

ESTIMATING THE RELATIVE SPATIAL AND SPECTRAL SENSOR RESPONSE FOR HYPERSPECTRAL AND MULTISPECTRAL IMAGE FUSION

Charis Lanaras^{*}, Emmanuel Baltsavias, Konrad Schindler

ETH Zurich, Institute of Geodesy and Photogrammetry
Stefano-Franscini-Platz 5, CH-8093 Zurich, Switzerland

Email: charis.lanaras@geod.baug.ethz.ch, manos@geod.baug.ethz.ch, schindler@geod.baug.ethz.ch

KEY WORDS: hyperspectral, image fusion, super-resolution, relative sensor response

ABSTRACT: In the framework of our research, we process jointly high spectral (hyperspectral) and high geometric (multispectral) resolution images with lower and higher spatial resolution respectively, exploiting their synergies, with aims to (a) generate a fused image of high spectral *and* geometric resolution, and (b) to improve spectral unmixing at sub-pixel level of hyperspectral images regarding estimation of endmembers and fractional abundances. To relate the two images much of the existing research work assumes that the spatial and spectral characteristics (responses) of the two sensors are known in advance. This assumption is true only for simulated data. When moving to real data, however, it is not obvious that the exact spatial and spectral sensor characteristics are accessible. The aim of this concrete work is to derive the relative sensor responses from the data, given (approximately) co-registered images. Recovering the relative spatial response amounts to reconstructing the 2D spatial blur kernel that integrates multispectral pixels into hyperspectral pixels. Conversely, to recover the relative spectral response we estimate the shapes and sizes of the 1D kernels that integrate hyperspectral bands into multispectral bands. The spectral and spatial response functions are coupled, in the sense that one must be known in order to directly solve for the other. In practice, we find that estimating them in two consecutive steps is sufficient. The proposed formulation includes non-negativity and other constraints, recovers remaining registration (translation) errors between the two images, and uses prior information to adjust to the shape of the spectral response (rectangular or ramp shaped) with either l_1 or l_2 norm regularization. The proposed method is tested with both real and simulated data, aerial, close-range and satellite, in the second case with ground truth. The results show that, also under real-world imaging conditions, it appears possible to compute the relative spatial and spectral responses in a data-driven manner.

1. INTRODUCTION

Hyperspectral images usually have lower spatial resolution than multispectral images. This is mainly because larger pixels are needed to achieve an acceptable signal-to-noise ratio (SNR) when the reflected energy is split into many narrow spectral bands. Recently, some researchers (Yokoya et al., 2012; Wycoff et al., 2013; Huang et al., 2014; Wei et al., 2015; Lanaras et al., 2015) use a multispectral image (MSI), with higher spatial resolution, to spatially enhance a hyperspectral image (HSI) with lower spatial resolution. This process is called *hyperspectral super-resolution* or *hyperspectral fusion*. To fuse the HSI and the MSI most of the existing work assumes that the relative spatial and spectral responses of the two sensors are known. In reality, virtually all work on hyperspectral super-resolution uses simulated data for evaluation, due to the difficulty of obtaining high-quality ground truth. Thus, the above assumption is always met. However, when moving to real data, the exact spatial and spectral sensor characteristics are not always known or may differ from the specifications due to aging of optomechanical and electronic components or a mismatch between the aerial/orbital environment and the laboratory (Wang et al., 2010). In addition, the HSI and MSI are often acquired with time difference and under different viewing angles, illumination and atmospheric conditions. The potential of the above mentioned hyperspectral fusion lies in a more accurate and detailed semantic interpretation of objects and their properties compared to original hyperspectral and multispectral images. Multispectral images are quasi standard, while hyperspectral images are increasingly acquired by several sensors and platforms (close-range, aerial and spaceborne), with applications often in environmental mapping, monitoring and change detection.

The aim of this research is to derive the relative spatial and spectral sensor response (relative characteristics) from the data, given approximately co-registered HSI and MSI. Estimating the spatial response amounts to finding the 2D spatial blur kernel that integrates MSI pixels into HSI pixels. In our approach, this also estimates remaining translational errors of the co-registration. To estimate the spectral response, we reconstruct the shapes and sizes of the kernels that integrate HSI bands into MSI bands. The spectral and spatial response functions are coupled, since one must be known in order to directly solve for the other. In practice, we found that estimating them in two consecutive steps is sufficient. Iterating the two steps is possible, but brings no significant improvement. Using images alone one can generally not recover absolute sensor characteristics, but only the relative response of one

^{*} Corresponding author.

sensor with respect to the other. However, approximate knowledge of absolute values can help initialize and guide the process, as explained below.

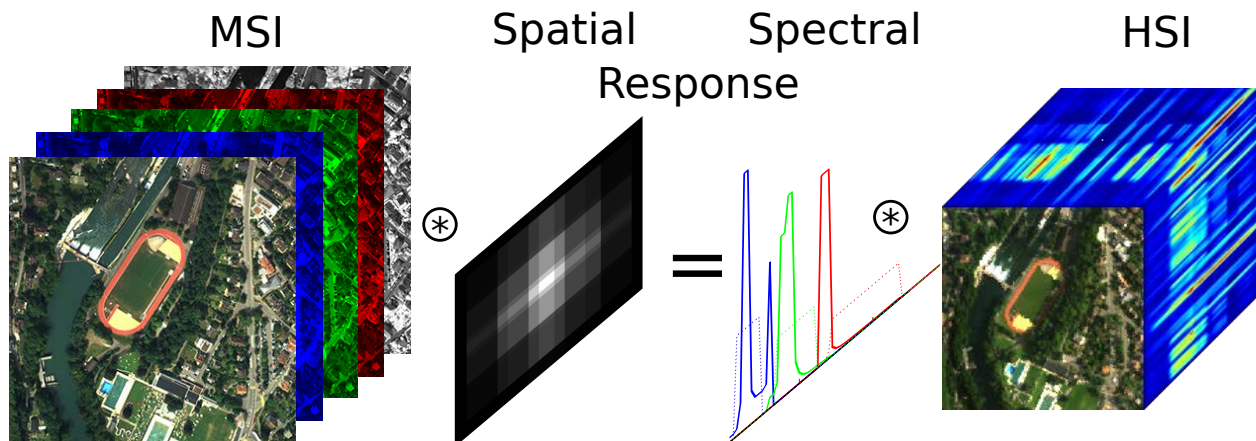


Figure 1. A visual representation of the underlying conceptual model used to estimate the relative spatial and spectral downsampling models of the MSI and HSI sensors.

1.1 Image Formation Model

We assume that both the spectral and spatial response functions are linear convolutions. A graphical overview of this model is shown in Figure 1. Below, we treat the two images not as 3D image cubes, but reordered to matrices with one pixel per column: $\mathbf{H} \in \mathbb{R}^{B \times p}$ for the HSI and $\mathbf{M} \in \mathbb{R}^{b \times P}$ for the MSI, with $B > b$ the respective numbers of spectral bands, and $P > p$ the number of image pixels. With the unknown relative spectral and spatial response functions \mathbf{R} and \mathbf{B} , the two images are thus, up to noise, related as $\mathbf{RH} = \mathbf{MS}$, where \mathbf{S} is a matrix that spatially downsamples \mathbf{M} according to the blur \mathbf{B} , so \mathbf{S} includes both blurring and downsampling.

1.2 Related Work

There is little work on estimating relative MSI/HSI sensor responses. As far as we know, this has been investigated only as a prerequisite for hyperspectral super-resolution, which itself is a quite recent topic. Simoes et al. (2015) formulate a regularized quadratic problem for estimating both response functions, assuming spectral response curves of known width in the MSI. Both spectral and spatial responses are regularized by penalizing their respective gradients. They also separately solve for spatial and spectral blur. In contrast to our work, non-negativity of the response is not enforced, and regularization is limited to quadratic penalizers, which in our experience tend to over-smooth the result. Yokoya et al. (2013) tightly constrain the spectral response to deviate only slightly from the known pre-launch values. The spatial response is modelled as a Gaussian blur, assuming perfect co-registration. Its variance is found by maximizing the cross-correlation between the gradients of the HSI and the downsampled MSI. Huang et al. (2014) have proposed an unconstrained solution for the spectral response, which has already been shown to lead to non-plausible results.

1.3 Own Contribution

Details of our previous work can be found in Lanaras et al., 2015. In this research, we focus on spatial and spectral sensor response issues, and propose a model with the following novel aspects. We consider non-negativity of the responses, which is physically reasonable and significantly stabilizes the estimation. For the spatial response we do not enforce a Gaussian blur, but rather only require it to be symmetric with respect to an unknown centre. We also refrain from regularization, which tends to affect the width of the blur kernel. Estimating the centre of the spatial response allows one to estimate small translational mis-registration errors between HSI and MSI. This is useful, since in practice accurate (sub-pixel) co-registration of images with different resolutions and possibly different geocoding is sometimes challenging. For the spectral response, apart from the l_2 norm we use also the l_1 norm for data fitting, to increase robustness to outliers and also because it is better suited for modelling steep, nearly rectangular spectral response functions as found in some MSI, such as ADS80 or Landsat-8 OLI (Cramer, 2011). We show in addition that, by applying the estimated relative sensor relations, one can detect areas, which deviate from the proposed model and may cause difficulties for a subsequent hyperspectral super-resolution. These areas are mainly due to multitemporal differences and could to a certain extent used for change detection.

2. RELATIVE SPATIAL RESPONSE

The spatial resolution difference (ratio) between the MSI and HSI is $S = \sqrt{P/p}$, here assumed that $S \in \mathbb{N}_+$ (though S theoretically and for the proposed method could be non-integer). Let $\mathbf{R}_0 \in \mathbb{R}^{b \times B}$ be an initial approximation of the relative spectral response, which reduces the number of spectral bands from B to b . Using the approximate spectral response does not change the sharpness of the image, thus we can start from estimating the spatial response.

Let $\mathbf{h} = \mathbf{R}_0 \mathbf{H} \in \mathbb{R}^{b \times p}$ denote the image created from \mathbf{H} having the spectral bands of \mathbf{M} . We seek to estimate the blur that will optimally fit \mathbf{M} to \mathbf{h} under some spatial subsampling. To do so, we split the estimation into the horizontal and vertical blurring kernels, and compute these for every MSI band i . By splitting in two 1D kernels, we assume the blur is separable. This reduces the number of unknown coefficients and gives more stable results, while still allowing for a flexible, anisotropic representation of the blur. The search window size for the blur coefficients in the MSI resolution is $W = (2k+1)S$, where $k \in \mathbb{N}$ is expressed in HSI pixels and is determined semi-empirically (i.e. usually the spatial resolution difference between HSI and MSI is known, and we choose k and thus W to be significantly larger). For pixels $j = 1, \dots, Z$ of \mathbf{h} we extract Z 1D patches of size W from \mathbf{M} , either horizontally or vertically. Z is the number of all HSI pixels for which the 1D kernel of size W (but now in HSI pixels) does not fall outside the HSI image (i.e. an image border of k HSI pixels is not processed). We solve the following optimization separately for each of the two 1D kernels and each MSI band i :

$$\arg \min_b \|\mathbf{h}_i - \mathbf{L}\mathbf{b}_i\|_2 \quad (1)$$

subject to $\mathbf{G}\mathbf{b}_i \geq 0$ (see Figure 2)

where $\mathbf{b}_i \in \mathbb{R}^W$ are the unknown 1D coefficients of the blur,
 $\mathbf{h}_i \in \mathbb{R}^Z$ are the Z values of band i at pixels j , and
 $\mathbf{L} \in \mathbb{R}^{Z \times W}$ is a matrix (including values of \mathbf{M}) with Z rows of W -dimensional patches.

$\mathbf{G} \in \mathbb{R}^{W \times W}$ is a matrix of which each row computes the difference between two elements in \mathbf{b}_i with successively increasing distances to the centre of gravity of the blur (closer minus further element) (see Figure 2). The row for the element with the largest distance has only the respective element set to 1, to ensure non-negativity of all elements in \mathbf{b}_i . Note that the sequence of the rows of \mathbf{G} can be arbitrary but we choose them such that the diagonal elements are 1. To find the centre of gravity (in horizontal and vertical direction), we first get a solution of Eq. (1) with $\mathbf{G} = \mathbf{I}_w$, with \mathbf{I} the identity matrix to enforce non-negativity. The final 2D blur for band i is given by the product of the vertical and horizontal kernels, $\mathbf{B} = \mathbf{b}_{\text{ver}} \mathbf{b}_{\text{hor}}^T$. Eq. (1) is solved as a quadratic problem, using the interior point method (Boyd and Vandenberghe, 2004). The constraint shown in Figure 2 also enforces symmetry for the 1D kernel coefficients and decreasing or equal blur kernel coefficients from the centre of the kernel. The offset of the blur kernel's centre of gravity from the centre of the window corresponds to the global, translational mis-registration between HSI and MSI. Note that we do not enforce that the sum of the blur coefficients is 1 (though in practice it is always close to 1). This is done to compensate for global radiometric differences between the HSI and MSI.

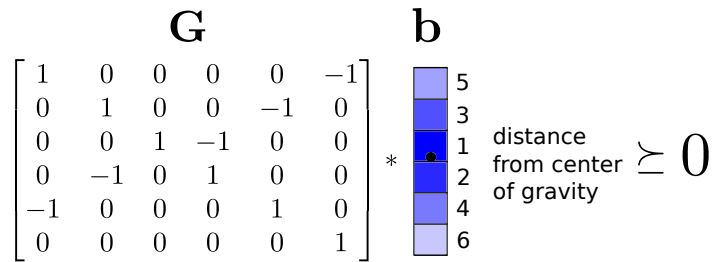


Figure 2. Visual explanation of the constraint of Eq. (1).

3. RELATIVE SPECTRAL RESPONSE

Given that the HSI bands are very narrow and assuming that they fully overlap with the MSI bands, we can express each MSI band i as a linear combination of HSI bands. The estimation of the spectral response $\mathbf{R} \in \mathbb{R}^{b \times B}$ is thus independent for each row of \mathbf{R} (i.e. each MSI band), which leads to the following optimization for the unknown spectral response $\mathbf{r}_i \in \mathbb{R}^B$ of band i :

$$\arg \min_{\mathbf{r}} \|\mathbf{F}(\mathbf{m}_i - \mathbf{H}^T \mathbf{r}_i)\|_1 + \lambda_i \|\mathbf{D} \mathbf{r}_i\|_a \quad (2)$$

subject to $\mathbf{r}_i \geq 0$

where $\mathbf{m}_i \in \mathbb{R}^p$ is the i^{th} MSI band, spatially downsampled with the blur \mathbf{b}_i and $\lambda_i \geq 0$ is a regularization parameter to enforce spectral response curve smoothness for band i , using the matrix operator $\mathbf{D} \in \mathbb{R}^{B-1 \times B}$ to compute the differences of each \mathbf{r}_i element of adjacent bands (current minus the right one in the electromagnetic spectrum). The diagonal matrix $\mathbf{F} = \text{diag}(m_{i1}^2, \dots, m_{ip}^2) \in \mathbb{R}^{p \times p}$ holds individual weights for the elements \mathbf{m}_i (individual pixels). Weights are selected such that pixels with higher intensity - and thus better SNR - in a band contribute more to the estimation of \mathbf{r}_i . Empirically, this weighting stabilizes the solution. The type of the norm a is selected to reflect prior knowledge about \mathbf{r}_i . Steep response curves require $a = 1$, whereas for smooth, gradually changing spectral response curves $a = 2$ is preferable. To solve Eq. (2), we again use the interior point method, starting from the approximate spectral response used in Section 2. In most cases, we know approximately the MSI and HSI spectral bandwidths, hence we can limit the search to a smaller number of bands $\hat{B} < B$. \hat{B} are the approximately known HSI spectral bands covering the spectral overlap between HSI and MSI in each MSI band extended on each side of the electromagnetic spectrum of the MSI band a bit more. The selection of the \hat{B} bands is done manually. Note, this only places soft lower and upper bounds on the spectral response curve, the exact spectral width of HSI and MSI bands need not be known. The parameters λ_i are set individually for each band, as discussed in Section 4.2.

4. TEST DATA, RESULTS AND DISCUSSION

4.1 Data

The first dataset includes HSI and MSI aerial images covering the city of Wettingen, Switzerland. The HSI was acquired by APEX (Schaeppman et al., 2015) with GSD ≈ 3 m on 12/07/2013. It has 284 bands covering the range 0.4-2.4 μm . The MSI was acquired by a Leica ADS80 with GSD ≈ 0.5 m on 27/07/2013. It has 4 spectral bands (RGB and NIR). ADS80 has very steep, almost perfectly rectangular spectral responses (Cramer, 2011). The two sensor image strips and their overlap are shown in Figure 3. Two regions (see Figure 3, blue and red rectangles) were used for testing. The blue rectangle is 876x870 MSI pixels and the red rectangle 1074x1590 MSI pixels. The HSI and MSI were co-registered using mutual information (Lanaras et al., 2014). HSI and MSI images were spectrally and spatially degraded respectively and the co-registration was done on both test regions independently. Only two global translations for each band were estimated. It was observed that the four bands were not perfectly co-registered with each other, implying mis-registration of the spectral bands of ADS80 and/or (most probably) the APEX ones. However, finally a common translation was used for all MSI image bands as an average of the individual band translations and these were applied to the ADS80 data for each test region.




Figure 3. APEX (left top) and ADS80 (left bottom) data strips with the two colour rectangles showing the test regions, and on the right the zoomed blue rectangle from ADS80 (top) and APEX (bottom).

The second dataset made use of the close-range CAVE image database (Yasuma et al., 2008). It includes many scenes with a large variety of colours and materials. The scenes were imaged in 31 hyperspectral bands in the range of 400-700 nm in 10nm increments using a liquid crystal tunable colour filter and also in RGB, whereby both HSI (called multispectral in the respective documentation) and RGB had the same spatial resolution. It is not clear how the delivered RGB images are generated. They could be real or synthetic, RGB or HDR RGB (HDR = High Dynamic Range). To the best of our knowledge, the RGB images are simulated, ground truth images, generated using the theoretical spectral responses of the 7 colour filters of the GAP sensor the authors designed and the 31 hyperspectral channels. They look spectrally very similar to the HDR RGB images of the GAP sensor. Since it was unknown how the RGB images were generated and whether there was a mis-registration with the HSI images, we also generated new RGB images, by using the HSI images and a spectral subsampling using as spectral response curves those of a typical commercial camera, the Nikon D7000. We call these two RGB images as original and new RGB CAVE data. We use the original RGB for the estimation of the relative spectral sensor response, and the new RGB for the relative spatial sensor response.

The third dataset came from the EO-1 satellite (see details at <https://eoportal.org/web/eoportal/satellite-missions/e/eo-1>). The test area (1100 x 190 pixels) is along the river Rhine, close to the French city of Strasbourg. Images were acquired on July 18th 2016. MSI comes from the ALI sensor (9 bands, covering the spectral range 433-2350 nm) with a 30m GSD and HSI from the Hyperion sensor with 198 useful bands (out of the 220, some bands did not contain any useful information), covering the spectral range 400-2500 nm with 10nm bandwidth and also 30m GSD. The two images were co-registered to less than a pixel, using the average translation derived from many match points using SIFT. For this dataset only the relative spectral response is estimated, because MSI and HSI have the same spatial resolution.

Image data in grey values (8- to 16-bit) were linearly transformed to the range 0-1.

4.2 Results and Discussion

The spatial response estimation was evaluated at the two test regions of the aerial images (Figure 3). Window widths with $k = 4,5,6$ were semi-empirically chosen (note that the spatial resolution difference between HSI and MSI was about 6) and showed the best numerical stability (in the first iteration, when $\mathbf{G} = \mathbf{I}_w$). Figure 4 shows the computed 1D and 2D blur kernels for the blue ADS80 band and $k = 4$. Table 1 shows the residual mis-registration estimated for various bands and k values. The estimated mis-registrations are up to ≈ 2.5 MSI pixels (≈ 0.4 HSI pixels), implying that the co-registration was not ideal (we believe that a good co-registration should not exceed errors of about 0.2-0.3 HSI pixels). The errors are quite stable for different k but vary significantly from band to band, which means that for one sensor (most probably APEX) the individual sensor bands have remaining co-registration errors.




Figure 4. Estimated relative spatial response for the blue band of ADS80 for $k = 4$. 1D kernels in colour, 2D kernel in grey. The pixels values at left and bottom are in MSI pixels. Note that the pixels of the 1D and 2D kernels as they appear in the figure actually often consist of multiple pixels.

Table 1. Mis-registration in pixels (in the MSI resolution) between HSI and MSI per image band of the aerial ADS80 and APEX images, for various k values.

k value	X-coordinate			Y-coordinate		
	4	5	6	4	5	6
Blue	-0.5	-0.4	-0.4	0.7	0.6	0.6
Green	0.6	0.6	0.7	1.0	1.1	1.0
Red	0.0	0.1	0.1	1.2	1.2	1.1
NIR	-2.3	-2.5	-2.6	1.9	1.9	1.9

In order to evaluate our method with known ground truth we used all CAVE HSI database images and the new RGB (MSI) images. Each HSI image was blurred with a 11×11 Gaussian kernel, where an x/y mis-registration of respectively 1.7/0.8 pixels was introduced (thus an interpolation was performed). The images were then downsampled with a local average filter by a factor 6, the same resolution difference as the aerial images. This new data were used as HSI images. MSI and HSI were also degraded by Gaussian noise of SNR = 30 dB. Table 2 shows that the estimated offsets of the blur kernel are almost identical to the introduced mis-registration errors and that the k value, as long as it is quite correct (big enough), has practically no influence. The differences between the bands are minimal.

Table 2. Mis-registration in pixels (in the MSI resolution) between the HSI and MSI (simulated data, using the new RGB) of the CAVE database for various spectral bands and k values.

k value	X-coordinate			Y-coordinate		
	4	5	6	4	5	6
Blue	-1.6	-1.6	-1.6	0.9	0.9	0.9
Green	-1.6	-1.6	-1.6	0.8	0.8	0.8
Red	-1.6	-1.6	-1.6	0.8	0.8	0.8




Figure 5. (a) Approximate specifications (dotted lines) and final estimates (solid lines) of the relative spectral response. (b) For the aerial data: influence of the regularizer $\|\mathbf{Dri}\|$ as a function of its weight λ_i . We chose flat regions between the saturated extremes, to ensure stable behaviour. (c) The estimated spectral response of the original RGB CAVE sensor (solid line). The spectral response of Nikon D7000 (taken from manufacturer specifications), which is assumed to be similar to the one used in the original RGB CAVE data (dotted line). The relative spectral response in (a) and (c) is unitless and is estimated using the best fit between each MSI band and the HSI bands used to reconstruct it.

The results of the spectral response estimation for the airborne and CAVE sensors are shown in Figure 5. In Figure 5b, we plot the influence of the regularization term $\|\mathbf{Dri}\|$ against its weight λ_i . One can clearly see flat areas for too low values of λ_i , where the solution is not affected at all, and for very high values, where oversmoothing occurs. In between, there are regions where the curves are flat, meaning that the balance between $\|\mathbf{Dri}\|$ and the data, and thus also the solution, is rather stable. We found empirically that selecting λ_i in flat areas (e.g. in Figure 5(b) and the red curve at $\lambda_i 10^{-1}$) yields plausible spectral responses. The selection of λ_i was manual for each band, whereby its value also depends on the selection of the l norm (1 or 2). For the aerial data we use the l_1 norm (due to the ADS80 steep and rectangular spectral characteristics), for the CAVE data the l_2 norm (due to the smoother spectral response curves). For the CAVE dataset, since not every image contains materials of all colours, we solve for a single spectral response over all images. For this data we do not restrict the solution to a part of the spectral range (so, no \hat{B} was used), and allow strong regularization with $\lambda = 10^3$. The results are plotted in Figure 5(c) with a solid line, while the dotted line corresponds to the spectral response of a Nikon D7000 camera. Although it is unknown how

the RGB images of the CAVE dataset are generated, we assume that this was done by a sensor (maybe synthetic) with a spectral response similar to that of the Nikon D7000. Figure 6 shows the relative spectral response results for the EO-1 ALI and Hyperion sensors. In the estimation, the l_1 norm was used.

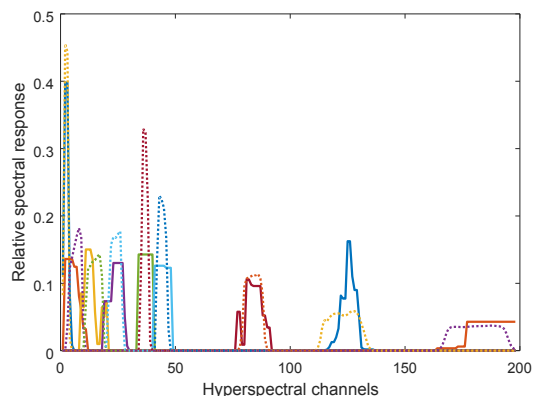


Figure 6. Solid line: Estimated spectral responses of the 9 spectral bands of ALI. Dotted line: The spectral responses according to the manufacturer.

In Figure 7 the residuals for one sub-region of the aerial dataset are shown, i.e. the differences between the spectrally downsampled HSI and the spatially downsampled MSI band 3 (red). The distribution of the residuals is plausible: the right part of the stadium tartan is shadowed in the HSI due to differences in image acquisition time. There are also some large residuals on buildings, which we attribute to differences in viewing angle and possibly different shadows due to multitemporal data acquisition. Figure 8 shows similar information for the CAVE dataset. We do not have an explanation for the high residuals on the lower right (background of the painting) of this CAVE image. Figure 9 shows the results for the EO-1 sensors ALI and Hyperion. The results are very good and this can certainly be also attributed to the fact that MSI and HSI have the same spatial resolution, are acquired simultaneously and from practically the same viewing angle, cover almost the same electromagnetic spectrum region and the image scale is small (thus small differences between HSI and MSI are negligible).




Figure 7. Residuals of the spectral response between the spatially subsampled MSI and the spectrally subsampled HSI for band 3 (red), measured in absolute differences between intensities in the range [0 - 1].



Figure 8. Original CAVE RGB (left), spectrally subsampled HSI (middle), residuals of the spectral response measured in absolute differences between intensities in the range $[0 - 1]$ (right).




Figure 9. Top: Bands MS-1', MS-2 and MS-3 (blue, green, red) of ALI. Middle: Estimated spectral subsampling of Hyperion (here the bands corresponding to RGB are shown). Bottom: residuals in the MS-2 band (green) of the spectral response measured in absolute differences between intensities in the range $[0 - 1]$. Note that the colour in the legend scale is different than in the figures above. Co-registration is visually judged (through subtracting ALI and Hyperion bands) to be sufficient and the model fits well to the data (residuals are small).

5. CONCLUSIONS

We have investigated the estimation of relative sensor characteristics between MSI and HSI, in the context of hyperspectral super-resolution for various datasets. Our main message is that, also under real-world imaging conditions, it appears possible to compute the spatial and spectral relative relations in a data-driven manner. We found that enforcing non-negativity for the spectral and spatial sensor response functions is important and stabilizes the estimation. Also important is to set the search window for the spatial response conservatively, to ensure it is larger than the actual blur. The estimation of the spectral response is a harder problem that needs to be strongly constrained to obtain plausible results. In particular, we find that prior knowledge about the shape of the response curves (steep rectangular vs. ramp-like flanks) is valuable and can be introduced by using an appropriate norm. While the results are encouraging, further work is needed to address situations with more image noise or multi-temporal differences between HSI and MSI. Moreover, the residuals of the data-driven estimation make it possible to identify regions that do not comply with the estimation model used here and would cause errors for a subsequent hyperspectral super-resolution, facilitating their elimination from the estimation.

ACKNOWLEDGEMENTS

We gratefully acknowledge the financial support of the Swiss Science National Foundation for this research and the provision of the aerial data by the Remote Sensing Labs, University of Zurich.

REFERENCES

Boyd, S. and Vandenberghe, L., 2004. Convex optimization. Cambridge University Press.

Cramer, M., 2011. Geometry perfect - radiometry unknown? In: Photogrammetric Week 2011, D. Fritsch (Ed.), Wichmann/VDE Verlag, Berlin & Offenbach. At <http://www.ifp.uni-stuttgart.de/publications/phowol1/080cramer.pdf> (accessed, 14 Sept. 2016)

Huang, B., Song, H., Cui, H., Peng, J. and Xu, Z., 2014. Spatial and spectral image fusion using sparse matrix factorization. *IEEE TGRS* 52(3), pp. 1693–1704.

Lanaras, C., Baltsavias, E. and Schindler, K., 2014. A comparison and combination of methods for co-registration of multi-modal images. In: Proc. of 35th Asian Conference on Remote Sensing, Nay Pyi Taw, Myanmar, 27-31 October. At https://www.ethz.ch/content/dam/ethz/special-interest/baug/igp/photogrammetry-remote-sensing-dam/documents/pdf/ACRS2014_Lanaras_et_al.pdf. (accessed, 15 Sept. 2016).

Lanaras, C., Baltsavias, E. and Schindler, K., 2015. Hyperspectral super-resolution by coupled spectral unmixing. In: Proc. IEEE ICCV, 7-13 Dec., Santiago de Chile. At https://www.ethz.ch/content/dam/ethz/special-interest/baug/igp/photogrammetry-remote-sensing-dam/documents/pdf/lanaras_et_al_iccv15.pdf (accessed, 15 Sept. 2016).

Schaepman, M.E., et al., 2015. Advanced radiometry measurements and earth science applications with the airborne prism experiment (APEX). *Remote Sensing of Environment* 158, pp. 207–219.

Simoes, M., Bioucas-Dias, J., Almeida, L. and Chanussot, J., 2015. A convex formulation for hyperspectral image superresolution via subspace-based regularization. *IEEE TGRS* 53(6), pp. 3373–3388.

Wang, T., Yan, G., Ren, H. and Mu, X., 2010. Improved methods for spectral calibration of on-orbit imaging spectrometers. *IEEE TGRS* 48(11), pp. 3924–3931.

Wei, Q., Bioucas-Dias, J., Dobigeon, N. and Tournier, J.-Y., 2015. Hyperspectral and multispectral image fusion based on a sparse representation. *IEEE TGRS* 53(7), pp. 3658–3668.

Wycoff, E., Chan, T.-H., Jia, K., Ma, W.-K. and Ma, Y., 2013. A non-negative sparse promoting algorithm for high resolution hyperspectral imaging. In: Proc. IEEE International Conference on Acoustics, Speech and Signal Processing, 26-31 May).

Yasuma, F., Mitsunaga, T., Iso, D. and Nayar, S., 2008. Generalized Assorted Pixel Camera: Post-Capture Control of Resolution, Dynamic Range and Spectrum. *IEEE TIP* 19(9), pp. 2241- 2253.

Yokoya, N., Yairi, T. and Iwasaki, A., 2012. Coupled nonnegative matrix factorization unmixing for hyperspectral and multispectral data fusion. *IEEE TGRS* 50(2), pp. 528–537.

Yokoya, N., Mayumi, N. and Iwasaki, A., 2013. Cross-calibration for data fusion of EO-1/Hyperion and TERRA/Aster. *IEEE JSTARS* 6(2), pp. 419–426.