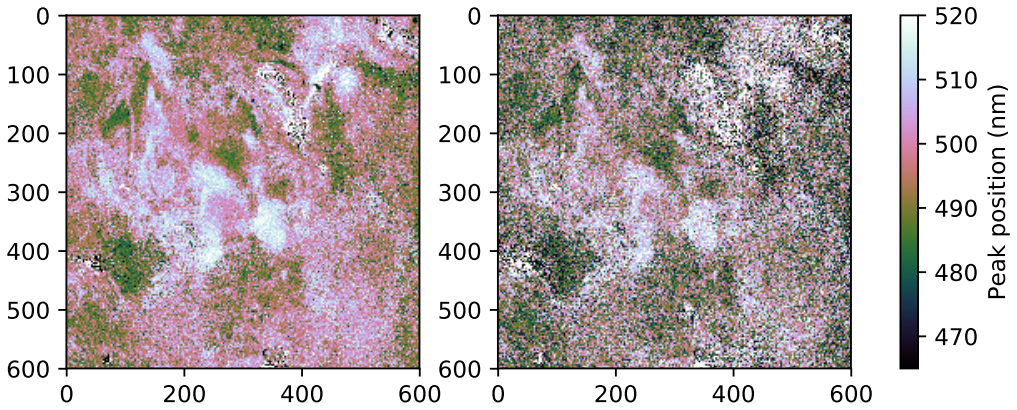

Figure 8. The rounded smoothing splines maximum estimate (left), compared to the maximum position as estimated using `max` on each raw spectrum (right).

Corresponding cross-validation error curves are shown in Fig. 7, along with examples of spectra smoothed at parameters outside the minimum. A parameter above the cross-validated parameter yields a spectrum which follows the noise. A parameter below the cross-validated parameter yields a smooth spectrum, but has bias. The smoothing parameter at the lowest cross-validation error yields a compromise between the two.

The peak position as estimated from the raw spectra by simply finding the index position for the max value throughout the noisy spectrum was compared against the rounded smoothing splines estimate in Fig. 8. Rounding the peak position as obtained from the smoothing splines method yields an index value similar to the simple max, but chooses a wavelength index closer to the sub-resolution estimate. This represents an estimate with less spatial noise. Working in the subresolution regime would especially be useful for characterization of gradual changes over time when the wavelength discretization is large. The use of cross-validation to estimate smoothness makes the results comparable across pixels and images.

5.2 Comparison between smoothing splines and MNF

The smoothing splines method was compared to MNF on the autofluorescence data and the simulations described in section 4.2.

5.2.1 Application of MNF to fluorescence data

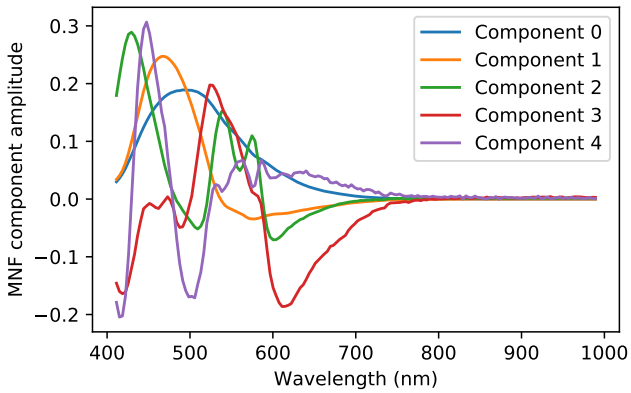

Figure 9. Components of the inverse MNF matrix.

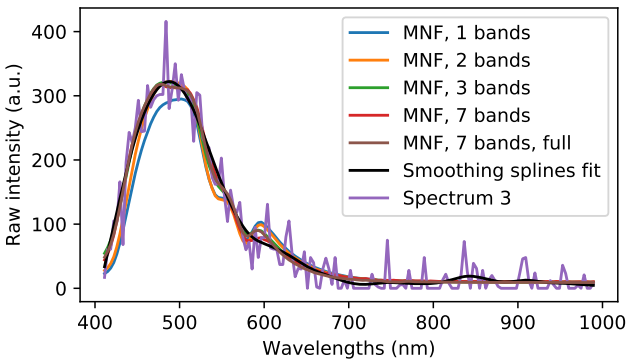

Figure 10. Inverse MNF spectra for different number of components in inverse, and statistics obtained from a larger subset of the image.

The first components in the inverse transform and the corresponding first bands in the forward transformed image are shown in Fig. 9 and 11, respectively. The inverse, denoised spectra will be a linear combination of

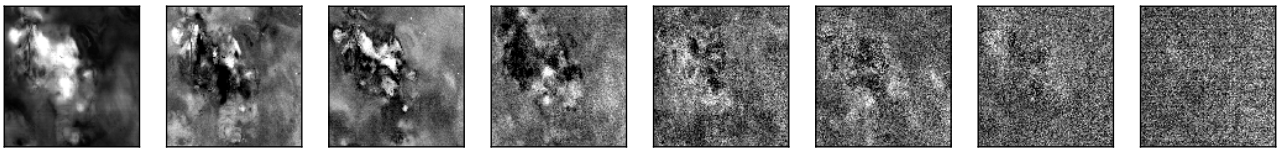

Figure 11. The first bands of the MNF transform.

the components in the inverse transform, with the pixel values in the transformed bands as the coefficients. It is seen that the components gradually contain more and more noise. Seven bands were taken to construct the inverse transform considered further on.

Example spectra are compared against corresponding splines in Fig. 6. This shows some deviations between the smoothing splines spectrum and the MNF spectrum: for spectrum 1, the denoised spectra do not overlap around 500 nm, and for spectrum 3, the MNF spectrum has oscillations not present in the smoothing splines spectrum. The true spectra are not known here, and the differences were therefore not evaluated for the measurements. Differences for simpler simulations are evaluated in section 5.2.3.

MNF has a tuning parameter similar to the smoothing parameter, the number of components to use in the inverse transform. Selecting a low number of components leads to bias in the inverse spectrum, which is demonstrated by the denoised spectra using 1 and 2 bands in the inverse in Fig. 10. Including all components in the inverse reproduces the original spectrum exactly. An intermediate value produces a smooth approximation, subject to the image statistics and the spectrum at hand.

Tuning number of components using cross-validation is not well-defined, as MNF would have to be fitted with one of the bands missing from the data. MNF does not interpolate missing bands, and calculating the prediction error on the left-out band is not well-defined. The number of components is therefore usually selected through visual inspection of the plot over the SNR as a function of the number of components, or by visual inspection of the noisiness of the transformed MNF bands or inverse spectra. In addition to subjectivity, this means that the parameter is chosen based on the bulk properties of the image. Individual spectra could require less or more components due to spatial variations in the SNR, but adjusment on pixel level is not possible due to the method not being interpolating.

MNF is expected to yield a less optimal estimate if the shape of the spectrum at hand is less represented in the statistics matrices. Essentially, pure components would be used to represent less represented spectra. Different components would be mixed where there is no mixture of such fluorescence components. Such mixture behaviour would be more prominent with less components, while more components leads to more noise, essentially being a bias-variance consideration.

5.2.2 Comparison of denoising matrices

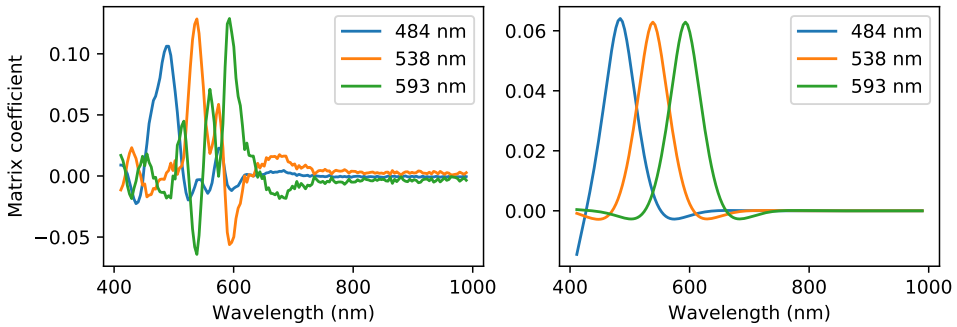

Figure 12. Components of the denoising matrices for MNF (left) and smoothing splines (right, $p = 0.001$) that yields denoised spectral values at three discrete bands.

Evaluating the smoothing splines at the training points x_i can be expressed as a matrix multiplication $\mathbf{S}_\alpha \boldsymbol{y}$, while denoising using MNF is also similarly expressed as a matrix multiplication $\mathbf{D}(\boldsymbol{y} - \bar{\mathbf{B}})$. A value at the band i is therefore calculated by linearly combining values of the raw spectrum using coefficients present along a row i of the denoising matrix. Linear combination coefficients yielding a specific band value are plotted in Fig. 12, and shows that the coefficients are similar for smoothing splines and MNF. There exists a relation between smoothing splines and ridge regression,[6] and ridge regression and PCA,[6] and MNF can be expressed as a PCA transform of the noise whitened image.[15] The similarity of the denoising matrices could indicate some relation between smoothing splines and MNF, but would have to be further investigated.

5.2.3 Application of MNF and smoothing splines to simulations

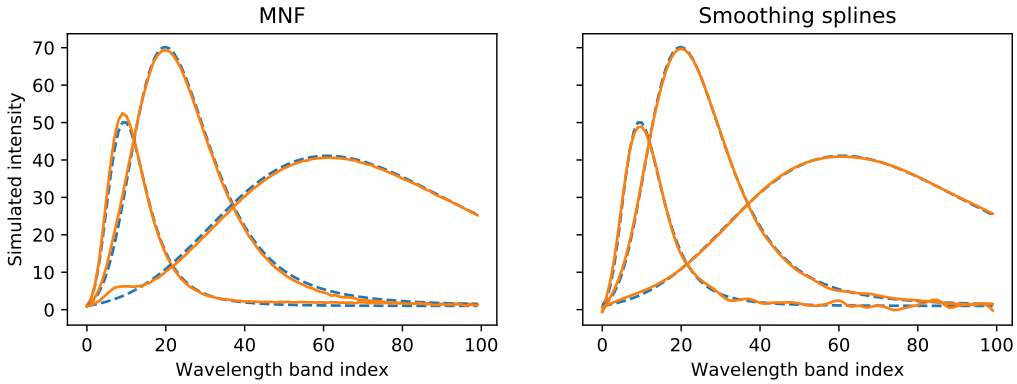

Figure 13. Comparison of MNF and smoothing splines as applied to simulated spectra.

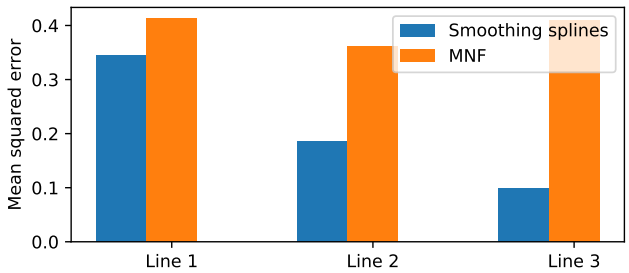

Figure 14. Mean squared error between smoothing splines and MNF for the three spectral types present in the simulated dataset.

MNF was run on the simulations with 6 components in inverse, with the seventh component consisting of pure noise. Examples of denoised spectra are shown in Fig. 13. The mean squared error between the original spectra before noise was added and the denoised spectra is shown in Fig. 13. The examples show alien spectral components in the MNF denoising of spectrum 3. This demonstrates component mixing, where components from from spectrum 1 are used to denoise spectrum 3. The spectra also show some bias. The bias is summed up as a higher bulk mean squared error as compared to the smoothing splines technique. The smoothing splines technique has lower mean squared error, and the examples show that the smoothing splines method well approximates the high-SNR parts of the original spectra. However, the smoothing splines estimate has high variance and follows the noise too closely for the low-SNR parts. The choice of α is low in order to accomodate the high-SNR parts, but results in a fit with too many degrees of freedom in the low-SNR parts. This is behavior that was also seen for the measured fluorescence spectra, especially evident in Fig. 10.

5.2.4 General discussion

Smoothing splines interpolates the trends of the noisy spectra, using an objective measure with a clear minimum to determine the amount of smoothness for each pixel. This makes the results comparable regardless of image

and noise statistic estimates, which is especially useful for characterization across images. The smoothing splines tuning parameter is continuous, allowing individual tuning down to an arbitrary numerical precision.

MNF is tuned by visual inspection of either components or SNR plots, and the results are not necessarily comparable across different images or even pixels. The MNF tuning parameter is discrete, restraining somewhat the possbility for tuning.

MNF takes image statistics from the image at large into account when smoothing the spectra, whereas smoothing splines considers only one spectrum at a time. This can have both advantages and disadvantages. It was seen in the simulations that using the full image to denoise single spectra can lead to component mixing, which can have negative ramifications for further processing. Different images would have different image statistics, and different components would be used to build the denoised spectra. Using components from the full image makes full use of the available information, however, and restrains the spectra in low-SNR parts properly, as seen in the simulations. The smoothing splines estimates of α mainly satisfies the high-SNR parts of the spectra, while the low-SNR parts will have an α which is not restrictive enough.

Thus, MNF and splines represent two extremes: MNF uses information from the image at large, but makes no individual tailoring, while smoothing splines considers each spectrum individually, but makes no use of prior information. Improvements could be made by centering and standardizing the image bands or applying weights to the method using the image covariance matrix. In addition, there exists modifications to the technique which have α dependent on x,[16–19] which would be appropriate here.

MNF would also be more adaptable for various types of noise distributions and image spectra, while the smoothing spline is more restrictive and sensitive to spectral behavior. Attempts to apply the method on reflectance data from human skin yielded unphysical results due to the high variation in noise variance as a function of wavelength, and absorption features being confused as undesired noise variance. The method was more suitable for fluorescence data, due to the simple polynomial-like behavior of these spectra. Centering and standardization of the image bands could improve applicability to reflectance data. It is expected that the method will be less suitable when the fluorescence spectra are strongly influenced by absorption features.

Inverse spectra obtained using the MNF transform can discard or mix important spectral features. Still, MNF is often the most suitable alternative, and for some data this might be a good enough approximation. The variant of MNF considered here is also the most extreme variant, in that a hard cut-off is used for the transformed bands. A more gradual variant could be to apply gradually more aggressive noise removal on the transformed bands before back-transformation, but objective evaluation of such improvements could be challenging. MNF is also valuable in dimensionality reduction.

Regardless of denoising method, the smoothing splines method is well suitable for investigating the spectra at subresolution positions and be used to estimate fluorescence emission peak positions or other properties that are easily extracted from spline coefficients. The results here also suggests that it is appropriate as a denoising method for fluorescence spectra.

6. CONCLUSION

The smoothing splines method has been adapted for use in hyperspectral images, with individual smoothing parameters that are appropriate for every single pixel in the image. The method has been applied to images of atherosclerosis plaques, and found to yield smooth spectra. The spline coefficients were used to find sub-resolution peak positions of the emission spectra. Comparison of MNF and smoothing splines on simulations showed the smoothing splines estimate to yield spectra with lower bias. While the appropriateness of MNF is subject to the availability of image statistics, smoothing splines represents a robust method with an objective choice between smoothness and closeness of the spectra, and which apparently has advantages over MNF for the type of data represented by autofluorescence images.

APPENDIX A. MANIPULATING SMOOTHING SPLINES SOFTWARE PACKAGES TO OBTAIN THE SMOOTHING MATRIX

The techniques presented in this paper assume the availability of $\mathbf{S}_\alpha$, the smoothing matrix. This is usually not made available in a given software implementation. Smoothing splines is implemented in MATLAB (version 9.2.0, The MathWorks Inc., Natick, Massachusetts, USA) through the function `csaps`, and a Python implementation can be found in the Python package `pywafo`.[20] Both implement smoothing splines according to the definitions given in de Boor,[7] which use p instead of α as the smoothing parameter.

A.1 Estimation using the identity matrix

This method requires no software modification, and obtains $\mathbf{S}_\alpha$ by manipulating the input. Given a spectrum $\boldsymbol{I_1} = [1, 0, \ldots, 0]^T$, the resulting fitted smoothing spline is

$$\hat{\boldsymbol{f}}_\alpha = \mathbf{S}_\alpha \boldsymbol{I_1} = [S_{11}, S_{12}, \ldots]^T. \tag{14}$$

Iteratively putting in columns of the identity matrix I reproduces $\mathbf{S}_\alpha$.

A.2 Estimation using intermediate results

This method requires software modification. As both MATLAB and pywafo implementations are based on de Boor,[7] a couple of intermediate results can be exploited to estimate $\mathbf{S}_\alpha$ directly. Below, this is outlined using de Boor's notation, usually translated into similar variable names in the code. de Boor's method also considers the covariance matrix of y, $\mathbf{D}$, but this is neglected below for simplicity.

The smoothing splines are here formulated in terms of their piece-wise polynomial form (ppform). One of the intermediate results is the vector $\boldsymbol{a}$, defined to be

$$f_p(x_i) = a_i, \tag{15}$$

i.e. the elements of $\boldsymbol{f_p}(\boldsymbol{x}) = \mathbf{S}_\alpha \boldsymbol{y}$. It is in the packages found from

$$\boldsymbol{a} = \boldsymbol{y} - 6(1-p)\mathbf{Q}\boldsymbol{u}. \tag{16}$$

The vector $\boldsymbol{u}$ is found by solving

$$\underbrace{\left(6(1-p)\mathbf{Q}^T\mathbf{Q} + p\mathbf{R}\right)}_{\stackrel{\text{def}}{=}\mathbf{B}_p}\boldsymbol{u} = \mathbf{Q}^T\boldsymbol{y} \tag{17}$$

Both $\boldsymbol{a}$ and $\boldsymbol{u}$ are used to estimate the other parameters defining the ppform of the smoothing spline, which can then be used for interpolation. The matrices $\mathbf{Q}$ and $\mathbf{R}$ are calculated from differences between data points x_i.

Inserting (17) into (16) then yields

$$\boldsymbol{a} = \underbrace{\left(\mathbf{I} - 6(1-p)\mathbf{Q}\mathbf{B}_p^{-1}\mathbf{Q}^T\right)}_{\mathbf{S}_\alpha}\boldsymbol{y}. \tag{18}$$

Thus, $\mathbf{S}_\alpha$ is estimated from $\mathbf{Q}$ and $\mathbf{B}_p$, which are all available in a given smoothing splines implementation.

ACKNOWLEDGMENTS

Thanks to Eivind L.P. Larsen for acquisition of the data.

REFERENCES

[1] 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), 026011 (2011).

[2] Noh, H. K. and Lu, R., "Hyperspectral laser-induced fluorescence imaging for assessing apple fruit quality," *Postharvest Biol. Technol.* (2007).

[3] Gowen, A. A., O'Donnell, C. P., Cullen, P. J., Downey, G., and Frias, J. M., "Hyperspectral imaging - an emerging process analytical tool for food quality and safety control," *Trends Food Sci. Tech.* **18**(12), 590–598 (2007).

[4] Lu, G. and Fei, B., "Medical hyperspectral imaging: a review," *J. Biomed. Opt.* **19**(1), 010901 (2014).

[5] Green, A. A., Berman, M., Switzer, P., and Craig, M. D., "A transform for ordering multispectral data in terms of image quality with implications for noise removal," *IEEE Trans. Geosci. Remote Sensing* **26**(1), 65–74 (1988).

[6] Hastie, T., Tibshirani, R., and Friedman, J., [*The elements of statistical learning*], Springer, second ed. (2008).

[7] de Boor, C., [*A practical guide to splines*], Springer-Verlag (1978).

[8] Marion, R., Michel, R., and Faye, C., "Measuring trace gases in plumes from hyperspectral remotely sensed data," *IEEE Trans. Geosci. Remote Sens.* **42**, 854 – 864 (2004).

[9] Berman, M., "Automated smoothing of image and other regularly spaced data," *IEEE Trans. Pattern Anal. Mach. Intell.* **16**, 460–468 (1994).

[10] Lee, J. S. and Cox, D. D., "Robust smoothing: Smoothing parameter selection and applications to fluorescence spectroscopy," *Comput. Stat. Data Anal.* **54**, 3131–3143 (2010).

[11] Dongarra, J., "Basic linear algebra subprograms technical forum standard," *Int. J. High Perform. Comput. Appl.* **16**(2), 1–111 (2002).

[12] Bjorgan, A. and Randeberg, L. L., "Real-time noise removal for line-scanning hyperspectral devices using a minimum noise fraction-based approach," *Sensors* **15**(2) (2015).

[13] Larsen, E. L., 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).

[14] Randeberg, L. L., Larsen, E. L. P., Haugen, O. A., and Svaasand, L. O., "Hyperspectral characterization of atherosclerotic plaques," *Proc. SPIE* **7368** (2009).

[15] Lee, J., Woodyatt, A. S., and Berman, M., "Enhancement of high spectral resolution remote sensing data by a noise-adjusted principal components transform," *IEEE Trans. Geosci. Remote Sens.* **28**(3) (1990).

[16] Mackay, D. J. C. and Takeuchi, R., "Interpolation models with multiple hyperparameters," *Stat. Comput.* **8**, 15–23 (1998).

[17] Abramovich, F. and Steinberg, D. M., "Improved inference in nonparametric regression using lk-smoothing splines," *J. Stat. Plan. Inference* **49**, 327–341 (1996).

[18] Pintore, A., Speckman, P., and Holmes, C. C., "Spatially adaptive smoothing splines," *Biometrika* **93**, 113–125 (2006).

[19] Ruppert, D. and Carroll, R. J., "Theory & methods: Spatially-adaptive penalties for spline fitting," *Aust N. Z. J. Stat.* **42** (2002).

[20] "pywafo source code repository." Available online on https://github.com/wafo-project/pywafo. Visited 2019-01-23.