


ISBN 978-82-326-6173-2 (printed ver.)
ISBN 978-82-326-6818-2 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (electronic ver.)



Doctoral theses at NTNU, 2021:225

João Fortuna

Hyperspectral Remote Sensing

Instrument design, field campaigns and
data analysis

Doctoral theses at NTNU, 2021:225

NTNU
Norwegian University of
Science and Technology
Thesis for the degree of
Philosophiae Doctor
Faculty of Information Technology
and Electrical Engineering
Department of Engineering Cybernetics

João Fortuna

Hyperspectral Remote Sensing

Instrument design, field campaigns and data analysis

Thesis for the degree of Philosophiae Doctor

Trondheim, June 2021

Norwegian University of Science and Technology

Faculty of Information Technology

and Electrical Engineering

Department of Engineering Cybernetics



Norwegian University of
Science and Technology

NTNU

Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Information Technology
and Electrical Engineering
Department of Engineering Cybernetics

© João Fortuna

ISBN 978-82-326-6173-2 (printed ver.)
ISBN 978-82-326-6818-2 (electronic ver.)
ISSN 1503-8181 (printed ver.)
ISSN 2703-8084 (electronic ver.)

ITK-report: 2021-3-W



Printed by Skipnes Kommunikasjon AS

*“I think we’ve got enough information now, don’t you?
–All we have is one "fact" you made up.
–That’s plenty. By the time we add an introduction, a few illustrations, and a
conclusion, it will look like a graduate thesis.”
– Calvin and Hobbes*

Summary

Hyperspectral remote sensing is still a young field of research, but with a lot of interest, justified by its promising ability to detect and quantify targets. Recent developments in camera technology and single-board computers have allowed for the appearance of extremely capable hyperspectral imaging systems, with low cost and small footprints.

This thesis explores several topics of relevance for hyperspectral remote sensing, always with a strong practical component. It is my intent that someone can use this document as a guide to their journey into the world of spectral image sensing, analysis, and understanding. The main contributions presented are:

Low-cost hyperspectral instrument design: A DIY instrument based on low-cost components, and additive manufacturing technology is described. Such instruments can be valuable both as research and education tools.

Integration of remote sensing instruments on UAVs: A practical guide focusing on a lightweight hyperspectral imaging payload with a push-broom imager, GPS, and an Inertial Measurement Unit (IMU) as well as data synchronization and acquisition systems.

Algorithms for data enhancement: First, a method for separating the effect of shadows (de-shadowing) and other partially known lighting condition changes from the actual targets of the analysis such as the effects of the physical, chemical or biological properties of the ground, which are of interest. Second, a method for fusing co-registered high spatial and low spectral resolution image data – e.g., RGB – with low spatial and high spectral resolution data – Hyperspectral. This is possible by exploiting the overlap in observed phenomena by the two cameras to create a model through least square projection.

Algorithms for understandable data compression: A novel method and software system for rational handling of time series of multichannel measurements.

Summary

This quantitative learning tool, the *On-The-Fly Processing* (OTFP), develops reduced-rank bilinear subspace models that summarise massive streams of multivariate responses, capturing the evolving covariation patterns among the many input variables over time and space. Thereby, a considerable data compression is possible without significant loss of useful systematic information.

Contents

Summary	iii
Contents	v
List of figures	ix
List of tables	xiii
Preface	xv
Glossary	xvii
1 Introduction	1
1.1 Hyperspectral Remote Sensing	1
1.2 The Hyperspectral Remote Sensing Panorama	3
1.3 Scope of work	7
2 Hyperspectral instrument design: <i>Do it yourself hyperspectral imager for handheld to airborne operations</i>	11
2.1 3D printing	11
2.2 Collimator-slit-front-optics assembly	12
2.3 Detectors	16
2.4 Alignment and calibration	16
2.5 Multi-purpose examples	19
2.6 Concluding remarks	26
3 Preparing the imager for field tests: <i>A lightweight payload for hyperspectral remote sensing using small UAVs</i>	29
3.1 System overview	29
3.2 Experimental results	34
3.3 Future work	35

3.4	Conclusions	36
4	Correcting for environmental disturbances: <i>Multivariate data modelling for de-shadowing of airborne hyperspectral imaging</i>	37
4.1	Introduction	37
4.2	The informative converse paradox	38
4.3	Method	40
4.4	Dataset	43
4.5	Results	44
4.6	Conclusions	45
5	Enhancing hyperspectral data: <i>Multivariate image fusion: A pipeline for hyperspectral data enhancement</i>	47
5.1	Introduction	48
5.2	Method - multivariate image fusion (MVIF)	50
5.3	Datasets	56
5.4	Results	58
5.5	Discussion	60
5.6	Conclusions	64
5.A	Companion methods	64
6	Dealing with big data streams: <i>On-The-Fly Processing of continuous high-dimensional data streams</i>	71
6.1	Introduction	72
6.2	System overview	76
6.3	Datasets	81
6.4	Results and discussion	82
6.5	Comparison with classical PCA	92
6.6	Discussion	93
6.7	Conclusions and perspectives	94
7	Hyperspectral time series analysis: <i>Hyperspectral image data streams interpreted by modelling known and unknown variations</i>	97
7.1	Introduction	98
7.2	Methods	100
7.3	Results	111
7.4	Conclusions	124
8	To infinity and beyond: <i>Hyperspectral remote sensing and the Small-Sat Lab</i>	127
8.1	Recapitulation	127

8.2	Future scope – The SmallSat Lab	130
8.3	A cubesat for ocean colour remote sensing	131
8.4	Final considerations	138
	References	141

List of figures

2.1	3D Model of transmission grating holders.	13
2.2	Optical diagram.	14
2.3	Exploded view of Hyperspectral Imager (HSI) components.	14
2.4	Model 2 snap together transmission grating holder.	15
2.5	Assembled prototypes.	17
2.6	Spectrogram.	19
2.7	Wavelength and sensitivity calibration of 3D printed Hyperspectral Imager.	20
2.8	Airborne data collection experiment.	22
2.9	Comparison of data from NDVI and HSI instruments.	23
2.10	Sample images from the 3D printed push-broom Hyperspectral Imager.	25
2.11	Handheld gimbal operation of the push-broom Hyperspectral Imager.	26
3.1	Prototype payload.	30
3.2	Data connection diagram.	30
3.3	Power distribution diagram.	30
3.4	Flash electrical wiring.	32
3.5	Fluorescent light spectrum.	33
3.6	Hyperspectral and RGB data collected during flight.	35
3.7	Spectra of interest.	36
4.1	The informative converse paradox.	39
4.2	AVIRIS dataset f140530t01p00r16 with region of study marked.	43
4.3	RGB reconstruction of original data.	44
4.4	Estimated spectra of known lighting effects (<i>S</i>): a constant "baseline" and a shadow signature (here estimated from sunny vs shaded mountain sides).	44
4.5	De-shadowed data, given as RGB.	45
4.6	Shadow data, given as RGB. 10 principal components were used for the estimation of the de-shadowed data.	45
4.7	Clustering in original data.	46

4.8	Clustering in de-shadowed data.	46
5.1	Example satellite images simulated with different resolutions.	48
5.2	Multivariate Image Fusion exploits the relative strengths of RGB and Hyperspectral data.	51
5.3	Pipeline overview diagram.	53
5.4	HICO dataset H2011145084342 with marked region of study.	57
5.5	RGB rendering of HICO scene used during tests.	57
5.6	RGB image from ZenMuse camera and RGB rendering of HS data from Hopavågen tests.	58
5.7	Non-negative Matrix Factorization of low resolution Y_L (HICO).	61
5.8	Non-negative Matrix Factorization of high resolution \hat{Y}_H (HICO).	61
5.9	Non-negative Matrix Factorization of low resolution Y_L (HSI).	62
5.10	Non-negative Matrix Factorization of high resolution \hat{Y}_H (HSI).	63
5.11	Indep. Component Analysis of low resolution residuals E_L (HICO).	64
5.12	Indep. Component Analysis of low resolution residuals E_L (HSI).	65
5.13	Noise level per band of low resolution HICO data.	66
5.14	RGB sensor simulated spectral sensitivity.	67
5.15	Spatial Maps (Scores) of a sample of factors from the HICO dataset. Higher factors are clearly more noisy than the lower/earlier.	68
5.16	Scree plot showing only the 10 first components of the low resolution residuals from the HICO dataset.	69
5.17	Example Spectral Signatures split into smooth or non-smooth according to our classifier.	70
5.18	Votes from each relevance classifier for the HICO dataset.	70
6.1	Schematic representation of the OTFP algorithm.	77
6.2	Vis-NIR light absorbance spectra from the B-Z reaction at three different points in time.	84
6.3	Vis-NIR light absorbance spectra from the B-Z reaction.	85
6.4	Hyperspectral NIR images – Compressed orange image	86
6.5	Hyperspectral NIR images – Modelling of orange image	87
6.6	Hyperspectral image from a push-broom camera installed on a flying drone.	89
6.7	Industrial process data.	91
6.8	Industrial process data analysis.	91
6.9	Hyperspectral image from a push-broom camera installed on a flying drone – PCA vs. OTFP.	93
7.1	Experimental setup.	100

7.2	Overview of experimental data acquisition and modelling of hyperspectral video.	102
7.3	Modelling the known: Spectral and temporal structure of the parameters from Extended Multiplicative Signal Correction (EMSC).	106
7.4	Weight vector used to assign weights to different wavelength regions during EMSC and OTFP.	108
7.5	Average moisture during the drying process.	112
7.6	Hyperspectral imaging input spectra.	114
7.7	Modelling the known: Spatial structure of EMSC parameters.	117
7.8	Modelling the unknown: Spectral and temporal structure of the parameters from adaptive bilinear modelling in the OTFP implementation.	119
7.9	Modelling the unknown: Spatial structure of OTFP parameters.	120
7.10	Variation at the different wavelengths explained by the sequence of modelling steps.	121
7.11	Kinetic modelling of the hyperspectral video developments.	123
8.1	Phytoplankton blooms observed from space.	132
8.2	Main components in a multiagent marine observation system.	133
8.3	Mapping of capabilities of various sensor platforms used for ocean colour.	135
8.4	Concept of operations for SmallSat.	136
8.5	Simplified data processing and control architecture.	137

List of tables

1.1	Remote sensing data processing levels.	2
2.1	Technical specification of the Collimator-slit-front-optics assembly. . .	15
2.2	Prototypes instrumental parameters.	16
3.1	Experiment parameters.	35
5.1	Performance indices of MVIF with HICO dataset.	59
5.2	Performance comparison in Moffett field dataset.	60
6.1	Vis-NIR light absorbance spectra from the B-Z reaction: values of the compression quality indices.	83
6.2	Hyperspectral NIR images: values of the compression quality indices.	88
6.3	Hyperspectral image from a push-broom camera installed on a flying drone: values of the compression quality indices.	88
6.4	Industrial process data: values of the compression quality indices. . .	92

Preface

This thesis is submitted in partial fulfilment of the requirements for the degree of Philosophiae Doctor (PhD) at the Norwegian University of Science and Technology (NTNU).

I have done this work at the Centre for Autonomous Marine Operations and Systems (NTNU AMOS) and the Department of Engineering Cybernetics. My supervisors have been Thor Inge Fossen, Tor Arne Johansen and Harald Martens.

During a period of my time at NTNU I was also employed at Idletechs.

This work was supported by the Norwegian Research Council (grant no. 223254) through the Centre of Autonomous Marine Operations and Systems (NTNU AMOS) at the Norwegian University of Science and Technology, the MASSIVE project (grant no. 270959), as well as the Norwegian Space Center.

Acknowledgements

This journey started in 2012, two years before I moved to Norway, when I first became a researcher at the Underwater Systems and Technology Lab (LSTS) at the University of Porto (FEUP). Being able to work with robotics in the field, be it aerial, surface or underwater, allowed me to develop my engineering skills, solving problems under stress, but also to experience the sweet taste of "mission accomplished" at the end of the day.

When a young, fresh out of high school, João gave up the dream of becoming a biologist, he could never have imagined he would end up working so closely with marine biologists and oceanographers aboard research vessels, or even neuroscientists looking at zebrafish brain activity.

Everyone has an "engineering problem", and that is something I grew very interested in solving. More than a decade later, I can definitely say I made the right choice.

It takes two to tango, but to get where I am today, it took a couple more...

I couldn't have made it this far without my supervisors, professors Thor Inge Fos-

sen, Tor Arne Johansen, and Harald Martens. Your guidance allowed me to get the most I could out of my PhD experience, more than exceeding my expectations. Thank you for letting me wander – just take a look at the breadth of topics in this thesis – but making sure I did not get lost. After all, I did get to the end, eventually! Thor Inge, I know this wasn't what you signed up for back in 2014, or me, for that matter, but even after I switched tracks on my research topic in 2015 you supported my decision. Tor Arne, to this day, I'm still not 100% convinced you don't secretly have a clone to get all that work done... Harald, I can't honestly say you were just a supervisor, you truly were a mentor both in and out of university.

My time in Trondheim brought me away from home, from family. It wouldn't have been the same without my "international family", my friends at NTNU, board game afternoons, *hytte* weekends, ski trips (that I just wasn't good at...), wednesday meetings, *julebords*, *pinnekjott*, *ribbe*, *akevitt*, and all the other Norwegian, Danish, Italian, German, Polish, American, Greek, Iranian et al., little things you shared with me. You made me more of a citizen of the world.

Thank you Kristian and Frederik, who planted the idea of a PhD when you visited Porto and then became my office mates, along with Kim, later Andreas and, honorarily, Mikkel. More than office mates, you were good friends.

Furthermore, coffee, or often tea breaks would simply not be the same without Siri, Konstanze and our philosophical discussions about every and anything. Thank you for all those moments!

It was only possible to do all this with the help of the team at the UAV-Lab, with special thanks to the pilots Lars and Pål, and the team at the SmallSat Lab, both of which I proudly say I was a part of. You made sure I didn't spend too much time locked in my office!

Collaboration is key when you want to do as many different things as I did. In that regard, I have to thank Fred Sigernes at The University Centre in Svalbard (UNIS), whom I had the great opportunity to visit and discuss hyperspectral camera designs with; Ingunn Burud at the Norwegian University of Life Sciences (NMBU), for introducing me to the world of spectral imaging; Geir and Axel at the Trondhjem Biological Station (TBS), for providing the scientific meaning and purpose to our engineering; Jo Arve and Alberto, for another opportunity to leave the office and join you on field trials of the Autonaut and the Otter; and of course, João Sousa and the team at the Underwater Systems and Technology Lab (LSTS) at the University of Porto (FEUP), not only for getting me hooked on marine robotics, but also for making me feel back at home any time I visited.

To my parents, Rosário and Luís, my brother, Pedro, and my family, for always encouraging me to take on new challenges, and believing I would be able to tackle them. And finally, I could not have pulled this off without the support of Joana, my "dance partner", especially during the home stretch, putting up with me during a long quarantine.

Glossary

2D/3D Two-/three-dimensional.

CubeSat A type of miniaturized satellite for space research that is made up of multiple cubic units – known as "Us" – of 10 × 10 × 10cm size.

Detector technologies Different detector (sensor) technologies exist, appropriate for different spectral ranges.

Charge-coupled device (CCD) and more recent Complementary metal-oxide-semiconductor (CMOS) sensors are silicon based. Their working spectral range is mostly VIS and a small portion of NIR: 400 to 1000nm. The difference between CCD and CMOS will not be discussed here.

Indium Gallium Arsenide (InGaAs) detectors are employed in NIR spectroscopy, typically at 1000 to 1700nm.

Diffraction grating An optical component with a periodic structure that splits and diffracts light into several beams travelling in different directions. Gratings may be of the reflective or transmissive type, analogous to a mirror or lens, respectively.

EO Earth observation is the gathering of information about the physical, chemical, and biological systems of the planet Earth. It can be performed via remote sensing technology and by ground-based techniques.

FWHM Full Width at Half Maximum. When applied to spectrometer resolution, is the difference between the two wavelengths at which the radiation intensity is equal to half of its maximum value, for a selected spectral band. In other words, it is the width of a spectrum curve measured by each spectral channel or band. A lower (narrower) value is associated with a more selective – higher spectral resolution – system.

HSI Hyperspectral imaging (system).

Hyper-/multispectral imaging Hyper- and multispectral imaging capture a scene in both spatial (position in space) and spectral (wavelength) dimensions, where each sampled spatial position contains multiple intensity values corresponding to different spectral bands. Going from multispectral to hyperspectral imaging (HSI), the spectral sampling is of much higher density, enabling the reconstruction of continuous spectra. This data of higher dimensionality allows both to identify materials or objects using the spectral data and to identify their location using the spatially resolved image plane, usually exhibiting a 2D spatial dimensionality.

IMU Inertial measurement unit. An electronic device that measures and reports a body's specific force, angular rate, and sometimes the orientation of the body, using a combination of accelerometers, gyroscopes, and sometimes magnetometers.

Ocean colour The "colour" of the ocean is determined by the interaction of incident light with substances or particles present in the water. The term is used in the context of remote-sensing observations. Using specialised sensors, such as hyperspectral cameras, one can quantify radiation emerging from the ocean with high spectral resolution. These measurements can be used to infer information such as phytoplankton biomass or concentrations of other living and nonliving materials that modify the characteristics of the incoming radiation. Monitoring the spatial and temporal variability of ocean colour over large regions up to the scale of the global ocean has been instrumental in characterising the dynamics of marine ecosystems, and is a key tool for research into how marine ecosystems respond to climate change and anthropogenic perturbations.

Remote sensing The acquisition of information about an object or phenomenon without making physical contact with the object and thus in contrast to on-site observation.

RGB Red Green Blue, a term used to refer to colour images and cameras.

SNR Signal-to-noise ratio is a measure that compares the level of a desired signal to the level of background noise.

Spectral ranges Electromagnetic radiation can be classified into different ranges, according to the wavelength. A commonly used scheme is:

UV Ultraviolet (200 - 400nm);

VIS Visible (400 - 750nm);

NIR Near-infrared (750 - 1400nm);

SWIR Short-wavelength infrared (1400 - 3000nm);

MWIR Mid-wavelength infrared (3000 - 8000nm);

LWIR Long-wavelength infrared (8000 - 15000nm), "thermal imaging" region or "thermal infrared".

UAV/UAS Unmanned Aerial Vehicle/System, commonly known as drone.

Chapter 1

Introduction

In this first chapter, I will lay out a short introduction to the topics that build up the foundation for this dissertation. With those concepts in mind, I go on to explore the current challenges in this field and the state of the art.

1.1 Hyperspectral Remote Sensing

Let me start by breaking down the title: "Hyperspectral Remote Sensing". What is remote sensing? What does hyperspectral mean?

1.1.1 Remote Sensing

The first term I should define is "remote sensing": to determine the bio/geo/chemical properties of a material or scene without physical contact. Usually, it is contrasted and complemented by *in-situ* measurements – more precise but limited in coverage.

Perhaps the most well-known examples of remote sensing data are aerial and satellite imagery. In addition to stunning views, the systematic collection of data over time and space allows for the monitoring of many atmospheric and ground variables. These variables give us an insight into the dynamics of phenomena transforming the Earth, either of natural – e.g., biological, geological – or human origin – e.g., political, economical.

Remote sensing is particularly interesting when compared to more localised data acquisition techniques because it allows gathering data with such a large coverage, one that is possible when using air- or space-borne sensing platforms. Furthermore, remote sensing data can provide information from areas that are not accessible for *in-situ* measurements.

Remote sensing data can have varying degrees of complexity and processing re-

quired after being acquired. While some represent direct physical characteristics measured by the sensors, e.g., reflected light intensity at 400nm by a certain area of the ocean, others are derived attributes, e.g., chlorophyll concentration on that same area. To help the remote sensing community better organize in terms of these different types of data, NASA's Earth Science Data Systems Program proposed a system of data processing levels – see Table 1.1.

Table 1.1: Remote sensing data processing levels. Adapted from [1].

Level	Description
0	Reconstructed, unprocessed instrument data at full resolution, with communications artifacts removed.
1A	Time-referenced Level 0 data, annotated with ancillary information (appended but not applied), including radiometric and geometric calibration coefficients and georeferencing parameters.
1B	Level 1A data processed to sensor units (e.g., radar back-scatter cross section, or brightness temperature).
2	Derived geophysical variables (e.g., ocean wave height, soil moisture, ice concentration) at the same resolution and location as Level 1 source data.
3	Variables mapped on uniform space-time grid scales, usually with some completeness and consistency (e.g., missing points interpolated, or complete regions mosaicked together from multiple orbits).
4	Model output or results from analyses of lower level data (e.g., variables that were not measured by the instruments but instead are derived from these measurements).

1.1.2 Hyperspectral Imaging

Imaging spectrometers, also known as hyperspectral cameras or imagers, are some of the instruments used in today's remote sensing satellites, and they are designed to measure spectral data across one or more spatial dimensions – pixels.

Information extracted from those data is invaluable for a growing number of fields. Hyperspectral (HS) instruments provide a richness of data that enables the classification and detection of such resources through passive and nondestructive measurements. Information on the distribution and abundance of natural resources is vital for science, education, policy-making, and management alike [2]. From agriculture [3] to ocean observation [4], aquaculture or even city planning [5], these

are just a few of the many areas currently using remote sensing hyperspectral data products to help in their work.

Generally speaking, an imaging spectrometer is composed of two main elements:

- The optical dispersion or filtering component, which allows us to split the total signal into several spectral bands.
- The electro-optical element, or sensor, which allows us to quantify each band and represent it as a digital value.

In terms of sensor technology, the most common electro-optical elements, i.e., the components that translate optical (light) signals into electrical (and then digital) signals, are Charge Coupled Device (CCD) or Complementary Metal Oxide Semiconductor (CMOS)[6] 2D pixel arrays, much like the ones present in conventional cameras. At each instant, it is possible to capture a 2D image, or a frame, which depending on the technique, can contain spatial, spectral, or both kinds of information.

The number of pixels in each frame is dependent on the size of the sensor and of the pixels. Larger pixels collect more photons, the signal we want to measure, leading to a higher signal-to-noise ratio (SNR). Smaller pixels, on the other hand, allow for a higher pixel density, so more pixels in the same sensor area. The number of pixels in each frame will set the resolution of all non-scanning dimensions in the datacube. The contents of each frame will depend on the optical elements used to sort the spectral information. To finally generate the 3D hyperspectral datacube, some digital postprocessing is required to "fold" the sequence of 2D frames.

1.2 The Hyperspectral Remote Sensing Panorama

1.2.1 Instrument Design

The push-broom hyperspectral technique has been used for the last three decades [7, 8] in industrial, airborne and satellite applications [9–12].

This method records a perpendicular cross-section of any target along the track direction. In spectroscopy, the cross-section area corresponds to the entrance slit image as a function of wavelength, forming an image known as a spectrogram. The arrival of electronic image sensors such as the CCD has made it possible to stack spectrograms at high speed to form a spectral datacube as we sweep the target. The number of reconstructed images from the datacube is only limited by the spectral bandpass and the range of the instruments. The ability to form images as a function of wavelength is a powerful tool in remote sensing applications like online spectral identification or target detection and classification.

However, the push-broom technique requires line scanning of the target with high precision. For example, on an airborne or satellite platform, the instrument needs

to be stabilized in attitude by a 3-axis gyro to sweep the target without spatial distortions along the direction of flight. Low cost Micro Electro Mechanical Systems (MEMS) Inertial Measurement Units (IMUs) with integrated 3-axis gyroscopes, accelerometers, and magnetometers are found to be accurate enough for unmanned aerial vehicle navigation [13] and for motorized gimbal stabilization.

Other hyperspectral techniques have emerged over the last decade, e.g., to apply Liquid Crystal Tunable Filters (LCTFs) [14, 15], and MEMS Fabry-Perot tunable interference filters [16] with direct imaging capability where no line scanning is necessary. The centre wavelength of such filters is controlled electronically. A drawback of these techniques is the low speed of spectral sampling, which can lead to a shift of the target position during acquisition. To minimize that effect, a precise attitude correction is required.

Recent developments in drone technology, CMOS image sensors, gimbals based on MEMS IMUs, and 3D printing open up new possibilities for the construction of low cost push-broom Hyperspectral Imagers (HSI).

1.2.2 Airborne Systems

Airborne hyperspectral systems are currently used in fields such as agriculture, environmental monitoring, or infrastructure inspection. For such applications, the spectral range of the sensors is usually 400-1000nm, commonly referred to as VNIR (visible and near-infrared).

Some manufacturers of hyperspectral sensors provide complete systems that can fit in a small unmanned aerial vehicle (UAV):

- Specim's AisaKESTREL10 [17] offers a 400-1000nm range, 2048 pixels spatial resolution at <5kg total system weight;
- Resonon's Airborne Hyperspectral Imaging System [18] can be customized with their different cameras. When using the Pika L model, the spectral range is also 400-1000nm, with 900 spatial pixels at <1.5kg for the total system;
- Headwall's Nano-Hyperspec sensor [19] with the same spectral range of 400-1000nm, spatial resolution of 640 pixels and <0.52kg without lens.

The Unmanned Aerial Vehicles Laboratory (UAV-Lab) at NTNU is focused on developing small, low cost, and open-source tools while still aiming at high performance.

1.2.3 Data Processing and Analysis

Correcting for environmental disturbances

Outdoors imaging in general, and airborne in particular, suffer from variations in apparent colour caused by illumination and not the reflecting object itself. While this variation may be difficult to correct in RGB images, hyperspectral sensors offer a wealth of redundant data that can be used to remove the effects of different light sources.

Enhancing Hyperspectral data

Using multiple sensors with different capabilities often creates a clearer picture of the environment when compared to a single sensor scenario. However, while our brains are good at fusing information from different sources, some work is required if the process is to be automated. Multivariate calibration, a term coined in the field of chemometrics, refers to the development of models to explain the properties of interest by combining different variables from multi-channel sensor measurements.

Fusing images from sensors with different spectral and spatial properties to generate a single, improved data product is a known and studied problem [20, 21]. Hyperspectral Image Super-Resolution methods can be grouped into four categories: Bayesian based approaches [22–28]; Tensor based approaches [29–34]; Matrix factorization based approaches [35–50] and; Deep Learning based approaches [51–58].

The mentioned methods assume that the images to be fused are co-located (registered). However, more recent methods drop that assumption and achieve simultaneous registration and super-resolution [59, 60].

Dealing with big data streams

Many modern measurement technologies generate massive amounts of data in a very short time – e.g., continuous streams of high-dimensional data via one-step analytical procedures¹. For instance:

- Modern spectrometers can deliver hundreds of informative, high-dimensional spectra per second.
- Hyperspectral cameras produce multivariate spatially resolved images. In addition, when configured in a time-lapse mode, they can yield continuous streams of high-dimensional spatiotemporal recordings.

¹Contrary to unstructured data from, e.g., free text, they are systematically recorded and are here referred to as quantitative data.

- Industrial monitoring for condition-based maintenance, as well as the control of complex dynamic processes, requires high-dimensional inputs to be sufficiently informative.
- Computer experiments, needed to study the behaviour of complex mathematical models, involve advanced workstations performing thousands of simulations, each one possibly characterized by just as many input and output properties.

Hence, a measurement revolution (recently termed *data tsunami* [61]) is currently taking place in numerous fields of applied science, ranging from analytical chemistry and medicine to environmental surveillance, informatics, and industrial *Internet of Things* (IoT). However, these incredibly quick advances run the risk of being practically useless for three reasons:

- The human ability to grasp the content of interest from data remains fairly constant, and data simplification is therefore desirable for interpretative purposes. Here, one possible solution could be the removal of irrelevant descriptors among the available ones. Nevertheless, for most applications their identification is not straightforward, which makes such a simplification risky and complicated.
- IoT threatens to flood both communication channels and users' cognitive capacity with overwhelming torrents of repetitive, more or less redundant data.
- Traditional computing systems are generally not capable of performing analytics on constantly streaming data, typical of today's world of multimedia communication [62].

In a scenario like this, if it were possible to simultaneously compress and model high-dimensional measurement series as they flow from, e.g., an analytical platform and without significant loss of useful information content, their storage, transfer, retrieval, visualization and interpretation would be radically eased.

Hyperspectral time series analysis

Hyperspectral time series and other multichannel spatiotemporal spectral measurement processes give overwhelming streams of Quantitative Big Data. The raw data are nonselective, in the sense that they are affected by many different sources of variation – variations in sample physics (e.g. light scattering) and sample chemistry (composition) as well as variations in the light source and the camera (systematic errors). To identify, separate, quantify and interpret the various sources of information in such data streams is thus a challenge. It is important to quantify phenomena already known for the given type of measurements, as well as

also discovering, quantifying, and displaying unexpected, but systematic variation patterns in the data.

1.3 Scope of work

This thesis documents my journey through the field of hyperspectral remote sensing, a journey that is by no means an extensive exploration of all there is to know about this field.

Following this introductory chapter, I will discuss the design of a hyperspectral camera in chapter 2, then move on to its integration with a small UAV for field testing in chapter 3. After that, I will look into data processing methods in the remaining chapters. Finally, in the last chapter, I bring all concepts together and introduce the work being done at the NTNU smallSat Lab.

1.3.1 Contributions

The aforementioned chapters are adapted from publications I co-authored, which are listed below chronologically.

Multivariate data modelling for de-shadowing of airborne hyperspectral imaging (2017)

J. Fortuna and H. Martens, “Multivariate data modelling for de-shadowing of airborne hyperspectral imaging,” *Journal of Spectral Imaging*, vol. 6, 2017.

In which we present an extension to the Informative Converse (IC) theorem, allowing for a simplified interpretation of hyperspectral images through de-shadowing. By imposing certain assumptions on the distribution of spectral information, we achieve an improved separation of light source variations from ground property variations.

This research was first presented at the International Association for Spectral Imaging (IASIM) 6th international conference in spectral imaging, in July 2016.

On-The-Fly Processing of continuous high-dimensional data streams (2017)

R. Vitale, A. Zhyrova, J. Fortuna, O. de Noord, A. Ferrer, and H. Martens, “On-The-Fly Processing of continuous high-dimensional data streams,” *Chemometrics and Intelligent Laboratory Systems*, vol. 161, 2017.

In which we present a novel method for compression and interpretation of streaming data. The On-The-Fly Processing (OTFP) is based on an evolving implementa-

tion of Principal Component Analysis (PCA) that updates, when necessary, both the preprocessing parameters and the principal components structure (whose changes and possible expansion can be optionally monitored in real time through intuitive graphical displays) as data arrives.

This method was patented and is a core technology of Idletechs AS².

Do it yourself hyperspectral imager for handheld to airborne operations (2018)

F. Sigernes, M. Syrjäsuo, R. Storvold, J. Fortuna, M. E. Grøtte, and T. A. Johansen, “Do it yourself hyperspectral imager for handheld to airborne operations,” *Optics Express*, vol. 26, p. 6021, 2018.

In which we present a do-it-yourself recipe for the construction of small low-cost hyperspectral push-broom imagers. Only off-the-shelf optical, mechanical, and electronic components are used. The instruments are assembled by the use of a 3D printer.

A lightweight payload for hyperspectral remote sensing using small UAVs (2018)

J. Fortuna and T. A. Johansen, “A lightweight payload for hyperspectral remote sensing using small UAVs,” in *2018 9th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, 2018.

In which we describe a lightweight system for Hyperspectral sensing, recording, and synchronization with navigation data. Software and hardware developed while preparing this payload system were then used repeatedly on field campaigns and as a blueprint for newer generations of airborne hyperspectral payloads.

Hyperspectral time series analysis: hyperspectral image data streams interpreted by modelling known and unknown variations (2020)

P. Stefansson, J. Fortuna, H. Rahmati, I. Burud, T. Konevskikh, and H. Martens, “Hyperspectral time series analysis: hyperspectral image data streams interpreted by modelling known and unknown variations,” *Data Handling in Science and Technology*, 2020.

In which we experimentally demonstrate a generic way to model a stream of hyperspectral time series data in terms of *a priori* known and unknown constituent

²www.idletechs.com

spectra, essentially enabling large dimensionality reduction of the data without loss of information. An additional benefit of the described methodology, apart from enabling substantial compression of the data, is that it autonomously highlights unidentified systematic spectral variations within the sample being studied, which aids in further exploration and understanding of the underlying chemical and physical processes causing variations.

Multivariate image fusion: A pipeline for hyperspectral data enhancement (2020)

J. Fortuna, H. Martens and T. A. Johansen, “Multivariate image fusion: A pipeline for hyperspectral data enhancement,” *Chemometrics and Intelligent Laboratory Systems*, 2020.

In which we describe a pipeline for enhancing the spatial resolution of hyperspectral data, by taking advantage of co-located RGB data.

Using multivariate dimensionality reduction techniques (like PCA), we find a reduced representation of the high spectral resolution hyperspectral (HS) data. The high spatial resolution RGB data is then resized to the same low spatial resolution of the HS data. When both datasets are reduced to similar sizes across all dimensions, we apply a regression method to find a representation of the reduced HS data through RGB. The RGB high spatial resolution information can then travel the reverse path of the reduced HS, providing a high spectral and spatial resolution dataset.

We provide a template for connecting functional blocks, with the aim of fusing multivariate datasets with different resolutions. This leaves room for changing the reduction and regression methods.

1.3.2 Other work

Cascaded line-of-sight path-following and sliding mode controllers for fixed-wing UAVs (2015)

J. Fortuna and T. I. Fossen, “Cascaded line-of-sight path-following and sliding mode controllers for fixed-wing UAVs,” in *2015 IEEE Conference on Control Applications (CCA)*, 2015.

In which we present a uniform semiglobal exponential stability (USGES) proof of a cascaded Line-of-Sight path-following and sliding mode control system for fixed-wing UAVs. The paper also contains simulation results that show good performance. Later, this algorithm was regularly used during flight campaigns.

It is part of the public repository of the DUNE project: <https://github.com/LST>

1. Introduction

[S/dune/tree/master/src/Control/Path/LOSnSMC](#)

Because its scope is not relevant for the topic of this thesis, this publication is not included.

Chapter 2

Hyperspectral instrument design

Do it yourself hyperspectral imager for handheld to airborne operations

This chapter describes the rapid prototype construction of small and lightweight push-broom Hyperspectral Imagers (HSI)¹. The dispersive element housings are printed by a thermoplastic 3D printer combined with S-mount optical components and commercial off-the-shelf camera heads. Four models with a mass less than 200g are presented with a spectral range from the visible to the near-infrared part of the electromagnetic spectrum. The bandpass is in the range from 1.4-5nm. Three test experiments with motorized gimbals to stabilize the attitude show that the instruments are capable of push-broom spectral imaging from various platforms, including airborne drone to handheld operations.

2.1 3D printing

When companies like MakerBot Industries started to mass produce 3D printers using thermoplastics as material back in 2009, the opportunity to create 3D objects became available to the public. Together with Computer Aided Design (CAD) software rapid prototyping of parts became possible. A number of studies have highlighted the use of 3D printing to produce optomechanical components [63] and spectrometers [64–66]. The latter studies did however not exploit the hyperspectral imaging capability.

The main constructional challenge in spectroscopy is the fact that the dispersive or

¹Based on the paper "Do it yourself hyperspectral imager for handheld to airborne operations", by Fred Sigernes, Mikko Syrjäsuo, Rune Storvold, João Fortuna, Mariusz Eivind Grøtte, and Tor Arne Johansen, published in Optics Express Volume 26, Issue 5, 2018.

refractive element, a grating or prism, bends the input light beam. A machine with multiple axes of freedom is needed in manufacturing the part that holds that element. In fact, a transmitting grating holder would require at least a 4-axes milling machine. A 3D printing machine uses Additive Manufacturing (AM) and can easily build parts with angular or more complex surfaces compared to conventional fabrication methods. In addition, the printing process produces parts by stacking honeycomb layers on top of each other, which gives high strength and low mass. The latter is a clear advantage for low mass payloads aimed at drone operations. Figure 2.1 panel (A) shows a 3D view of a simple transmission grating holder (Model 1). The solid modelling prior to 3D printing was carried out in Tinkercad™, which is an online CAD program by the company Autodesk Inc. The design is based on a $\beta = 19.36^\circ$ wedge using a 600 lines/mm transmission grating with the collimated input beam entering at the bottom. The optical diagram is shown in Figure 2.2. This design's centre output wavelength $\lambda = 552.5$ nm, parallel to the wedge normal, follows the grating equation [67] for the first spectral order ($k = 1$):

$$\lambda = \left(\frac{a}{k}\right) \cdot \sin \beta \quad (2.1)$$

Where a , the groove spacing, is 1666.67nm. The grating is fixed by glue and the collimator lens holder hole is threaded by an M12 \times 0.5 tap to fit S-mount lens holder components.

To avoid both gluing and threading, the holder can be split into two halves that encompass the optical element. Figure 1 panel (B) shows Model 2, which is a snap-fit or push-together design. The grating is half-way embedded into the holder. The same applies to the collimator-slit-front-optics assembly and the detector lens. Due to the fact that thermoplastic materials tend to shrink during 3D printing, there is a need to scale up all parts 1-2% of their original designed size to make the embedded components frictional snap-fit fixed.

One obvious advantage is that parts may now be replaced without risk of damage and destruction. Nevertheless, the most brittle part is the grating, and to avoid damaging it, the slot height is designed 0.2mm higher than the grating. A too close fit tends to break the grating when pushing the parts together. Thin slices of 0.2mm thick rubber bands are used at the top and bottom of the slot to secure and fix the grating. Two straight through 45mm long machine screws of diameter 3mm hold the design together. All parts are frictional press locked and require no glue or epoxy when assembled.

2.2 Collimator-slit-front-optics assembly

The grating is collimated using the mix and match S-mount components of the company Edmund Optics (EO) Ltd. The basic optical parts comprise one high pre-

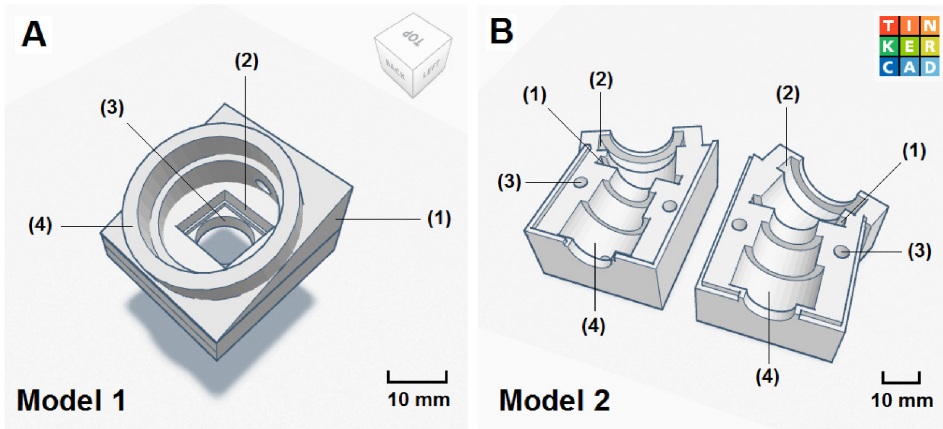


Figure 2.1: A 3D view of transmission grating holders using Tinkercad™. Panel (A) Model 1: (1) 19.36° wedge, (2) square 12.5 × 12.5mm² mount chamber for a 600 grooves/mm transmission grating, (3) 12mm diameter hole for the Collimator-slit-front-optics assembly, and (4) cylindrical 25mm diameter detector lens holder. Panel (B) Model 2: (1) 25 × 25mm² square grating holder, (2) detector lens holder, (3) straight through mount holes, and (4) Collimator-slit-front-optics assembly holder.

cision slit, a collimator lens, and a front lens. This requires three thin lens holders, one focus tube, and one locking ring. An exploded view including the detector lens is shown in Figure 2.3. Details of each part are listed in Table 2.1.

The thin lens holders of EO are designed for 10mm diameter optical components. The useful optical aperture is 9mm. The slit and the field lens are mounted together on one of the thin lens mounts. The slit is in contact with the flat side of the field lens. A second thin lens mount is used for the collimator lens that illuminates the grating. Both lens holders are mounted on an S-mount focus tube. The input f /value to the grating is $f/3.3$, which requires a $f/2.5$ camera lens to capture all the first spectral order ($k = 1$) light diffracted by the grating. The grating is larger than required in order to only use the centre part that is not affected by frictional damage when sliding it into its holder. The separation between the centre of the grating and the centre of the two lenses is only 5mm. The third lens holder is used in front of the slit to provide sufficient thread and focus distance for the front lens and its lock nut. The $f/4$ front lens is the main aperture stop of the system, which prevents field-of-view overfill of the collimator lens.

Figure 2.4 shows the assembly of the Collimator-slit-front-optics, grating and detector lens embedded into one half of the holder. We used the MakerBot's Replicator 2 armed with PLA (Polylactic Acid) thermoplastics filament to print the part.

2. Hyperspectral instrument design: Do it yourself hyperspectral imager for handheld to airborne operations




Figure 2.2: Optical diagram: (1) front lens, (2) entrance slit, (3) field lens, (4) collimator lens, (5) 600 lines/mm transmission grating, (6) detector lens, (7) exit focus plane. Centre detector diffraction angle $\beta = 19.36^\circ$ for wavelength $\lambda = 552.5\text{nm}$ for first spectral order ($k = 1$).




Figure 2.3: Exploded view of Hyperspectral Imager (HSI) components: (1) front lens, (2) lock nut, (3) air slit, (4) field lens, (5) three thin lens mounts, (6) focus tube, (7) collimator lens, (8) transmission grating, (9) detector lens, (10) focus spacers, and (11) C-mount lens adapter.

Table 2.1: Technical specification of the Collimator-slit-front-optics assembly.

Item	Description	Part# (EO)	~Cost [USD]
1	Front lens f/4 Focal Length (FL) 16mm	#83-107	50
2	M12 lock nut for μ -video lenses	#64-102	10
3	Precision air slit $25\ \mu\text{m} \times 3\text{mm}$	#38-558	101
4	Field lens FL = 10mm	#63-519	38
5	3 \times S-mount thin lens mounts	#63-943	81
6	S-mount focus tube	#63-953	49
7	Collimator lens FL = 30mm	#63-523	37
8	600 grooves/mm transmission grating 25mm^2	#49-580	105
9	Detector lens f/2.5 FL = 25mm	#56-776	60
10	S-mount brass spacer rings	#54-461	70
11	C-mount to μ -video lens adapter	#53-675	25
			626




Figure 2.4: Model 2 snap together transmission grating holder: (1) detector lens, (2) $25 \times 25\text{mm}^2$ square 600 grooves/mm transmission grating, (3) 3D printed grating holder, and (4) Collimator-slit-front-optics assembly.

2. Hyperspectral instrument design: Do it yourself hyperspectral imager for handheld to airborne operations

Table 2.2: Prototypes instrumental parameters.

Instr.	Sensor type	Sensor size [mm ²]	Spectral range [nm]	FWHM [nm]	Mass [g]	Sensor cost [USD]
A	Sony Super HAD Colour CCD	4.800× 3.600	435.8 - 733.6	1.4	106	34
B	5M pixel Colour CMOS	4.300× 3.200	434.8 - 701.3	1.4	128	280
C	Monochrome CMOS	11.264× 5.986	297.5 - 1005.5	1.4	152	1300
D	Monochrome CMOS	11.251× 7.032	281.8 - 966.1	5.0	168	920

Slit height magnification at the sensor focus plane is 1.28 due to the field lens in the slit entrance plane and the spectral bandpass is FWHM = 1.4nm (Full Width at Half Maximum). The visible part of the electromagnetic spectrum (400-700nm) illuminates an area in the focus plane of approximately $4.8 \times 3.2\text{mm}^2$, which means that a 1/3" sized image sensor can be used.

2.3 Detectors

Several prototypes have been assembled with different low-cost camera heads. Four of them are shown in Figure 2.5. A miniature colour CCD (Charge Coupled Device) video camera head (Turnigy 600 TVL) used by the Radio Control (RC) community, a CamOne Infinity action camera, and two monochrome industrial CMOS camera heads are tested out as potential detectors.

The main differences between the prototypes are the type of detectors and how they are mounted on the grating house. Key instrumental parameters are summarized in Table 2.2. The optical parts are identical, except for instrument (D) where the slit width is increased to 75 μm and the collimator lens is decreased in focal length to 25mm. The latter is done to increase throughput and sensitivity.

The camera head of instrument (A) is snap fitted by side plates that are added to the 3D print model. A wireless video transmitter holder is also included. For instruments (B) and (C), 3D printed side plates are mounted to press fit the camera heads. The two 45mm long bolts that go straight through the grating holder (Model 2) are used for this purpose. A metal strip is added to the instrument (C) to improve the design. The same technique is used on instrument (D), except that the side plates are aluminum. The camera head is now mounted to the plates by $4 \times 3\text{mm}$ diameter machine screws. The aluminum plates also act as a heat sink to the camera head.

2.4 Alignment and calibration

The instrument is focused by first mounting the camera lens to the detector. Brass spacer rings are used to obtain a fixed back focal length of 17.5mm, required for




Figure 2.5: Assembled prototypes. Panel (A): (1) Collimator-slit-front-optics assembly, (2) Model 1 grating holder with camera side holder plates embedded, (3) Turnigy 1/3" Sony Super HAD CCD camera, and (4) wireless video transmitter. Panel (B): (5) Snap together Model 2 grating holder, (6) camera side holder plates, (7) CamOne Infinity action camera, and (8) 3-axis motorized gimbal stabilizer by Feiyu Tech, Inc. (model MG). Panel (C): (9) uEye UI-3360CP-NIR-GL industrial camera head by Imaging Development Systems (IDS) GmbH, (10) metal strip, and (11) two straight through mounting bolts with nuts. Panel (D): (12) aluminum side plates, (13) IMX174LL CMOS camera head by The Imaging Source, LLC, and (14) USB 3 connector.

2. Hyperspectral instrument design: Do it yourself hyperspectral imager for handheld to airborne operations

C-mount camera heads to be focused at infinity. Secondly, after mounting all parts together, the spectral focus is achieved by tuning the S-mount focus tube. A spectral line source such as a fluorescence tube illuminating a diffuse surface is used to identify spectral lines as sharp as possible. Thirdly, the front lens is focused by identifying sharp structures in the horizontal direction of the spectrograms, which originate from any sharp edged target object located at infinity as seen by the front lens.

A spectrogram of a white paper edge illuminated by an OSRAM fluorescent tube is shown in Figure 2.6. The horizontal edge across the spectrum should be as sharp as possible to obtain focus of the front lens. Note that the edge is sharper at the wavelength centre of the spectrogram. This is believed to be due to chromatic aberrations in the collimator lens. An achromatic lens should be tested in future designs. The spectral lines are also slightly curved and tilted ($\pm 0.45\text{nm}$) due to astigmatism and a small misalignment of the slit angle, respectively. The wavelength calibration is therefore conducted by a polynomial fit of the known emission line peaks in the centre row spectrum of a mercury vapour tube spectrogram. As noted above, any misalignment of the slit and astigmatism effects will produce errors in the order of $\pm 0.5\text{nm}$.

It should be emphasized that the use of colour sensors is not necessary in spectroscopy, but the low cost and availability make them ideal for testing on prototypes. In addition, coloured spectrograms are useful for pedagogic and demonstrational reasons. For further processing, the spectrograms are simply converted to grayscale by adding the response from each colour channel and scaling the net result down to 8-bit. This is not ideal for target intensity work since we do not demosaic the raw response from the Colour Filter Array (CFA) in front of the sensor surface. Instrument (A) and (B) are as a consequence not intensity calibrated.

Instruments (C) and (D) use monochrome CMOS sensors that are also sensitive in the Near-Infrared (NIR) region of the spectrum. However, based on our calibration results, the effective or useful spectral range is only from 400 to 800nm. These instruments are not sensitive to Ultra Violet (UV) light below 400nm due to absorption by the glass lenses, and the second spectral order ($k = 2$) appears or blends in above 800nm. The effect is visualized in Figure 2.7 in the wavelength calibration of instrument (D).

The first-order emission lines of Mercury (Hg) at wavelength 404.7 and 435.8nm reappear above 800nm as second-order lines. This opens the opportunity to increase the readout time by cropping the useful part of each spectrogram, and cutting out the second-order part in NIR and the nonilluminated areas in UV. The latter could also be used to detect the background level of the sensor.

Sensitivity calibration was conducted using a standard certified irradiance tungsten lamp and a Lambertian screen as target [68]. The raw 8-bit count response and source spectrum in units of $\text{mW m}^{-2}\text{nm}^{-1}$ of the screen are shown in Figure

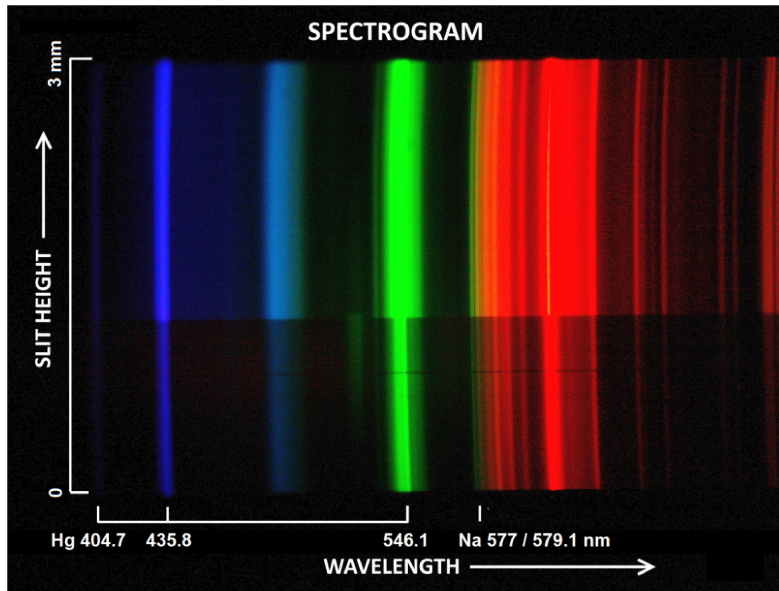


Figure 2.6: Spectrogram from instrument (B). Target is a white paper illuminated by a fluorescent tube (OSRAM FQ 54W/830 HO). The emission lines of mercury (Hg) at wavelengths 404.7, 435.8 and 546.1nm are marked. The doublet at ~ 580 nm is Sodium (Na). The upper part of the spectrogram is the white paper, and the lower part is a light gray coloured surface (office bench).

2.7. Certification above 800nm requires an order blocking filter with cutoff wavelengths above 400nm. A motorized filter wheel could be installed in front of the grating to handle this issue, but it would complicate the design, introduce moving parts, and add mass. Another solution would be to sacrifice spectral resolution and use a prism instead of the grating with no overlapping spectral orders.

2.5 Multi-purpose examples

The following three test applications are setup to demonstrate that the constructed HSIs are capable of forming multispectral images where the target objects are visually identified. Lessons learned are also included.

2.5.1 Octocopter drone experiment

The Octocopter was equipped with a motorized 2-axis gyro platform for the payload. The gyro stabilized pitch and roll attitude angles during the flight. Note that the yaw or heading of the payload was not stabilized. In addition to the HSI -

2. Hyperspectral instrument design: Do it yourself hyperspectral imager for handheld to airborne operations




Figure 2.7: Wavelength and sensitivity calibration of 3D printed Hyperspectral Imager (HSI) - instrument (D). Panel (A): The spectra are sampled from the centre horizontal row of the detector. The gain was set to zero. The blue spectrum is from a Mercury (Hg) vapour tube supplied by Edmund Optics Ltd. (SN K60-908). The red curve represents the spectrum of a fluorescent tube (OSRAM FQ 54W/830 HO). Each mercury emission line is marked according to wavelength and spectral order k . The green spectrum is a 30 second exposure of a Lambertian screen (Lab-sphere SRT-99-180) illuminated by a 1000W Tungsten lamp (ORIEL SN7-1275) located 8.54m away from the screen. Black coloured spectrum is the irradiance of the screen in absolute units of $\text{mW m}^{-2} \text{nm}^{-1}$. Panel (B): The spectrogram of the fluorescent tube.

instrument (A), a Normalized Difference Vegetation Index (NDVI) camera was installed to map the terrain and the vegetation [69]. The camera is a modified Canon PowerShot SX260. The original CFA is replaced by a Near Infrared (NIR) sensitive CFA by the workshop MaxMax.com - Llewellyn Data Processing (LDP LLC). The NDVI camera stored 1/1000 second exposures on an internal memory card every 2 seconds. Each image was tagged with GPS position data. The video from the HSI was transmitted and received by a video recorder at ground level. The exposure time was set to 1/50 second. The live stream was 25 frames per second. The automatic gain of the camera was enabled.

The main objective of the experiment was to determine whether the drone and the gyro together could act as a line scanning platform for the HSI instruments. A stable attitude flight where the image in the slit entrance plane is moved parallel along the slit height axis and perpendicular to the flight direction, should in theory

sample the target in a linear manner without spatial disturbance.

A geo-corrected orthomosaic image constructed from the NDVI image sequence is used to compare and detect similar spatial features obtained from the hyperspectral datacube. The mosaic is visualized in Figure 2.8. Five white diffuse coloured surfaces with a square area of 1m^2 are used as ground control points for calibration of the image position. Each control point position is measured by a Differential Global Positioning System (DGPS) prior to the airborne campaign. Geolocation and rectification were carried out by the students as part of their training.

The pilot operated the drone close to the white van seen at the parking lot to the left in the middle of Figure 2.8. The flight pattern is shaped like a fan. Yellow boxes mark the sampling area of the HSI. For each track, the drone was flown out into the field in different directions in order to return with low velocity back towards the position of the pilot. The pilot manually controlled the aircraft in altitude hold mode. More advanced intelligent modes like automatic GPS waypoint navigation were then not an option. The altitude ranged from 100 to 300m with ground track speed velocities from 1.7 to 3.1m/s. Note that the weather conditions were not optimal due to wind gusts up to 5m/s, especially at the high altitude range. As a consequence, data from only 4 out of 25 tracks were selected for further analysis based on stable flight performance of the drone.

The comparison between the NDVI and the HSI camera is shown in Figure 2.9. The NDVI images are first scaled down to 30% of their original size. The HSI images are colour RGB composites. Each colour channel is constructed by stacking the accumulated response in wavelength across a vertical sliced box of the video frames (spectrograms). The position and width of the box define the centre wavelength and the image bandpass, respectively. The blue, green and red channels of the composite HSI images are centred at 470, 550, and 630nm, all with a bandpass of 10nm. Secondly, the colour images are scaled down to match the spatial scale of the NDVI images. The adjustment of the scale was done by visual examination of the resulting match.

It is clear that the HSI is capable of reproducing the same features and objects that the NDVI camera detects. The main errors are occasional horizontal stripes associated with wind gusts and scaling in the flight direction. The latter is due to variations in aircraft velocity and flight direction (yaw).

2.5.2 Handheld gyro stabilized experiment

The next experiment was conducted with HSI-instrument (B). The imager was mounted on a 3-axis gimbal made by the company Feiyu Tech, model MG. The gimbal is designed for handheld operation of Digital Single Lens Reflex (DSLR) cameras. Integrated 3-axis accelerometers and gyroscopes feed the three brushless motors with attitude data (tilt, roll, and pan). The HSI was mounted on an

2. Hyperspectral instrument design: Do it yourself hyperspectral imager for handheld to airborne operations




Figure 2.8: Drone experiment in Skarsteindalen at Andøya, Norway on the 9th of August 2016. Panel (A) shows the Octocopter operated by the Remotely Piloted Aircraft (RPA) group at Andøya Space (AS). Panel (B): Hyperspectral Imager (HSI) - instrument (A) and the NDVI (Normalized Difference Vegetation Index) camera mounted on the 2-axis stabilized gyro platform of the drone. Background panel (C): NDVI orthomosaic photo. Yellow numbered boxes mark sampled areas of the HSI. The label (5) marks the upper left ground control point.




Figure 2.9: Side by side comparison between the NDVI (Normalized Difference Vegetation Index) camera and the Hyperspectral Imager (HSI) - Instrument (A) from the drone experiment conducted at Skarsteindalen on the 9th of August 2016. The bottom axis groups the 4 recorded image scenes. The HSI (colour) and the NDVI images (pinkish scaled) are tagged at the top. The HSI RGB composites are constructed by combining 10nm bandpass images at centre wavelengths 470nm (blue), 550nm (green) and 630nm (red). The time of flight is at the top of each image bar in seconds.

2. Hyperspectral instrument design: Do it yourself hyperspectral imager for handheld to airborne operations

L-shaped aluminum mount bar that increased the mass with 90g. It was necessary to add mass to the gimbal in order to balance all 3-axis correctly. The HSI recorded the spectral movie on two internal 32GB memory cards installed in the CamOne Infinity action camera head and the gimbal was powered by its own batteries, making the assembled system self-contained with no auxiliary connections or computers.

A panoramic sweep using the gimbal is shown in Figure 2.10. The target is the houses of Longyearbyen on the 25th of September 2017. The HSI and the gimbal were stationary mounted on the roof of The University Centre in Svalbard (UNIS) with the yaw axis panning from East to West up the Longyearbyen valley. Nine images were extracted from the hyperspectral datacube every 30nm from 470 to 670nm centre wavelength. All images have a bandpass of 10nm. The total sweep took 152 seconds. The RGB composite was constructed using centre wavelengths at 480nm (blue), 550nm (green), and 620nm (red).

As expected, the HSI produced a colour image close to visual perception for all targets in the scenario. The images seem to be more blurry in the blue part compared to the green-red part of the spectrum. The effect is mainly due to chromatic aberration as explained in Section 2.4. Note that the images are scaled down a factor of 3 in the vertical direction and equalized in intensity to increase contrast and brightness by the program paint.net from dotPDN LLC.

2.5.3 Computer screen test experiment

In the last experiment, the gimbal and the HSI instrument (B) were used handheld to sweep a computer screen displaying an image of a fruit collection. 571 spectrograms were recorded in 22s at a distance of 1.5m from the computer screen. The target screen colour photo and the HSI RGB composite are shown in Figure 2.11. Centre wavelengths at 490nm (blue), 552nm (green), and 620nm (red) were used to construct the image. Image bandpass is 10nm for all channels. The image is resized to match the target size. The vertical and horizontal scale factors are 0.22 and 1.38, respectively. The horizontal black line close to the top in the image is due to dust on the slit.

There are two advantages of using a computer screen as a test target. First, there is no need for an additional reference camera to document the target. Secondly, the target is illuminated by the screen itself.

The experiment shows that the gimbal and the HSI together are capable of reproducing both colour and shape of the target objects in the scene. The focus of the HSI RGB composite is not as sharp as the test image. The size of the slit width requires a downscaling in the vertical direction to obtain equal sized or square image pixels. In addition, multicolored and non-smooth transitions are seen at the edges of the objects. The effects are mainly due to chromatic aberration and operator




Figure 2.10: Sample images from the 3D printed push-broom Hyperspectral Imager (HSI) - instrument (B). Location is at the roof of the University Centre in Svalbard (UNIS) on the 25th of September 2017. Vertical centre line of the images is towards South - up the Longyearbyen valley. Each image is labeled with the centre wavelength to the right. The individual images have a bandpass of 10nm. The bottom RGB composite is constructed by combining images at centre wavelengths 480nm (blue), 550nm (green) and 620nm (red).



Figure 2.11: Handheld gimbal operation of the push-broom Hyperspectral Imager (HSI) - instrument (B). Panel (A): Target computer screen image. Panel (B): HSI RGB composite constructed by combining centre wavelength images at 490nm (blue), 552nm (green) and 620nm (red). Bandpass is 10nm for each colour channel.

shake, respectively.

The sensitivity of the HSI prototypes has not been a limiting factor in the study. As mentioned, higher throughput may be achieved by increasing the slit width and decreasing the collimator focal length to 25mm. A faster front lens at $f/2.8$ may then be used. As a result, both the spectral and spatial resolution would be degraded. Future studies will be conducted to optimize performance bases on the level of target illumination and spectral calibration.

2.6 Concluding remarks

Together with off-the-shelf optical components, low-cost camera heads, and 3D printing, it is possible to construct a push-broom hyperspectral camera with a mass less than 200g and cost below US\$700. Laboratory calibrations supplemented with field experiments have proven the conceptual design based on visual comparison of colour and shape with the RGB camera images recorded simultaneously, which

is believed to be sufficient proof at this stage of the development.

The principal results obtained in this study can be summarized as follows:

1. 3D printing is a powerful tool for construction of rapid prototype parts with sufficient accuracy for optical systems, like a transmission grating house presented in this study. The lightweight and strength of the printed thermo-plastic material is an advantage when optimizing payload mass and size, especially for drone operations.
2. The accessibility of mass produced low cost camera sensors developed for the industry, the Radio Control (RC) model community and action camera users, makes it possible to record visible and even the near infrared part of the electromagnetic spectrum at high frame rates.
3. Line scanning or push-broom hyperspectral imaging from an airborne platform or simply handheld can be done with a commercial off-the-shelf motorized gyro stabilized gimbal.
4. Self-contained motorized gyro gimbals and internal camera head recording reduce auxiliary device support requirements and complexity of field operations.
5. The push-broom hyperspectral imaging technique is revitalized.

Acknowledgments

We would like to thank the students Joy Muraszko, Henna-Reetta Hannula, Stefan Ram, Alistair McConnel, Elena Berndt, Daniel Kepski, Sara Marie Strand, Birgitte Madsen, Daniel Nilsson, Eivind Kolås, Leif Erik Andersson, Josh Marks, Nicole Nagy, Anne Kathrine Wenaas Ribe, Caelia Gardiner, Christopher Cosgrove, Conor Bolas, Maximilian Fuermann and Yitayew Temesgen Gebrie of the Arctic Earth Observational summer course which was successfully hosted by Christoffer Stausland from the Centre for Space-related Education (NAROM) at the Andøya Space (AS) in 2016. Also thanks to Mats Mikalsen Kristensen and Jostein Sveen from the Remotely Piloted Aircraft (RPA) group at AS.

Chapter 3

Preparing the imager for field tests

A lightweight payload for hyperspectral remote sensing using small UAVs

In this chapter we describe a lightweight hyperspectral imaging payload consisting of a push-broom imager, GPS, and IMU sensors as well as data synchronization and acquisition systems¹. The payload was developed in a modular and customizable way, making it a flexible UAV payload for research activities.

First we present the hardware and software set up, with details on each of the components and their interaction. Then we show the results from field tests.

3.1 System overview

The described payload, Figure 3.1a, is a highly modular system: sensors can be added, removed, or swapped. All changes have of course an impact on the quality of the final data product. E.g., having a lower performance, but cheaper GPS will produce less accurate georeferencing. In this paper, we present the current version of the payload's set-up.

Figures 3.2 and 3.3 show the data and power flow, respectively. In dark colour are

¹Based on the paper "A lightweight payload for hyperspectral remote sensing using small UAVs", by João Fortuna and Tor Arne Johansen, presented at the 9th Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing (WHISPERS), Amsterdam, Netherlands, 2018. This work was supported by the Norwegian Research Council (grant nos. 223254 and 221666) through the Centre of Autonomous Marine Operations and Systems (NTNU AMOS) at the Norwegian University of Science and Technology, the MASSIVE project (grant no. 270959), as well as the Norwegian Space Agency.

3. Preparing the imager for field tests: A lightweight payload for hyperspectral remote sensing using small UAVs



(a) Standalone.

(b) Mounted on Maritime Robotics' hexacopter.

Figure 3.1: Prototype payload.




Figure 3.2: Data connection diagram.




Figure 3.3: Power distribution diagram.

data or power sources, in white the conversion or routing elements and finally in a light colour are data storage/transmission or power consuming components. The following sections will explain each component interaction in more detail.

3.1.1 SenTiBoard

The **S**ensor **T**iming **B**oard provides accurate timestamps with 10ns resolution for all sensors. This component is described in [70][71]. By using data from the IMU

and GPS sensors and a navigation filter [72], we can generate georeferenced maps of hyperspectral data.

3.1.2 Hyperspectral instrument

At the heart of our system is the push-broom [73] hyperspectral instrument. This model was developed using 3D printed parts and commercial off-the-shelf (COTS) optical components and the optics are discussed in detail in Chapter 2. By developing the system in such way that the camera can be chosen with some degree of freedom, it allowed us to select a model with the required spectral sensitivity as well as easy Linux integration.

Camera

The camera model used in the current version of the system is the UI-3360CP-NIR-GL, with enhanced NIR sensitivity, produced by IDS Imaging Development Systems GmbH [74]. This corresponds to model C, refer to Section 2.3, Figure 2.5. Raw data depth is 12bit, it provides a USB 3.0 interface for data transfer, and an 8-pin I/O connector with flash signal, trigger and GPIO pins.

Optics

The principle used for hyperspectral imaging is a push-broom instrument. This is achieved through the use of a slit, a dispersing element (grating), and required focusing lenses, see Section 2.2. Because this camera was developed using COTS components, the available slits did not allow taking advantage of the full size of the sensor. While the sensor has 1088 pixels in the spatial axis, we can only use 580.

The spectral range of the instrument is theoretically 300-1000nm, however wavelengths below 400nm are blocked by the glass lenses, and in addition, frequencies above 800nm suffer from blending with the second spectral order as previously discussed in Section 2.4. Despite these limitations we record the 350-950nm range so we can later decouple some of these effects in post-processing.

The ground resolution is given by:

$$d_x = \frac{z \times w}{f_1} \quad [\text{m}] \quad (3.1)$$

$$\Delta_x = d_x + v \times \Delta_t \quad [\text{m}] \quad (3.2)$$

$$\Delta_y = \frac{z \times h}{f_1 \times N} \quad [\text{m}] \quad (3.3)$$

Where $f_1 = 16\text{mm}$ (front lens focal distance), $w = 25\mu\text{m}$ (slit width) and $h = 3\text{mm}$ (slit height), Table 2.1. The remaining variables: $N = 580$ (usable pixels in




Figure 3.4: Flash electrical wiring.

spatial axis) was discussed before, then z (flight altitude), v (ground speed) and Δ_t (exposure time) can be changed in each flight.

Data acquisition

To enable the synchronization of the camera frames to the rest of the data, the flash output signal is fed to the SenTiBoard. Figure 3.4 shows the wiring schematic used to capture the flash signal, as recommended in the camera manual [75]. The signal in the Time of Validity (TOV) node will be high (5V) whenever the flash is active, and low (GND) otherwise. When initializing the camera, we configure the flash to be active for the entire duration of the exposure time. The SenTiBoard will record the time of every falling edge (transition from 5V to GND) in TOV, thus timestamping the end of each exposure.

Due to limitations in processing power and writing speed in the on-board computer, we perform *off-chip* spectral binning [73] on the frames before saving to disk. The binning operation reduces the size of the image by adding columns together. After experimenting with different settings, a value of $10\times$ seems to allow recording data at 30fps. Since the raw data is captured at 12bit depth, it is possible to bin up to $16\times$ without risking overflow when using a 16bit image container file. The chosen binning factor results in a spectral resolution of approximately 4nm per band. Because of the limited usable spatial and spectral ranges caused by the optics, we read a cropped area of the sensor, that area is 1800×600 pixels. When applying $10\times$ spectral binning to those frames we end up with 16bit images of 180×600 pixels that should be stored.

Calibration

Our instrument was wavelength calibrated to match pixels in the spectral axis to wavelengths. The procedure used was to point the camera to a fluorescent light tube with known sharp emission bands [76] and capture data. Then, visually match the peaks to those bands as seen in Figure 3.5. Finally, fit a second degree polynomial on the pixel-wavelength pairs discovered, as mentioned in Section 5 of [77].




Figure 3.5: Fluorescent light spectrum.

Radiometric calibration was not performed, as we are still testing the sensor to find optimal settings.

3.1.3 Other hardware

Navigation sensors

For high positioning accuracy, we used an RTK (Real Time Kinematic) capable GPS module. The model selected was uBlox-NEO-M8T, assembled on a breakout board by CSG Shop [78]. We time and log the following messages: RXM-RAWX, RXM-SFRBX, NAV-PVT as per the instructions in [71].

The chosen IMU was the ADIS16490 3-axis gyroscope and accelerometer.

There is no heading sensor installed in this version of the payload, but heading information can be retrieved from the UAV autopilot logs. Timing is synchronized through GPS times both on the payload and autopilot.

On-board computer

As this system is to be used for small-scale UAV-based airborne surveys, there was a need to have an on-board computer that can acquire and log data. We chose the Odroid-XU4 board [79]. This board is also recommended by the camera manufac-

turer for embedded applications.

To accommodate for the high data rate coming from the camera, we opted for a USB 3.0 SSD disk. Initial tests with USB flash drives and μ SD cards showed that those devices were not fast enough.

Communications

In its stand-alone version, the payload includes a Wi-Fi radio for communication. During the field experiments this radio was configured as an Access Point (hotspot), this way we could connect to it using the integrated network card on a laptop. Data rate requirements are low, as external communication is used only for status monitoring and acquisition control.

3.1.4 Software

One of our goals when developing this payload was to make it as easy to use as possible. By simplifying and automatising the acquisition process, field experiments can be made more efficient. We achieve this by using the *LSTS Toolchain* [80], an open-source project with a suite of tools for autonomous systems that include on-board control and logging (DUNE), communication protocol (IMC) and a graphical user interface (Neptus) for mission control and log analysis. The UAV-Lab has contributed to and used this toolchain in the past [81], thus making it the prime candidate for powering our system.

DUNE [82] is the software running on the on-board computer. It manages the activation and logging of the sensor data, as well as communication with the ground station. To support the chosen hardware, there was a need to develop two "Tasks" (software modules) in DUNE: one for interfacing the camera and another for the SenTiBoard.

3.2 Experimental results

On March 22nd, 2018, we flew over Hopavågen, a salt water pool in the Trøndelag region, Norway. This is an area of interest for the Department of Biology at NTNU, who provided us with the experiment goal of mapping different types of habitats in shallow water such as seagrass, rocks, sand, or sea urchins. In the past, this has been performed manually.

The target to observe was underwater, and in addition, on that day the sky was very cloudy. That resulted in nonideal conditions. Before flight, ground tests were made to find the best settings for image capture. Table 3.1 shows the flight parameters and sensor configuration.

The payload was mounted on a hexacopter as seen in Figure 3.1b. During the

Table 3.1: Experiment parameters.

Parameter	Value	Unit
Frame rate	20	frames/s
Exposure time	49.9 (max)	ms
Sensor Gain	4	-
Flight Altitude	10	m
Flight Speed	1.5	m/s




Figure 3.6: Hyperspectral data (RGB rendering) overlaid on RGB image taken during flight. Swath width of HSI is 1.875m.

flights an additional RGB camera was also capturing images. Data from both RGB and HS cameras are seen in Figure 3.6. The two images were matched manually based on timing information and ground features. Spectral data from each of the marked points of interest can be seen in Figure 3.7.

Using the values from Table 3.1 in Equations 1-3, we get ground resolutions per pixel of $\Delta_x \approx 90mm$ (along track) and $\Delta_y \approx 3mm$ (across track). Ground sampling distance is $v/fps = 75mm$.

3.3 Future work

We are currently working on the final piece of the HS data acquisition and pre-processing pipeline: automatic georeferencing, map generation, and GIS software compatibility. The first step towards it is the processing of raw navigation data into a full navigation solution [72].

Other features that should be added in the near future are the ability to change settings more easily, detection of relevant target signatures during flight, stream-

3. Preparing the imager for field tests: A lightweight payload for hyperspectral remote sensing using small UAVs

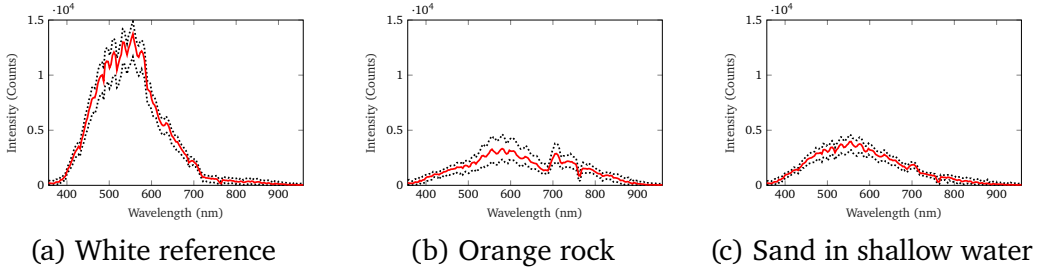


Figure 3.7: Mean spectrum (red solid line) and range (dotted black line) for each point of interest.

ing of more detailed capture status, and data quality diagnostics. In addition, we should perform accurate radiometric and geometric calibrations. Regarding the mechanical assembly of the payload, the current solution is experimental. Once we decide on a fixed set of components, we can design a custom enclosure with minimal weight and volume, while keeping the required mechanical and thermal stability.

3.4 Conclusions

Our system proved to be easy to use in the field, requiring little interaction beyond the initial parameter setting.

Because of the grating-slit design of the push-broom camera, the amount of light captured by the sensor is very limited, this sets some requirements on light and/or weather conditions during experiments. Changing flight parameters or the slit width can increase the amount of light, however that comes at the expense of reduced area coverage, spatial and/or spectral resolution.

Chapter 4

Correcting for environmental disturbances

Multivariate data modelling for de-shadowing of airborne hyperspectral imaging

Airborne hyperspectral imaging is a powerful technique for high-resolution classification of large areas of ground, applied today in fields like agriculture and environmental monitoring. Even though many classification algorithms are capable of handling shadows without a decrease in performance, visual inspection can be made easier if shadows are removed.

In this chapter, we present a method for separating the effect of shadows (de-shadowing) and other partially known lighting condition changes from the effects due to the physical, chemical or biological properties of the ground, which are of interest¹.

4.1 Introduction

Effects of shadows usually represent a multiplicative gain change in reflectance (R) data, traditional de-shadowing techniques solve the task by finding the more or less complex gain to compensate for it [83].

In this chapter, we present an approach that relies on the multiplicative effect of illumination in reflectance mode, that is, additive in the apparent absorbance

¹Based on the paper "Multivariate data modelling for de-shadowing of airborne hyperspectral imaging", by João Fortuna and Harald Martens. First presented at the International Association for Spectral Imaging (IASIM) sixth international conference in spectral imaging, held in Chamonix-Mont-Blanc (France), July 2016. Later published in the Journal of Spectral Imaging (JSI), Volume 6, 2017, per invitation of the conference committee.

$(\log_{10}(1/\mathbf{R}))$.

The next section presents the theory behind the method, with all necessary steps. Following that are the results of the application of the method to airborne hyperspectral data. The last section contains some conclusions and open discussion topics.

The methodology presented here represents an extension of the theory of the informative converse (IC) [84]. This theory concerns inadequate mathematical modelling of variations in data from multichannel measurements (Q variables in N samples), and points out how multivariate data modelling can give surprising new insights about unexpected – and thus unmodelled – phenomena in the system measured: it is well known that unmodelled phenomena can lead to alias ("multivariate bias") problems in modelled phenomena's estimated parameters. However, the IC theory shows that a subsequent multivariate analysis of the resulting ($N \times Q$) table of unmodelled residuals can give a surprisingly good characterization of some aspects of the elements missing in the oversimplified modelling. The present chapter refines the IC theory by bringing in additional knowledge about the system at hand.

As mentioned, established approaches try to estimate changes in intensity gain to compensate for shadows in hyperspectral image data. Let the data matrix $\mathbf{Y}(N \times Q)$ be a table of transformed spectra (N pixels, Q wavelength channels), obtained by treating the multichannel spectral data in the N pixels as if they represented Q -dimensional spectra from a spectrophotometer: reflected intensity from a scene is observed by the camera (\mathbf{I}). For gain control, each pixel's measured \mathbf{I} is element-wise divided by I_0 , the reflected intensity of a white body with the same light source, to convert \mathbf{I} to reflectance (\mathbf{R}). Finally, the linearized data are given by the apparent absorbance $\mathbf{Y} = \log_{10}(1/\mathbf{R})$. Hence, to ensure that systematic variations in I_0 can be modelled additively instead of multiplicatively, we replace the camera's intensity measurements \mathbf{I} by the apparent absorbance, $\mathbf{Y} = \log_{10}(1/\mathbf{R})$ in the subsequent multivariate spectral modelling.

4.2 The informative converse paradox

Figure 4.1 illustrates the original IC paradox. The basic problem is illustrated in Figure 4.1a): the measured data \mathbf{Y} often represent the sum of two causal sources – expected and unexpected ones – plus random errors (\mathbf{F}). The unexpected variation sources can seriously interfere with the modelling of \mathbf{Y} . For instance, assume that the spectral data \mathbf{Y} have contributions from an expected phenomenon with spectral profile \mathbf{s} and levels \mathbf{c} in n pixels, plus contributions from an unexpected phenomenon with spectral profile \mathbf{z} and levels \mathbf{d} in the n pixels:

$$\mathbf{Y} = \mathbf{c}\mathbf{s}^\top + \mathbf{d}\mathbf{z}^\top + \mathbf{F} \quad (4.1)$$

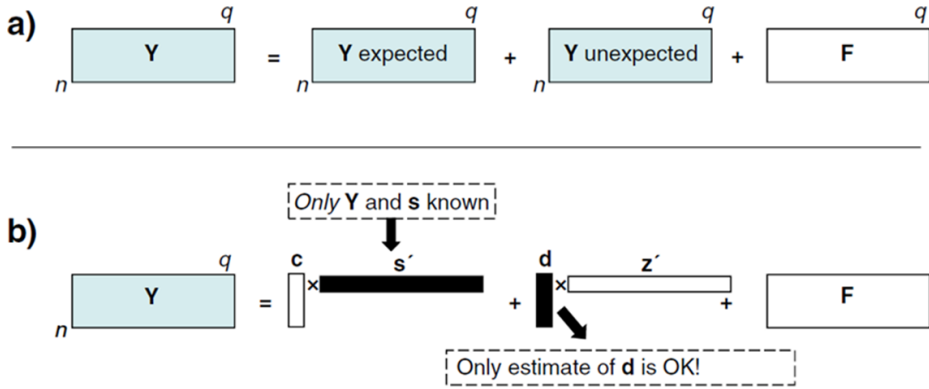


Figure 4.1: The informative converse paradox, illustrated for a linear model of the two constituents, observed in $Y(N \times Q)$ at Q wavelengths in N pixels; $Y(n \times q)$ is the original notation from [84]. a) The measurements Y have both an expected and an unexpected source contribution, in addition to random measurement errors F . b) The expected and unexpected source contributions may be modelled as cs^T and dz^T respectively. When only measurements Y and the expected source's known spectrum s are known, while its concentration c and unexpected concentration d and spectrum z are unknown, the original IC method gives a surprisingly good estimate of the unexpected d , but erroneous (aliased) estimates both for the expected c and the unexpected z . (Adapted from [84] and reproduced with the permission of the publisher).

where $Y(N \times Q)$: hyperspectral reflectance measurements for N pixels (rows) at Q different wavelength channels (columns), linearized as $\log_{10}(1/R)$; $cs^T = c(N \times 1)s^T(1 \times Q)$: contributions from a partially known cause; $dz^T = d(N \times 1)z^T(1 \times Q)$: contributions from a totally unknown; $F(N \times Q)$: measurement and modelling errors (assumed, for simplicity, to be random, normally distributed).

If the spectral data in Y are modelled inadequately Figure 4.1b) in terms of known spectrum s only, while z is ignored, the estimate of c (e.g. $\hat{c} = Ys(s^T s)^{-1}$) will be contaminated by contributions from d and thus "unnatural". The IC *paradox*, Figure 4.1b), is that a subsequent multivariate PCA analysis of $Y - \hat{c}s^T$ yields estimates of d with the correct "natural" expectancy (except for a trivial scaling factor). The corresponding estimate of z from that PCA will be "unnatural" due to its orthogonality to s .

The present chapter modifies this IC methodology by correcting the estimates of c and z with the use of additional assumptions (such as non-negativity), extends it to multicomponent models and applies it to hyperspectral airborne images of ground scenes in order to separate illumination changes (e.g., shadows) from ground properties (e.g., geology, vegetation).

4. Correcting for environmental disturbances: Multivariate data modelling for de-shadowing of airborne hyperspectral imaging

In the following, we demonstrate that shadow effects and illumination changes may be expected to give additive phenomena, allowing us to write the following multivariate linear model:

$$Y = CS^T + DZ^T + F \quad (4.2)$$

For hyperspectral data Y each matrix takes the meaning below:

- $Y(N \times Q)$ – hyperspectral absorbance for N pixels (rows) at Q different wavelength channels (columns), linearized as $\log_{10}(1/R)$
- $CS^T = C(N \times J)S^T(J \times Q)$ – contributions from partially known causes (e.g., light sources)
- $DZ^T = D(N \times M)Z^T(M \times Q)$ – contributions from totally unknown causes (e.g., ground components)
- $F(N \times Q)$ – measurement and modelling errors (assumed, for simplicity, to be random, normally distributed)

It is, of course, possible to estimate the expected sources' unknown concentration parameters $C(N \times J)$ of J constituents in N pixels from their known spectra $S(Q \times J)$, using the over-simplified linear model ignoring the unexpected variation source(s) DZ^T :

$$Y = CS^T + F \quad (4.3)$$

But since this represents an over-simplified mathematical model of reality, it can give serious errors in the parameters identified. In this case, the estimates of C from S only. The purpose of the present chapter is to give good estimates of the expected source's pixel levels C and the unexpected sources' contribution levels D and spectra Z as well as residuals F in such cases.

4.3 Method

The IC modelling extension may be applied in different ways. For illustration, we here assume that Z , the spectral properties of the main ground constituents (e.g., plants, rocks, water) are unknown, and so are their pixel-by-pixel level ("concentration") variations D , while S , the spectra of possible lighting conditions (e.g., sunlight, overcast), is known, while their pixel-by-pixel level variations C are unknown. However, the presented method also works in the converse case, when only C is known, while their spectra S are unknown (see [84]). Likewise, the method may be used to estimate and correct for the spectral effects of unknown light sources, particularly if the spectral properties or the pixel-by-pixel level variations of some or all of the main ground components are known. Necessary modifications to the algorithm are trivial and will not be explained here.

4.3.1 Relating the unknowns to the known or easily estimated parameters

We can always express the spectra Z and levels C as function of the spectra S and levels D respectively:

$$Z = SA + Z_{\perp S} \quad (4.4)$$

$$C = DB + C_{\perp D} \quad (4.5)$$

where $A(J \times M)$ and $B(M \times J)$ represent the coupling of unexpected spectra Z to known spectra S and expected but unknown levels C to unexpected levels D , respectively. The M -dimensional subspace of D is easy to estimate by IC analysis [84]. If all elements in A (or B) are 0, Z and S (or C and D) are orthogonal, the alias problem disappears. If not, then they represent spectral- and/or level- overlaps that will create alias errors unless estimated and compensated for.

4.3.2 Summary model

$$Y = CS^T + DZ^T + F \quad (4.6)$$

$$= (DB + C_{\perp D})S^T + D(A^T S^T + Z_{\perp S}^T) + F \quad (4.7)$$

$$= D(B + A^T)S^T + C_{\perp D}S^T + DZ_{\perp S}^T + F \quad (4.8)$$

$$= DHS^T + C_{\perp D}S^T + DZ_{\perp S}^T + F \quad (4.9)$$

where $H(M \times J)$ is the sum of the unknown ambiguity matrices in Z (w.r.t. S) and C (w.r.t. D):

$$H = B + A^T \quad (4.10)$$

4.3.3 Estimation methodology

In the following steps, ordinary least squares (OLS) regression is used in all estimations. If different input variables (columns in Y) are known to have different relevance or reliability, then the OLS may be replaced by weighted least squares (WLS) or generalized least squares (GLS) in all the regressions over channels. Likewise, if different pixels (rows in Y) need different weights, then WLS or GLS regression may be used in all regressions over observations. Alternatives to OLS, WLS or GLS, such as best linear unbiased predictors (BLUP) or robust statistical estimators, may also be used.

4.3.4 Estimation of model parameters

- \mathbf{S} is assumed known
- \mathbf{D} and $\mathbf{Z}_{\perp S}$ are easy to estimate by the IC analysis [84]:

$$\widehat{\mathbf{C}}^* = \mathbf{Y}\mathbf{S}(\mathbf{S}^T\mathbf{S})^{-1} \quad (4.11)$$

$$\mathbf{E}(\widehat{\mathbf{C}}^*) = \mathbf{D} + \mathbf{D}\mathbf{A}^T \quad (4.12)$$

$$\mathbf{F}_{IC} = \mathbf{Y} - \widehat{\mathbf{C}}^*\mathbf{S}^T \quad (4.13)$$

$$\mathbf{E}(\mathbf{F}_{IC}) = \mathbf{D}\mathbf{Z}_{\perp S}^T \quad (4.14)$$

Singular Value Decomposition (SVD) of the IC residual \mathbf{F}_{IC} gives, for its q most significant components:

$$[\mathbf{U}_{IC}, \mathbf{s}_{IC}, \mathbf{V}_{IC}] = \text{svd}(\mathbf{F}_{IC}, q) \quad (4.15)$$

which yields estimates of the bilinear structure $\mathbf{D}\mathbf{Z}_{\perp S}^T$. With:

$$\widehat{\mathbf{D}} = \mathbf{U}_{IC}\text{diag}(\mathbf{s}_{IC}) \quad (4.16)$$

$$\widehat{\mathbf{Z}}_{\perp S} = \mathbf{V}_{IC} \quad (4.17)$$

- $\mathbf{C}_{\perp D}$ is estimated by:

$$\widehat{\mathbf{C}}_{\perp D} = [\mathbf{I} - \widehat{\mathbf{D}}(\widehat{\mathbf{D}}^T\widehat{\mathbf{D}})^{-1}\widehat{\mathbf{D}}^T]\widehat{\mathbf{C}}^* \quad (4.18)$$

4.3.5 Estimating the ambiguity matrix \mathbf{H}

$$\mathbf{Y} = \mathbf{D}\mathbf{H}\mathbf{S}^T + \mathbf{C}_{\perp D}\mathbf{S}^T + \mathbf{D}\mathbf{Z}_{\perp S}^T + \mathbf{F} \quad (4.19)$$

$$\mathbf{D}\mathbf{H}\mathbf{S}^T = \mathbf{Y} - \mathbf{C}_{\perp D}\mathbf{S}^T - \mathbf{D}\mathbf{Z}_{\perp S}^T - \mathbf{F} \quad (4.20)$$

Insert estimates:

$$\widehat{\mathbf{D}}\mathbf{H}\mathbf{S}^T \approx \mathbf{Y} - \widehat{\mathbf{C}}_{\perp D}\mathbf{S}^T - \widehat{\mathbf{D}}\widehat{\mathbf{Z}}_{\perp S}^T \quad (4.21)$$

Solve to estimate \mathbf{H} :

$$\widehat{\mathbf{H}} = (\widehat{\mathbf{D}}^T\widehat{\mathbf{D}})^{-1}\widehat{\mathbf{D}}^T(\mathbf{Y} - \widehat{\mathbf{C}}_{\perp D}\mathbf{S}^T - \widehat{\mathbf{D}}\widehat{\mathbf{Z}}_{\perp S}^T)\mathbf{S}(\mathbf{S}^T\mathbf{S})^{-1} \quad (4.22)$$




Figure 4.2: AVIRIS dataset f140530t01p00r16 with region of study marked.

4.3.6 Estimation of expected source levels C and unexpected spectra Z

From Equations 4.4, 4.5 and 4.10:

$$\widehat{C} = \widehat{D}(\widehat{H} - A^T) + \widehat{C}_{\perp D} \quad (4.23)$$

$$= \widehat{D}\widehat{H} + \widehat{C}_{\perp D} - \widehat{D}A^T \quad (4.24)$$

$$\widehat{Z} = S(\widehat{H}^T - B^T) + \widehat{Z}_{\perp S} \quad (4.25)$$

$$= S\widehat{H}^T + \widehat{Z}_{\perp S} - SB^T \quad (4.26)$$

Hence, \widehat{C} (and \widehat{Z}) have unknown matrices $A(J \times M)$ (and $B(M \times J)$), which describe how Z depends on S (how C depends on D), and that generates ambiguity w.r.t. \widehat{D} (and S). This ambiguity can only be eliminated by additional information about C (and Z) (e.g. non-negativity, unimodality, smoothness, ICA based on entropy).

4.3.7 Estimation of spectral overlap A and level overlap B

By applying previous knowledge about C and Z such as smoothness, A and B can be estimated by an optimizer, for instance YALMIP [85].

Since only the product DZ^T can be estimated from Y and S , the individual matrices D and Z need to be determined by methods such as MCR-ALS [86].

4.4 Dataset

We applied our method to data from NASA's AVIRIS programme²[87]. AVIRIS, the Airborne Visible InfraRed Imaging Spectrometer, is a push-broom system that delivers calibrated images of the upwelling spectral radiance in 224 spectral bands with wavelengths from 400 to 2500nm (VIS-NIR).

This particular scene was extracted from flight f140530t01, run ID p00r16, just

4. Correcting for environmental disturbances: Multivariate data modelling for de-shadowing of airborne hyperspectral imaging



Figure 4.3: RGB reconstruction of original data.




Figure 4.4: Estimated spectra of known lighting effects (S): a constant "baseline" and a shadow signature (here estimated from sunny vs shaded mountain sides).

outside of Las Vegas, Nevada, USA. Data was collected on 30 May 2014. The full dataset has 850 samples over 3337 lines, however, in this study we consider a 500×500 region only, check Figures 4.2 and 4.3.

4.5 Results

Applying the method to airborne hyperspectral images yields promising results. RGB images were generated using a weighted sum of spectral channels in the visible range, in a similar fashion to the spectral response of the human eye. For this, we ensured the data were in reflectance mode. The weights for the weighted sum were calculated using Gaussian curves centred on the peak wavelengths of the cone cells in the human eye, such wavelengths are well known and easily available [88].

S was simplified to two spectra, given in Figure 4.4, the skylight spectrum was estimated by taking the difference of spectra between pixels in the shadow and light regions in small neighborhoods (so that the ground effects are expected to be the same).

²At the time of development and writing of the work on which this chapter is based (2015-2016), the instrument described in Chapter 2 was not ready or tested aboard our UAV. That is why we did not use our own data.




Figure 4.5: De-shadowed data, given as RGB.



Figure 4.6: Shadow data, given as RGB. 10 principal components were used for the estimation of the de-shadowed data.

Figure 4.5 depicts the de-shadowed data. At first glance, it seems as if all the terrain information was removed, like a flattened version of the original data. On the other hand, the "shadow map" in Figure 4.6 appears to contain only terrain information and nothing about the different constituents of the ground. Both de-shadowed and shadow data hold valuable information, chemistry/geology and physics/topography, respectively.

By comparing the performance of a simple *k-means* clustering (chosen number of clusters = 3) on the original (Figure 4.7) and de-shadowed (Figure 4.8) data, we can see how the effect of shadows and morphology is removed and clusters are based on ground "chemistry". Increasing the number of clusters beyond three revealed more interesting ground features (not shown here).

4.6 Conclusions

Theoretically, the extended IC method presented here seems able to simplify the interpretation of hyperspectral images by better "de-shadowing": improved separation of light source variations from ground property variations. Using additional assumptions, such as spatial ground topology or other known ground or light source features, both expected and unexpected signal sources can thus be quantified rather completely, or at least with less alias errors, when compared to the original IC method without such assumptions.

4. Correcting for environmental disturbances: Multivariate data modelling for de-shadowing of airborne hyperspectral imaging



Figure 4.7: Clustering in original data.




Figure 4.8: Clustering in de-shadowed data.

For instance, in the hyperspectral aerial image illustrated, the extended IC attained good "de-shadowing" using additional assumptions about where to discover the light source spectrum s (in this case the contrast between the sunny and shaded sides of mountains). Some features that appear in the de-shadowed RGB representation of the image (Figure 4.5) are not apparent in Figure 4.3. A dark patch in the lower right area of the picture is particularly noticeable. One may hypothesize that the information was contained in the NIR bands which was not apparent in the original RGB reconstruction, but when removing the effects of illumination it was made more evident.

Hence, the extended IC method resolved some of the paradox in the original IC. We believe that the new method can simplify the human interpretation as well as the automated quantitative use of hyperspectral imaging. Work is in progress to improve the method's optimization step further.

Chapter 5


Enhancing hyperspectral data

Multivariate image fusion: A pipeline for hyperspectral data enhancement

Hyperspectral cameras provide high spectral resolution data, but their usual low spatial resolution, when compared to colour (RGB) instruments, is still a limitation for more detailed studies. This chapter presents a simple yet powerful method for fusing co-registered high spatial and low spectral resolution image data – e.g., RGB – with low spatial and high spectral resolution data – Hyperspectral¹. The proposed method exploits the overlap in observed phenomena by the two cameras to create a model through least square projection. This yields two images: 1) A high-resolution image spatially correlated with the input RGB image but with more spectral information than just the 3 RGB bands. 2) A low-resolution image showing the spectral information that is spatially uncorrelated with the RGB image. We show the results for semi-artificial benchmark datasets and a real-world application. Performance metrics indicate the method is well suited for data enhancement.

¹Based on the paper "Multivariate image fusion: A pipeline for hyperspectral data enhancement", by João Fortuna, Harald Martens and Tor Arne Johansen, published in *Chemometrics and Intelligent Laboratory Systems*, Volume 205, 2020.

This work was supported by the Norwegian Research Council (grant no. 223254) through the Centre of Autonomous Marine Operations and Systems (NTNU AMOS) at the Norwegian University of Science and Technology, the MASSIVE project (grant no. 270959), as well as the Norwegian Space Agency.



(a) High-resolution ($500 \times 500\text{px}$) (b) Low-resolution ($100 \times 100\text{px}$)

Figure 5.1: Example satellite images with simulated different resolutions. Sensors have the same area but a different number of pixels (different pixel size and density). The higher number of pixels in (a) means less light (fewer photons) per pixel, hence more noise. Photo by NASA on Unsplash [89].

5.1 Introduction

5.1.1 There is no such thing as a free lunch, or photons

Even though **number of pixels** is a different concept from **spatial resolution**, they are tightly coupled. Say we have two monochrome focal plane array (FPA) cameras, *A* and *B*, on a satellite and both image the same area on the ground, see Figure 5.1. If *A* has five times the amount of pixels in both axis, then the spatial resolution – the ability to differentiate between two close objects, or in this case the size of one pixel on the ground – will be five times better than the one of camera *B*. Here we assume all other properties of the cameras, such as optics, and image capture to be the same and that they do not limit the resolution, i.e. the sensor is the bottleneck.

A higher resolution has the clear benefit of allowing us to see finer details, however increasing the number of pixels is not always possible, in particular as we increase the number of spectral bands. To understand why the number of bands affects the spatial resolution, we need to think of light reaching a camera as a stream of a finite number of photons. These particles travel through the optical elements of the camera before being distributed by the pixels in the sensor. There, they generate an electrical current which translates into digital data. Higher pixel density means fewer photons per pixel, as we divide the same finite amount of photons by a higher number of pixels, and fewer photons per pixel lead to a weaker signal

(low SNR). When the signal is fainter, the noise contribution becomes apparent if we increase the sensor gain, compare (a) and (b) in Figure 5.1. If we now try to sort the photons into many spectral bands, they become even scarcer. We need to compromise on the number of spatial pixels to have a usable signal.

Throughout this chapter, we will mention high and low-resolution data, we mean both the number of pixels and the ability to resolve a smaller object in the image. Higher-resolution data will have more pixels and conversely, lower resolution, fewer pixels.

Panchromatic images contain information from a broad spectrum in a single band, hence can more easily have a high spatial resolution. Pansharpening methods were initially developed in the mid-1980s for air- and space-borne multispectral imagers with a low resolution that could be improved with high-resolution panchromatic images. With the ever-growing availability of hyperspectral instruments, some of those methods were adapted to hyperspectral data and others developed anew. The literature describes several methods to achieve such sharpening [20].

Thanks to advances in sensor technology, we now have spatially high-resolution colour cameras (Red-Green-Blue – RGB) and even some multispectral cameras (with few, but more than 3, bands). We can use those data instead of panchromatic when enhancing hyperspectral data. The advantages of using colour cameras are clear: Even with only three bands, we now have multivariate – as opposed to univariate – spectral data in high spatial resolution. Such multivariate data gives much better selectivity, as it adds colour information to the simple measurement of light intensity provided by panchromatic sensors. This richness of data is relevant when combining a spatially high-resolution RGB camera with a spatially low-resolution hyperspectral camera. The higher the number of bands with high spatial resolution, the more we can improve the spatial resolution of hyperspectral data, assuming that both cameras have recorded the same spatial scene and therefore can be correlated.

Some of the most commonly used pansharpening methods are useful when the intended goal is to produce a high-resolution RGB image from multi-/hyperspectral data, having only a high-resolution monochromatic image. However, those enhancements are not visible when the goal is, for example, to find high-resolution estimated abundance maps of geological or biological resources of interest.

5.1.2 Motivation

For the past years, our research group has been working on a lightweight hyperspectral imaging system for Unmanned Aerial Systems (UAS), see Chapters 2 and 3. Because of weight, cost and complexity limitations, and operating conditions, the spatial resolution has been the most limiting factor when it comes to generating high-quality data products. RGB cameras are very often already a part of the

payload carried by such Unmanned Systems, and if not, they are a simple addition. Taking that into consideration, trying to improve HS spatial resolution with RGB data was the goal of the present research.

5.1.3 Contribution

In this chapter we describe a generic framework for multivariate image fusion, building on the ideas of pansharpening while also trying to enhance the output for further processing, instead of just visual representation. Hence, the purpose of the present methodology is to use the high spatial resolution of the RGB measurements to yield an equally high-resolution representation of the low-resolution HS measurements. Other studies [90, 91] have pursued somewhat related strategies, although, the method proposed here is to the authors' knowledge, a new development. It is fast, when compared to other methods in the literature [20], and requires very little knowledge of the calibration parameters or relationship between the two datasets to be fused, only assuming that they were spatially registered beforehand. In summary, the present method combines the input consisting of a low-resolution multi-channel HSI image and a high-resolution 3-channel RGB image of the same scene, into two output images: 1) A high-resolution multi-channel image showing what is spatially correlated to the RGB image, and; 2) A low-resolution multi-channel image showing what is spatially uncorrelated to the RGB image. The following section describes each step of the method. We then show some results of applying the method to both artificially degraded real data and real low-resolution data. To conclude, a discussion on possible improvements and future work.

5.2 Method - multivariate image fusion (MVIF)

In this section we present a method for fusing RGB and HS data in order to get a data product that takes advantage of the relative strengths of both, Figure 5.2. Such fusion is possible because HS data is in most real world applications typically very rank deficient - the number of HS wavelength bands is much higher than the number of statistically independent spectral variation types in the image. That means it is possible to learn all relevant patterns of variation in HS and replicate them through a low rank, but high spatial resolution, approximation based on RGB.

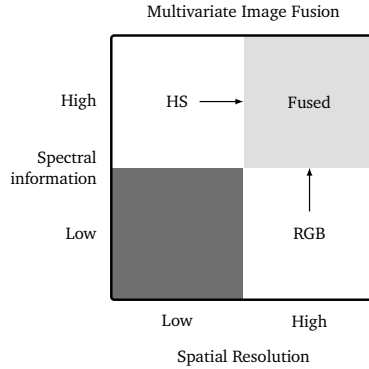


Figure 5.2: Multivariate Image Fusion exploits the relative strengths of RGB and Hyperspectral data.

5.2.1 Data Model

Before continuing, it is useful to write down the data model we will be working with. We take a similar approach to Chapter 4, based on the hypothetical model:

$$Y = CS^T + DZ^T + F \quad (5.1)$$

In this model we assume that the data cubes are unfolded, so all elements are 2D matrices.

Y is a high spectral resolution (hyperspectral, here called full-spectral) image dataset, with high or low spatial resolution, depending on the context.

C is the RGB data. The product CS^T contains the known spatial pattern from the RGB camera (C), based on initially unknown, but estimated full-spectral (S^T) information.

In analogy to the previous element, DZ^T contains spatial (D) and spectral (Z^T) information, but now from phenomena that are not seen by the RGB camera. Initially unknown, (D) and spectral (Z^T) have to be estimated.

F will, ideally, contain only noise.

When dealing with remote sensing spectral data, it is helpful to think of the total signal in each pixel of Y as a sum of contributions from all the phenomena that were observed in that single pixel. These contributions have 2 properties: concentration/abundance C and D and spectral signature S or Z . When we consider all the pixels in Y , the concentrations become spatial distribution maps of each of those phenomena. Adding all the spectral signatures – S and Z – weighted according to their respective concentration per pixel – C and D – yields the signal in Y , aside from noise – F :

$$Y = [C, D][S, Z]^T + F \quad (5.2)$$

Even though the number of bands (spectral resolution) can be of several tens or even hundreds, hyperspectral data is typically rank deficient, which means we can obtain a much lower dimension representation with less noise, while still keeping the relevant information. On the other hand, RGB data usually has a full rank of 3 in the spectral domain. When the spectral range of both instruments is overlapping, they observe the same phenomena and we can obtain a low rank representation of HS using RGB data.

If we use RGB data as is – with 3 bands – we are limited to a 3 dimension low rank representation, and while this may be enough for some datasets, it will prove inadequate for more complex scenes. This is a limitation of using a linear projection method as we have done here, other methods may not face this problem. Fortunately, it is possible to artificially expand the number of high resolution bands by appending the result of nonlinear operations on the original RGB data, thus increasing the rank of the high resolution data. Such operations are for example: interaction terms (product of different bands) or square terms:

$$\mathbf{C}_{RGB} = [\mathbf{C}_R, \mathbf{C}_G, \mathbf{C}_B] \quad (5.3)$$

$$\mathbf{C}_{int} = [\mathbf{C}_R \circ \mathbf{C}_G, \mathbf{C}_R \circ \mathbf{C}_B, \mathbf{C}_G \circ \mathbf{C}_B] \quad (5.4)$$

$$\mathbf{C}_{sqr} = [\mathbf{C}_R \circ \mathbf{C}_R, \mathbf{C}_G \circ \mathbf{C}_G, \mathbf{C}_B \circ \mathbf{C}_B] \quad (5.5)$$

$$\mathbf{C}_{ext} = [\mathbf{C}_{RGB}, \mathbf{C}_{int}, \mathbf{C}_{sqr}] \quad (5.6)$$

If we consider \mathbf{C}_{RGB} – Eq. 5.3 – to be the original RGB-only data, unfolded, where each column represents a colour channel, then we can define a matrix of first-degree interactions, \mathbf{C}_{int} , and another of square terms, \mathbf{C}_{sqr} , respectively, by element-wise multiplying each band by another, Eq. 5.4, or by itself, Eq. 5.5. The operator \circ represents the element-wise multiplication, also known as Hadamard product. When composing the \mathbf{C} matrix to input to the algorithm, we could use \mathbf{C}_{ext} if we wanted to include interaction and square terms.

Fundamentally, the interaction and square terms do not add new information, however, by providing these nonlinear terms to the linear algorithm, it allows it to find nonlinear spectral variations, which are expected to exist. This is analogous to how different wavelength channels have different nonlinear relations to the chemical sample composition in NIR multichannel reflectance measurements. These unknown but different nonlinear relations may be regarded as a special type of unknown interference. Using the pragmatic but incorrect $\log(1/R)$ transform allows linear multivariate calibration modelling, e.g., by PLSR to utilize additional subspace dimensions spanned by the channels' unknown differences in nonlinearity, to pick up and correct for these unknown interferences, as described in [92]. For our present RGB data we do not know the detailed camera properties. In addition, the light signal is affected by the atmospheric absorbance and light scattering

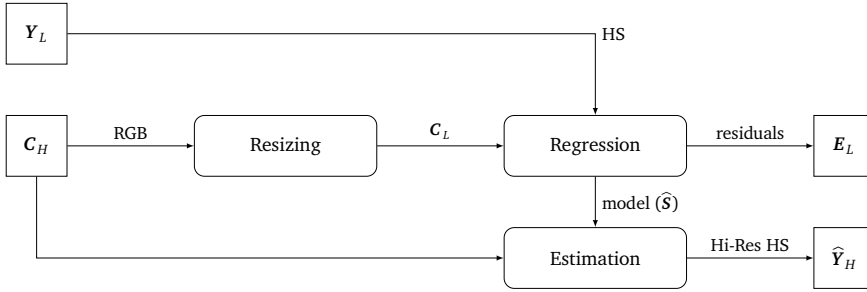


Figure 5.3: Pipeline overview diagram.

effects in the water phase that the photons have to go through, on their way from the light source, the Sun, via the bottom object and back to the camera. We do not know the ideal mathematical transform from chemical and physical properties of the objects on the bottom to the RGB signal of the camera, but the transform $\log(1/R)$ is probably too simplistic. However, by adding new "wavelength channels" by nonlinear combinations, e.g., interactions and square terms of the original RGB channels, the linear multivariate calibration has a better chance of finding a subspace that spans both the chemical and physical signal variations and their nonlinearities.

Later, when we discuss the performance of our proposed method, we show how different combinations of nonlinear terms affect it.

5.2.2 Notation

Throughout the chapter we will use the following notation:

- Unfolded HS data cube with k_H bands and low spatial resolution $n_L^h \times n_L^w$ — $Y_L \in \mathbb{R}^{n_L \times k_H}$
- Unfolded RGB data (and appended artificially generated terms) with k_L bands and high spatial resolution $n_H^h \times n_H^w$ — $C_H \in \mathbb{R}^{n_H \times k_L}$
- Enhanced HS data with high spectral and spatial resolutions — $\hat{Y}_H \in \mathbb{R}^{n_H \times k_H}$

Where $n_L = n_L^h \times n_L^w$ and $n_H = n_H^h \times n_H^w$ are respectively the total number of spatial pixels – height (h) and width (w) – in low and high resolution data.

5.2.3 Method overview

The algorithm can be summarized in the following steps, also visible in Figure 5.3:

1. **Resizing** (shrinking) high-resolution RGB data to low-resolution HS size, with image registration, to ensure that the pixels in both images represent the same ground positions.

5. *Enhancing hyperspectral data: Multivariate image fusion: A pipeline for hyperspectral data enhancement*

2. Noise weighted modelling to estimate HS from low-resolution RGB by regression over the low resolution pixels, through **Regression**. Lack-of-fit HS residuals are kept for further analysis of spectral patterns not seen in RGB.
3. **Estimation** of HS using high-resolution RGB.

5.2.4 Resizing

As mentioned before, Y_L and C_H have different image dimensions, so we start by shrinking C_H for it to coincide C_L . Resizing should take into account the properties of the HS instrument that resulted in such low resolution – i.e., if sampling frequency is low but exposure time is also low, resulting in a subsampled target, then we should resample the high resolution RGB data. If on the other hand, the exposure time is long and the target (ground) is fully sampled/observed, but the information is mixed/convolved due to motion blur in each pixel, then we should apply a similar convolution to mix the RGB pixels. This way we ensure that the ground contributions are similarly represented between the HS and RGB data.

5.2.5 Noise-balancing wavelength weights

Hyperspectral instruments have varying levels of noise for each band. The proposed methodology involves least-squares based estimation of parameters in a reduced rank regression model. For such methods, it is important to balance the noise level of the different wavelength channels. For the most common pushbroom slit-grating design, the noise is generally worse as we move away from the centre of the sensor and the operation range of the (electro-)optical components. In VIS-NIR instruments, usually with CMOS sensors, the performance degrades quickly for bands below 400nm and above 900nm, but even inside that range, the noise level varies from channel to channel. Knowing how the noise varies improves the modelling performance, by downweighing the noisy bands, we reduce the risk of over fitting noise. We estimated the noise according to the method described in 5.A.1.

5.2.6 Regression and Estimation

Once preprocessing is done, we can proceed to the core of our method, the regression step. Here we estimate S , knowing Y and C . Generically speaking, we establish a projection model:

$$Y = CS^T + E \tag{5.7}$$

Then we can apply the previous equation to our data:

$$\mathbf{Y}_L = \mathbf{C}_L \mathbf{S}^\top + \mathbf{E}_L \quad (5.8)$$

$$\widehat{\mathbf{S}}^\top = (\mathbf{C}_L^\top \mathbf{C}_L)^{-1} \mathbf{C}_L^\top \mathbf{Y}_L \quad (5.9)$$

$$\widehat{\mathbf{Y}}_L = \mathbf{C}_L \widehat{\mathbf{S}}^\top \quad (5.10)$$

$$\mathbf{E}_L = \mathbf{Y}_L - \widehat{\mathbf{Y}}_L \quad (5.11)$$

The matrix $\widehat{\mathbf{S}}$ contains the estimate of a dictionary that translates the variations in \mathbf{C}_L into variations in \mathbf{Y}_L . Furthermore, we can now use \mathbf{C}_H with that same dictionary and compute:

$$\widehat{\mathbf{Y}}_H = \mathbf{C}_H \widehat{\mathbf{S}}^\top \quad (5.12)$$

which gives a high spatial resolution estimation of \mathbf{Y} .

5.2.7 Low resolution residuals analysis

When estimating $\widehat{\mathbf{S}}$ and $\widehat{\mathbf{Y}}_L$, we are left with unmodelled low-resolution residuals – \mathbf{E}_L . These can be analysed to give us some insight into what could not be enhanced to a higher resolution – systematic information not captured by the model – and estimate how much of it was random independent noise. To do that, we need to further decompose \mathbf{E}_L , through some bilinear matrix decomposition techniques, according to the model:

$$\mathbf{E}_L = \mathbf{D}_L \mathbf{Z}^\top + \mathbf{F}_L \quad (5.13)$$

$$\mathbf{D}_L \in \mathbb{R}^{n_L \times A} \quad (5.14)$$

$$\mathbf{Z} \in \mathbb{R}^{k_H \times A} \quad (5.15)$$

where A is the number of factors, or components extracted.

Referring back to our overall model in Eq. 5.1, then \mathbf{D}_L and \mathbf{Z} contain low-resolution spatial and respective spectral information of phenomena that are not measured by the RGB camera. \mathbf{F}_L contains unmodelled noise, in low-resolution.

Matrix decomposition

Matrix decomposition or factorization, also called unmixing in the context of spectral data, is a family of methods that split a matrix into a product of other matrices. For hyperspectral data, those resulting matrices usually correspond to some type of spectral signatures and the respective spatial distribution and/or concentration. A simple yet useful factorization method is Singular Value Decomposition (SVD) [93][92]. However, a property of the resulting spectral features is that they are

orthonormal, hence not directly representative of bio/geo/chemical spectral signatures.

On the other hand, with Non-Negative Matrix Factorization (NNMF) [94], Multivariate Curve Resolution (MCR) [95] or Independent Component Analysis (ICA) [96], the spectral components (loadings) are often related to the actual spectral signatures of phenomena seen in the captured scene. This comes at the expense of more complex computation.

Hyperspectral data is notoriously rank deficient, meaning that factorization methods will model a limited number of meaningful components and many noise components.

For the implementation of the MVIF pipeline here described, we first use SVD to estimate the number of non-noise components (A) in the residuals, then we can use one of the more complex methods knowing how many components to expect. There is not a consensus among specialists regarding which is the best method to select the appropriate number of relevant components when using SVD [97], furthermore, many of them require visual inspection of plots. We propose a solution based on a voting system: three methods evaluate different metrics and vote on whether a component is relevant or not. If a component gets all of the votes, it is deemed relevant. Find more details about this method in Appendix 5.A.3.

Once we have all the votes from the 3 classifiers, we decide how many factors to keep – A . Then it is simply a matter of running the unmixing method of our choice, to obtain the factorization as in Equation 5.13.

$$E_L \xrightarrow[\text{decomp}]{\text{matrix}} \widehat{D}_{L,A}, \widehat{Z}_A \quad (5.16)$$

$$\widehat{E}_{L,A} = \widehat{D}_{L,A} \widehat{Z}_A^T \quad (5.17)$$

$$F_{L,A} = E_L - \widehat{E}_{L,A} \quad (5.18)$$

5.3 Datasets

Control Dataset

We used data from the HICO instrument, available at [98]. HICO (Hyperspectral Imager for the Coastal Ocean) was a hyperspectral imager installed on the International Space Station (ISS) that captured data from 2009 to 2014. HICO datasets have 87 bands (400-900nm), cross-track resolution of 500 pixels, with ground sample distance (GSD) of 90m.

Both the RGB representation of the HS data and the reference RGB image were, for this dataset, extracted from the HS data – the usual procedure for benchmarking



Figure 5.4: HICO dataset H2011145084342 with marked region of study.



(a) RGB representation of Y_H , which in this case is also C_H .



(b) RGB representation of Y_L .


Figure 5.5: RGB rendering of HICO scene used during tests. The procedure described in 5.A.2 was used to create both the high ($500 \times 500\text{px}$) and low resolution ($100 \times 100\text{px}$) images. Some colour adjustment was applied to make the images more aesthetically pleasing.

these types of algorithms. See 5.A.2 for more details.

We used a sample of HICO data at full resolution as the high-resolution HS reference – dataset ID H2011145084342, check Figure 5.4. This data was captured over the Strait of Dover, 25 May 2011. From this dataset we extracted a 500×500 region of interest, see Figure 5.5. The low resolution HS data is 100×100 spatial pixels and 87 bands.

Test/Field Dataset

We also include the results of applying our method to data obtained in a field trial with drone mounted HS and RGB cameras, see Figure 5.6. This flight was con-



(a) High resolution RGB – C_H .

(b) RGB representation of low resolution HSI – Y_L – stretched to same ratio of C_H .

Figure 5.6: RGB image from ZenMuse camera (1500×560 px) and RGB rendering of HS data (50×560 px) from Hopavågen tests.

ducted in Hopavågen, Norway, in March 2018. More details on the data capture and experiment are available in Chapter 3.

In this dataset, the loss of spatial resolution was due to suboptimal flight conditions and instrument design limitations. The across-track resolution (here seen as the horizontal axis) is equivalent in both RGB and HS, but the along-track direction (vertical axis) is much lower for the HS camera. In fact, for each HS pixel, there are 30 RGB pixels.

Details about the spatial registration of the two images are outside of the scope of this work. In broad strokes, a first coarse registration was possible due to timestamp synchronization between the two cameras, then fine-tuned through an image registration method – available in MATLAB as `imregister()`.

5.4 Results

In this section, we show and analyse the results of applying the presented method to the two distinct datasets. In addition, we also compare the metrics for a third benchmark dataset.

First, we generate low resolution hyperspectral data by degrading a high resolution hyperspectral data cube. RGB data is also extracted from the high resolution HS data. As we have a high resolution reference, we can quantify the performance of the algorithm using the performance metrics described in [20].

Second, we use a sample of data from a UAV field campaigns, for which this method was conceived for. RGB comes from a separate camera. Since there is no high resolution reference, performance can only be evaluated visually.

Finally, the benchmark dataset is just briefly analysed to compare the performance of MVIF with that of another method from literature on the same dataset.

5.4.1 Performance metrics and benchmark dataset

As described in [20], there is a number of metrics commonly used to evaluate the performance of enhancement methods. Those metrics are: Cross Correlation (CC), measuring the spatial enhancement, with 1 as optimal value; Spectral Angle Mapper (SAM), as the name suggests indicates spectral fidelity, 0 is ideal; Root Mean Square Error (RMSE) and Erreur Relative Globale Adimensionnelle de Synthèse (ERGAS), both global quality indices, with 0 as ideal value. In Table 5.1 we compare how the values change for different choices of additional terms in \mathbf{C} – both \mathbf{C}_L and \mathbf{C}_H .

Table 5.2 shows the performance of MVIF in a benchmark dataset – Moffet field, also studied in [20] – compared to the best performing method mentioned in that publication. According to its authors, that performance is achieved in a machine with an Intel Core i5 3230M 2.6GHz with 8GB RAM. Our method was running on an Intel Core i7 4510U 2.0GHz with 8GB RAM, so an equivalent performance is expected. Even though the results seem better – slightly better values for the quality indices, and an extreme time reduction – we do have to say that our method uses an RGB reference, while the methods discussed in that publication are using a univariate – panchromatic – high resolution reference.

Table 5.1: Performance indices of MVIF with HICO dataset. Cross Correlation (CC), 1 is ideal. Root Mean Square Error (RMSE), Erreur Relative Globale Adimensionnelle de Synthèse (ERGAS) and Spectral Angle Mapper (SAM), 0 is ideal.

Terms in \mathbf{C}	CC	RMSE	ERGAS	SAM (deg)	Time (s)
RGB	0.823	0.081	10.339	6.888	1.751
RGB, Square	0.964	0.050	6.337	5.009	1.767
RGB, Square root	0.971	0.045	4.878	5.852	1.854
RGB, Interaction	0.970	0.043	5.485	4.686	1.849
RGB, Square, Square root	0.978	0.040	4.413	5.109	1.809
RGB, Interaction, Square, Square root	0.981	0.036	4.679	3.868	1.895

5. Enhancing hyperspectral data: Multivariate image fusion: A pipeline for hyperspectral data enhancement

Table 5.2: Performance comparison in Moffett field dataset. Cross Correlation (CC), 1 is ideal. Root Mean Square Error (RMSE), Erreur Relative Globale Adimensionnelle de Synthèse (ERGAS) and Spectral Angle Mapper (SAM), 0 is ideal. Values for Bayesian Sparse method extracted from [20]. To note that Bayesian Sparse enhanced HS data using only panchromatic high resolution data, while we used RGB.

Method	CC	RMSE	ERGAS	SAM (deg)	Time (s)
Bayesian Sparse	0.982	200.158	3.426	6.625	133.61
MVIF	0.985	164.861	3.424	5.427	0.95

5.4.2 Plots

Results shown here were obtained taking into consideration the values in Table 5.1, meaning we opted for adding interaction, square and square root terms to \mathbf{C} before submitting it to the regression step, as it gave the overall best performance. In addition, we used Non-Negative Matrix Factorization (NNMF) to unmix the hyperspectral data cubes – \mathbf{Y}_L and $\hat{\mathbf{Y}}_H$ – and make it possible to represent in low dimension. Again, we estimated the number of factors using the method in 5.A.3. The plots with the spectral signatures for each component are matched in colour with the most similar between low – Figures 5.7 and 5.9 – and high – Figures 5.8 and 5.10 – resolution for each dataset.

Notice that each of the low resolution abundance maps are enhanced to high resolution, leaving no trace of low spatial resolution artifacts.

In Figures 5.11 and 5.12 we show a factorization of the residuals that could not be enhanced. Here we opted for using Independent Component Analysis (ICA), implemented as FastICA [96], instead of NNMF since the residuals are not non-negative. ICA gives a more interpretable factorization than SVD/PCA, while dealing well with possible, or in this case likely, negative concentrations. Another reason for using this method is that it is fast and can give us a clue regarding whether there is relevant data that we overlooked, or not. If there is some indication that we should further analyse the residuals, other methods such as MCR can also be applied. We stress the importance of doing such complementary analysis on the residuals. This creates awareness regarding the limitations of the method, and even if low resolution, these are still relevant data.

5.5 Discussion

Correlation in noise for artificial datasets When RGB and low resolution HS data are artificially generated for benchmarking, they both originate from the same

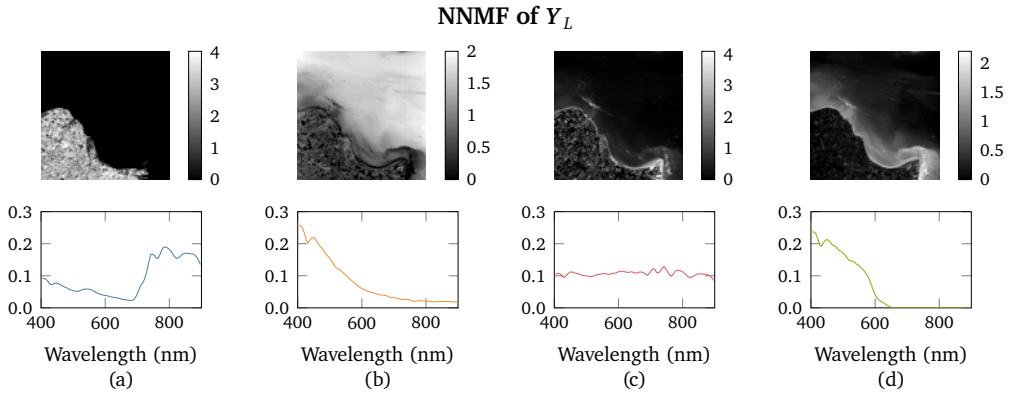


Figure 5.7: Non-negative Matrix Factorization (NNMF) of low resolution Y_L . Spatial coefficients in the top row were refolded to a 2D map. The corresponding Spectral feature – systematic radiance pattern – is shown below each map. The components shown here can be interpreted as land based vegetation (a), some combination of CDOM (colour dissolved organic matter) and phytoplankton (b and d), and an albedo-like property of pixel (c) - almost flat spectrum.

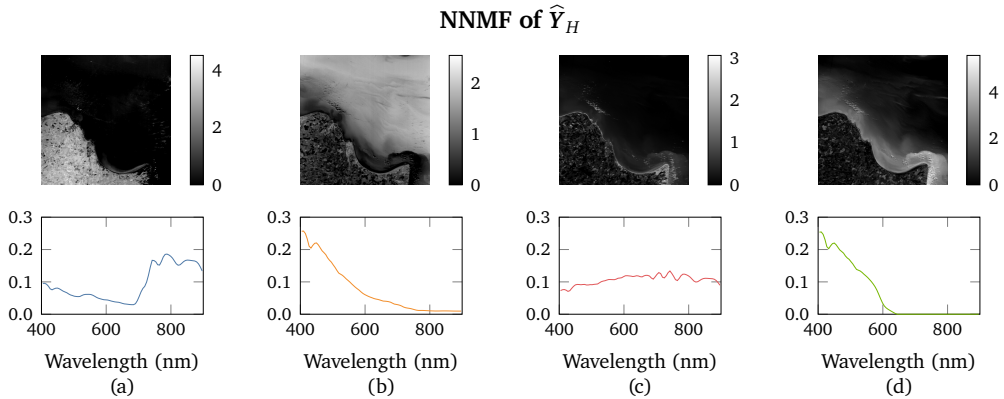


Figure 5.8: Non-negative Matrix Factorization (NNMF) of high resolution \hat{Y}_H . Spatial coefficients in the top row were refolded to a 2D map. The corresponding Spectral feature – systematic radiance pattern – is shown below each map. The components shown here can be interpreted as land based vegetation (a), some combination of CDOM (colour dissolved organic matter) and phytoplankton (b and d), and an albedo-like property of pixel (c) - almost flat spectrum.

high resolution HS reference. Naturally, there is a concern that noise in both is correlated. We have tried to minimize this issue by following the convolution, blurring

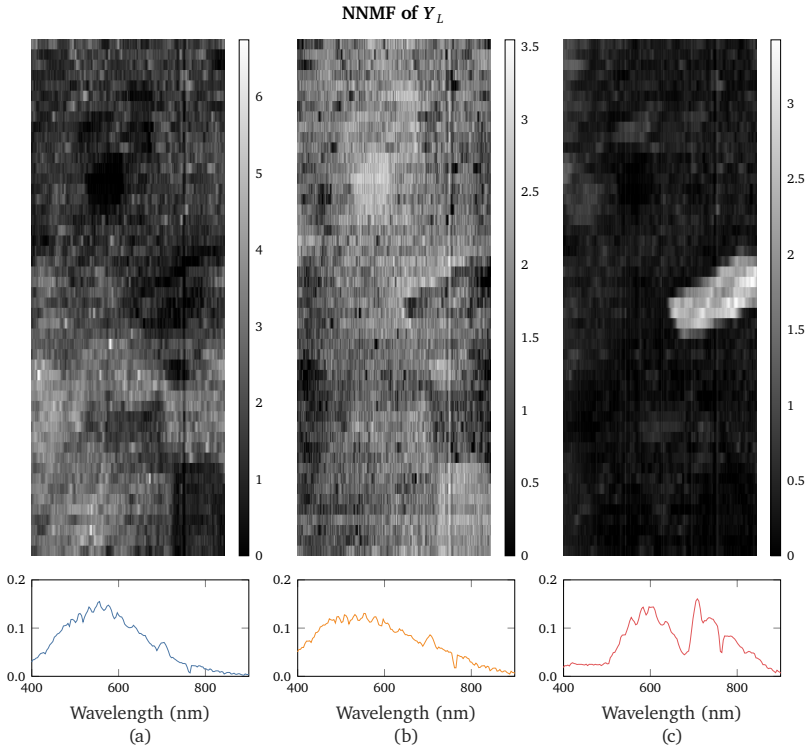


Figure 5.9: Non-negative Matrix Factorization (NNMF) of low resolution Y_L . Spatial coefficients in the top row were refolded to a 2D map. The corresponding Spectral feature – systematic radiance pattern – is shown below each map. In this dataset, raw data is not corrected for solar radiance. Component (a) is strongly influenced by the solar spectrum, showing what looks like intensity of reflected solar spectrum, affected both by the albedo of different materials and in-water path length – i.e. depth, which increases from bottom to top of image, due to ground slope – while (b) seems to pick out the darker rocks. Component (c) is isolating the orange/yellow rock, but the spectrum indicates some influence of chlorophyll, most likely due to the algae covering that rock. Looking at Figure 5.6 might help understand these components.

and downsample procedure used in [20], available at [99].

This is not an issue for our field trial dataset, as the data comes from two distinct instruments.

Image registration The work here discussed focuses on the fusion of data with different resolutions, originating from separate sensors. Here we ignore the reg-




Figure 5.10: Non-negative Matrix Factorization (NNMF) of high resolution \hat{Y}_H . Spatial coefficients in the top row were refolded to a 2D map. The corresponding Spectral feature – systematic radiance pattern – is shown below each map. Notice that all the low resolution artifacts and bad pixels disappear. In this dataset, raw data is not corrected for solar radiance. Component (a) is strongly influenced by the solar spectrum, showing what looks like intensity of reflected solar spectrum, affected both by the albedo of different materials and in-water path length – i.e. depth, which increases from bottom to top of image, due to ground slope – while (b) seems to pick out the darker rocks. Component (c) is isolating the orange/yellow rock, but the spectrum indicates some influence of chlorophyll, most likely due to the algae covering that rock. Looking at Figure 5.6 might help understand these components.

istration problem, assuming data is previously aligned and matched. A future development direction would be to integrate data matching and registration as a preprocessing step.

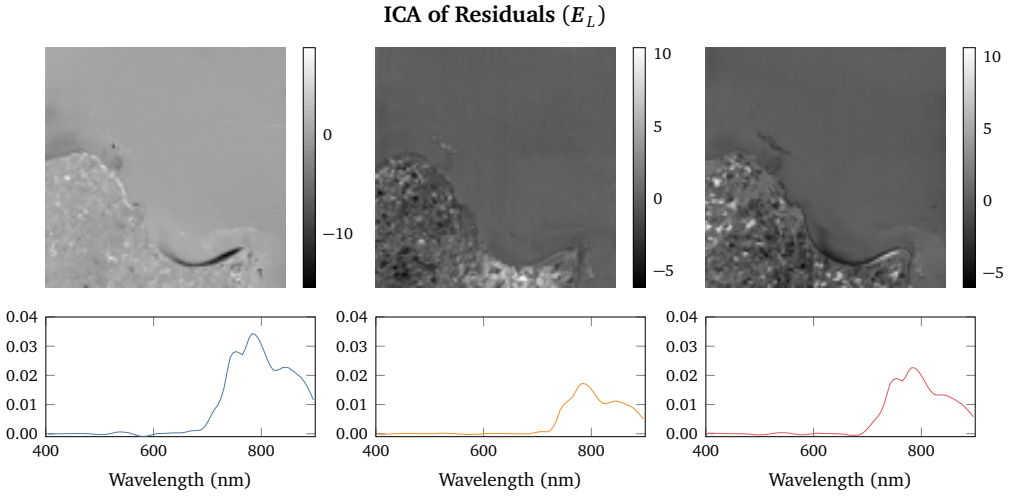


Figure 5.11: Independent Component Analysis (ICA) of low resolution residuals E_L . Spatial coefficients in the top row were refolded to a 2D map. The corresponding Spectral feature – systematic radiance pattern – is shown below each map.

5.6 Conclusions

In this chapter, we describe a pipeline for enhancing the spatial resolution of HS data, by taking advantage of co-located RGB data. The method is simple and fast, while giving good quality results.

Furthermore, as the epithet pipeline indicates, the method is composed by a sequence of steps. Here we describe a possible pipeline, where we use a simple projection in the regression/estimation step, and expand the number of variables of the RGB data in a certain way. However, the reader may find that for their application, a nonlinear regression method and different variations of the high resolution data could work better.

The main contribution of this research is to provide a template for connecting functional blocks, with the aim of fusing multivariate datasets with different resolutions.

5.A Companion methods

5.A.1 Noise Estimation

The method for estimating the noise level per channel is very simple and intuitive, nonetheless the results match the expected instrument performance.

In essence, we check how rough, i.e. non-smooth, each band image is. Gaussian

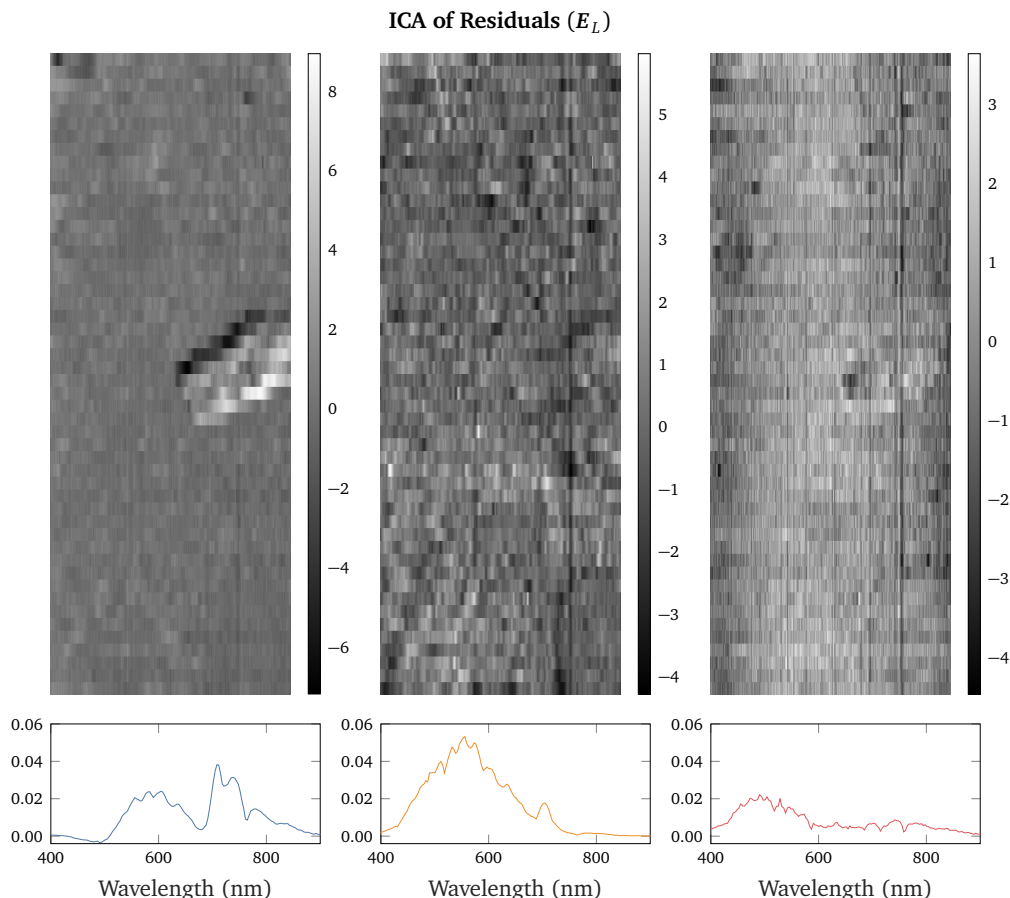


Figure 5.12: Independent Component Analysis (ICA) of low resolution residuals E_L . Spatial coefficients in the top row were refolded to a 2D map. The corresponding Spectral feature – systematic radiance pattern – is shown below each map.

noise will cause sharp peaks very visible when taking the second difference along the horizontal and vertical axes. By looking at the absolute values, we get a direct indication of how noisy a pixel is.

```
# Diff twice along each direction (vertical and horizontal), for each band image
diff0 = abs(diff(Y_3d, n=2, axis=0))
diff1 = abs(diff(Y_3d, n=2, axis=1))

# Flatten diff results
diff0_flat = reshape(diff0, [diff0.shape[0] * diff0.shape[1], diff0.shape[2]])
```

5. Enhancing hyperspectral data: Multivariate image fusion: A pipeline for hyperspectral data enhancement

```
diff1_flat = reshape(diff1, [diff1.shape[0] * diff1.shape[1], diff1.shape[2]])
```

To get an overall value per band we can either use the median of all values, horizontal and vertical together:

```
# Vertical and Horizontal all together  
diff = concatenate((diff0_flat, diff1_flat), axis=0)  
noise = median(diff, axis=0)
```

Or average vertical and horizontal noise:

```
# Average of Vertical and Horizontal noise levels  
median0 = median(diff0_flat, axis=0)  
median1 = median(diff1_flat, axis=0)  
noise = (median0 + median1) / 2
```

Once the noise level per band is known, see Figure 5.13, the inverse of that is used as weights.




Figure 5.13: Noise level per band of low resolution HICO data.

5.A.2 Artificial data generation

Here we describe our process to generate RGB and low resolution data from a high resolution HS reference.

RGB data

Unlike common practice, instead of picking 3 bands out of the HS data cube, we simulate the sensitivity of an RGB sensor. RGB sensors do not have very narrow

band-pass filters for Red-Green-Blue wavelengths, so picking one single band to represent each channel would not give us a realistic dataset. Therefore, a weighted sum of different bands extracted from the HS data, with the weights based on the specification sheet of a CMOS RGB sensor, was used for each channel instead, see Figure 5.14. This way we ensure, through sensible assumptions, that the RGB data




Figure 5.14: Weights per band. Weights are designed to simulate an RGB sensor. The black dashed line represents the NIR cut-off filter, usually present in RGB sensors. The vertical scale is unimportant here, we merely want to show the shape of the curves.

are realistic.

Low resolution HS data

For validation, we generated low resolution data using the same method as the study in [20], available as MATLAB code in [99].

5.A.3 Relevant Components

When analysing the residuals before unmixing, the first step is to decompose them through SVD. Then U (Scores), Σ (singular values), and V (Loadings) are evaluated by different methods:

1. the noise level of the spatial distribution maps (U)
2. the slope of the scree plot [100] (Σ)
3. the smoothness of the spectral signatures (V)

Noise level of the spatial distribution maps The same noise per channel routine that was used to find the noise in the raw data is re-used, now on the refolded distribution maps (U). Relevant components are expected to have little noise in

5. Enhancing hyperspectral data: Multivariate image fusion: A pipeline for hyperspectral data enhancement

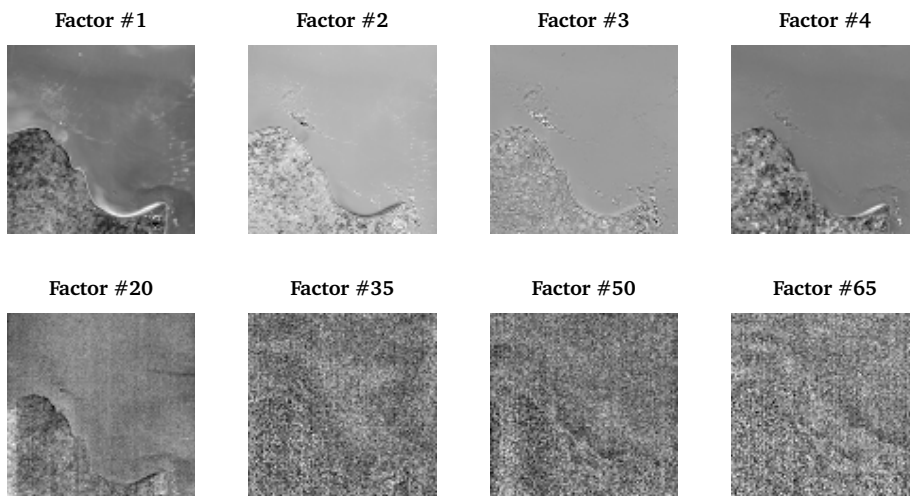


Figure 5.15: Spatial Maps (Scores) of a sample of factors from the HICO dataset. Higher factors are clearly more noisy than the lower/earlier.

the spatial domain. Note that U is not scaled with Σ , so every band has a similar range, and the same threshold can be applied. See Figure 5.15

Slope of the scree plot This is an implementation of an autonomous scree test [100], which is usually a visual inspection test. The scree plot will have an "elbow", which represents the boundary between relevant and non-relevant factors. Through a linear fit, we find the slope of the plateau that corresponds to the noise components, then we start checking lower numbered components until the slope of the linear fit starts to change, indicating we have reached the elbow. See Figure 5.16.

```
# Maximum value is always 1,  
# this way we can use same threshold values for different data  
s /= s.max()  
nc = len(s)  
  
# More than half the components are usually noise in remote sensing data,  
# so we start from the middle  
curr_idx = nc // 2  
z = polyfit(x=range(curr_idx, nc), y=s[curr_idx:], deg=1)  
  
while curr_idx > 1:
```

```

curr_idx -= 1
z_new = polyfit(x=range(curr_idx, nc), y=s[curr_idx:], deg=1)
fit_chg = abs(z - z_new)
# If slope changes too much from previous fit, we are past the elbow
if fit_chg[0] > threshold_fit:
    break
else:
    z = z_new

p = poly1d(z)
fit_p = p(range(nc))

serr_all = (s - fit_p)**2
serr_fit = serr_all[curr_idx:]

# Relevant components will have large error to fitted line (not in flat region)
vote_slope = serr_all > threshold_error

```




Figure 5.16: Scree plot showing only the 10 first components of the low resolution residuals from the HICO dataset. Red dashed line was fitted to the flat section. When the dashed and solid lines diverge we have the relevant factors (marked with filled circles).

Smoothness of the spectral signatures Even though the loadings matrix (V) resulting from an SVD of hyperspectral images cannot be directly interpreted as spectral signatures of natural phenomena, which are usually smooth, they are linear combinations of smooth signatures. For sufficiently high spectral resolution sensors, which we assume a hyperspectral camera has, this means that informa-

5. Enhancing hyperspectral data: Multivariate image fusion: A pipeline for hyperspectral data enhancement

tion is smooth in the spectral domain, and noise is not. A smooth spectrum will have small variations in slope. See Figure 5.17.

```
# Absolute sum of diff along wavelength axis
# large values will show for rough spectra
v_diff = diff(v, n=2, axis=1)
sum_diff = sum(abs(v_diff), axis=1)

# Relevant components have smooth spectra
vote_smooth = sum_diff < threshold_smooth
```

Once we have all the votes from the 3 classifiers, we decide how many factors to




Figure 5.17: Example Spectral Signatures (Loadings) split into smooth or non-smooth according to our classifier. The high frequency variations seen in the spectra on the right plot are often correlated with noise. From the HICO dataset.

keep. Figure 5.18 shows the result of such voting.




Figure 5.18: Votes from each relevance classifier for the HICO dataset. Factors with 3 votes are considered relevant. Here, 3 relevant factors should be possible to extract from residuals.

Chapter 6

Dealing with big data streams

On-The-Fly Processing of continuous high-dimensional data streams

This chapter presents a novel method and software system for rational handling of time series of multichannel measurements¹. This quantitative learning tool, the *On-The-Fly Processing* (OTFP), develops reduced-rank bilinear subspace models that summarize massive streams of multivariate responses, capturing the evolving covariation patterns among the many input variables over time and space. Thereby, a considerable data compression can be achieved without significant loss of useful systematic information.

The underlying OTFP methodology is relatively fast and simple: it is linear/bilinear and does not require a lot of raw data or large cross-correlation matrices to be kept in memory. Unlike conventional compression methods, the approach allows the high-dimensional data stream to be graphically interpreted and quantitatively utilized in its compressed state. Unlike adaptive moving window methods, it allows all past and recent time points to be reconstructed and displayed simultaneously. This new approach is applied to four different case studies: i) multichannel Vis-NIR spectroscopy of the Belousov-Zhabotinsky reaction, a complex, ill-understood chemical process; ii) quality control of oranges by hyperspectral imaging; iii) environmental monitoring by airborne hyperspectral imaging; iv) multi-sensor process analysis in the petrochemical industry. These examples demonstrate that the OTFP can automatically develop high-fidelity subspace data models, which simplify the storage/transmission and interpretation of more or less continuous time series of high-dimensional measurements – to the extent that there are covariations among

¹Based on the paper "On-The-Fly Processing of continuous high-dimensional data streams", by Raffaele Vitale, Anna Zhyrova, João Fortuna, Onno E. de Noord, Alberto Ferrer and Harald Martens, published in *Chemometrics and Intelligent Laboratory Systems*, Volume 161, 2017.

the measured variables.

6.1 Introduction

6.1.1 Data compression strategies

Data compression plays a central role in telecommunications and many other scientific and technological branches of interest [101]. According to the nature and features of the algorithmic procedure through which it is performed, it can be defined as either *lossless* or *lossy*. Lossless methods utilize statistical distribution properties and simple patterns in the data for compression, converting the inputs into compressed bit series².

Lossy compression techniques – e.g., the various dedicated versions of JPEG and MPEG methods used for digital image, video, and sound compression – approximate the main, perceptible variations in the input data by local *ad hoc* patterns, filtering out less perceptible variation types and noise. Lossy approaches are commonly much more efficient (in terms of compression rate) than lossless ones, like *algebraic* zipping, but allow the original input to be only roughly restored. Moreover, when set to compress too much, they not only cause loss of valid information (resulting in, e.g., image blurring or loss of high-frequency sound), but can also introduce undesired decoding artefacts (e.g., visible block effects or audible errors). Whether lossless or lossy compression methods are used, the compressed data are represented by *per se* meaningless streams that cannot be directly used for quantitative calculations, mathematical modelling, or graphical representation.

The novelty of the developed *On-The-Fly Processing* (OTFP) tool is represented by the fact that a hitherto under-utilized source of redundancy (the intercorrelation usually evolving in multi-channel data streams) is mathematically modelled to prevent significant loss of useful systematic information carried by the original measurements. Based on the model's automatically estimated parameters, the data stream may be interpreted and utilized for prediction, forecasting, and fault detection in the compressed state. The idea behind this strategy was outlined in [102]. Here, more algorithmic details will be given and its applicability to different types of high-dimensional data streams is demonstrated.

Conceptually, the OTFP system may be motivated by the following thought experiment: assume that a space probe should be constructed and sent out to explore, for the first time, the unknown geological properties of the hidden back side of a remote planet, using a hyperspectral camera. Prior to the launch, scarce knowledge about this planet is available to design the ideal instrument, and after the probe

²Most of the lossless compression approaches, such as standard file *zipping*, recodes the original input by using shorter bit sequences for often encountered data and larger ones for rare data.

has landed, it is too late to change anything. Which wavelengths should be chosen, and how should the imaging data be transmitted back to Earth? Some individual wavelengths distinguishing between already known, earthly rock types might be included. However, possible geological *surprises* should also be taken into account. Therefore, it is decided to equip the probe with a wide spectral range camera, capable of measuring, say, 1000 different wavelength channels. However, the limited communication bandwidth then becomes a problem: the probe cannot transmit all those measurements for every point in time and space.

What would be the best way to send spectral data back to Earth? Perhaps, could that be automatically decided on-the-fly by the space probe's computer itself, based on what its camera measures? The on-board computer could be programmed to discover, compress and transmit the essence of all the recorded images, in a continuous learning-and-communicating process that never sends the same information twice.

But how to quantify this compact spectral essence comprehensively? To understand the unknown geological landscape, a reliable approximation of the spectral profile of every pixel in every image, with as many spectral and spatial details and as few artefacts as possible, is needed. A lossless multivariate spectral preprocessing followed by a continuously developing bilinear compression/classification model could deliver a compact summary of the sequence of hyperspectral image data, which would yield maximal insight here on Earth from the limited quantity of received data. The first three application examples described below will illustrate this, albeit in more mundane settings.

6.1.2 Subspace compression

The OTFP is based on evolving bilinear subspace modelling. The software automatically detects systematic patterns of covariation in the data and uses these to model the data mathematically. Subspace projection and dimensionality reduction techniques based on bilinear models, e.g., Principal Component Analysis (PCA), constitute one of the possible ways to compress and approximate a certain set of data, removing simultaneously both statistical redundancy and uninformative noise. Their basic principles can be summarized as follows: let $j = 1, 2, \dots, J$ be the number of input channels (J wavelengths of light per pixel in a hyperspectral camera, J sensor variables monitored during a dynamic industrial process or even J metabolites quantified in biological samples) recorded for each of $n = 1, 2, \dots, N$ measurements performed, for instance, on N objects on a conveyor belt, at N spatial locations, N time steps or N different experimental conditions. In the present-day instrumental context, where J might be very large, the useful information carried by such data structures ($N \times J$ matrices) is usually intercorrelated among various input channels over the continuously growing set of registered measurements.

In these circumstances, for a chosen degree of acceptable accuracy (depending on the amount of data variance explained), it is possible to reduce the J -dimensional space of the original descriptors to an A -dimensional subspace, onto which all the N objects under study can be projected and represented as new points. *Prima facie*, as $A < J$, this projection can be regarded as a compression operation, whose efficiency is related to the ratio $\frac{A}{J}$.

6.1.3 PCA bilinear structure model

The well known PCA bilinear approximation of a generic $N \times J$ matrix of observed data, X , can be described by the following structure model:

$$X = \mathbf{1}\mathbf{m}^T + TP^T + E \quad (6.1)$$

where $\mathbf{1}$ ($N \times 1$) is a vector of ones, \mathbf{m} ($J \times 1$) contains a typical profile, e.g. the mean values of the J input variables in X , P ($J \times A$) is a matrix of *loadings* associated to such input variables, which determine the A basis vectors or *components* of the PCA subspace, T ($N \times A$) defines the projection coordinates or *scores* of all the N considered objects (locations, time points or experimental conditions) on this lower-dimensional space and E ($N \times J$) represents the matrix of unmodelled residuals, i.e. the portion of X not *explained* by the model at the chosen rank, A [103].

The PCA solution may be formulated in different, equivalent ways. Here, it is assumed to show the following properties:

$$P^T P = I \quad (6.2)$$

$$T^T T = \text{diag}(\lambda_A) \quad (6.3)$$

where I is an identity matrix of dimensions $A \times A$, while the a -th element of λ_A ($A \times 1$) corresponds to the eigenvalue of the a -th PCA component.

One of the most critical points when deriving the PCA approximation of a set of data is how to choose the A components of its subspace to prevent losing important portions of useful information and to filter out uniquely statistical redundancy and uninteresting noise. Some of these A dimensions may sometimes be defined according to prior knowledge of the investigated system. For instance, the number of known chemical constituents of mixtures characterized by spectroscopic methods might be appealed to for this purpose. However, in cases like this, also more or less *unexpected* constituents and/or physical phenomena may affect the performed measurements, generating new patterns of variation and thus new subspace dimensions which need to be retained for a proper data approximation and interpretation. Therefore, at least to a certain extent, the identification of the new basis vectors associated to these unforeseen sources of variability has to be carried

out through a preliminary exploratory analysis of the available empirical records. If a continuous data stream is dealt with and N rapidly grows over time, correctly determining new possible subspace dimensions is even more complex: new, unexpected patterns of covariation may spring up in the information flow. Therefore, in such situations, it becomes crucial to automatically recognize when the set of initial basis vectors needs to be reestimated and extended, and to address this task in a statistically valid and computationally efficient way.

6.1.4 PCA as a multivariate series expansion of the underlying data generation mechanism

As outlined in [104–106], the bilinear PCA model can be thought of as a Taylor expansion of the function f defining how the measurement descriptors are jointly related to their common structure. For instance, for each of the J aforementioned input channels, one can envision a local linear approximation of the underlying (unknown) causal phenomena driving their evolution. Mathematical summary modelling of such J local approximations (achieved by PCA or related methods like Partial Least Squares Regression – PLSR – Independent Component Analysis – ICA – or non-linear versions of these) can detect and display their main patterns of covariation. This can unveil the underlying causality of the data generation mechanism.

6.1.5 Algorithms for PCA decomposition

The PCA approximation of a certain dataset can be efficiently attained by a variety of algorithms, among which the most widespread and popular one is certainly Singular Value Decomposition (SVD) [107]. However, if N is very high, classical SVD may be very demanding in terms of both CPU load and memory requirements. In the last few years, several variants of classical SVD have been proposed for performing PCA on very large matrices without entirely keeping them in the computer memory (*out-of-core*) [108–112]. *Out-of-core* PCA can be carried out by different procedures:

- a $J \times 1$ cumulative sum vector and a $J \times J$ cross-product matrix may be accumulated over time, combined and used for eigen-analysis of the covariance in X , which yields the PCA loadings. That is appropriate for parallelization, but then the scores for the past time or space samples are lost;
- if also J is very high (e.g., thousands of wavelengths in an hyperspectral camera monitoring a certain scene or process), the $J \times J$ covariance matrix cannot be easily handled. Evolving moving-window/recursive PCA approaches may then be used instead, working on the most recent subset of observations.

However, that gives problems when comparing past and present records in graphical *scores plots*.

In the attempt of overcoming all these limitations, the OTFP tool is here proposed. The purpose of the OTFP is to identify systematic trends and patterns in high-dimensional data flows, compress these and display them graphically, in addition to automatically detect outliers – key points to be addressed when continuous quantitative data streams are dealt with [113]. Based on what was detailed before, it represents an extension of classical bilinear PCA, specifically developed for processing multi-channel records as soon as they are collected. It extracts patterns of covariation between the input variables by comparing previous and new observations and thereby identifying and modelling new variation phenomena, without needing large amounts of data or parameters to be retained in memory. Given, for instance, a continuously growing stream of high-dimensional data, the OTFP modelling system gradually develops a minimal bilinear summary model of the input data stream. For each point in space and/or time, already established components are quantified as spatiotemporal scores by projection of their multi-channel loadings. Furthermore, new, unmodelled patterns of covariation are automatically detected, refined, and quantified in terms of additional spatiotemporal scores and multi-channel loadings, then appended to the OTFP model. Hence, unlike bilinear moving-window solutions, this dynamic model extension is executed so that the system preserves the quantitative connection between all past and present records. Yet it does not need to retain all past inputs or bilinear scores in memory – for long-lasting processes the memory usage would grow prohibitively high. Besides, the OTFP system does not require to hold and update a huge $J \times J$ covariance matrix – for many applications that would also be of a prohibitive size. Instead, it repeatedly stores the necessary scores and loadings, avoiding excessive memory usage during the process.

6.2 System overview

The present OTFP algorithm (schematically outlined in Figure 6.1) is characterized by three fundamental aspects: i) its self-learning and ii) adaptive nature and iii) its stabilising modelling principles. It allows massive amounts of data collected along time to be compressed and modelled with minimal loss of significant information content. The algorithm is initiated with the preliminary choice of a typical input vector, \mathbf{m} , and the best guess of which weights to give to the different input channels for balancing their variances, \mathbf{c} . In addition, a set of predefined component loadings, \mathbf{P} , derived for instance from an initial exploratory investigation of the system under study and representing systematic variation patterns expected to affect the incoming data stream, may or may not be supplied. Then, various system

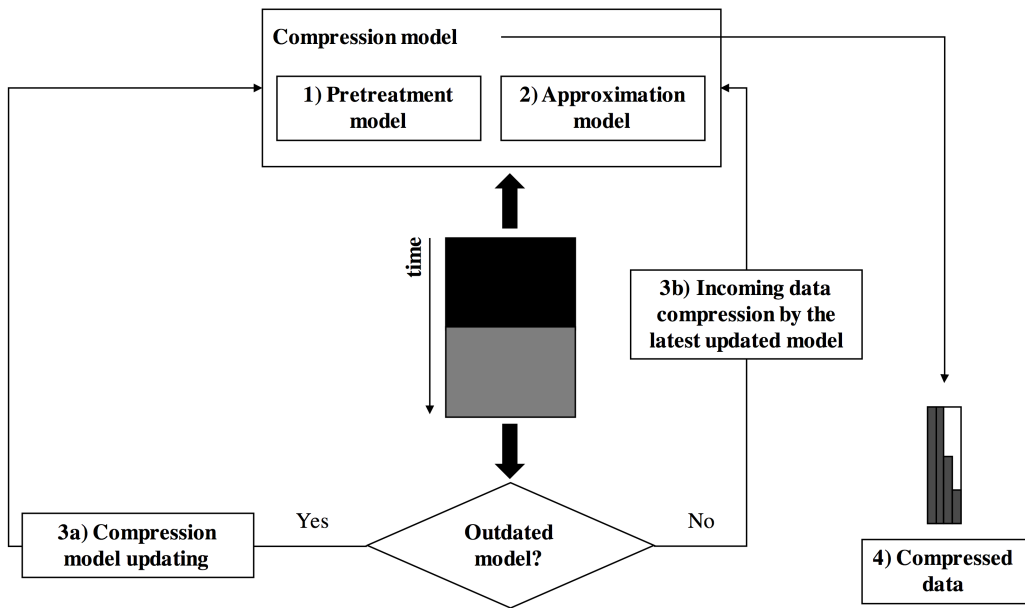


Figure 6.1: Schematic representation of the OTFP algorithm: a first set of data (black block) is input to 1) a pretreatment and 2) a PCA-based dimensionality reduction stage. As new measurements are recorded (grey block), they can be either 3a) exploited for the reparametrization of the compression model, if it is found to be outdated, or 3b) just approximated by its latest version. 4) Bilinear approximation loadings and preprocessing parameters are saved by keeping track of how they have been initially defined and/or changed during model updating. The time series of bilinear approximation scores are more or less continuously stored and deleted from memory to subsequently process new input data.

parameters such as the desired modelling fidelity – the fraction of data variance the OTFP model has to capture, also known as the amount of *explained* data variance – need to be specified. As the multi-channel data starts to flow, it may deliver a more or less continuous stream of individual J -dimensional input records, e.g., a set of measurements collected by the same set of J simple channels or sensors during the evolution of an industrial process. Alternatively, it may deliver a sequence of input data blocks – batches – each containing N_g records ($g = 1, 2, \dots, G$) and the same set of J channels, e.g., N_g spectral profiles, constituted of absorbance values measured at J wavelengths and associated to individual pixels of an hyperspectral image. Such records are then treated by the following procedure:

1. The J -dimensional data are (optionally) submitted to a lossless preprocessing, linearising the responses and balancing the variable variances to ease

the subsequent bilinear modelling. This step is domain-specific and the way it is executed has to be set *a priori*. For this reason, the best pretreatment strategy should always be selected based on both the nature of the instrumental equipment and technical knowledge;

2. The preprocessed data are projected onto the subspace defined by the bilinear loadings, P , already established at this point in time, to estimate the scores for the respective components;
3. The residuals left in the data after the projection on known components are input to a bilinear (here PCA-like) modelling stage to detect new unknown components and isolate outliers. If new components are found, they are quantified in terms of new scores and loadings. Thus, the statistically redundant J original variables are replaced by a smaller number (A) of principal components (PCs). The number of such components determines the degree of fidelity initially specified by the user. The algorithm automatically learns to identify and quantify all the systematic types of covariation in the data stream as it flows, while most of the random measurement errors and individual or irrelevant outliers are removed, provided the latter do not constitute a new pattern of variation. This compressed representation is suitable for graphical interpretation and quantitative use, and from it the pretreated data can be reconstructed at any time;
4. At regular intervals, the OTFP model may be refined and reorthogonalized in a linear updating stage;
5. The pretreatment information associated to the different blocks is stored as output together with the approximation model scores and loadings.

As specified before, the OTFP algorithm detects all the systematic types of covariation in the data stream – be it from the flow of observed objects (expected information) or from the measuring process itself (unexpected information, anyway needed for reliable interpretation and quantitative use of the data). Phenomena considered irrelevant during preprocessing, as well as individual outliers discovered by the OTFP algorithm, are noted and then excluded from the self-modelling process. So is much of the random, independent measurement error, since it does not represent a systematic pattern of covariation.

At any time, the systematic part of the data stream can be reconstructed from the data model, e.g., for visualization. However, this reconstruction is not mandatory; the compressed data model parameters, representing the known and/or unknown types of systematic phenomena in the data stream, are themselves suitable for efficient storage and transmission, human graphical interpretation and applied quantitative usage.

These steps will now be described. For further details, the reader may contact either the corresponding or the last author.

6.2.1 Input

The *ever-lasting* raw data stream, \mathbf{X} , divided into a sequence of blocks, \mathbf{X}_g ($N_g \times J$, $g = 1, 2, \dots, G$), is submitted to the optional preprocessing stage, which includes a linearization and a signal-conditioning step, and then to the OTFP self-modelling. The number of observations encompassed by these blocks can be freely set by the user. Unless the preprocessing parameters and the OTFP centre and scaling vectors (\mathbf{m} and \mathbf{c}) are established *a priori*, the start of the modelling process (i.e. for \mathbf{X}_1) requires sufficient observations to enable a precise and relevant initialization of them.

Linearization

The linearization of the input data in \mathbf{X}_g is domain-specific. For instance, non-linearities in light spectroscopy data may be reduced by transformation of the recorded light intensity, I , at each wavelength, first to transmittance, $T = \frac{I}{I_0}$, where I_0 represent the blank signal, and then to absorbance, $A = \log \frac{1}{T}$, to better conform to Beer's law of linear chemical responses.

Another aspect of the linearization is to convert non-additive variation types (e.g., multiplicative light scattering in absorbance spectra, motions in RGB or hyperspectral videos) into additive signal contributions or preprocessing parameters. For instance, multi-channel pretreatments such as Standard Normal Variate (SNV) [114], Multiplicative Scatter Correction (MSC) [115, 116] and Extended Multiplicative Signal Correction (EMSC) [117] can reparametrize multiplicative effects. Two-domain IDLE modelling [102] can convert confusing motion effects into nicely additive motion flow fields. Domain transforms, like Fast Fourier Transform (FFT) and wavelet analysis can change data locally from time to frequency domain. This more or less lossless, model-based preprocessing, may produce additional parameters which may be highly informative and must be stored for later data reconstruction.

Weighing the variables for better signal conditioning

In general, for an optimal data approximation, the J originally measured descriptors in \mathbf{X}_g are approximately centred, e.g., by subtraction of their mean values estimated from the data digested up to the current step. They may then be weighed to ensure a better balance among their variances so that:

$$\mathbf{X}_{g,p} = (\mathbf{X}_g - \mathbf{1}\mathbf{m}^T) \circ \mathbf{1}\mathbf{c}^T \quad (6.4)$$

where $\mathbf{1}$ ($N_g \times 1$) is a vector of ones, \mathbf{m} ($J \times 1$) and \mathbf{c} ($J \times 1$) contain the model centre and the input weighing factors (these weighing factors could be defined as the

inverse of the standard deviation values of the J recorded variables at the current step), respectively, while \circ identifies the element-wise (Hadamard) product. The same pretreatment is applied to all consecutive data blocks until m and/or c are readjusted as part of the model updating operation (see below).

6.2.2 Fit to already established model subspace

The linearized, centred and weighed records in $X_{g,p}$ are now projected onto the already established loadings P (if they exist at the current step), according to the linear structure model:

$$X_{g,p} = T_{g,p}P^T + E_{g,p} \quad (6.5)$$

Clearly, the frequency at which such a projection step is carried out depends on the number of observations in $X_{g,p}$, that is, as aforementioned, a user-defined parameter³.

6.2.3 Bilinear model expansion

In the present implementation of the OTFP, once calculated, the residual vectors in $E_{g,p}$ are examined: if they are deemed small enough to be considered uninteresting noise, the respective original records are simply discarded and their scores gathered in $T_{g,p}$. If this is not the case such residual vectors are introduced into a temporary repository to check whether they represent a new systematic trend in the data stream or not. At regular intervals or when its size or variance exceeds a specific user-defined threshold, this temporary repository is used for the estimation of a new set of loadings and scores. If their respective factors are found to explain a sufficiently high amount of the repository variation⁴, these new scores and loadings are appended to those of the already established PCs in P and $T_{g,p}$, respectively. Otherwise, if leverage analysis of the new scores points out that only scattered objects have contributed to them, these are dismissed as incidental outliers, their scores are stored, and the original model is retained.

Since the size of the entire scores matrix can become very large as the information flows, the scores are saved to the local disk at regular intervals and then deleted from memory along with $X_{g,p}$ and $E_{g,p}$ ⁵.

³In the case-studies described in Section 6.4, the projection frequency was found to affect only the computational time of the algorithmic procedure (as it increases, the number of data blocks the OTFP has to consecutively handle becomes larger) but not its final outcomes.

⁴The scores for these new PCs are – implicitly – defined to be zero for all the previous observations.

⁵In the case-studies described in Section 6.4, the storage of the scores on the local disk proved not to constitute a limiting step for the execution of the OTFP algorithm.

6.2.4 Model updating

Whenever necessary (e.g., if the model is characterized by a relatively high bias), preprocessing parameters, loadings and scores for both old and new observations are readjusted to ensure PCA-like orthogonality and thus a more efficient compression of the data. For such an updating, the OTFP does not need to recall the whole array of scores stored on the local disk, but directly operates on two summary indices of such an array, which are kept in memory in place of it (namely its column-wise cumulative sum vector and its cross-product matrix). The dimensionality of the reestimated model is automatically established according to the user's desired optimization criterion. Here, for simplicity, the percentage of data variance that has to be captured is used. This allows the original information stream to be retrieved with a predetermined reconstruction accuracy. Other criteria, based on the statistical significance of the eigenvalues associated with the single components [118, 119], may also be exploited.

6.3 Datasets

To evaluate the potential of the proposed method, 4 different sets of time series data were compressed and modelled as detailed before and reconstructed afterwards:

- High-speed multi-channel monitoring of a chemical reaction: 4329 multi-channel Vis-NIR spectra were measured in-line between 400 and 1098nm (350 wavelengths) via a NIRS 6500 spectrophotometer, equipped with a fibre-optic bundle, during several replicates of the self-oscillating Belousov-Zabhotinsky (B-Z) reaction [120]. The final matrix had dimensions 4329×350 . This example is intended to illustrate a new way to handle more or less continuous, high-dimensional measurements of a complex dynamic system not yet fully understood from a scientific point of view;
- Detailed remote characterization of a set of related, complex objects: three 245×210 -sized hyperspectral NIR images of three oranges were registered within the near-infrared spectral range 898-1690nm (247 wavelengths) by a XEVA-FPA-1.7-320 line-scanner camera (Xenics, Belgium). To enable their handling, such three-way arrays need to be unfolded into a unique matrix, so that a single pixel spectrum is contained in each one of its rows. After background removal, its dimensions were 72365×247 . This example was chosen to illustrate how non-invasive bio-spectroscopy can reveal hidden aspects of related complex biological samples;
- Airborne environmental surveillance: an hyperspectral image was recorded by a push-broom device installed on an Unmanned Aerial Vehicle (Drone-

Spex, Norut AS – University Centre in Svalbard – Norwegian University of Science and Technology, Norway [121]), flying over Faial (Azores, Portugal). At each accumulation step, the optical sensor collected the absorbance values at 450 wavelengths in the visible light range between 420 and 640nm for a strip of 245 pixels. A total number of 1000 consecutive snapshots were captured, which led to a three-way array of dimensions $1000 \times 245 \times 450$. Moreover, in this case, it was unfolded into a 245000×450 matrix. This example is intended to show how data from a modern environmental monitoring instrument, a drone, can be automatically compressed for efficient storage and transmission and interpreted in the compressed state;

- Traditional industrial process analysis: 76 engineering variables, mainly including temperatures, pressures and flow rates, were recorded at hourly intervals to follow the evolution of a continuous industrial process. The complete data structure had dimensions 14561×76 . This example illustrates the application of the OTFP to records measured over time by a relatively small set of conventional sensors.

6.4 Results and discussion

The power of the OTFP approach and the quality of the initial data retrieval were assessed in all case-studies at hand according to the following indices:

- A : number of extracted PCs;
- EV_{raw} : percentage of explained raw data variance;
- EV_p : percentage of explained preprocessed data variance;
- $RMSRE$: Root Mean Square Reconstruction Error defined as $\sqrt{\frac{\sum_{n=1}^N \sum_{j=1}^J (x_{n,j} - \hat{x}_{n,j})^2}{NJ}}$, where $x_{n,j}$ is the (n, j) -th element of X and $\hat{x}_{n,j}$ refers to its respective reconstructed value;
- t_c : compression time expressed in seconds⁶;
- CR : compression ratio⁷.

EV_{raw} , EV_p and $RMSRE$ are strictly related to the OTFP approximation accuracy degree, while A , t_c and CR can be considered measures of computational speed and efficiency.

Calculations were executed by using Idletechs' prototype software⁸ in a Matlab

⁶ t_c is computed as the time needed to compress the entire concerned dataset.

⁷ CR is computed as the ratio between the memory usage of the uncompressed and compressed (preprocessing parameters, scores and loadings matrices) data structures, both saved as double precision .mat files.

⁸www.idletechs.com

Table 6.1: Vis-NIR light absorbance spectra from the B-Z reaction: values of the compression quality indices. The number of original measured variables is reported in the first column.

J	A	EV_{raw}	EV_p	$RMSRE$	t_c	CR
350	10	99.93	99.61	0.0019	12.5	26.81 ($\frac{9768173 \text{ bytes}}{364370 \text{ bytes}}$)

R2012b⁹ environment, set up on a MacBook Pro equipped with a 2.3GHz Intel Core i7 and 8GB, 1600MHz DDR3 RAM.

6.4.1 High-speed multi-channel monitoring of the Belousov-Zhabotinsky reaction

Table 6.1 lists the values of the aforementioned parameters related to the Vis-NIR data compression. The initialization measurements were centred and weighed ($\mathbf{c} = \frac{1}{\mathbf{m}+0.05}$) after baseline correction¹⁰. The model centre vector, \mathbf{m} , was updated at regular intervals as new spectroscopic details were encountered in the process, while the variable weighing vector, \mathbf{c} , was kept constant for simplicity.

In order to more clearly appreciate the performance of the OTFP, 3 uncompressed and reconstructed spectra associated to different reaction stages are displayed in Figure 6.2. The full approximation model is sketched in Figure 6.3, in terms of final model mean (Figure 6.3a), chosen weighing factors (Figure 6.3b), de-weighted and scaled loadings (Figure 6.3c) and lack-of-fit residuals (Figure 6.3d). This example has shown that the OTFP automatically discovered and quantified various systematic variation patterns in the complex, ill understood B-Z reaction. At our chosen fidelity fraction (relative reconstruction error variance < 0.01%, resulting in 10 PCs), only very slight differences between the original and reconstructed profiles are detectable to the naked eye. Had we demanded higher fidelity fractions, more PCs would have been included. Conversely, had we demanded fewer PCs, that would have given higher reconstruction error variance. When submitting this high-dimensional data stream to the automatic model-based data compression, the main patterns of systematic variability in the data were automatically found and extracted. In this example, each high-dimensional spectrum was measured at a single space point only. The next example will show how an overwhelming data stream that arises when thousands of such high-dimensional spectra are measured in parallel by a hyperspectral camera can be dealt with by the OTFP.

⁹The MathWorks, Inc., Natick, Massachusetts, United States

¹⁰The reported results refer to the baseline-subtracted spectra for better illustration.




Figure 6.2: Vis-NIR light absorbance spectra from the B-Z reaction at three different points in time: input (black solid lines) and OTFP modelled and reconstructed (red dotted lines) spectra.

6.4.2 Detailed remote characterization of orange samples

This example concerns efficient quality control of physical objects – in this case oranges. The individual pixel NIR spectra were submitted to a model-based pre-treatment, MSC, to remove the undesired light scattering effects and prevent actual chemical signals, often of lesser magnitude [122], from being overlooked. They were subsequently centred and weighed to down-scale noisy spectral regions. Here, the model centre was continuously updated, the variable weighing factors kept constant all over the processing and the MSC parameters additionally stored along with all the other retained information.

As indicated in Table 6.2, the compression of the orange hyperspectral images also yielded satisfactory outcomes. In addition to a very precise data retrieval, since noise is partly filtered out, various imperfections, probably due to instrumental problems, are apparently removed (see Figure 6.4).

¹¹The mean vector closely resembles the lower-absorbance spectral profiles, due to their high abundance in the Vis-NIR dataset.

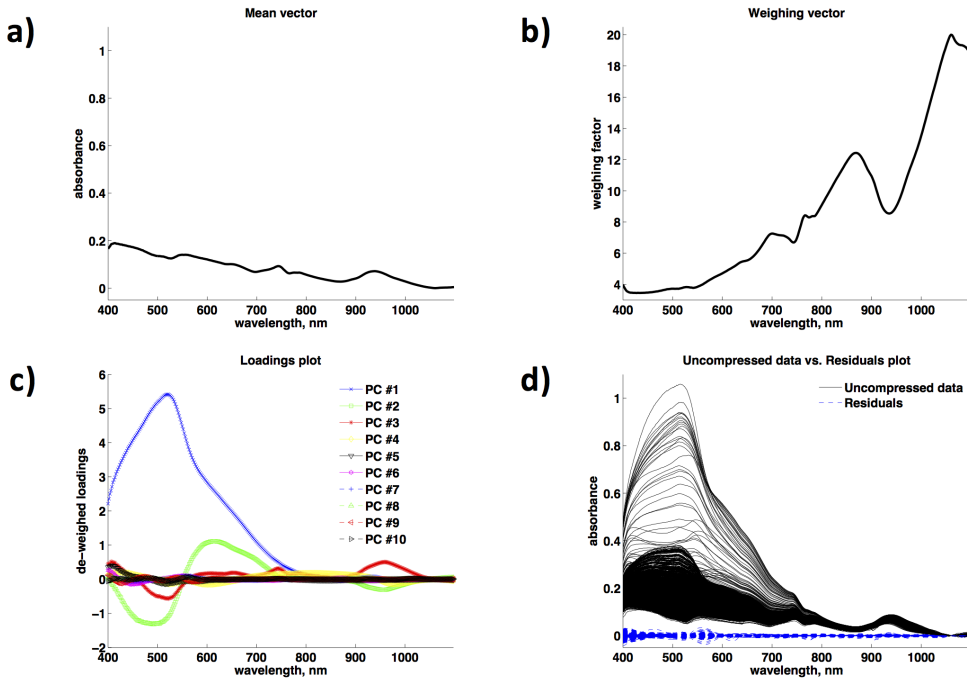


Figure 6.3: Vis-NIR light absorbance spectra from the B-Z reaction: representation of the full compression model. a) Final mean vector, b) variable weighing factors (kept constant throughout the algorithmic procedure), c) loadings profiles (divided by the channel weights, c, and scaled by their respective singular values) and d) input absorbance spectra (black solid lines) and lack-of-fit residuals (blue dotted lines)¹¹.

As an example of the added value the bilinear modelling offers unlike conventional compression methods in terms of understanding and interpretability, the scores distribution maps¹² (or scores plots) of image #2 related to the first three extracted PCs and their corresponding loadings profiles are displayed in Figure 6.5 along with the MSC preprocessing parameters used to correct the spectra of the individual pixels, the corresponding root weighed Residuals Sum-of-Squares (RSS) image (after the extraction of five PCs), the final mean vector and the variable weighing factors resulting from the OTFP. PC #1 seems to reflect an overall lighting variation on the 3D orange. The texture of the orange peel is partly captured by PC #2, along with a particular defect located on the bottom-left area of its surface and a 3D illumination and/or penetration effect generating a gradual

¹²The darkness of the pixels is proportional to the value of their scores on the respective components.




Figure 6.4: Hyperspectral NIR images: a-c) Uncompressed and d-f) OTFP modelled and reconstructed grey-scale orange image #1, #2 and #3 at 1675nm

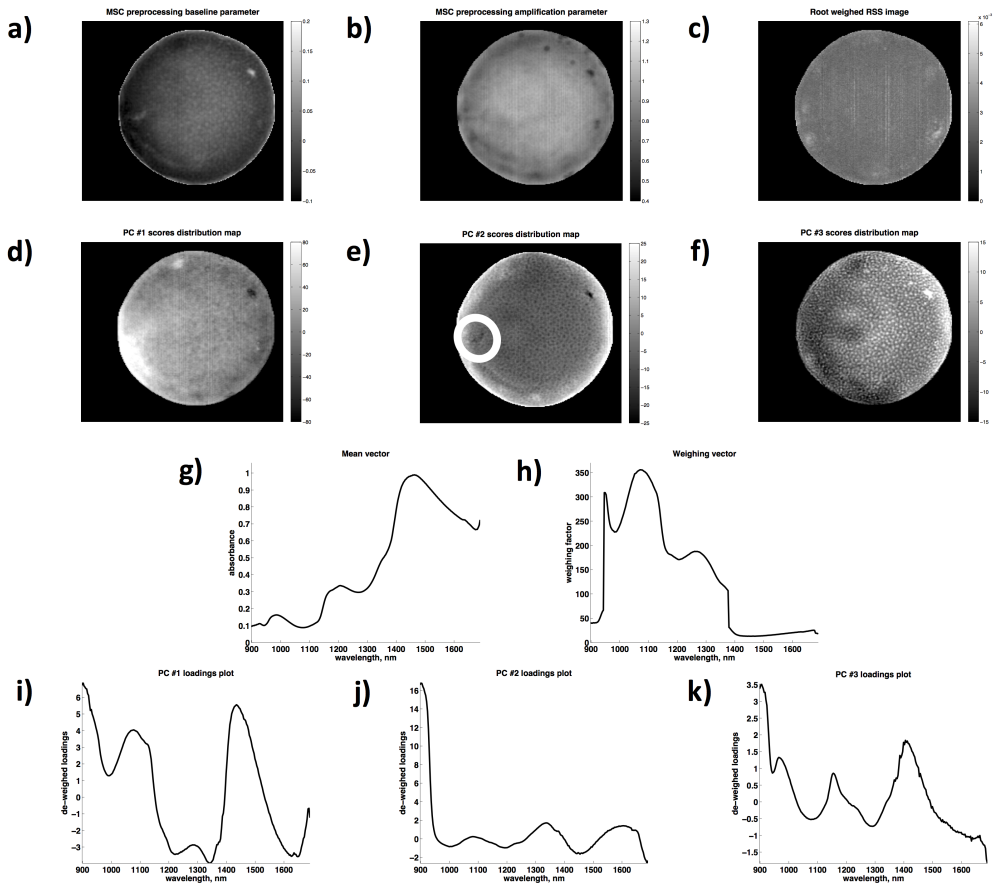


Figure 6.5: Hyperspectral NIR images – Modelling of orange image #2: a) baseline variations and b) amplification variations, estimated by MSC preprocessing and used to correct the spectra of the individual pixels, c) summary of the unmodelled residuals (root Residuals Sum-of-Squares, RSS, of the weighed wavelength channels after the extraction of 5 OTFP PCs), d) PC #1, e) PC #2 and f) PC #3 grey-scale scores distribution maps, g) final wavelength mean vector and h) wavelength weighing factors (kept constant throughout the algorithmic procedure) i) PC #1, j) PC #2 and k) PC #3 loadings profiles (divided by the channel weights, c, and scaled by their respective singular values). The white circle in e) highlights a particular defect on the surface of the orange sample

6. Dealing with big data streams: On-The-Fly Processing of continuous high-dimensional data streams

Table 6.2: Hyperspectral NIR images: values of the compression quality indices. The number of original measured variables is reported in the first column.

J	A	EV_{raw}	EV_p	$RMSRE$	t_c	CR
247	5	99.93	93.27	0.0096	43.8	33.29 ($\frac{129235545 \text{ bytes}}{3882254 \text{ bytes}}$)

Table 6.3: Hyperspectral image from a push-broom camera installed on a flying drone: values of the compression quality indices. The number of original measured variables is reported in the first column.

J	A	EV_{raw}	EV_p	$RMSRE$	t_c	CR
450	3	99.82	99.02	0.015	300.2	45.02 ($\frac{241451269 \text{ bytes}}{5363455 \text{ bytes}}$)

decrease in the scores values from the border to the centre of the sample. PC #3 seems to represent a purely textural component.

This example has shown that the self-modelling process simplified the interpretation and usage of the enormous amounts of data from a hyperspectral camera recording a series of similar objects. The model parameters gave high compression as well as interesting graphical insights. The next case-study will illustrate an even more overwhelming data stream from a continuously measuring hyperspectral camera installed on a flying drone.

6.4.3 Environmental surveillance by airborne hyperspectral imaging

The high compression of the hyperspectral push-broom image is proven by both Table 6.3 and Figures 6.6a and b. In this case, the spectrum of each pixel at each point in time was just centred. Specifically, the model centre was continuously updated, while the variable weighing factors were set to 1 and kept constant all over the processing.

Despite the notable reduction in the memory usage, the uncompressed and reconstructed pseudo-RGB pictures, constructed by selecting only three of the available wavelength channels¹³, exhibit barely perceptible discrepancies.

It is well known that while bilinear models from orthogonal subspace estimation methods (including PCA and the present OTFP) capture the essential variation information in data, the individual components are not intended to be meaningful

¹³Around 445nm, 535nm and 575nm, where the eye cones have their maximum sensitivity to blue, green and red light, respectively.

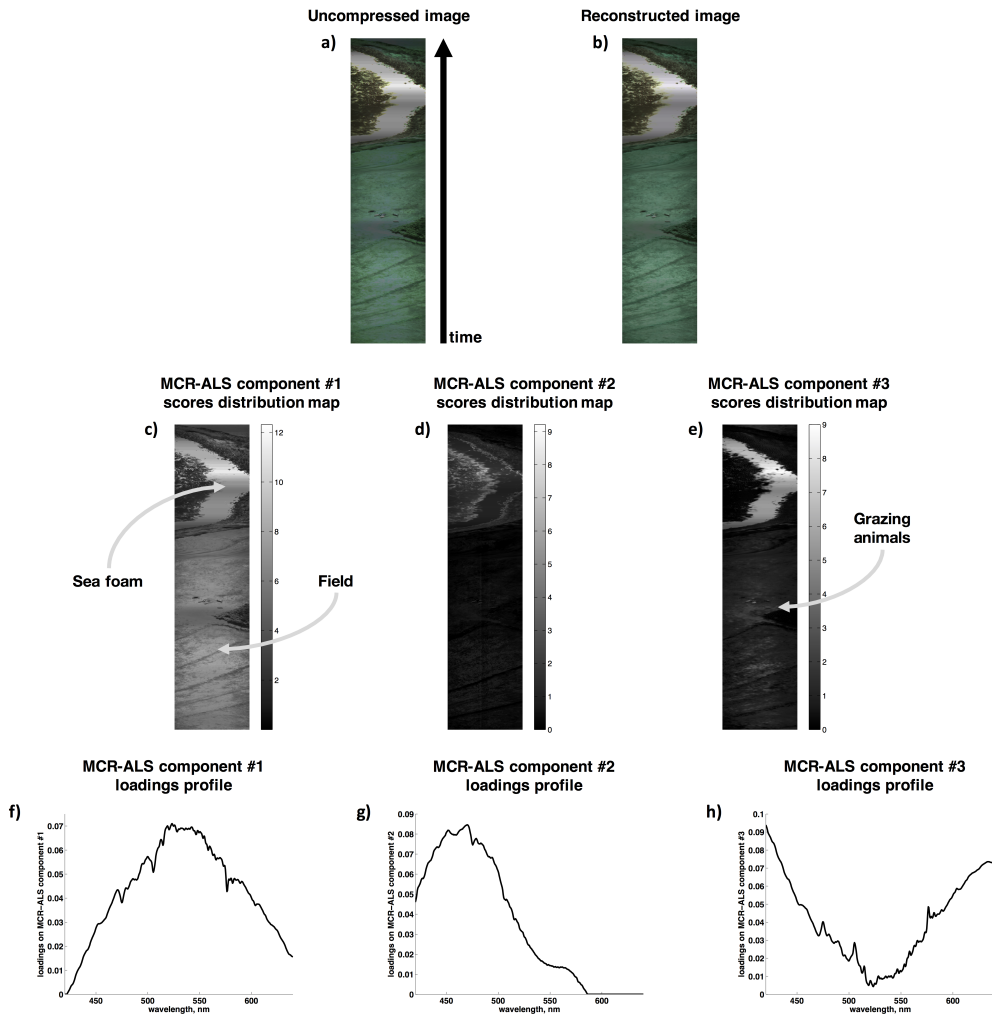


Figure 6.6: Hyperspectral image from a push-broom camera installed on a flying drone: a) uncompressed and b) OTFP modelled and reconstructed images in pseudo-RGB colours, c) MCR-ALS component #1, d) MCR-ALS component #2 and e) MCR-ALS component #3 grey-scale scores distribution maps, f) MCR-ALS component #1, g) MCR-ALS component #2 and h) MCR-ALS component #3 loadings profiles

from a physicochemical perspective, due to their mutual orthogonality (see Equations 6.2 and 6.3). Relaxing these orthogonality constraints and possibly adding other criteria, such as non-negativity in loadings and scores, may give more meaningful individual component plots. For example, Figure 6.6 also includes three different component scores distribution maps and loadings profiles (Figures 6.6c, 6.6d, 6.6e, 6.6f, 6.6g and 6.6h), obtained by a Multivariate Curve Resolution-Alternating Least Squares (MCR-ALS) [123] transformation of the global OTFP model. MCR-ALS is a soft bilinear-modelling technique, analogous to PCA, originally conceived for the resolution of multicomponent evolving chemical systems into pure individual contributions, not necessarily completely uncorrelated. It is based on an iterative sequence of optimization steps, but requires appropriate initial guesses of these contributions to achieve a reliable solution. Here, the scores and the loadings represented in Figures 6.6c, 6.6d, 6.6e, 6.6f, 6.6g and 6.6h were reestimated by executing MCR-ALS on the OTFP reconstructed data, appealing to the final OTFP loadings as input.

Although MCR-ALS components #1 and #3 are seemingly dominated by the sea foam pixels (whose corresponding signal was found to be saturated in a large spectral range), three distinct patterns are visibly recognisable: the field pixels in the first scores distribution map, the pixels surrounding the sea foam in the second scores distribution map and those capturing several animals grazing at the centre of the image in the third scores distribution map. Therefore, *ça va sans dire*, the OTFP may be employed for preliminary image treatment before further handling or segmentation.

Independent Component Analysis (ICA) [96, 124] or Parallel Factor Analysis (PARAFAC) [125, 126] and extensions of these coupled with various pixel clustering methods also belong to the rich flora of post-processing methods that can be applied to bilinear models like those coming from the OTFP.

The three first illustrations have shown how broad data streams from multi-channel sensors can be handled by the OTFP. The last example concerns a very different kind of data – a more or less random collection of individual, single-channel sensors. Traditional process industry is often extensively equipped with temperature-and-pressure sensors. Often, each new sensor gets its own display screen with its own alarm limits. How can the burden for the process operators be reduced as well as the number of false alarms? Perhaps by finding common patterns of covariation among the many sensors?

6.4.4 Analysis of an industrial manufacturing process

This example illustrates how the OTFP may be used for more rational handling of traditional industrial process data.

According to the quality indices reported in Table 6.4, the general performance of

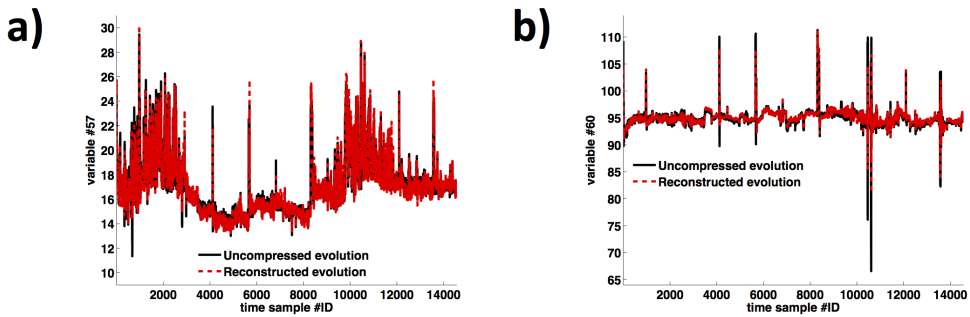


Figure 6.7: Industrial process data: Uncompressed (black solid line) and OTFP modelled and reconstructed (red dotted line) temporal evolution of a) variable #57 and b) variable #60




Figure 6.8: Industrial process data: a) PC #1/PC #2 scores (blue dots and red squares refer to Normal Operating Conditions and shut-down time samples, respectively) and b) loadings plots (the numbers correspond to the #IDs of the original variables and are represented according to their respective PC #1/PC #2 loadings values)

the OTFP when applied to this rather low-dimensional stream of industrial process data was found to be slightly worse than in the previous case-studies. This is not unexpected, given the low number of variables under study and their widely varying nature, and is a consequence of the fact that their correlation structure is not so strong that just few PCs can practically summarize all their significant variation. Nevertheless, for most of them an acceptable reconstruction was achieved, as Figure 6.7 confirms.

Besides, examining both scores and loadings can provide remarkable insights into the process behaviour, particularly if meaning can be ascribed to the input records – or at least to some of them – by human expert characterization. This is illustrated

6. Dealing with big data streams: On-The-Fly Processing of continuous high-dimensional data streams

Table 6.4: Industrial process data: values of the compression quality indices. The number of original measured variables is reported in the first column.

J	A	EV_{raw}	EV_p	$RMSRE$	t_c	CR
76	13	99.47	81.33	0.4640 ¹⁴	49.5	3.35 (^{4895674 bytes} _{1459315 bytes})

by the scores plot in Figure 6.8a. PC #1 separates two groups of observations: blue dots and red squares refer in fact to Normal Operating Conditions (NOC) and shut-down time samples, respectively. As the latter present negative projection coordinates on this component, they will be characterized by lower-than-average values of all the measured variables featuring a relatively large positive PC #1 loading (which actually assumed a nearly 0-level during shut-down periods) and *vice versa* (see Figure 6.8b). On the other hand, within-cluster differences seem to be mainly spotted by PC #2.

6.5 Comparison with classical PCA

To what extent does the OTFP mimic the corresponding traditional data modelling strategy? The present implementation of the OTFP employs similar criteria in the model updating stage to those of standard PCA, so it is natural to compare both the approaches. While the OTFP needs to hold only a small part of the data in memory at a time, traditional PCA requires all the data to be held in memory at the same time, at least if both loadings and scores are to be assessed. The time span of the hyperspectral drone imaging example (Figure 6.6 and Table 6.3) was chosen short enough to allow conventional PCA to be run and its solution to be compared to the final OTFP model.

Figure 6.9 permits to appraise the performance of the two methods for the same dataset. Figure 6.9a shows that the mean spectrum used for model centring in PCA is more or less identical to the model centre vector, m , gradually developed by the OTFP. The outcomes of the two techniques are also virtually indistinguishable if looking at the plot of the cumulative percentage of explained preprocessed data variance (Figure 6.9b) and the variable-wise RMSRE (after the extraction of three PCs, Figure 6.9c) as well as at the loadings profiles of PC #1 (Figure 6.9d), PC #2 (Figure 6.9e) and PC #3 (Figure 6.9f). The corresponding spatiotemporal OTFP scores (not displayed due to the high number of data points) were also found to be very similar to the PCA ones. Consequently, both PCA and the OTFP led to practi-

¹⁴As the original variables were characterized by different units of measurements, the reported RMSRE value concerns the final centred and weighed data array.

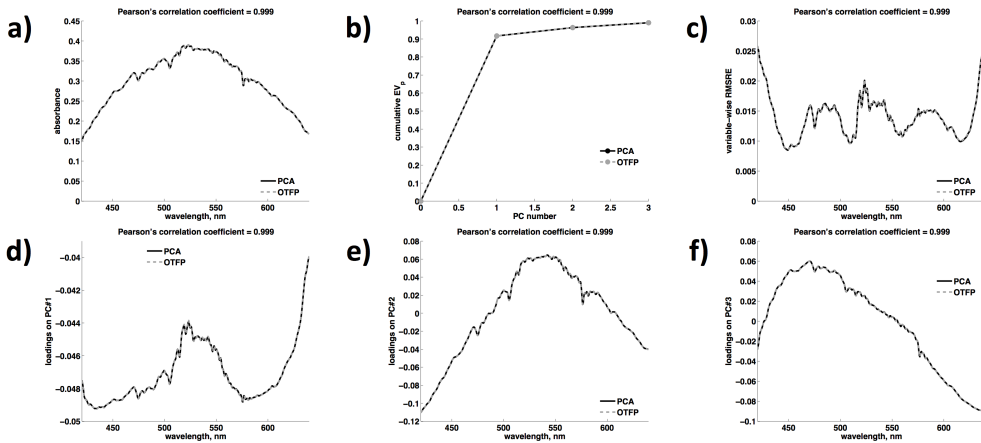


Figure 6.9: Hyperspectral image from a push-broom camera installed on a flying drone – Classical PCA (black solid line) vs. OTFP (grey dotted line): a) Mean vectors, b) cumulative percentages of explained preprocessed data variance, c) lack-of-fit (root mean square error) for the individual variables after the extraction of 3 OTFP PCs (negligible if compared with the original signal magnitude) d) PC #1, e) PC #2 and f) PC #3 loadings. Variable weighing factors (not shown) were set to 1 for all the spectral wavelengths and kept constant all over the OTFP

cally identical values of the diagnostic indices listed at the beginning of Section 6.4 except for t_c . The decomposition was achieved faster by PCA, which had simultaneous access to all the available information. On the other hand, the OTFP had to handle it by evolving its bilinear model on-the-fly as the data flowed. Nevertheless, the comparison highlighted that the OTFP can be considered a feasible alternative to standard PCA, when this latter is not applicable (e.g., when the measurements are collected in real time or the size of the analysed matrix is prohibitively large).

6.6 Discussion

The OTFP treats the incoming data one record or one batch of records at a time, and gradually develops a compact quantitative model of this data stream from the covariation patterns that it discovers. Still, Figure 6.9 illustrates the OTFP behaved quite similarly – at least for the first three components – to the corresponding conventional *global* multivariate data modelling method (in this case PCA), which analyses all the data simultaneously. Their results are almost identical even if the OTFP repeatedly has to make sense out of small chunks of data as they arrive. Therefore, it has to make many temporary decisions about what to throw away as random noise, while the global PCA has access to all data at once. On the other

hand, this is precisely the motivation behind the development of the OTFP tool, that is to always maintain an updated, compact summary of all the systematic changes, which have taken place in an otherwise overwhelming, *ever-lasting* high-dimensional data stream, with low computational or memory requirements.

The OTFP uses a multivariate data driven approximation model (here, PCA-like) as a generic Taylor expansion around a chosen set point or model centre, to summarize whatever known or unknown phenomena has caused the systematic co-variation patterns in the input data stream. The OTFP data model has a linear, additive structure and therefore gives the best approximation performance when non-additive and/or non-linear effects in the input data have been corrected for in the preprocessing step. Preprocessing is then helpful for reducing response curvature and other types of non-linearities¹⁵. Response linearity was improved in the first example (Figures 6.2 and 6.3) by converting the fibre-optic transfection data to absorbance values. Patterns known to be non-interesting may be removed during preprocessing, as illustrated by the simple baseline correction in the same case-study. In the second example (see Figure 6.4), unknown additive baseline variations and multiplicative amplification variations were estimated, parametrized and removed jointly by MSC. On the other hand, when the input variables are given in very different units, preprocessing should also scale them to balance their uncertainty levels – or, if that is unknown, to balance their total variances as shown for the industrial process data (Figures 6.7 and 6.8).

The OTFP components are mathematical basis vectors that characterize the data stream. When plotted in combination they give useful insights into the main patterns of data variation, as illustrated in Figure 6.8. But such orthogonal basis vectors are not meant to be interpreted individually. Therefore, the OTFP solution may be at any time readjusted for better visualization and more causal interpretation. This was shown by the conversion of the orthogonal, PCA-like OTFP component profiles into more graphically distinct ones by requiring non-negative scores and loadings in an MCR-ALS-based post-processing.

6.7 Conclusions and perspectives

In the near future, a drastic increase in the collection and use of high-dimensional, continuous measurements is expected. Rational use of such data streams requires generic data modelling tools that not only give good prediction and classification as the information flow evolves, but that also reveal its essential structure for human

¹⁵In case preprocessing is not of help, complex non-linearities and system heterogeneities may be handled by automatically splitting the data stream under study into two or more disjoint OTFP models (in a similar way as for the well-known static Soft Independent Modelling of Class Analogy – SIMCA – approach [127, 128]).

interpretation and efficient compression. In this article, the On-The-Fly Processing (OTFP) tool for the on-the-fly gradual modelling and compression of continuous quantitative data streams was proposed. It is based on an evolving implementation of PCA that updates on-line, when necessary, both preprocessing parameters and principal component structure (whose changes and possible expansion can be optionally monitored in real time through intuitive graphical displays). It combines the advantages of three different ways of attaining PCA or PCA-like bilinear decompositions, while avoiding their disadvantages:

- repeated use of conventional PCA, each time bringing increasing amounts of data into memory for simultaneous analysis, which yields bilinear models relevant for both past and present observations, but becomes prohibitively slow and memory-demanding for *ever-lasting* data streams;
- moving-window PCA, which repeatedly merges new observations with a bilinear subspace loading summarising past observations, ensuring that the bilinear model is up-to-date and thus relevant for the latest observations at any given moment, while losing relevance for older observations;
- eigen-analysis of the cumulative $J \times J$ cross-product matrix, a simple and fast computation as long as J is not too large, which is suitable for parallelization and *out-of-core* estimation of the PCA loadings with relevance also for past observations, but without quantitative scores for them.

The OTFP discovers new systematic patterns of covariation in multi-channel data streams, and thereby extends its current bilinear model with new dimensions in a computationally efficient way. Over time, the observation scores are stored to disk in packets and then deleted from memory. The model is continuously updated to be as PCA-like as possible, but in such a way that past scores can always be recalled and compared to present ones.

The algorithmic procedure exhibited very satisfactory performance in terms of compression rate and time and quality of the input reconstruction, especially if measurement series underlain by strong correlation structures (e.g., in Vis-NIR spectra or hyperspectral images) were dealt with. On the other hand, in the industrial process example, its power was not as prominent, probably due to a lower degree of intercorrelation in this data stream. Still, the retrieval of the temporal evolution of the original variables was reasonably precise. This could represent an important cross-road for manufacturing companies, whose modern information storage systems are commonly based on univariate calculations, not taking into account the possible interdependences among various instrumental responses, destroying their essential multivariate nature and eliminating much of their meaningful content [129]. Finally, the scores and loadings estimated through the PCA-based dimensionality reduction feature distinctive interpretability properties, extremely helpful for data understanding, utilization and further exploration by

complementary statistical approaches, such as MCR-ALS. As far as the authors are concerned, no available compressor guarantees such a noteworthy added value. In the near future, this strategy for continuous, automatic model development based on multi-channel measurements, may become useful also for processing a wider range of data stream types. For instance, the Internet-of-Things (IoT) will result in an enormous increase in technical measurements in many fields of interest like medicine, industry, or communications. Many of these IoT sensors will be multi-channel (cameras, spectrometers, etc.). Others will be univariate, but even these will generate multi-channel data: the time series from one single, more or less continuous data source will lead to high-dimensional frequency spectra (spectrograms), after suitable domain transforms (e.g., by FFT or wavelet analysis). Since the methodology here relies solely on linear algebra, it is expected to work properly also within the more general BIG DATA context.

Acknowledgments

We would like to thank Prof. Bjørn Alsberg for providing the Vis-NIR equipment and the Underwater Systems and Technology Lab (LSTS) at the University of Porto, the Hydrographic Institute of the Portuguese Navy and the University of the Azores for carrying out the REP15 exercise, during which the hyperspectral push-broom image was collected.

Chapter 7

Hyperspectral time series analysis

Hyperspectral image data streams interpreted by modelling known and unknown variations

In this chapter, we experimentally demonstrate a generic method for compact quantification and interpretation of multichannel spatiotemporal data ("hyperspectral video") in terms of known and unknown variation types¹. The process of drying a wet wood sample was characterized by a series of high-resolution hyperspectral images taken at 150 consecutive time steps over a period of 21 hours. Each pixel intensity was measured at 159 wavelength channels in the vis-NIR region, which resulted in a time series of approximately half a terabyte of raw spectral data.

Passing the massive stream of data through a four-stage data modelling procedure resulted in a substantially compressed 10 component bilinear model comprised of five *a priori* known and five newly discovered spectral components. From this compressed subspace model, a filtered version of the original data stream could be reconstructed.

First, the measured intensity spectra were normalized by transformation into reflectance and linearized into apparent absorbance units. These absorbance spectra were secondly submitted to simplified causal modelling of known phenomena by Extended Multiplicative Signal Correction (EMSC) – to identify variations in

¹Based on the paper "Hyperspectral image data streams interpreted by modelling known and unknown variations", by Petter Stefansson, João Fortuna, Hodjat Rahmati, Ingunn Burud, Tatiana Konevskikh and Harald Martens. Published as a chapter in the book "Hyperspectral Imaging", Data Handling in Science and Technology, Elsevier, Volume 32, 2020. This work was partially funded by the Norwegian Research Council in the project "WOOD/BE/BETTER" code 225345.

the main light scattering and light absorption variation types. Thirdly, the high-dimensional streams of lack-of-fit residual spectra from EMSC were analysed for possible remaining systematic structures – e.g., due to unknown and hence unmodelled variation types – by an adaptive bilinear modelling method (ABLM). In the final modelling stage, the dynamics of the various known and unknown physical and chemical variations of the drying process were assessed from their temporal developments.

7.1 Introduction

Diffuse multichannel reflectance spectroscopy in, e.g., the visible (vis) and near-infrared (NIR) wavelength range is a fast and informative methodology for simultaneous measurement of a range of chemical and physical properties in complex biological samples [130, 131]. However, to resolve the selectivity problems of chemical and physical variations with similar spectral effects, multivariate calibration is required [92]. In hyperspectral imaging, where each pixel in an image is represented by a spectrum of reflected light, the spatial distribution of these properties can be quantified in a heterogeneous sample such as a piece of wood, using mathematical modelling and Multivariate Image Analysis (MIA) [132].

When a given sample is measured repeatedly *over time* by hyperspectral imaging, the resulting "hyperspectral video" also provides information about temporal developments. The multichannel spatiotemporal measurements in hyperspectral video generate really Big Data. In this work, a piece of Scots pine was monitored using hyperspectral imaging in the vis-NIR wavelength region for a period of 21 hours as the wood underwent desorption from a moisture saturated state to a dry state. We use statistical data-driven models to estimate the temporal development of properties within the wood as it dries and demonstrate how a seemingly overwhelming stream of hyperspectral video data can be converted into relatively simple quantitative spatiotemporal information by a combination of various pragmatic mathematical modelling techniques.

Chemical absorption and physical scattering are the two phenomena dominating the interaction between light and complex materials like wood. Both of them are well understood: Due to electronic or molecular resonances in chemical compounds, those complex materials absorb light at a given wavelength and convert it into heat, or emit it at other wavelengths. Changes in the degree of absorbed light between samples are useful for the quantification of their chemical composition. Variations in the physical properties of the samples (such as the amount, size, shape and refractive index of particles) cause variations in several light scattering phenomena (angular distribution of reflected or transmitted light, effective optical path length, specular surface reflectance, etc.); measuring these is useful

for quantification of the samples' physical properties.

The present work evaluates a stream of hyperspectral images representing a hyperspectral time series, which we analysed with respect to the common underlying spectral variation patterns. In addition, since the wood sample is monitored over time in a fixed position, the spatial and temporal structure of these variation patterns can be studied qualitatively and quantitatively.

In theory, causal mathematical modelling of how these phenomena affect light measurements is therefore possible: for instance, light absorption and light scattering can be modelled via the complex refractive index [133]. However, such causal modelling requires detailed information about the chemical and physical structures of the measured samples, and this information may not be available. Moreover, it may require a measurement set-up and/or extensive tests involving simultaneous transmittance/reflectance measurements in an integrating sphere, etc. For routine analysis, such quantification is usually too cumbersome, too slow, or too expensive.

Instead, this work employs a combination of four simple mathematical approximation modelling stages:

1. Normalization and linearization of the raw data to facilitate subsequent linear modelling according to Beer-Lambert's Law.
2. Modelling known structures: Linear additive/multiplicative model of an extended version of Beer-Lambert's Law (*Extended Multiplicative Signal Correction*, EMSC [117, 134]) to quantify and extract known chemical and physical variations.
3. Modelling unknown structures: Bilinear search for unknown and hence unmodelled variation types in the spectral residuals after the EMSC, based on Adaptive Bilinear Modelling (ABLM) using the On-The-Fly-Processing (OTFP)² tool presented in Chapter 6.
4. Temporal kinetics modelling of the known and unknown state variables (component time series) obtained from the EMSC and ABLM based data-models of the hyperspectral video measurements.

The outline of the chapter is as follows: first, we describe the actual experiment and the mathematical modelling methods employed to analyse the results. Then we summarize the results, starting with a preliminary, overall assessment of the complexity of the drying process, based on the reduction of weight of the wood sample. Next, the entirety of the hyperspectral measurements are analysed with respect to known and unknown phenomena, based on the EMSC and OTFP methods, respectively. Finally, we assess the kinetics of the EMSC and OTFP temporal parameters.

²www.idletechs.com

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations

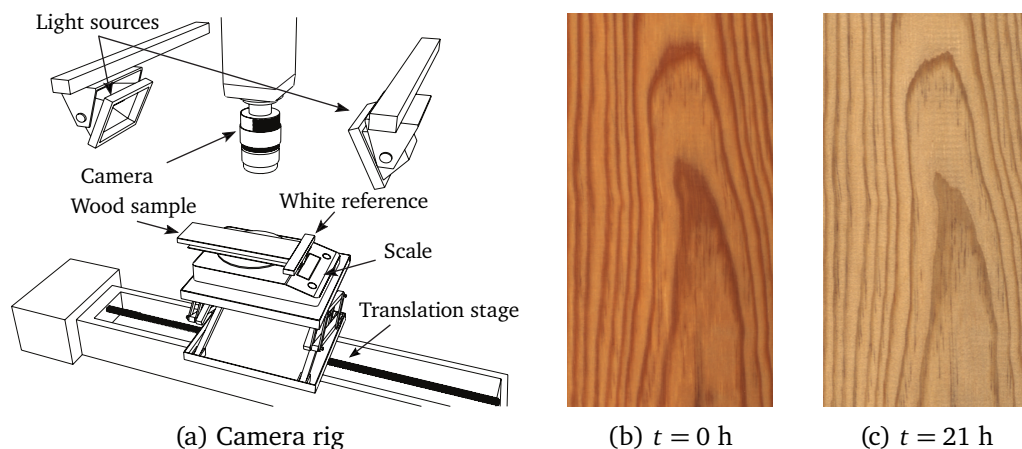


Figure 7.1: The experiment. Illustration of experimental setup used to measure the spectral reflectance and weight of a drying wood sample (a). RGB renderings of wood sample in wet (b) and dry (c) states.

7.2 Methods

7.2.1 The experiment

Drying of a wood sample

A wood sample of the species Scots pine (*Pinus sylvestris*) originating from a forest in Hobøl, roughly 30km south of Oslo, Norway, was cut into dimensions $18 \times 100 \times 280$ mm. The sample was then placed in a drying oven where it was dried at a temperature of 103°C for 48 hours until it was ensured, through repeated weight measurements, that as much as possible of the water in the material had been evaporated. The sample was then taken from the oven and immediately weighed. Its dry weight was determined to be 245.46g. Once the sample's dry weight was established, the sample was submerged in water and left to soak for approximately 24 hours. After the soaking period, the sample was removed from the water and placed on a digital scale. Its initial wet weight was 336.51g. The digital scale with the wood sample and a Spectralon white reference were attached to a translation stage as illustrated in Figure 7.1a.

The wood sample was artificially illuminated by two halogen spotlights positioned on either side of a hyperspectral camera, which monitored the sample for a period of about 21 hours. From the weight loss of the wood sample registered by the scale, the relative moisture content could be calculated at different time points and used for a preliminary assessment of the kinetic complexity of the drying process.

Monitoring the drying process by hyperspectral time-lapse camera data ("hyperspectral video")

The hyperspectral push-broom camera (Specim, Oulu, Finland) captured images with a spatial resolution of 4480×1312 pixels (pixel dimension $69.2 \times 68.6 \mu\text{m}^2$), each characterized by 200 wavelength bands in the 392-1022nm range. Because bands at the edge of the camera's detection limit were found to suffer from a low signal-to-noise ratio, the wavelength range was cropped after the signal acquisition, down to 159 bands covering the 500-1005nm region. The software controlling the experimental setup was programmed to automatically repeat the image acquisition of the sample every eight minutes over a period of about 21 hours as the sample and the setup remained completely untouched, resulting in 150 hyperspectral frames of the sample at various moisture contents. The spatial resolution of the region of interest within each frame (area excluding the white reference plate, wood edges, etc.) was 2200×1070 pixels, resulting in a four-dimensional dataset of size $2200 \times 1070 \times 159 \times 150$ (rows \times columns \times spectral bands \times time). Altogether, this corresponds to roughly 418GB of data when stored in double precision floating point (fp64) format. The sample was not moved during the 21 hours, therefore the pixels of each hyperspectral image are expected to correspond to the same location of the sample, apart from a minor offset caused by the contraction of the wood sample as it dries. Thus, the resulting data set can be seen as four-dimensional, with each one-dimensional pixel spectrum being a function of both space and time as illustrated in Figure 7.2 a).

The recorded photon count ("light intensity" I) in each of the 150×159 images has 2D spatial information about the physical and chemical structure of the wood sample; the distinction between *early-* and *latewood* growth zones (lighter and darker regions of the wood, respectively) is clearly visible. These zones reflect seasonal fluctuations in the growth rate of the tree during its lifetime due to variations in temperature and precipitation, resulting in chemical and physical differences in wood structure. This sample heterogeneity is shown in the two wood images in Figure 7.1b and 7.1c.

To use the massive stream of high-resolution hyperspectral images to study the physics and chemistry of the wood drying process, these measured intensity data I were passed through several mathematical modelling stages:

1. Response linearization;
2. Semi-mechanistic multivariate modelling of known effects;
3. Data-driven multivariate modelling of the remaining unknown effects; and finally,
4. Statistical summary and kinetic modelling of the dynamics of the known and unknown effects.

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations

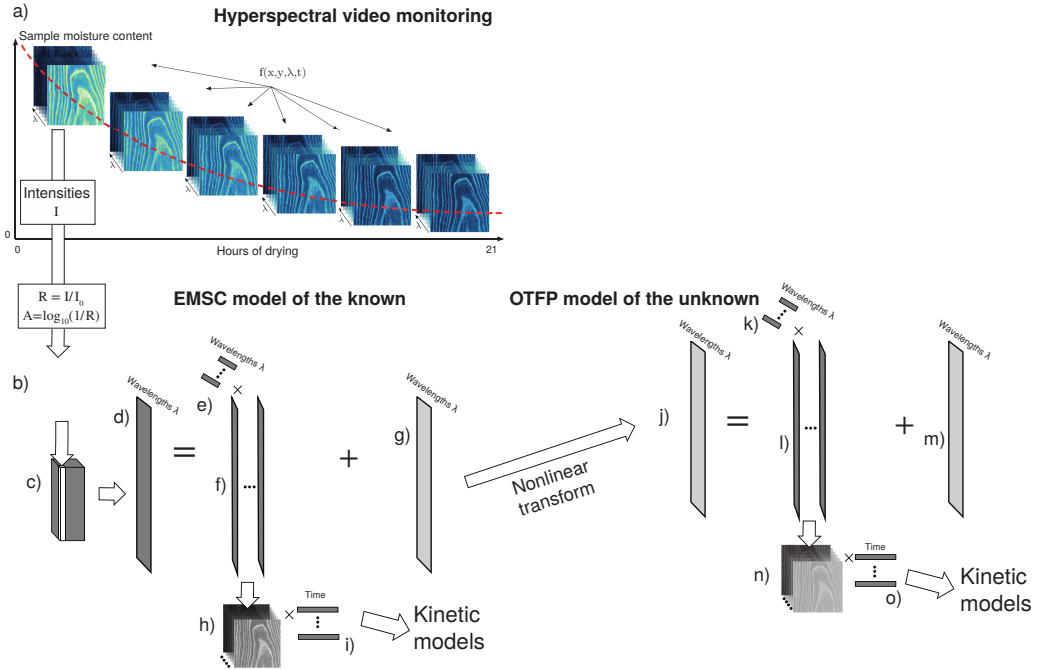


Figure 7.2: Overview of experimental data acquisition and modelling of hyperspectral video. a) Input data: (2200×1070) pixels \times 159 wavelength channels \times 150 time points. b) - i) Model what is known about input data: EMSC modelling of 2-way input data for 159 wavelength channels at $353\,100\,000$ pixels $(2200 \times 1070 \times 150) \times 159$ wavelengths, and spatiotemporal averaging. j) - o) Model what is unknown: Adaptive bilinear modelling of 2-way residual data of $353\,100\,000$ pixels \times 159 wavelengths.

7.2.2 Mathematical modelling stages

Weight-based assessment of the kinetic complexity of the drying process

The weight-based water saturation (sat_{H_2O} (%)) in the wood sample was calculated as

$$sat_{H_2O}(t) = \frac{w_{wood}(t) - w_{dry\ wood}}{w_{wood}(t)} \cdot 100 \quad (\%), \quad t = 1, 2, \dots, 150 \quad (7.1)$$

Where $w_{dry\ wood}$ ($= 245.46$ g) represents the weight of the wood sample dried at 103°C for 48 hours and $w_{wood}(t)$ represents the varying wood weight during the drying period. To assess the overall complexity of the drying process, the temporal derivative of $sat_{H_2O}(t)$ was calculated: For $y(t) = sat_{H_2O}(t)$ at time step t , its

temporal derivative $dy(t)/dt$ was calculated as

$$\frac{dy(t)}{dt} = \frac{y(t+1) - y(t-1)}{h(t+1) - h(t-1)} \quad (7.2)$$

where $h(t)$ is the actual time in hours at time step t . The derivative $dy(t)/dt$ was then plotted against $y(t)$ itself, to assess the kinetics of the weight-based drying process.

Since the drying process appeared to be quite complex, a simpler alternative was also employed, namely, a graphical assessment of how well the weight loss dynamics followed the simple first-order reaction:

$$\frac{dy(t)}{dt} = -k \cdot y(t) \quad (7.3)$$

where k is the rate constant. Integrating this equation over time gives the following expression in natural logarithms (ln):

$$\ln(y(t)) = -k \cdot t + \ln(y(0)) \quad (7.4)$$

Hence, a plot of $\ln(y(t))$ vs time t (in hours) should give a straight line if the process follows first-order kinetics. Observed deviations from a single straight line hint towards a more complex drying process.

Hyperspectral modelling

A hyperspectral video is a 3-way data set (pixels \times wavelength channels \times time points) that, in principle, could have been modelled by a tensor-algebraic 3-way model, e.g. a PARAFAC-model [126] of the type

$$A = B \otimes C \otimes D + E \quad (7.5)$$

where A is the 3-way video input data, B , C and D represent a low-dimensional model with vectors in the pixel-, wavelength- and time-domains, respectively, and E represents measurement noise.

However, the sheer amount of input data makes this N-way modelling too computation intensive. Moreover, the light scattering varies with the drying time (conf. Figure 7.1b - 7.1c), and this is likely to involve changes in the effective optical path length. Hence, since path length variations give non-additive effects, this purely additive 3-way model was discarded.

Instead, the following sequence of theory- and data-driven modelling steps (Figure 7.2) were chosen here.

Response linearization The light intensities I from the wood sample, measured at each of the wavelengths for each of the pixels at each of the points in time (Figure 7.2a) were converted to reflectance units by $\mathbf{R} = (I - I_d)/(I_0 - I_d)$, where I_0 represents the intensity measured on a Spectralon white reference plate and I_d is the dark signal (image taken with the shutter closed). The reflectance data were in turn linearized with respect to chemical response by the conventional transformation to apparent absorbance $\mathbf{A} = -\log_{10}(\mathbf{R})$. After this response normalization and linearization of each of the $> 5 \times 10^{10}$ individual light intensity readings, the 2D image of the apparent absorbance at each wavelength channel at each point in time (Figure 7.2 b) is unfolded into a 1D column vector. This resulted in a virtual 3-way array of absorbance (Figure 7.2 c) with 2 354 000 pixels per image and 159 wavelengths \times 150 time points.

This array was unfolded into a virtual 2-way absorbance matrix (Figure 7.2 d) of 353 100 000 pixels \times 159 wavelengths, which was submitted to three more stages of modelling:

1. **Modelling the known:** The 159-dimensional spectrum of each of the over 353 million pixels was submitted to a semi-causal modelling of what is known about how light interacts with matter in complex samples such as wood, based on the Extended Multiplicative Signal Correction (EMSC) model [117, 134].

In total, five different phenomena with known spectra were modelled, in order to estimate their unknown spatial and temporal distributions. The purpose of this stage is to model the linearized absorbance spectra in terms of a sum of variation types whose spectral profiles (Figure 7.2 e) are known but actual levels in each pixel (Figure 7.2 f) is unknown. After subtracting these five estimated effects, the residual spectra (Figure 7.2 g) should only contain random measurement noise, if the chosen mechanistic model had been perfect.

The chosen EMSC model spectra (Figure 7.2 e) represent three known physical and two known chemical variation patterns that are expected to affect the apparent absorbance spectrum \mathbf{z}_i of each pixel $i = 1, 2, \dots, n_{\text{pixels}}$ relative to a chosen, "typical" reference pixel spectrum \mathbf{m} (in our case chosen to be the mean of all pixels in the image taken after 21 hours of drying), according to the model:

$$\mathbf{z}_i = b_i \times \left(\mathbf{m}^T + \Delta c_{i,\text{Water}} \times \mathbf{s}_{\text{Water}}^T + \Delta c_{i,\text{WoodPigment}} \times \mathbf{s}_{\text{WoodPigment}}^T \right) + a_i \times \mathbf{1}^T + d_i \times \mathbf{f}^T + \boldsymbol{\varepsilon}_i \quad (7.6)$$

In this equation \mathbf{m} , $\mathbf{s}_{\text{Water}}$, $\mathbf{s}_{\text{WoodPigment}}$, $\mathbf{1}$, \mathbf{f} , \mathbf{z}_i , $\boldsymbol{\varepsilon}_i$ are column-wise vectors with the same length as the number of wavelength channels. b_i , $\Delta c_{i,\text{Water}}$, $\Delta c_{i,\text{WoodPigment}}$, a_i and d_i are scalars. The reference spectrum \mathbf{m} (Figure 7.3 a)

was chosen in the modelling for estimating the effective relative path length parameter b_i , in each pixel i and it is corrected for by division. The reason why it is important to estimate the relative optical path length is that according to Beer-Lambert's law, the absorbance effects of path length variations can be very big, and are multiplicative, while e.g. chemical pigment variations can be very small and give additive absorbance effects.

This parameter, the relative optical path length, is intended as a pragmatic measure of the "diffuse thickness" [135] of the wood sample at different states of drying, relative to that of the dry reference sample which defined the reference spectrum \mathbf{m} . Popularly speaking, b_i should thus show how far, on "average", the photons travel inside the wood after hitting the wood surface before they emerge again at the surface to be detected by the camera. The longer a photon travels in the chemically absorbing environment, the higher the probability is that it will be lost by chemical absorption and converted into heat or lost by other means. This relative optical path length estimate b_i is expected to vary more or less inversely proportional to the scattering coefficient S in sample i : $b_i = 1/S$, but with the simplifying assumption that S is the same at all measured wavelengths.

In order to attain robust estimates of the relative path length variations by projection of input spectra \mathbf{z}_i on reference spectrum \mathbf{m} , the other variation types that also affect the input spectra must be modelled so as to avoid alias errors:

A flat baseline $\mathbf{1} = [1, 1, \dots, 1]^T$ of length 159 (Figure 7.3 b), and a straight line \mathbf{f} with monotonically increasing values evenly spaced between -1 and 1 (Figure 7.3 c) are chosen for estimating the pixel's spectral baseline-offset a_i and baseline slope d_i , respectively, again due to physical light scattering variations, and are corrected for by subtraction.

A wood colour spectrum $\mathbf{s}_{\text{WoodPigment}}$ (Figure 7.3 d) was chosen for estimating and subtracting spatially visible wood structure variations $\Delta c_{i,\text{WoodPigment}}$. It was defined as the average scatter-corrected difference between the absorbance spectra of *early*- and *latewood* pixels in the last of the images (the driest state).

The absorbance spectrum of water, $\mathbf{s}_{\text{Water}}$ (Figure 7.3 e) was chosen in order to quantify how the concentration of water differs from the presumed concentration of water in the reference pixel, $\Delta c_{i,\text{Water}}$, and subtract this effect. The water spectrum, $\mathbf{s}_{\text{Water}}$, was defined as the wavelength dependent specific absorption coefficient of water within the 500-1005nm range (data from [136]). Residual spectrum $\boldsymbol{\varepsilon}_i$ represents any variation in the input absorbance spectrum \mathbf{z}_i that is not described by this EMSC model, once the unknown parameter values b_i , $\Delta c_{i,\text{WoodPigment}}$, $\Delta c_{i,\text{Water}}$, a_i , d_i in each pixel i of these known spectral variation types have been estimated. To estimate

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations

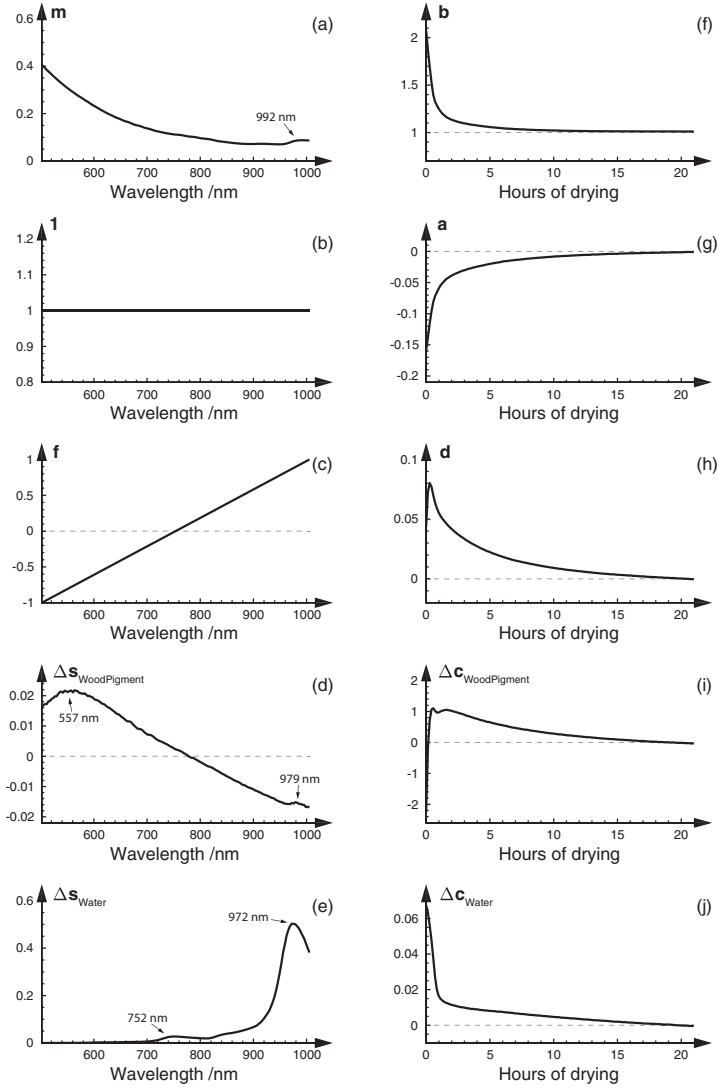


Figure 7.3: Modelling the known: Spectral and temporal structure of the parameters from Extended Multiplicative Signal Correction (EMSC). Left column shows EMSC model spectra chosen for modelling apparent absorbance; (a) Reference spectrum m for estimating optical path length, calculated as the average spectra of the last (driest) image in the series. (b) Constant "spectrum" for estimating baseline offset. (c) Linear "spectrum" for estimating baseline slope. (d) Dominant pigment spectrum $s_{\text{WoodPigment}}$, defined as the average difference between *early*- and *latewood* pixels in the last (driest) image in the series. (e) Water spectrum s_{Water} . Right column shows the temporal development of all EMSC parameters (estimated at each point in time by averaging over all image pixels).

these parameters simultaneously in a linear regression model, the products $b_i \times \Delta c_{i,\text{Water}}$ and $b_i \times \Delta c_{i,\text{WoodPigment}}$ were redefined as $g_{i,\text{Water}}$ and $h_{i,\text{WoodPigment}}$ respectively. Hence, the EMSC model is rewritten:

$$\begin{aligned} \mathbf{z}_i = & b_i \times \mathbf{m}^\top + g_{i,\text{Water}} \times \mathbf{s}_{\text{Water}}^\top + h_{i,\text{WoodPigment}} \times \mathbf{s}_{\text{WoodPigment}}^\top \\ & + a_i \times \mathbf{1}^\top + d_i \times \mathbf{f}^\top + \boldsymbol{\varepsilon}_i \end{aligned} \quad (7.7)$$

Defining the set of known model spectra:

$$\mathbf{M} = [\mathbf{m}, \mathbf{s}_{\text{Water}}, \mathbf{s}_{\text{WoodPigment}}, \mathbf{1}, \mathbf{f}]^\top \quad (7.8)$$

And the corresponding set of unknown parameters (pixel properties):

$$\mathbf{p}_i = [b_i, g_{i,\text{Water}}, h_{i,\text{WoodPigment}}, a_i, d_i] \quad (7.9)$$

the description for each pixel i becomes a simple linear model:

$$\mathbf{z}_i = \mathbf{p}_i \times \mathbf{M} + \boldsymbol{\varepsilon}_i \quad (7.10)$$

This allows the unknown parameter values in \mathbf{p}_i to be estimated by ordinary least squares (OLS) regression of each pixel spectrum \mathbf{z}_i on \mathbf{M} :

$$\hat{\mathbf{p}}_i = \mathbf{z}_i \times \mathbf{M}^\top \times (\mathbf{M} \times \mathbf{M}^\top)^{-1} \quad (7.11)$$

The multivariate modelling methods used here, EMSC and OTFP, are both based on weighted least squares. Therefore, it is important to balance the presumed relevance of the 159 wavelength channels against their estimated noise levels. Uncertainties in the measured signal from the hyperspectral camera vary across different wavelengths due to e.g. the spectral response of the camera detector. To account for this in the modelling stage, wavelengths were weighted with a vector of weights $\mathbf{v}_\lambda, \lambda = 1, 2, \dots, \Lambda$ according to their signal-to-noise ratio (SNR) so that the signal of wavelengths associated with greater uncertainties were down weighted, whilst signal originating from wavelengths with low noise level were to a greater extent preserved. The SNR of an image was approximated by dividing the average raw signal intensity in the white reference region of an image with the average signal intensity in the dark reference region of the same image:

$$\text{SNR} = \frac{I_{\text{signal}}}{I_{\text{noise}}} = \frac{I_{\text{white}}}{I_{\text{dark}}} \quad (7.12)$$

As the detector and the halogen lights used to artificially illuminate the wood




Figure 7.4: Weight vector used to assign weights to different wavelength regions during EMSC and OTFP. Red dotted line represents measured signal-to-noise ratio. Dark solid line represents the smoothed curve used as weight vector \mathbf{v} in both EMSC and OTFP.

sample become warmer over time, slight changes in their characteristics occur which causes minor changes in the SNR between images taken at different points in time during an extended image acquisition period such as 21 hours. The SNR was in the present study therefore calculated for nine different images, evenly sampled from the first to the last image of the acquisition period, and then averaged into one SNR curve. The average SNR curve was then smoothed and normalized to lie in the 0.4-0.9 range before used as a weight vector in both EMSC and OTFP. Figure 7.4 shows the final, smoothed, weight vector \mathbf{v}_λ together with the SNR curve.

The ordinary least squares solution therefore is replaced by a weighted least squares (WLS) solution:

$$\hat{\mathbf{p}}_i = \mathbf{z}_i \times \mathbf{V} \times \mathbf{M}^T \times (\mathbf{M} \times \mathbf{V} \times \mathbf{M}^T)^{-1} \quad (7.13)$$

where weights $\mathbf{V} = \text{diag}(\mathbf{v}_\lambda)$ balances the sum-of-squares contributions from the different wavelength channels with respect to their relevance and noise levels. Once the relative effective optical path length b_i had been estimated for every pixel, the chemical parameters were estimated by division:

$$\Delta c_{i,\text{Water}} = b_i^{-1} \times g_{i,\text{Water}} \quad (7.14)$$

$$\Delta c_{i,\text{WoodPigment}} = b_i^{-1} \times h_{i,\text{WoodPigment}} \quad (7.15)$$

Then, for each of these five EMSC model elements, the many spatiotemporal pixel parameters were summarized in terms of spatial structure (Figure 7.2

h), comparison of the first and the last image in the drying sequence) and their temporal developments (Figure 7.2 i). The obtained time series were finally analysed with respect to their kinetic properties.

To correct for the physical light scattering variations while retaining the chemical light absorbance variations, the EMSC post-processing of the input spectra is defined as

$$\mathbf{z}_{i,\text{corr}} = b_i^{-1} \times (\mathbf{z}_i - a_i \times \mathbf{1}^\top - d_i \times \mathbf{f}^\top) \quad (7.16)$$

which corresponds to

$$\mathbf{z}_{i,\text{corr}} = \Delta c_{i,\text{Water}} \times \mathbf{s}_{\text{Water}}^\top + \Delta c_{i,\text{WoodPigment}} \times \mathbf{s}_{\text{WoodPigment}}^\top + b_i^{-1} \times \boldsymbol{\varepsilon}_i \quad (7.17)$$

The estimated residual spectrum of each pixel $\boldsymbol{\varepsilon}_i$ (Figure 7.2 g), is obtained by:

$$\boldsymbol{\varepsilon}_i = \mathbf{z}_i - b_i \times \mathbf{m}^\top - g_{i,\text{Water}} \times \mathbf{s}_{\text{Water}}^\top - h_{i,\text{WoodPigment}} \times \mathbf{s}_{\text{WoodPigment}}^\top - a_i \times \mathbf{1}^\top - d_i \times \mathbf{f}^\top \quad (7.18)$$

This residual spectrum is expected to contain random measurement noise as well as possible unmodelled spectral structures remaining after the EMSC modelling. Before the residual spectra $\boldsymbol{\varepsilon}_i$ were submitted to further scrutiny to search for unknown features, they were scaled in order to remove the non-additive effects of varying path length:

$$\boldsymbol{\varepsilon}_{i,\text{scaled}} = b_i^{-1} \times \boldsymbol{\varepsilon}_i \quad (7.19)$$

This EMSC modelling yielded re-scaled and weighted 159-dimensional residual spectra $\boldsymbol{\varepsilon}_{i,\text{scaled}}$ (Figure 7.2 j) for more than 353 million 159-dimensional spectra:

$$\mathbf{E}_{\text{scaled}} = [\boldsymbol{\varepsilon}_{i,\text{scaled}}, 1, 2, \dots, 353\ 100\ 000] \quad (7.20)$$

2. **Modelling the unknown:** In order to look for unknown, and hence unmodelled patterns of variations in the hyperspectral video data, the residuals $\mathbf{E}_{\text{scaled}}$ were passed to the next stage in the modelling: A joint analysis to discover, quantify and display unknown, unmodelled spectral variation patterns and to separate these from the background of (presumably random) measurement noise.

The estimated residual spectra $\mathbf{E}_{\text{scaled}}$ are expected to contain not only "random" measurement noise, but also systematic – but unknown – variation structures. These may be due to errors in the *shape* of the mechanistic model (here: normalization, linearization, linear EMSC model), *errors* in the employed model elements (here: estimated spectra $\mathbf{s}_{\text{Water}}$ and $\mathbf{s}_{\text{WoodPigment}}$) or

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations

unmodelled phenomena (here: e.g. physical specular reflection from wood surface, or chemical variation in the cellulose/lignin ratio in different parts of the wood).

Bilinear data modelling by principal component analysis (PCA) and PCA-like methods is useful for discovering, quantifying and graphically displaying non-random covariation structures in data.

For "path length"-scaled, mean-centred variables and weighted residuals the model may thus be described as:

$$\mathbf{E}_{\text{ABLM,scaled,weighted}} = \sum_{a=1}^A \mathbf{c}_a \times \mathbf{s}_a^T + \mathbf{E}_{\text{ABLM,final,weighted}} \quad (7.21)$$

where the bilinear contribution from each principal component (PC) a in this case consists of the product of the spatiotemporal scores \mathbf{c}_a and the loading spectra \mathbf{s}_a .

Ideally, the first PCs are included in the final model because they represent non-random covariation structures, while $\mathbf{E}_{\text{ABLM,final,weighted}}$ represents the remaining random noise.

In the present case, the number of spectra in $\mathbf{E}_{\text{scaled}}$ is very high - for different 2D pixel positions at different times. To simplify the exploratory adaptive bilinear modelling, the On-The-Fly-Processing (OTFP) software from Idletechs AS³ was employed, in order to be able to handle this stream of Quantitative Big Data on a regular PC within a reasonable computation time. The OTFP (Chapter 6) yields results very similar to a weighted principal component analysis (PCA), but develops this bilinear model sequentially and therefore allows PCA-like modelling of more or less continuous streams of high-dimensional data such as spectra from hyperspectral imaging. The same statistical weight vector of the wavelength channels (Figure 7.4) was used in the OTFP as in the EMSC.

The OTFP analysis will gradually, but automatically, discover, extract and quantify any clear systematic covariation pattern remaining in the residuals. The components in the bilinear model consist of the product of unknown spectral loading profiles (Figure 7.2 k) and unknown spatiotemporal score vectors (Figure 7.2 l), and represent a succession of "unexpected patterns". In the present case, four patterns from the residuals were found to have smooth spectral loadings. To be conservative, a fifth principal component was also included in the final model.

As with the spatiotemporal EMSC scores, the estimated spatiotemporal OTFP score patterns of each of these five PCA components was refolded with respect to its 2D wood picture (Figure 7.2 n) and its temporal development

³www.idletechs.com

(Figure 7.2 o) during the drying process.

The residuals (Figure 7.2 m) left after this two-stage modelling were briefly interpreted graphically. Since no important patterns were found, they were summarized statistically and then discarded.

3. **Temporal kinetics modelling:** Each of the average temporal score vectors from the EMSC modelling of known phenomena (Figure 7.2 i) and OTFP modelling of unknown phenomena (Figure 7.2 o) were finally fitted to the simplified first-order kinetic model. The purpose of this was to shed some light on the temporal development of the hyperspectral video measurements, and hopefully also on the mechanisms governing how the optical properties of the wood sample changed during the drying process.

For each of the EMSC and OTFP components the average score was computed at each of the 150 time points by averaging over all pixels at time t . After defining the component's score vector as \mathbf{y}_t^* , $t = 1, 2, \dots, 150$, a linearly transformed score vector $\mathbf{y}_t = \mathbf{y}_t^* \times a + b$ was defined so as to ensure that \mathbf{y}_t is positive and falling gradually towards zero. The scaling factor a was defined to ensure that one unit of change in \mathbf{y}_t corresponded to one unit change in apparent absorbance of the input data. Offset b was then defined so that the minimum value \mathbf{y}_{150} , at the end of the drying process, was 0.001 absorbance units or more. Then the temporal derivative of \mathbf{y}_t was computed using Eq. 7.2.

The dynamic complexity of the component time series \mathbf{y}_t^* could then be assessed by plotting $d\mathbf{y}_t/dt$ vs time t : If the processes affecting component \mathbf{y}_t^* had followed simple first-order kinetics, then the 150 data points would fall approximately along a straight line (conf. Eq. 7.4). Hence, deviations from the straight lines indicate more complex kinetics.

7.3 Results

7.3.1 Average moisture in the drying process

The total weight of the wood sample ($w_{\text{wood}}(t)$) is shown in Figure 7.5 a) for the drying period 0-21h, in total consisting of 150 time points. The weight-based average moisture content ($\text{sat}_{\text{H}_2\text{O}}(t)$) in the wood sample is shown in Figure 7.5 b). The figure shows that the wood sample has 27.09% moisture at time $t = 0$, and after $t = 21$ h of low-temperature drying, the wood sample still contains 6.20% moisture, compared to its oven-dried state.

In order to study the overall kinetics of the drying process, the average moisture content $\text{sat}_{\text{H}_2\text{O}}(t)$, its temporal derivative $d\text{sat}_{\text{H}_2\text{O}}(t)/dt$ (Eq. 7.2) is plotted against drying time (Figure 7.5 c) and against $\text{sat}_{\text{H}_2\text{O}}(t)$ itself (Figure 7.5 d). In theory, if

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations

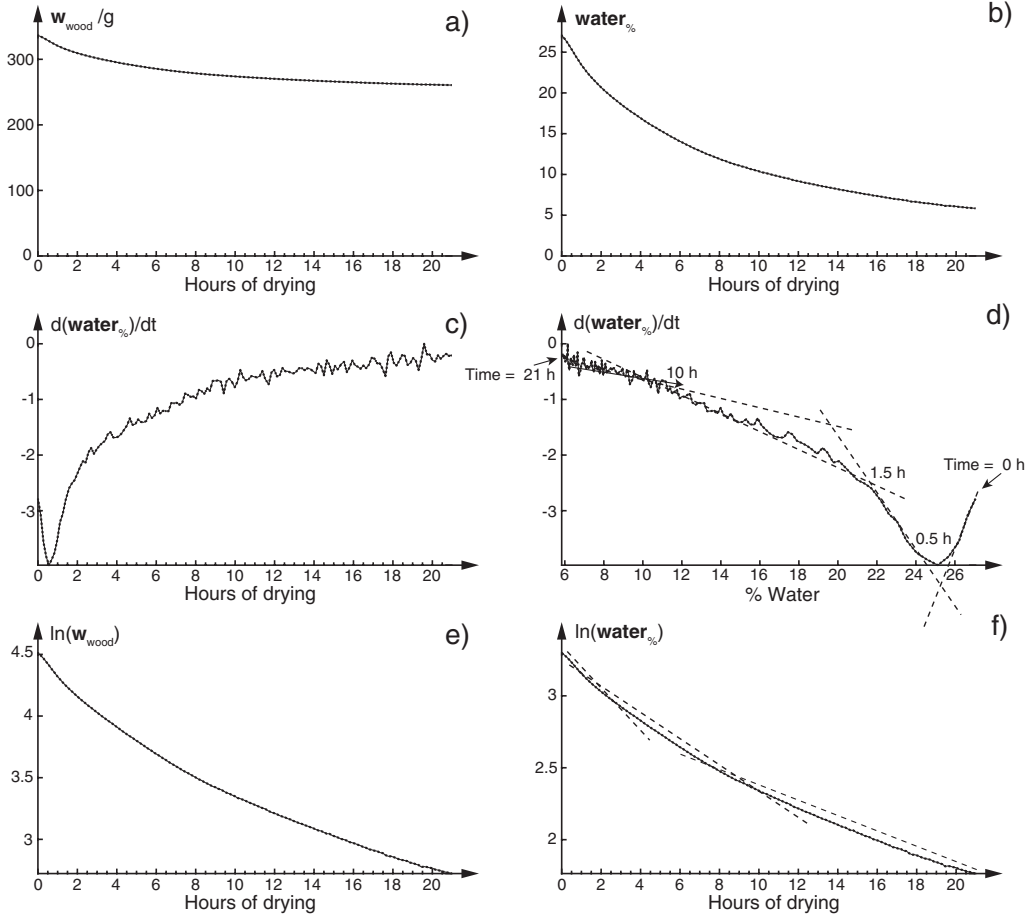


Figure 7.5: a) Weight of wood sample as function of drying time ($w_{\text{wood}}(t)$). b) Percentage of water in wood sample as function of drying time ($\text{sat}_{\text{H}_2\text{O}}(t)$). c) Rate of change $d\text{sat}_{\text{H}_2\text{O}}(t)/dt$ as function of drying time. d) Rate of change $d\text{sat}_{\text{H}_2\text{O}}(t)/dt$ as function of $\text{sat}_{\text{H}_2\text{O}}$, with four local approximation lines. e) $\ln(w_{\text{wood}}(t) - w_{\text{dry wood}})$ as function of drying time. f) $\ln(\text{sat}_{\text{H}_2\text{O}}(t))$ as function of drying time, with three local approximation lines.

$dsat_{H_2O}(t)/dt$ vs. $sat_{H_2O}(t)$ in Figure 7.5 d) had displayed a simple linear relationship, this would have meant that the drying process follows simple first-order kinetics. But at least four different near-linear relationships at distinct drying stages may be observed: 0-0.5h, 0.5-1.5h, 1.5-10h, and 10-21h. This indicates that the drying of wood represents a rather complex dynamic process, where especially the initial hour or two follow very different kinetics.

A more noise robust but also less sensitive way to study the observed dynamics is illustrated by plotting the development of $y = w_{wood}$ (Figure 7.5 e) and $y = sat_{H_2O}$ (Figure 7.5 f) on logarithmic scales against drying time. While a simple first-order kinetics process would have yielded a single straight line, $\ln(y(t)) \approx -k \times t + \ln(y(0))$, clear curvatures are observed. In Figure 7.5 f) three local approximation line segments are drawn, to illustrate the complex drying dynamics. Hence, while the weight loss of the wood sample gives important insight into the overall drying process, it does not expose the different mechanisms involved in the process. Better methods are necessary in order to monitor and study the different mechanisms leading to water evaporation in the wood sample.

7.3.2 Hyperspectral imaging input spectra

Figure 7.6 illustrates the modelling of the vis-NIR spectra in terms of the physical and chemical phenomena expected to dominate the data, for ten typical pixels. The top left and right subplots represent the absorbance spectra of pixels at $t = 0$ (i.e. very wet wood sample) and $t = 21h$ (i.e. dry wood sample). The inserted windows magnifies the wavelength region were water absorbance is seen most clearly. As expected, the absorbance spectra of wet wood (Figure 7.6, top left) show a clear peak in the 940-1005nm region, which the specific absorption coefficient spectrum (Figure 7.3 e) suggests is heavily associated with water absorption. After drying for 21 hours, it decreased, but not disappeared.

Overall, the apparent absorbance spectra seem to be dominated by brown wood pigments showing their strongest light absorbance in the blue and green wavelength regions (<550nm). Moreover, the shorter wavelengths in the visible region appear to have a significantly higher absorbance in a wet wood sample than in a dry one. There are also considerable absorbance variations in the water absorbance range. This is not counterintuitive, considering that a wet piece of wood is perceived as darker than a dry one even to the human eye, and probably represents variations in physical light scattering properties. Less intuitive, however, is the fact that the absorbance decreases, i.e. the wood is perceived as lighter, in the 650-900nm range when saturating the sample with water. This difference in absorbance is likely an effect of different scatter properties between the two states of the sample. Introducing water into the wood pores reduces lateral light scattering [137], which in turn allows more light to be reflected back to the camera.

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations

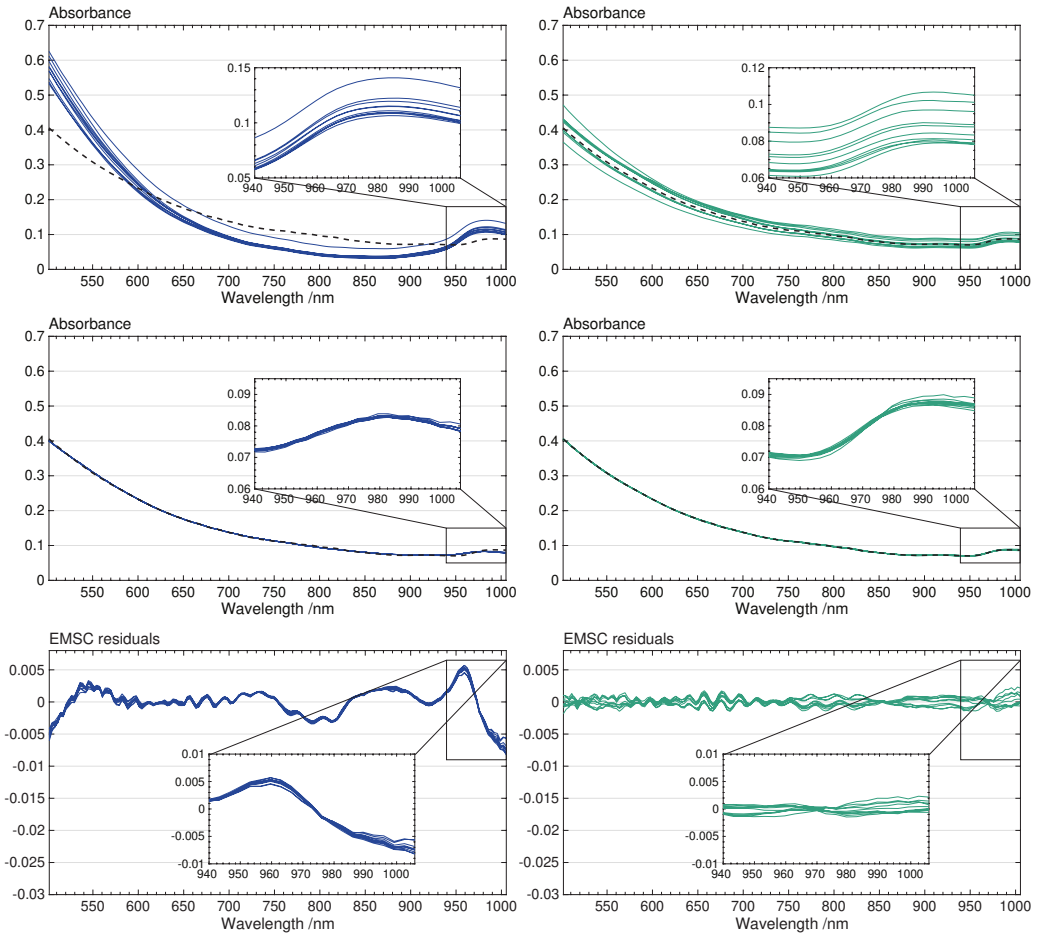


Figure 7.6: Apparent absorbance spectra from 10 typical pixels of the wood sample in wet condition (left) at $t = 0$ h and dry condition (right) at $t = 21$ h. The black dotted line represents the chosen reference spectrum m , which is the average of all pixels in the image taken after 21 hours of drying. Top figures show spectra before EMSC pre-processing. Middle figures show spectra after EMSC pre-processing. Bottom figures show unmodelled spectral residuals after the EMSC modelling. Windows within the figures show a magnification of the 940-1005nm region strongly associated with water absorption.

Hence, the water-related physical properties of the wood, as seen from its light scattering properties (as well as its drying kinetics in Figure 7.5d and 7.5f) have considerable complexity.

7.3.3 Modelling known structures by EMSC

The purpose of the EMSC is to quantify and remove effects from the spectra that are assumed to be understood. The middle row of curves in Figure 7.6 shows the spectra of the same subset of pixels from the wood sample in both wet ($t = 0\text{h}$, left) and dry condition ($t = 21\text{h}$, right) after the EMSC correction:

$$\mathbf{z}_{i,\text{corr}} = \mathbf{m} + \boldsymbol{\varepsilon}_{i,\text{scaled}} \quad (7.22)$$

As can be seen in the middle row of Figure 7.6, clearly the EMSC model largely succeeds in modelling the "known" types of spectral variations, and then removing them, both by subtractions (baseline variations, water and wood pigment variations) and by division (optical path length variations). The EMSC-corrected spectra $\mathbf{z}_{i,\text{corr}}$ are brought together around the chosen reference spectrum \mathbf{m} . By comparing the spectra of the sample in dry and wet states in the 940-1005nm region – a region heavily associated with water absorption – it is clear that the characteristics of the spectra are different in the two states even after the spectral correction. In this region, the dry sample exhibits an S-shaped absorbance while the wet version of the same sample has a more blunt and flat absorbance spectrum. Because the sample underwent desorption (drying) in a laboratory with a room temperature of about 20°C, the sample will certainly still contain chemically bound water after 21 hours of drying. It is possible that the differences in absorbance characteristics seen in the 940-1005nm range can be accredited to the different absorbance properties of free vs. chemically bound water in the wood.

For more graphical resolution, the bottom curves of Figure 7.6 shows the estimated residual spectra $\boldsymbol{\varepsilon}_{i,\text{scaled}}$ (without reference spectrum \mathbf{m}) for the same pixels in sample in the wet ($t = 0\text{ h}$) and dry ($t = 21\text{ h}$) condition. Because most of the residuals are non-zero, there are undoubtedly unknown spectral phenomena taking place within the sample which are not completely captured by the chosen EMSC model spectra shown in Figure 7.3 (left side). The EMSC accounted for more than 99% of the total variance in the 159 weighted wavelength channels. However, it is clear that the residuals are not just random measurement noise. Unmodelled spectral structures are clearly visible, in particular in the water absorbance region and at the shortest wavelengths. Before moving on to analyse these unknown variations, the spatiotemporal properties of the EMSC parameter estimates will be discussed.

7.3.4 Temporal development of fitted EMSC parameters

Figure 7.3 (right side) shows the temporal development of the fitted EMSC parameters associated with the EMSC model vectors (left side), averaged across all pixels at each of the 150 points in time $t = 0, \dots, 21\text{h}$.

Within the first hour of drying, all modelled properties of the wood sample appear

to change rapidly. It is important to note that the hyperspectral monitoring process of the wood sample started immediately after the sample had been taken from a state of full submersion in water. As a consequence of this, there was an amount of liquid water on top of the wood during the initial period of drying, effectively forming a film of water covering the wood and partially cloaking the spectral properties of the wood itself. This probably caused the particularly complex drying behavior at the beginning of the drying process, as also seen in Figure 7.5 d).

Figure 7.6 (bottom left) showed a lot of unmodelled water absorption at $t = 0$ h. It is obvious that the chosen EMSC model spectra alone do not allow an adequate description for the spectra of wood with the highest moisture content. Since the absorbance spectra of bound water probably overlap with the EMSC model of free water (Figure 7.3c), so-called alias errors [84] are expected in the EMSC estimates. For this reason, some of the initial changes seen in the modelled properties of the sample could be misleading during the initial stages of drying, as they have large uncertainties associated with them.

After one to two hours, the average of the multiplicative optical path length score b_i (Figure 7.3f) seems to approach the level of reference. The average additive baseline scores a_i and d_i (Figure 7.3g and h) likewise level off. Hence, at first glance the different physical light scattering parameters show similar behavior.

Concerning the chemical absorbance effects, the average scores for $\Delta c_{i, \text{WoodPigment}}$ should ideally have been constant over time, to the extent the structure of the wood sample itself is constant. As can be seen in subplot (e) of Figure 7.3, this is not the case. While $\Delta c_{i, \text{WoodPigment}}$ does not vary as much as the other EMSC model parameters, some variation is evident. This could indicate that the *in situ* estimated spectrum describing the pigment differences within the wood is too crude to adequately model the inhomogeneity of the wood. The origin of this is not clear, but it could be related to the small spatial contraction expected when the wood dries. Because less energy is required to evaporate free water within a wood sample compared to bound water [137], the drying process of a water saturated piece of wood happens in at least two phases. First, a rapid evaporation of free water occurs, which is then followed by a slower evaporation of bound water. As such, the temporal development of the average water saturation suggested by the EMSC model in Figure 7.3j appears realistic; the curve quickly drops during the first hour, suggesting the evaporation of free water and surface water, which is then followed by a significantly slower drying process for the remaining 20 hours of drying.

The EMSC-based estimate of $\Delta c_{i, \text{Water}}$, is based on a projection on the spectrum of free water, s_{Water} , so it is intended to be proportional to the concentration of free water. This behavior of the estimated free water is what might be expected if the free water were primarily situated on or close to the wood surface. It is distinctly different from that of the weight-based overall moisture content (Figure 7.5b), which represents the loss of both free and bound water.




Figure 7.7: Modelling the known: Spatial structure of EMSC parameters. 2D visualization of fitted EMSC parameters in wet wood sample, i.e. $t = 0$ h, (upper row) and dry wood sample, i.e. $t = 21$ h, (lower row) for all parameters used in the EMSC model.

7.3.5 Spatial development of fitted EMSC parameters

Figure 7.7 shows the corresponding spatial distribution of the EMSC parameters for two of the 150 time points: $t = 0$ h (wet; top) and $t = 21$ h (dry, corresponding to the reference time point for spectrum m ; bottom). In the figure, lighter colors mean high while darker mean low values. All the EMSC parameters show clear spatial patterns of *early-* vs *latewood*, in addition to the differences between wet and dry wood.

In summary, the effective optical path length is lower in dry wood than in wet wood. This is probably due to an increase in light scattering as air replaces water in the wood pores and therefore increases the variability in the refractive index inside the wood material. On the other hand, the absorbance baseline offset increases upon drying. This may also be due to the increasing light scattering, e.g. by increasing the angular distribution of the reflected light and thus reducing the fraction of light reaching the narrow angle of the camera's sensor. The consequence of the path length reduction is apparently stronger than the consequence of the increased spectral absorbance level, since the overall effect is to render dry wood visually lighter than wet wood. A more detailed spatial comparison of *early-* and *latewood* pixels confirms the opposite trend of these two effects of light scattering. The slightly decreasing baseline slope indicates a slightly increased wavelength

dependency of this absorbance baseline. The pigment concentration estimate – $\Delta c_{\text{WoodPigment}}$ – increases slightly; this is probably due to insufficient model detail. Finally, the estimate of free water concentration, Δc_{Water} , falls distinctly upon drying, as expected. The EMSC modelling accounted for 98.04% of the weighted variance.

7.3.6 Modelling unknown structures with OTFP

Each of the residual spectra remaining after the theory-driven EMSC modelling of known phenomena, illustrated for 10 of the pixels at the bottom of Figure 7.6, were divided by the estimates of their relative effective optical path length b_i . These more than 350 million residual spectra, accounting for 1.06% of the total variance of the input spectra, were then submitted to the OTFP software system, one small batch at a time.

Since the OTFP is a weighted least squares procedure that describes the stream of incoming spectra with as few principal components as possible, the first step in the OTFP is to multiply each residual spectrum by the same relevance-vs-reliability weights (Figure 7.4) that were used in the EMSC modelling. Then the bilinear PCA-like model was gradually developed, to describe as much of the variation in the incoming stream of residual data with as few principal components as possible. The sequence of the five first OTFP components representing unknown, but systematic spectral structures, decreased this variance further, from 1.06% to 0.52%, 0.15%, 0.05%, 0.02% and 0.01%, respectively.

7.3.7 Spectral and temporal development of estimated OTFP parameters

Figure 7.8 shows the results for the first five PCs, for the deweighted spectral loadings for the 159 wavelength channels (left), and the temporal scores, averaged over all the pixels in each of the 150 points in time $t = 0, \dots, 21$ h (right). The first four PCs display clear, smooth, loadings and scores. Even though the fifth PC shows rather noisy loadings, its score vector is smooth. Comparing their relative sum-of-squares contributions based on weighted loadings, PCs #1 and #2 accounted for 95% and 4.8%, while PCs #3, #4 and #5 together accounted for 0.2% only. Hence, the first two PCs dominate the weighted spectral residuals.

All five OTFP PCs show very different behavior for the first hour of drying, compared to the rest of the 21 hour drying period. Moreover, PCs #1, #2 and #3 show clear signatures in the water absorbance region above 900nm. PCs #2, #3, #4 likewise show clear spectral features at the lowest wavelengths. PCs #1 and #2 also show some structure at intermediate wavelengths.

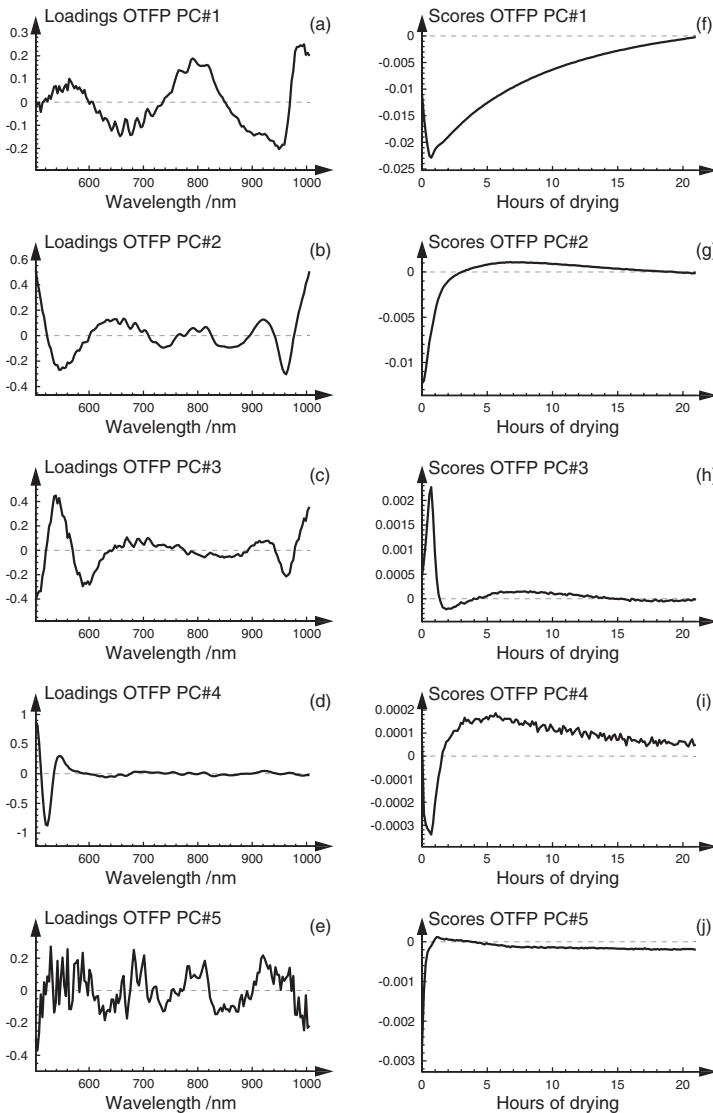


Figure 7.8: Modelling the unknown: Spectral and temporal structure of the parameters from adaptive bilinear modelling in the On-The-Fly-Processing (OTFP) implementation. Left column shows the OTFP model spectra estimated for modelling of apparent absorbance. De-weighted loadings for component 1-5. Right column shows the temporal development of the ABLM parameters (estimated at each point in time by averaging over all image pixels).

7. Hyperspectral time series analysis: Hyperspectral image data streams interpreted by modelling known and unknown variations




Figure 7.9: Modelling the unknown: Spatial structure of OTFP parameters. 2D visualization of reconstructed OTFP scores of wet the wood sample, i.e. $t = 0$ h, (upper row) and dry wood sample, i.e. $t = 21$ h, (lower row) for the five first PCs.

7.3.8 Spatial structure of estimated OTFP parameters

Figure 7.9 shows the corresponding spatial distribution of the OTFP parameters for two of the 150 time points: $t = 0$ h (wet; top) and $t = 21$ h (dry; bottom). Again, in the figure, lighter colors mean high while darker mean low values. In particular, the first, second and fifth components show spatial differences between wet and dry wood.

For instance, PCs #1 and #2 reveal a local spatial structure at the bottom of the wood sample that is not visible in the EMSC parameters. However, in general, the temporal and spatial structures of the OTFP components were not as clear as for the EMSC parameters in Figure 7.3 and 7.7.

A subsequent axis rotation from the basic, orthogonal PCA score and loading representation of the OTFP to a simpler bilinear structure (by e.g. varimax or independent component analysis ICA) may be expected to give an easier wood-related interpretation. However, that is beyond the scope of the present work, which is primarily intended to demonstrate how massive streams of hyperspectral video data can be efficiently modelled in terms of known and unknown features.




Figure 7.10: How the variation at the different wavelengths was explained by the sequence of modelling steps: In the input data, after EMSC and after OTFP PCs #1, #2, #3, #4 and #5. Left: Residual standard deviations, statistically weighted. Right: Residual standard deviations, de-weighted.

7.3.9 Overall model assessment

Statistical summary

Figure 7.10 (left) shows how the standard deviation of each of the 159 wavelength channels is reduced by the succession of modelling stages, for the weighted wavelength channels used in the weighted least squares modelling in EMSC and OTFP. It is clear that the EMSC modelling of known spectral phenomena (dashed green curve) has removed a very large part of the variation. However, substantial residuals remain in the visible wavelength range (610-720nm) and in two NIR ranges (750-840nm and 870-1010nm). Interestingly, the water absorption region 920-970nm shows clear, unmodelled spectral features.

Most of the residuals after the EMSC, including the remaining water absorption, are accounted for by the first OTFP component. The effects of OTFP components 2, 3, 4, and 5 are hardly visible in the plot. Figure 7.10 (right) show the same residual standard deviations after having removed the statistical weights (Figure 7.4).

For statistically optimal EMSC and OTFP modelling from a weighted least squares perspective, the inverse of the residual standard deviation after five OTFP components might have been used as wavelength weights. However, we have chosen to retain the original weights as shown in Figure 7.4, in order to emphasize the chemical relevance in the longer wavelengths (water absorption) and the shortest wavelengths (wood pigment effects).

The following statistics summarize the weighted residual standard deviations (left, Figure 7.10) over the 159 weighted wavelengths: The EMSC modelling of known spectral phenomena accounted for as much as 98.94% of the total initial variance,

leaving 1.06% unexplained. The sequence of five OTFP components representing unknown, but systematic spectral structures, decreased this further to 0.52%, 0.15%, 0.05%, 0.02% and 0.01% of the total initial variance, respectively. Hence, in total, about 99.99% of the input variance was explained by the modelling of both known and unknown spectral phenomena.

It should be noted that due to partial overlap between the known and unknown spectral variation phenomena in the drying wood, alias errors are to be expected in the EMSC scores, which may affect the subsequent OTFP modelling due to the multiplicative (rather than purely additive) residual definition. Moreover, the estimated OTFP loading spectra have the unnatural property of being orthogonal to the chosen EMSC spectra [84]. In Chapter 4 we presented a method for resolving alias problems in purely additive systems. This approach might also have been useful here. However, that would require special attention to the non-additive estimation and correction for the optical path length, so that is not pursued here.

Temporal kinetics modelling

The kinetic analysis of the overall drying process (Figure 7.5) showed that the weight-based estimates of the moisture content dynamics were rather complex. The phase space of the moisture percentage displayed at least four types of kinetics. The EMSC and OTFP models have shown that several different variation phenomena affect the spectroscopic properties of the wood sample during the drying process.

The final modelling step is an attempt at assessing the complexity of the processes happening during the drying process of the wood sample. Each of the individual temporal score averages from the EMSC model (Figure 7.3, right) and the OTFP model (Figure 7.8, right) were modified and fitted to the decay model corresponding to a simple first-order reaction. In each case, the fit included a rescaling, a sign change if needed, and an offset correction, followed by a logarithmic transformation.

Figure 7.11 shows that for the first hour of drying, the simple first-order kinetic model was not suitable. However, in the period 1-20h drying, the log transforms for four of the five EMSC parameters as well as the first (and dominant) OTFP parameter fitted well to the linear model expected for first-order process dynamics, with different rate constants.

Overall model summary

A time series of hyperspectral images, comprising over 350 million spectra, was logarithmically linearized, weighted statistically and then decomposed by mathematical subspace modelling into "known" and "unknown" components.

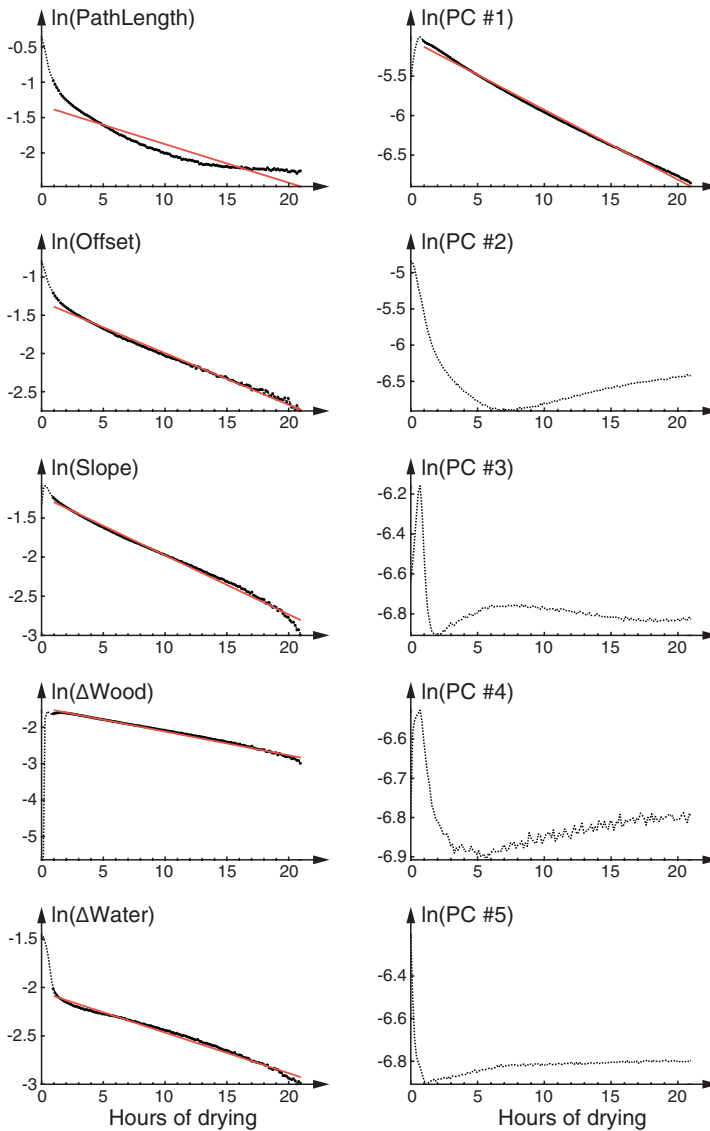


Figure 7.11: Kinetic modelling of the hyperspectral video developments: Analysis of the known and unknown temporal developments of the parameters from EMSC (left) and OTFP (right), averaged over all pixels, as first order dynamic processes. Dotted: $\ln(\text{normalized parameters})$; densely dotted represent the data points used in the least squares estimation of the kinetic parameters. Straight red lines: Model fitted.

The "known" components represented three physical light scattering effects and two chemical absorbance effects expected to affect the measured spectra. Together, these five "known" components accounted for 98.94% of the variance of all statistically weighted spectra, leaving 1.06% variance unexplained. Then a sequence of "unknown" components, i.e., unexpected, but systematic variation patterns, were discovered and extracted from the spectral residuals. The five first components reduced the remaining unexplained variance to 0.52%, 0.15%, 0.05%, 0.02% and 0.01% respectively. Hence, by compressing the 159 wavelength channels into these 10 modelled components, more than 99.98% of the variance of all absorbance data was thereby accounted for.

Linear and bilinear modelling was used in the component estimation, and compensated for by subtraction, in analogy to Beer's law. Moreover, one of the estimated physical components, the relative effective optical path length, was corrected for by division, also in analogy to Lambert's law.

The dynamics of the wood drying process was further assessed by tentatively fitting each of the component averages, after a suitable linear transformation, to the loglinear kinetic model. Some of the known and unknown components fit quite well to this model of a first-order reaction, but with different reaction rates, while other components displayed more complex time developments.

Compared to the initial assessment of the overall weight loss during the wood drying process, the hyperspectral video appeared to give much more detailed, quantitative account of the complex water diffusion and evaporation processes along with their kinetics. It also revealed how their spatial distribution in the wood changed from wet to dry wood.

7.4 Conclusions

By monitoring a drying wood sample, we experimentally demonstrated a generic way in which a stream of hyperspectral time series data can be modelled in terms of *a priori* known and unknown constituent spectra to enable large dimensionality reduction of the data, essentially without any loss of information. An additional benefit of the described methodology, apart from enabling substantial compression of the data, is that it autonomously highlights unidentified systematic spectral variations within the sample being studied, which aids in further exploration and understanding of the underlying chemical and physical processes causing the variations. The kinetic analysis of the weight loss curve of the drying wood sample indicated that the overall drying process was rather complex. This complexity was addressed by resolving the hyperspectral time series data in terms of a multivariate, mixed multiplicative-additive EMSC modelling, involving three known physical variation phenomena related to varying light scattering and two known chem-

ical variation phenomena related to changing wood composition. After removal of these known effects, the hyperspectral imaging data had clear unmodelled spectral variations, particularly in images taken during the first hour of drying.

Most of these residual variations were picked up by the subsequent data-driven OTFP modelling, which revealed two major and a couple of minor unexpected variation patterns. Four of the known variation phenomena and one of the unexpected variation patterns seemed to follow relatively simple first-order kinetics. Most notably, the effective optical path length did clearly not follow first-order kinetics.

The RGB images of the wood sample in Figure 7.1b and 7.1c showed the wood to be darker and more yellow-brown when in a wet state compared to that when dry. This corresponds well to how the drying seemed to affect the light scattering, causing several types of variations, the most dominant one being a strong reduction in the effective optical path length.

Chapter 8

To infinity and beyond

Hyperspectral remote sensing and the SmallSat Lab

8.1 Recapitulation

As we come to the conclusion of this thesis, it is adequate to discuss how the results presented so far contribute to the field of hyperspectral remote sensing.

Instrument design, integration, and testing

In chapter 2, we present our work on the development of a low cost and low mass imaging spectrometer which represents a step towards the democratization of this field of research. Traditionally, hyperspectral instruments are expensive, barring some research groups from using them.

The instrument we describe is affordable enough to be used in high-school level STEM education. Nevertheless, the same design is also useful for professional operations using imaging spectrometers.

Chapter 3 builds on top of the previous chapter, describing a complete payload for hyperspectral remote sensing. This includes the instrument itself, GNSS (Global Navigation Satellite System) receiver, INS (inertial navigation system), and a data timing and synchronization system. Synchronized ancillary position and attitude data is required for producing geographically accurate observations and data products. The complete payload was tested in several field campaigns onboard small rotary- and fixed-wing unmanned systems. These results can be transferred to other moving platforms, such as satellites.

Data enhancement methods

Starting in chapter 4, we move away from hardware design, integration, and testing, towards algorithm development. The first problem we focused on was the removal of shadows (or other undesired variations) from air- and spaceborne hyperspectral images. An important application of (hyperspectral) remote sensing is to monitor changes over time – days, months, or even years. To be able to monitor changes in the target area, we need to be able to separate between changes in illumination – such as sunlight at different times of day, shadows from terrain, or shadows from clouds – from actual changes of interest – i.e., plant health, ocean colour, mineral composition of the surface.

Results of previous work [84] show that if we approximate a hyperspectral observation by a linear model having contributions – spectra \times concentrations – from illumination and ground, and we have an estimate of the spectra of illumination in the scene, we can estimate with great accuracy the spatial distribution – concentrations – of ground components. Spectra of illumination (sunlight) can be estimated using bright pixels in an image and using the difference between shaded and lit areas in close vicinity, where we assume that the ground constitution is the same. However, this leaves us with some ambiguity regarding the spectra of the ground constituents and spatial distribution of illumination. Our method proposes a way to solve that ambiguity and fully resolve both the known and unknown contributions by imposing assumptions derived from domain knowledge on the spatial variation of illumination. In addition, these results can be extended to other problems, where a component in a mixture is well known, but unknown components need to be precisely quantified. E.g., when we need to extract the concentrations and signatures of algae or plankton from a spaceborne hyperspectral image of an area of the ocean, but need to decouple that signal from the well-known signature of the water itself.

In chapter 5 we looked at another method for data enhancement, tackling a different issue: the usual low (spatial) resolution of hyperspectral imaging when compared to RGB (colour) imaging. We proposed a method for fusing co-located hyperspectral and RGB datasets, taking advantage of both the high spatial resolution of RGB and the high spectral resolution of hyperspectral data. Since it is common for remote sensing platforms to have both types of cameras, the relevance of such type of method is clear.

In an oversimplified way, we are able to do this by reducing the spectral dimensionality of HS data and the spatial resolution of RGB data so that we end up with two similarly sized datasets, in all dimensions. Since the observed target is the same, it is fair to assume that the spatial information sensed by both instruments is almost the same. This assumption makes it so that through a regression method, we can find a common representation for both datasets, which in turn means we can apply

the "inverse" of both dimensionality reductions to this common representation, to get a new high -spatial and -spectral resolution dataset.

Furthermore, the discussed method is presented as a pipeline, a way to connect families of submethods of dimensionality reduction and regression. This means the results are generalizable beyond the shown examples, not only in terms of the choice of submethods, but also in the types of data to fuse.

Interpretable data compression

Hyperspectral data is inherently large: for each spatial pixel, we can have several hundreds of bands – wavelengths. However, as reasearch of many dimensionality reduction methods has shown before, the rank of hyperspectral datasets – the number of actual "information" dimensions – is much lower than the number of bands. It is usual that the Principal Component Analysis (PCA) – a reduction method – of a hyperspectral remote sensing dataset results in fewer than 10 components that explain almost all (>99%) of the variance in data. Supposing the original data had 200 bands, we would get a compression rate of about 20 to 1 – impressive! Even though this is a lossy compression, the lost part is mostly noise, so we are able to both compress and clean data. Another advantage of PCA-like methods is that in its compressed form, the dataset is interpretable, usually even more easily than in its original form.

Still, a particularity of many such methods is that they require all data at once to generate a model. This can impose heavy requirements on computational and memory capabilities, so heavy that it quickly becomes unfeasible to run these algorithms on personal computers.

In chapter 6 we present a novel method, the *On-The-Fly Processing* (OTFP) tool, a PCA-like modelling tool for streaming data. The underlying OTFP methodology is relatively fast and simple: it is linear/bilinear and does not require a lot of raw data or large cross-correlation matrices to be kept in memory. Additionally, it allows the high-dimensional data stream to be graphically interpreted and quantitatively used in its compressed state. Unlike adaptive moving window methods, it allows all past and recent time points to be reconstructed and displayed simultaneously.

The OTFP has great potential for remote sensing applications, where both data processing and transmission are limited. By smartly reducing the amount of data to transmit, we ensure that the most relevant data is sent and that bandwidth is not wasted on noise or redundant data.

Towards the end, in chapter 7, we extend the results of OTFP to high-dimensional spatiotemporal data – hyperspectral video.

First, a preprocessing step, where we linearize the spectra to facilitate further processing with linear methods.

Second, we extract the effects of the expected phenomena, giving us an estimation of their distribution over space – i.e., in each frame – and time – i.e., the evolution from frame to frame.

Then, once we remove the *a priori* known variation, we use the OTFP tool to model the systematic effects left in the residuals. This model will have a reduced number of components, each with a spectral signature – loadings – and a temporal evolution from frame to frame, for every spatial pixel – scores.

Finally, the changes over time of the scores of known and unknown effects were modelled to find what type of dynamics they represent.

These results were shown for a sample of wood during the process of drying, starting at a soaked state. However, the extrapolation to remote sensing data, in particular to satellite data of a target area collected over long periods of time is very straightforward.

8.2 Future scope – The SmallSat Lab

Towards the end of my period at NTNU, in 2017-2018, I helped start the NTNU SmallSat Lab¹. In the following sections, I will discuss our ambitions, based on our understanding of matters at the time. Since then, my involvement with the SmallSat Lab decreased, so, most likely, some of the points I mention are now outdated. In the context of my research, the Hyperspectral SmallSat was a capstone project of sorts, mashing together the concepts I presented over the previous chapters.

Our goal was to use a hyperspectral instrument – see Chapter 2 – as a CubeSat payload, to observe and measure ocean phenomena like algal blooms, phytoplankton, or even river plumes, off the coast of Norway.

To enhance data quality, particularly the ground resolution, we required sufficient overlapping exposures. We proposed to solve this problem by having the satellite perform an optimized slewing manoeuvre while pointing at the target of interest – this would require accurate attitude control. We also discussed using an additional RGB camera for resolution improvement, as suggested by the work in Chapter 5. The imaging phase would take about one minute and precede the processing phase. The onboard computer would perform geometric, spectral, and spatio-temporal corrections of the raw data. Furthermore, using advanced onboard algorithms to extract relevant information – Chapter 6 – the downlink requirements would be lowered many-fold when compared to sending raw data. Mission design and preliminary sensitivity analysis of the HSI payload indicated the concept as feasible for dedicated remote sensing off the coast of Norway, with the main limiting factors being cloud cover and a short peak season during summer.

¹<https://www.ntnu.edu/ie/smallsat>


SmallSats could directly augment a coordinated robotic platform containing unmanned and autonomous underwater, surface and aerial systems. Succeeding on the first flight, we at the SmallSat Lab hoped to develop a pipeline of small satellites. Future missions would continue to focus on oceanography and remote sensing, providing higher spatial, temporal, and spectral resolution of a target area than current ocean colour satellites.

8.3 A cubesat for ocean colour remote sensing

From sinks for greenhouse gases to a home for marine life, or even a source of valuable resources, the world's oceans role is undeniable. As such, many scientists, technologists, policy-makers, and the general public understand the oceans' importance for the well-being of those who inhabit the planet, both present and future. The influence of the changing climate and its impact on 70% of our planet covered by water, needs to be studied at resolutions ranging from fine-scale (microbiology) to large-scale (atmospheric phenomena such as hurricanes, the extent of the global ice melt, harmful algal blooms, and fronts).

Traditional methods of taking measurements in the ocean are neither sustainable nor flexible. On the one hand, ship-based measurements require extensive engineering and science infrastructure and support. They also subject people to harsh sea-faring conditions. Additionally, those methods typically result in point measurements for phenomena spread over large spatial extents. On the other hand, multi-national remote sensing capabilities require consensus-driven management, years of development, and substantial ground support to manage the inherent risks involved in building, flying and operating a complex spacecraft, all of which amount to a massive investment. Both space and ocean offer hostile environments for exploration, yet provide critical clues to our origins and continuing well-being [138]. An alternative approach is a network of autonomous underwater vehicles (AUVs), autonomous surface vehicles (ASVs), and unmanned aerial vehicles (UAVs). This network is capable of coordinated missions, executed in concert with conventional vehicles, buoys, and fixed sensors, as envisioned in the Autonomous Ocean Sampling Network (AOSN) [139, 140]. Such systems enable not only a significant reduction in costs, and increased safety, but even more importantly provide substantially more and continuous information about the observed targets and features of scientific interest. They do so by taking advantage of the complementary and coordinated capabilities of such autonomous assets related to position, range, endurance, mobility, sensors, and across large spatiotemporal scales for observing oceanographic phenomena.

A variety of phenomena of oceanographic interest can be detected from space. Most sensors are either optical or radar-based. Our work is currently focused on the



(a) Bloom observed off the coast of Norway on June 10, 2006.

(b) Bloom observed in the Baltic Sea on July 11, 2010.

Figure 8.1: Phytoplankton blooms observed from space, by Envisat's MERIS instrument. *Image courtesy: ESA*

former, typically in the context of looking at processes with a large spatiotemporal extent, like algal blooms. Algal blooms are a relevant case for these technologies, as they have clear surface signatures (Figure 8.1), and some blooms generate neurotoxins with a significant impact on both marine and human populations [141], making their detection important.

A primary light-absorbing substance in the oceans is chlorophyll, which is involved in phytoplankton photosynthesis [142]. Phytoplankton absorbs the red and blue components of the light spectrum and reflects primarily green light. Other substances which absorb light and are composed of organic matter, are often referred to as colour dissolved organic matter (CDOM). A push-broom hyperspectral imager (HSI) operating in the visible and near-infrared (VNIR) spectrum can sweep over the ocean surface to observe both organisms and matter in the upper water column. From this data, the goal is to detect and characterize the spatial extent of algal blooms [143], measure primary productivity, and other substances resulting from aquatic activity and pollution, to support environmental monitoring, climate research and marine resource management. Using the proposed HSI payload on a small satellite platform (SmallSat), in concert with multiple aerial, surface, and underwater vehicles, we propose an autonomous multiagent system for marine observations. Figure 8.2 shows an operational view of such a system.

There are several motivations for the need of such a combined and complementary capability:

- Traditional multipurpose Earth observation (EO) satellites operated by agencies such as NASA and ESA provide excellent data covering the whole




Figure 8.2: Main components in a multiagent marine observation system: a CubeSat equipped with an HSI payload ③; an unmanned aerial vehicle (UAV) ④; an autonomous surface vessel (ASV) ⑤; an autonomous underwater vehicle (AUV) ⑥; and a drifting buoy ⑦. Lines --- represent sensing, while lines - - - represent communications.

Earth. They may have medium to high spatial resolution, but usually have low temporal resolution – revisit times can be days to weeks – and low to medium spectral resolution – generally multispectral rather than hyperspectral [144]. We need to augment upper water column process studies where observing large tracts of the ocean and tracking changes over space and time requires a finer-grained revisit time – typically, this would be in the order of hours to a day.

- Dedicated nano- and micro-satellites [145] can be operated to provide high spatial resolution by pointing an HSI with a small field of view (FOV) to narrower (but still relatively large) target areas of interest, with shorter revisit times and potentially high spectral resolution. They may be using onboard remote sensor data analysis to detect features of special interest while obeying communication bandwidth constraints.
- Simultaneous observations from one or more UAVs with an HSI payload – ④ – and ground truth provided by *in-situ* platforms near and on the ocean surface – ⑤, ⑥, ⑦ – can then be used to accurately measure and validate the features of interest at a much finer scale in the upper water column, and without distortion from the atmosphere or being limited by cloud cover, see Figure 8.2.

- Finally, we hope to leverage the excitement of space flight, as a means to engage young researchers towards ocean monitoring, marine robotics, and the systematic design and build of complex systems with the need to leverage advances in Artificial Intelligence and robotics.

The components of such an autonomous multiagent system must be tightly knit together by communication technology in combination with intelligent information processing and coordinated control as well as mission planning where tasks are dynamically allocated amongst the available assets and systems. Coordination in such a context would involve observing the same patch of the ocean co-temporally across diverse assets with a range of sensing techniques to piece together a comprehensive and cogent view.

While the use of an HSI for earth observation is not new (see [145, 146]), our focus exclusively as a tool for both oceanographic measurements and synoptic *in-situ* field measurements is a novel approach with a significant potential for reducing cost and improving data quality in oceanography. Furthermore, event-response capabilities can be finely tailored to specific events so that impact on the public is limited, e.g., Harmful Algal Blooms (HABs).

The main focus of the SmallSat Lab work is the design of a SmallSat with a push-broom HSI payload, its mission design, and the concept of operation as part of an autonomous multiagent observation system. Both the HSI and SmallSat are under development and several versions of the SmallSat and multiagent system are expected to be operational during the next years.

8.3.1 Conceptual Overview

Figure 8.3 summarizes the capabilities of different agents in an autonomous marine observation system dedicated to ocean colour observations such as the one illustrated in Figure 8.2.

The right end of the spatial scale axis defines the spatial coverage of the sensor system, while the left end defines the smallest spatial scale that can be captured by an ocean colour sensor on the platform. The upper end of the temporal capability defines the endurance of the platform, while the lower end illustrates the revisit time or the fastest temporal scale that can be observed. It should be noted that the endurance of UAVs, ASVs, and AUVs may typically be easily extended by fast relaunch after re-fueling. Furthermore, other important dimensions exist beyond the temporal and spatial scale that also characterize the differences between platforms. One important dimension is related to the maximum speed of UAVs, ASVs, and AUVs, implying that high temporal and spatial resolutions can only be achieved within their range.

This means that these types of vehicles will need adaptive sampling and to be guided towards "interesting" parts of the target area where they can be expected




Figure 8.3: Mapping of capabilities of various sensor platforms used for ocean colour.

to make informative observations [139]. Other important dimensions are related to weather sensitivity, payload weight limitations, and operational complexity. The figure illustrates the complementary capabilities of the different sensor platforms and shows the motivation for a coordinated observation system that exploits the advantages of each platform type.

Based on Figure 8.3, the HSI SmallSat capability is differentiated from traditional optical EO satellites (such as Sentinel 3) by having a better temporal, spectral and spatial resolution, but within a much smaller area and shorter lifetime. A constellation with multiple HSI SmallSat will extend the capability in all dimensions compared to a single SmallSat. We also observe that an HSI SmallSat is a useful and complementary platform to AUVs, ASVs, UAVs, and buoys/drifters in particular given the limited mobility and speed of the mentioned platforms. In particular, the capabilities covered by SmallSat, UAV and ASV are very well aligned with the requirements for observing phytoplankton blooms [147].

To achieve high spectral, spatial, and temporal resolution of an HSI in a SmallSat system with low cost and low weight, our design assumes that the observation target area is limited to a small number of patches of the Earth (possibly only one), as illustrated in Figure 8.3.

We emphasize that our system objective is not to map the entire Earth, but a tiny fraction corresponding to our specific target area(s) of interest. This enables the use of an imaging system with a relatively narrow FoV. With accurate attitude control and a slewing motion, the push-broom HSI will sweep over a small area. See




Figure 8.4: Concept of operations for SmallSat:

- ① Wake-up and Uplink: The SmallSat will spend most of its orbit in standby, harvesting solar energy. When nearing the ground station ③, it wakes up, to receive an updated mission plan – target area, cloud coverage information, observations by other unmanned assets ④ to be used for calibration, etc.
- ② Preparation: The SmallSat activates attitude control, starting to point the HSI towards the target area.
- ③ Start observation: The SmallSat enables the HSI and starts recording data to onboard storage.
- ④ Target tracking: Accurate pointing and slewing motion control during acquisition to maximize data quality.
- ⑤ End observation and Data processing: After the target area is scanned, the SmallSat disables the HSI. Then, it starts processing the datacube, extracting relevant information to be sent to the ground.
- ⑥ Downlink and Sleep: Depending on the location of the next ground antenna ④ that can communicate with the SmallSat, it might directly downlink the results of its tasks, or go to standby mode and wake up for communication at a later time.

Figure 8.4 for an overview of the different phases of the SmallSat orbit.

8.3.2 Scientific Objectives

The main driver for this mission is oceanography, specifically dedicated to narrow field-of-view monitoring and mapping of ocean colour phenomena. Key science objectives are to:

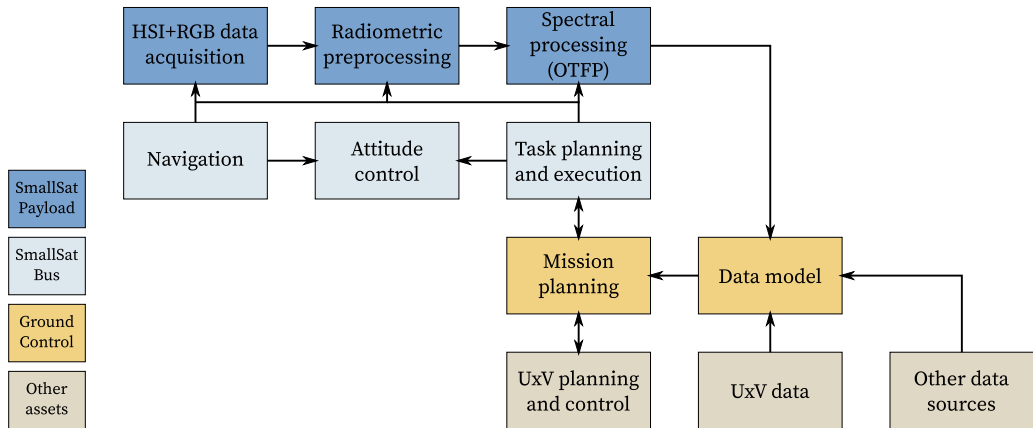


Figure 8.5: Simplified data processing and control architecture.

- detect algal blooms and phytoplankton in both Case 1 and Case 2 waters² and determine chl-a, -b, -c₃, -c₄ concentrations
- detect other discolourations of oceans such as oil spills and river plumes
- perform *in-situ* validation of remote sensing data with UAVs, ASVs, AUVs, manual sampling

8.3.3 HSI Data Processing and Control Architecture

Figure 8.5 illustrates the information flow, integrated processing and control architecture. The layers from top to bottom correspond to the SmallSat HSI payload with onboard processing, SmallSat bus onboard task planning and control, ground control/operations with data and model management, and payloads and control of other assets such as UAV/AUV/USV and traditional satellite data products.

A brief description of the main functionality of the blocks in Figure 8.5 is given below:

- The SmallSat HSI payload and its On-the-Fly-Processing (OTFP) follow a processing chain that is similar to what is found in other small satellite systems, e.g. [148]:
 - The HSI+RGB data acquisition block captures image frames, both hyperspectral and RGB. The hyperspectral data is composed of several thousands of frames that build up a datacube – this is the push-broom concept explained in previous chapters. The RGB data is orders of magnitude smaller in terms of the number of frames, but should have better

²Case 1: Open clear water; Case 2: Coastal turbid waters

spatial resolution, so that it can be used for data enhancement as presented in Chapter 5. Prediction and observations of illumination (sun position and atmospheric conditions such as clouds) can be used to optimize the attitude control to get the best possible data results within the target area.

- Radiometric preprocessing transforms each pixel value into absolute reflectance values. This involves applying calibration parameters obtained from models and measurements of the atmosphere, solar radiation and/or geo-referenced reference measurements near the ground.
- Spectral processing is based on multivariate data modelling that projects the data onto expected and unexpected/observed spectral signatures corresponding to the features of interest, see Chapter 6. This acts as a smart data compression algorithm before data downlink.
- The SmallSat bus system has functions for attitude control (reaction wheels and magnetic torquing), navigation (GPS, INS, star-tracker), communication datalink that can be partly used by the data payload, resource/energy/power management and resilience/redundancy/fault tolerance.
- The ground control segment includes overall mission planning and execution, data management, oceanographic model with data assimilation and simulation capabilities, operator interfaces, etc. The overall mission management will optimize the use of available resources and assets by providing information and tasks to each one of them. Data received from the SmallSat and other assets will be used to update the oceanographic data model. To improve this model, the incoming data will be subjected to several algorithms such as detection of features, removal of the effects of specular reflections in the water, shadows due to clouds, and other undesired optical features, resolution enhancement, see Chapters 4 and 5.
- Other assets include buoys, manned and unmanned aerial/surface/underwater vehicles with payload sensors, traditional satellite data products, and meta-ocean services.

8.4 Final considerations

By applying techniques like OTFP (see Chapter 6), we are able to reduce the size of data to be downloaded from remote sensing systems many-fold. Even though this is a lossy compression, ideally, the loss of relevant data is minimal and we are able to "shave off" noise and unwanted information.

Furthermore, by using data enhancement algorithms like de-shadowing (see Chapter 4) and MVIF (see Chapter 5) we improve the quality of the models, and sub-

sequently the knowledge gained with low-cost remote sensing platforms. A big challenge with some of these algorithms is optimizing them for embedded computational systems, such as the ones that would preferably be used in a Cubesat or a drone. Embedded systems have limited resources when compared to personal computers where us researchers develop and test these algorithms. Mixed-architecture systems with ARM and FPGA cores present new opportunities for enhancing performance with strict power restrictions.

There is no doubt that we need tools to handle this data tsunami, be it from remote sensing, the Internet-of-Things (IoT), or one of the many sources of enormous amounts of data. My time at NTNU allowed me to explore just a small subset of these tools.

However, even more importantly than the opportunity to learn specific algorithms, it exposed me to many different data problems: from airborne hyperspectral data to thermal camera images of a ship's engine room, vibration data from a failing bearing, or even zebra-fish neural activity. Most of these problems can, in the end, be solved by using just one of the many different algorithms, so the key thing is knowing how to transform the raw data in such a way that it will fit nicely into the input box of your algorithm of choice. I believe that figuring out how to do this is one of the last things machines will learn by themselves.

As someone once said, "*Sometimes science is more art than science. A lot of people don't get that.*". "Getting it" is one of the most useful things I learned.

References

- [1] NASA, “Data Processing Levels | Earthdata,” <https://earthdata.nasa.gov/collaborate/open-data-services-and-software/data-information-policy/data-levels>, 2019.
- [2] C. K. Singh, Ed., *Geospatial Applications for Natural Resources Management*, CRC Press, 2018.
- [3] D. Caballero, R. Calvini, and J. M. Amigo, “Hyperspectral imaging in crop fields: precision agriculture,” in *Hyperspectral Imaging*, J. M. Amigo, Ed., vol. 32 of *Data Handling in Science and Technology*. Elsevier, 2020.
- [4] D. J. Keith, B. A. Schaeffer, R. S. Lunetta, R. W. Gould Jr., K. Rocha, and D. J. Cobb, “Remote sensing of selected water-quality indicators with the hyperspectral imager for the coastal ocean (HICO) sensor,” *International Journal of Remote Sensing*, vol. 35, no. 9, 2014.
- [5] U. Heiden, W. Heldens, S. Roessner, K. Segl, T. Esch, and A. Mueller, “Urban structure type characterization using hyperspectral remote sensing and height information,” *Landscape and Urban Planning*, vol. 105, no. 4, 2012.
- [6] P. A. Jerram, M. Fryer, J. Pratlong, A. Pike, A. Walker, B. Dierickx, B. Dupont, and A. Defernez, “Hyperspectral CMOS imager,” in *International Conference on Space Optics ICSO 2010*, E. Armandillo, B. Cugny, and N. Karafolas, Eds. International Society for Optics and Photonics, 2017, vol. 10565, SPIE.
- [7] G. Vane, Ed., *Proceedings of SPIE Imaging Spectrometry II*, vol. 0834, 1987.
- [8] W. L. Wolfe, *Introduction to Imaging Spectrometers*, SPIE, 1997.
- [9] E. Herrala and J. Okkonen, “Imaging spectrograph and camera solutions for industrial applications,” *International Journal of Pattern Recognition and Artificial Intelligence*, vol. 10, no. 01, 1996.

- [10] T. S. Hyvärinen, E. Herrala, and A. Dall’Ava, “Direct sight imaging spectrograph: a unique add-in component brings spectral imaging to industrial applications,” in *Digital Solid State Cameras: Designs and Applications*. 1998, vol. 3302, SPIE.
- [11] F. Sigernes, K. Heia, H. Nilsen, and T. Svenøe, “Imaging spectroscopy applied in the fish industry?,” in *Norwegian Society for Image Processing and Pattern Recognition 2*, 1998.
- [12] F. Sigernes, D. A. Lorentzen, K. Heia, and T. Svenøe, “Multipurpose spectral imager,” *Applied Optics*, vol. 39, no. 18, 2000.
- [13] J. Hosen, H. H. Helgesen, L. Fusini, T. I. Fossen, and T. A. Johansen, “Vision-Aided Nonlinear Observer for Fixed-Wing Unmanned Aerial Vehicle Navigation,” *Journal of Guidance, Control, and Dynamics*, vol. 39, no. 8, 2016.
- [14] P. J. Miller, “Use of Tunable Liquid Crystal Filters to Link Radiometric and Photometric Standards,” *Metrologia*, vol. 28, no. 3, 1991.
- [15] F. Sigernes, Y. Ivanov, S. Chernouss, T. Trondsen, A. Roldugin, Y. Fedorenko, B. Kozelov, A. Kirillov, I. Kornilov, V. Safargaleev, S. Holmen, M. Dyrland, D. Lorentzen, and L. Baddeley, “Hyperspectral all-sky imaging of auroras,” *Optics Express*, vol. 20, no. 25, 2012.
- [16] J. Antila, R. Mannila, U. Kantojärvi, C. Holmlund, A. Rissanen, I. Näkki, J. Ollila, and H. Saari, “Spectral imaging device based on a tuneable MEMS Fabry-Perot interferometer,” in *Next-Generation Spectroscopic Technologies V*. 2012, vol. 8374, SPIE.
- [17] Specim, “AisaKESTREL,” <http://www.specim.fi/products/aisakestrel-hyperspectral-imaging-system/>, Accessed March 2018.
- [18] Resonon, “Airborne Hyperspectral Imaging Systems,” <https://www.resonon.com/Products/airborne.html>, Accessed March 2018.
- [19] Headwall, “Nano-Hyperspec,” <http://www.headwallphotonics.com/spectral-imaging/hyperspectral/nano-hyperspec>, Accessed March 2018.
- [20] L. Loncan, L. B. De Almeida, J. M. Biucas-Dias, X. Briottet, J. Chanussot, N. Dobigeon, S. Fabre, W. Liao, G. A. Licciardi, M. Simões, J. Y. Tournet, M. A. Veganzones, G. Vivone, Q. Wei, and N. Yokoya, “Hyperspectral Pan-sharpening: A Review,” *IEEE Geoscience and Remote Sensing Magazine*, vol. 3, no. 3, 2015.

-
- [21] N. Yokoya, C. Grohnfeldt, and J. Chanussot, "Hyperspectral and Multispectral Data Fusion: A comparative review of the recent literature," *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, no. 2, 2017.
- [22] L. Bungert, D. A. Coomes, M. J. Ehrhardt, J. Rasch, R. Reisenhofer, and C.-B. Schönlieb, "Blind image fusion for hyperspectral imaging with the directional total variation," *Inverse Problems*, vol. 34, no. 4, 2018.
- [23] N. Akhtar, F. Shafait, and A. Mian, "Bayesian sparse representation for hyperspectral image super resolution," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [24] M. Simões, J. Bioucas-Dias, L. B. Almeida, and J. Chanussot, "A Convex Formulation for Hyperspectral Image Superresolution via Subspace-Based Regularization," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 6, 2015.
- [25] Q. Wei, J. Bioucas-Dias, N. Dobigeon, and J. Tourneret, "Hyperspectral and Multispectral Image Fusion Based on a Sparse Representation," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 53, no. 7, 2015.
- [26] Q. Wei, N. Dobigeon, and J. Tourneret, "Bayesian Fusion of Multi-Band Images," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 6, 2015.
- [27] Y. Zhang, S. D. Backer, and P. Scheunders, "Noise-Resistant Wavelet-Based Bayesian Fusion of Multispectral and Hyperspectral Images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 47, no. 11, 2009.
- [28] Y. Chang, L. Yan, H. Fang, S. Zhong, and Z. Zhang, "Weighted Low-rank Tensor Recovery for Hyperspectral Image Restoration," 2017.
- [29] R. Dian, L. Fang, and S. Li, "Hyperspectral Image Super-Resolution via Non-local Sparse Tensor Factorization," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [30] K. Zhang, M. Wang, S. Yang, and L. Jiao, "Spatial-Spectral-Graph-Regularized Low-Rank Tensor Decomposition for Multispectral and Hyperspectral Image Fusion," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 11, no. 4, 2018.
- [31] H. Li, L. Jing, Y. Tang, and H. Ding, "An Improved Pansharpening Method for Misaligned Panchromatic and Multispectral Data," *Sensors*, vol. 18, no. 2, 2018.

- [32] C. I. Kanatsoulis, X. Fu, N. D. Sidiropoulos, and W. Ma, “Hyperspectral Super-Resolution: A Coupled Tensor Factorization Approach,” *IEEE Transactions on Signal Processing*, vol. 66, no. 24, 2018.
- [33] Y. Xu, Z. Wu, J. Chanussot, and Z. Wei, “Nonlocal Patch Tensor Sparse Representation for Hyperspectral Image Super-Resolution,” *IEEE Transactions on Image Processing*, vol. 28, no. 6, 2019.
- [34] R. Dian, S. Li, and L. Fang, “Learning a Low Tensor-Train Rank Representation for Hyperspectral Image Super-Resolution,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 30, no. 9, 2019.
- [35] R. Kawakami, Y. Matsushita, J. Wright, M. Ben-Ezra, Y. W. Tai, and K. Ikeuchi, “High-resolution hyperspectral imaging via matrix factorization,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. 2011, IEEE.
- [36] N. Yokoya, T. Yairi, and A. Iwasaki, “Coupled Nonnegative Matrix Factorization Unmixing for Hyperspectral and Multispectral Data Fusion,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, no. 2, 2012.
- [37] N. Akhtar, F. Shafait, and A. Mian, “Sparse Spatio-spectral Representation for Hyperspectral Image Super-resolution,” in *Computer Vision – ECCV 2014*, D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds. 2014, Springer International Publishing.
- [38] M. Selva, B. Aiazzi, F. Butera, L. Chiarantini, and S. Baronti, “Hyper-Sharpening: A First Approach on SIM-GA Data,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 8, no. 6, 2015.
- [39] C. Lanaras, E. Baltsavias, and K. Schindler, “Hyperspectral Super-Resolution by Coupled Spectral Unmixing,” in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [40] H. Kwon and Y. Tai, “RGB-Guided Hyperspectral Image Upsampling,” in *2015 IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [41] Q. Wei, J. Bioucas-Dias, N. Dobigeon, J. Tourneret, M. Chen, and S. Godsill, “Multiband Image Fusion Based on Spectral Unmixing,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 54, no. 12, 2016.
- [42] W. Dong, F. Fu, G. Shi, X. Cao, J. Wu, G. Li, and X. Li, “Hyperspectral Image Super-Resolution via Non-Negative Structured Sparse Representation,” *IEEE Transactions on Image Processing*, vol. 25, no. 5, 2016.

-
- [43] M. A. Veganzones, M. Simões, G. Licciardi, N. Yokoya, J. M. Bioucas-Dias, and J. Chanussot, "Hyperspectral super-resolution of locally low rank images from complementary multisource data," *IEEE Transactions on Image Processing*, vol. 25, no. 1, 2016.
- [44] K. Zhang, M. Wang, and S. Yang, "Multispectral and Hyperspectral Image Fusion Based on Group Spectral Embedding and Low-Rank Factorization," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 55, no. 3, 2017.
- [45] C. Yi, Y. Zhao, and J. C. Chan, "Hyperspectral Image Super-Resolution Based on Spatial and Spectral Correlation Fusion," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 56, no. 7, 2018.
- [46] X. Han, B. Shi, and Y. Zheng, "Self-Similarity Constrained Sparse Representation for Hyperspectral Image Super-Resolution," *IEEE Transactions on Image Processing*, vol. 27, no. 11, 2018.
- [47] L. Zhang, W. Wei, C. Bai, Y. Gao, and Y. Zhang, "Exploiting Clustering Manifold Structure for Hyperspectral Imagery Super-Resolution," *IEEE Transactions on Image Processing*, vol. 27, no. 12, 2018.
- [48] Y. Fu, Y. Zheng, H. Huang, I. Sato, and Y. Sato, "Hyperspectral Image Super-Resolution With a Mosaic RGB Image," *IEEE Transactions on Image Processing*, vol. 27, no. 11, 2018.
- [49] Z. Pan and H. Shen, "Multispectral Image Super-Resolution via RGB Image Fusion and Radiometric Calibration," *IEEE Transactions on Image Processing*, vol. 28, no. 4, 2019.
- [50] R. Dian and S. Li, "Hyperspectral Image Super-Resolution via Subspace-Based Low Tensor Multi-Rank Regularization," *IEEE Transactions on Image Processing*, vol. 28, no. 10, 2019.
- [51] C. Wang, Y. Liu, X. Bai, W. Tang, P. Lei, and J. Zhou, "Deep Residual Convolutional Neural Network for Hyperspectral Image Super-Resolution," in *Image and Graphics*, Y. Zhao, X. Kong, and D. Taubman, Eds. 2017, Springer International Publishing.
- [52] X. Han, B. Shi, and Y. Zheng, "SSF-CNN: Spatial and Spectral Fusion with CNN for Hyperspectral Image Super-Resolution," in *2018 25th IEEE International Conference on Image Processing (ICIP)*, 2018.
- [53] R. Dian, S. Li, A. Guo, and L. Fang, "Deep Hyperspectral Image Sharpening," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 29, no. 11, 2018.

- [54] Y. Chang, L. Yan, H. Fang, S. Zhong, and W. Liao, "HSI-DeNet: Hyperspectral Image Restoration via Convolutional Neural Network," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 57, no. 2, 2019.
- [55] Y. Qu, H. Qi, and C. Kwan, "Unsupervised Sparse Dirichlet-Net for Hyperspectral Image Super-Resolution," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [56] O. Sidorov and J. Y. Hardeberg, "Deep Hyperspectral Prior: Denoising, Inpainting, Super-Resolution," 2019.
- [57] Q. Xie, M. Zhou, Q. Zhao, D. Meng, W. Zuo, and Z. Xu, "Multispectral and Hyperspectral Image Fusion by MS/HS Fusion Net," 2019.
- [58] X. Han, Y. Zheng, and Y. Chen, "Multi-Level and Multi-Scale Spatial and Spectral Fusion CNN for Hyperspectral Image Super-Resolution," in *The IEEE International Conference on Computer Vision (ICCV) Workshops*, 2019.
- [59] Y. Zhou, A. Rangarajan, and P. D. Gader, "An Integrated Approach to Registration and Fusion of Hyperspectral and Multispectral Images," *IEEE Transactions on Geoscience and Remote Sensing*, 2019.
- [60] W. Wang, W. Zeng, Y. Huang, X. Ding, and J. Paisley, "Deep Blind Hyperspectral Image Fusion," in *The IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [61] L. Buydens, "Towards tsunami-resistant chemometrics," *The Analytical Scientist*, vol. 0813, 2013.
- [62] A. Katal, M. Wazid, and R. Goudar, "Big Data: issues, challenges, tools and good practices," in *Sixth International Conference on Contemporary Computing (IC3)*. 2013, IEEE.
- [63] C. Zhang, N. C. Anzalone, R. P. Faria, and J. M. Pearce, "Open-Source 3D-Printable Optics Equipment," *PLOS ONE*, vol. 8, no. 3, 2013.
- [64] M. A. Hossain, J. Canning, S. Ast, K. Cook, P. J. Rutledge, and A. Jamalipour, "Combined "dual" absorption and fluorescence smartphone spectrometers," *Optics Letters*, vol. 40, no. 8, 2015.
- [65] T. C. Wilkes, A. J. S. McGonigle, J. R. Willmott, T. D. Pering, and J. M. Cook, "Low-cost 3D printed 1 nm resolution smartphone sensor-based spectrometer: instrument design and application in ultraviolet spectroscopy," *Optics Letters*, vol. 42, no. 21, 2017.

- [66] L. Wang, Y. Chang, R. Sun, and L. Li, "A multichannel smartphone optical biosensor for high-throughput point-of-care diagnostics," *Biosensors and Bioelectronics*, vol. 87, 2017.
- [67] C. Palmer and E. Loewen, *Diffraction Grating Handbook*, Thermo RGL, 2000.
- [68] F. Sigernes, J. M. Holmes, M. Dyrland, D. A. Lorentzen, S. A. Chernous, T. Svinyu, J. Moen, and C. S. Deehr, "Absolute calibration of optical devices with a small field of view," *Journal of Optical Technology*, vol. 74, no. 10, 2007.
- [69] W. Nijland, R. de Jong, S. M. de Jong, M. A. Wulder, C. W. Bater, and N. C. Coops, "Monitoring plant condition and phenology using infrared sensitive consumer grade digital cameras," *Agricultural and Forest Meteorology*, vol. 184, 2014.
- [70] S. M. Albrektsen and T. A. Johansen, "User-Configurable Timing and Navigation for UAVs," *Sensors*, vol. 18, no. 8, 2018.
- [71] "SenTiBoard," <https://sentiboard.com/>, Accessed March 2018.
- [72] T. H. Bryne, J. M. Hansen, R. H. Rogne, N. Sokolova, T. I. Fossen, and T. A. Johansen, "Nonlinear Observers for Integrated INS/GNSS Navigation: Implementation Aspects," *IEEE Control Systems Magazine*, vol. 37, no. 3, 2017.
- [73] F. Dell'Endice, J. Nieke, B. Koetz, M. E. Schaepman, and K. Itten, "Improving radiometry of imaging spectrometers by using programmable spectral regions of interest," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 64, no. 6, 2009.
- [74] IDS, "UI-3360CP Rev. 2," <https://en.ids-imaging.com/store/ui-3360cp-rev-2.html>, Accessed March 2018.
- [75] IDS, "uEye Camera Manual," https://en.ids-imaging.com/manuals/uEye_SDK/EN/uEye_Manual_4.90.3/index.html, Accessed March 2018.
- [76] "Fluorescent lighting spectrum peaks labelled — Wikipedia, The Free Encyclopedia," https://commons.wikimedia.org/wiki/File:Fluorescent_lighting_spectrum_peaks_labelled.gif, Accessed March 2018.
- [77] F. Sigernes, "Airborne Hyperspectral Imaging," Tech. Rep., University Centre in Svalbard, 2002.

- [78] C. Shop, “InCase PIN series NEO-M8T TIME & RAW receiver board (RTK ready),” http://www.csgshop.com/product.php?id_product=240, Accessed March 2018.
- [79] Hardkernel, “Odroid-XU4,” http://www.hardkernel.com/main/products/prdt_info.php?g_code=G143452239825, Accessed March 2018.
- [80] J. Pinto, P. S. Dias, R. Martins, J. Fortuna, E. Marques, and J. Sousa, “The LSTS toolchain for networked vehicle systems,” in *2013 MTS/IEEE OCEANS - Bergen*, 2013.
- [81] A. Zolich, T. A. Johansen, K. Cisek, and K. Klausen, “Unmanned aerial system architecture for maritime missions,” in *2015 Workshop on Research, Education and Development of Unmanned Aerial Systems (RED-UAS)*, 2015.
- [82] “DUNE: Unified Navigation Environment, GitHub repository,” <https://github.com/lsts/dune>, Accessed March 2018.
- [83] S. M. Adler-Golden, M. W. Matthew, G. P. Anderson, G. W. Felde, and J. A. Gardner, “Algorithm for de-shadowing spectral imagery,” in *Proc.SPIE Imaging Spectrometry VIII*, 2002, vol. 4816.
- [84] H. Martens, “The informative converse paradox: Windows into the unknown,” *Chemometrics and Intelligent Laboratory Systems*, vol. 107, no. 1, 2011.
- [85] J. Lofberg, “YALMIP : a toolbox for modeling and optimization in MATLAB,” in *2004 IEEE International Conference on Robotics and Automation (IEEE Cat. No.04CH37508)*, 2004.
- [86] J. Jaumot, A. de Juan, and R. Tauler, “MCR-ALS GUI 2.0: New features and applications,” *Chemometrics and Intelligent Laboratory Systems*, vol. 140, 2015.
- [87] NASA JPL, “AVIRIS Data Portal,” https://aviris.jpl.nasa.gov/alt_locator/, Accessed June 2016.
- [88] “Color Vision — Wikipedia, The Free Encyclopedia,” http://en.wikipedia.org/w/index.php?title=Estimation_lemma&oldid=375747928, Accessed June 2016.
- [89] NASA (@nasa) on Unsplash, “Water, nature, ocean and reef,” https://unsplash.com/photos/6-jTZysYY_U, Accessed August 2019.

- [90] C. Bedia, À. Sierra, and R. Tauler, "Application of chemometric methods to the analysis of multimodal chemical images of biological tissues," *Analytical and Bioanalytical Chemistry*, 2020.
- [91] S. Piqueras, C. Bedia, C. Beleites, C. Krafft, J. Popp, M. Maeder, R. Tauler, and A. de Juan, "Handling Different Spatial Resolutions in Image Fusion by Multivariate Curve Resolution-Alternating Least Squares for Incomplete Image Multisets," *Analytical Chemistry*, vol. 90, no. 11, 2018.
- [92] H. Martens and T. Næs, *Multivariate calibration*, John Wiley & Sons, 1992.
- [93] G. H. Golub and C. Reinsch, "Singular value decomposition and least squares solutions," in *Linear Algebra*. Springer, 1971.
- [94] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, 1999.
- [95] A. de Juan and R. Tauler, "Multivariate Curve Resolution (MCR) from 2000: Progress in concepts and applications," *Critical Reviews in Analytical Chemistry*, vol. 36, no. 3-4, 2006.
- [96] A. Hyvärinen and E. Oja, "Independent component analysis: algorithms and applications," *Neural networks*, vol. 13, no. 4-5, 2000.
- [97] I. T. Jolliffe, *Principal Component Analysis*, Springer Series in Statistics. Springer-Verlag, New York, 2 edition, 2002.
- [98] Oregon State University, "HICO - Hyperspectral Imager for the Coastal Ocean," <http://hico.coas.oregonstate.edu/>, Accessed August 2019.
- [99] Open Remote Sensing, "Hyperspectral Pansharpening: A review," <https://openremotesensing.net/knowledgebase/hyperspectral-pansharpening-a-review/>, 2015, Accessed April 2019.
- [100] R. B. D'agostino Sr and H. K. Russell, "Scree Test," in *Wiley StatsRef: Statistics Reference Online*. American Cancer Society, 2014.
- [101] D. Salomon and G. Motta, *Handbook of Data Compression*, Springer-Verlag Inc., London, UK, fifth edition, 2010.
- [102] H. Martens, "Quantitative Big Data: where chemometrics can contribute," *Journal of Chemometrics*, vol. 29, 2015.
- [103] I. Jolliffe, *Principal Component Analysis*, Springer-Verlag Inc., New York, USA, second edition, 2002.

- [104] S. Wold, "A theoretical foundation of extrathermodynamic relationships (Linear Free Energy relationships)," *Chemica Scripta*, vol. 5, 1974.
- [105] S. Wold and M. Sjöström, "Chemometrics and its roots in physical organic chemistry," *Acta Chemica Scandinavica*, vol. 52, 1998.
- [106] H. Martens, K. Tøndel, V. Tafintseva, A. Kohler, E. Plahte, J. Vik, A. Gjuvsland, and S. Omholt, *New perspectives in Partial Least Squares and related methods*, vol. 56, chapter PLS-based multivariate metamodeling of dynamic systems, Springer-Verlag Inc., New York, USA, first edition, 2013.
- [107] C. Eckart and G. Young, "The approximation of one matrix by another of lower rank," *Psychometrika*, vol. 1, 1936.
- [108] A. Balsubramani, S. Dasgupta, and Y. Freund, "The fast convergence of incremental PCA," in *Advances in Neural Information Processing Systems 26*. 2013, Curran Associates, Inc.
- [109] N. Halko, P. Martinsson, Y. Shkolnisky, and M. Tygert, "An algorithm for the Principal Component Analysis of large data sets," *SIAM Journal on Scientific Computing*, vol. 33, 2011.
- [110] N. Kettaneh, A. Berglund, and S. Wold, "PCA and PLS with very large data sets," *Computational Statistics and Data Analysis*, vol. 48, 2005.
- [111] E. Rabani and S. Toledo, "Out-of-core SVD and QR decompositions," in *10th SIAM Conference on Parallel Processing for Scientific Computing*. 2001, Society for Industrial and Applied Mathematics.
- [112] F. Vogt and M. Tacke, "Fast Principal Component Analysis of large data sets," *Chemometrics and Intelligent Laboratory Systems*, vol. 59, 2001.
- [113] J. Camacho, "Visualizing Big Data with Compressed Score Plots: approach and research challenges," *Chemometrics and Intelligent Laboratory Systems*, vol. 135, 2014.
- [114] R. Barnes, M. Dhanoa, and S. Lister, "Standard Normal Variate transformation and de-trending of near-infrared diffuse reflectance spectra," *Applied Spectroscopy*, vol. 43, 1989.
- [115] H. Martens, S. Jensen, and P. Geladi, "Multivariate linearity transformation for near-infrared reflectance spectrometry," in *Proc. Nordic Symp. on Applied Statistics*, O. Christie, Ed. 1983, Stokkand Forlag Publ., Stavanger, Norway.

-
- [116] P. Geladi, D. MacDougall, and H. Martens, "Linearization and scatter-correction for near-infrared reflectance spectra of meat," *Applied Spectroscopy*, vol. 39, 1985.
- [117] H. Martens, J. Nielsen, and S. Engelsen, "Light scattering and light absorbance separated by Extended Multiplicative Signal Correction. Application to near-infrared transmission analysis of powder mixtures," *Analytical Chemistry*, vol. 75, 2003.
- [118] I. Endrizzi, F. Gasperi, M. Rødbotten, and T. Næs, "Interpretation, validation and segmentation of preference mapping models," *Food Quality and Preference*, vol. 32, 2014.
- [119] R. Vitale, J. A. Westerhuis, T. Næs, A. K. Smilde, O. E. de Noord, and A. Ferrer, "Selecting the number of factors in principal component analysis by permutation testing Numerical and practical aspects," *Journal of Chemometrics*, vol. 31, no. 12, 2017, e2937 cem.2937.
- [120] A. Zaikin and A. Zhabotinsky, "Concentration wave propagation in two-dimensional liquid-phase self-oscillating system," *Nature*, vol. 225, 1970.
- [121] K. Nordkvist, "Ocean color retrieval using DroneSpex - A miniature imaging spectrometer," M.S. thesis, Department of Space Science, MSc Programmes in Engineering, Space Engineering, Luleå University of Technology, 2007.
- [122] K. Esbensen, *Multivariate Data Analysis - in practice*, CAMO Process AS, Oslo, Norway, fifth edition, 2002.
- [123] J. Jaumot, R. Gargallo, A. de Juan, and R. Tauler, "A graphical user-friendly interface for MCR-ALS: a new tool for multivariate curve resolution in MATLAB," *Chemometrics and Intelligent Laboratory Systems*, vol. 76, 2005.
- [124] P. Comon, "Independent component analysis, a new concept?," *Signal Processing*, vol. 36, 1994.
- [125] F. Hitchcock, "The expression of a tensor or a polyadic as a sum of products," *Cambridge Journal of Mathematics and Physics*, vol. 6, 1927.
- [126] R. Bro, "PARAFAC. Tutorial and applications," *Chemometrics and Intelligent Laboratory Systems*, vol. 38, 1997.
- [127] S. Wold, "Pattern recognition by means of disjoint principal components models," *Pattern Recognition*, vol. 8, 1976.

- [128] S. Wold and M. Sjöström, *Chemometrics: Theory and Application*, vol. 52, chapter SIMCA: a method for analyzing chemical data in terms of similarity and analogy, American Chemical Society, Washington D.C., USA, first edition, 1977.
- [129] T. Kourti, “Application of latent variable methods to process control and multivariate statistical process control in industry,” *International Journal of Adaptive Control and Signal Processing*, vol. 19, 2005.
- [130] P. Williams and K. Norris, *Near-infrared technology in the agricultural and food industries*, St. Paul, Minn.: American Association of Cereal Chemists, 2004.
- [131] M. Manley, “Near-infrared spectroscopy and hyperspectral imaging: non-destructive analysis of biological materials,” *Chemical Society Reviews*, vol. 43, 2014.
- [132] H. Grahn and P. Geladi, *Techniques and Applications of Hyperspectral Image Analysis*, John Wiley & Sons, Ltd, 2007.
- [133] A. Kohler, J. Sulé-Suso, G. D. Sockalingum, M. Tobin, F. Bahrami, Y. Yang, J. Pijanka, P. Dumas, M. Cotte, D. G. van Pittius, G. Parkes, and H. Martens, “Estimating and correcting mie scattering in synchrotron-based microscopic fourier transform infrared spectra by extended multiplicative signal correction,” *Applied Spectroscopy*, vol. 62, 2008.
- [134] H. Martens and E. Stark, “Extended multiplicative signal correction and spectral interference subtraction: New preprocessing methods for near infrared spectroscopy,” *Journal of Pharmaceutical and Biomedical Analysis*, vol. 9, no. 8, 1991.
- [135] G. S. Birth, “Diffuse Thickness as a Measure of Light Scattering,” *Applied Spectroscopy*, vol. 36, no. 6, 1982.
- [136] Natural Phenomena Simulation Group – University of Waterloo, “Water and Ice Data,” <http://www.npsg.uwaterloo.ca/data/water.php>, Accessed February 2018.
- [137] M. Andersson, L. Persson, M. Sjöholm, and S. Svanberg, “Spectroscopic studies of wood-drying processes,” *Optics Express*, vol. 14, no. 8, 2006.
- [138] J. Bellingham and K. Rajan, “Robotics in Remote and Hostile Environments,” *Science*, vol. 318, no. 5853, 2007.

-
- [139] T. B. Curtin, J. G. Bellingham, J. Catipovic, and D. Webb, "Autonomous Oceanographic Sampling Networks," *Oceanography*, vol. 6, 1993.
- [140] M. Faria, J. Pinto, F. Py, J. Fortuna, H. Dias, R. Martins, F. Leira, T. A. Johansen, J. Sousa, and K. Rajan, "Coordinating UAVs and AUVs for Oceanographic Field Experiments: Challenges and Lessons Learned Experiments in UAV and AUV control for Coastal Oceanography," in *IEEE Int. Conf. Robotics and Automation, Hong Kong*, 2014.
- [141] D. A. Jessup, M. A. Miller, J. P. Ryan, H. M. Nevins, H. A. Kerkering, A. Mekebri, D. B. Crane, T. A. Johnson, and R. M. Kudela, "Mass Stranding of Marine Birds Caused by a Surfactant-Producing Red Tide," *PLOS ONE*, vol. 4, no. 2, 2009.
- [142] G. Johnsen, M. A. Moline, L. H. Pettersson, J. Pinckney, D. V. Pozdnayakov, E. S. Egeland, and O. M. Schofield, *Optical monitoring of phytoplankton bloom pigment signatures*, chapter 14, Cambridge University Press, 2011.
- [143] H. Dierssen, G. B. McManus, A. Chlus, D. Qiu, B.-C. Gao, and S. Lin, "Space station image captures a red tide ciliate bloom at high spectral and spatial resolution," *Proceedings of the National Academy of Sciences*, vol. 112, no. 48, 2015.
- [144] K. R. Turpie, S. Ackelson, T. Bell, H. Dierssen, J. Goodman, R. O. Green, L. Guild, E. Hochberg, V. V. Klemas, S. Lavender, C. M. Lee, P. Minnett, T. Moisan, F. Muller-Karger, J. Ortiz, S. Palacios, D. R. Thompson, and R. Zimmerman, "Global Observations of Coastal and Inland Aquatic Habitats," Tech. Rep., NASA, 2016.
- [145] A. G. C. Guerra, F. Francisco, J. Villate, F. Agelet, O. Bertolami, and K. Rajan, "On small satellites for oceanography: A survey," *Acta Astronautica*, vol. 127, 2016.
- [146] M. Guelman and F. Ortenberg, "Small satellite's role in future hyperspectral Earth observation missions," *Acta Astronautica*, vol. 64, no. 11, 2009.
- [147] T. D. Dickey and R. R. Bidigare, "Interdisciplinary oceanographic observations: the wave of the future," *Scientia Marina*, vol. 69, 2005.
- [148] M. Soukup, J. Gailas, D. Fantin, A. Jochemsen, C. Aas, P. J. Baeck, L. Benhadj, S. Livens, B. Delauré, M. Menenti, B. G. H. Gorte, S. E. Aria Hosseini, M. Esposito, and C. N. van Dijk, "HyperScout: Onboard Processing of Hyperspectral Imaging Data on a Nanosatellite," in *Small Satellites, System and Services Symposium (4S), Valletta, Malta*, 2016.

-
- [139] T. B. Curtin, J. G. Bellingham, J. Catipovic, and D. Webb, "Autonomous Oceanographic Sampling Networks," *Oceanography*, vol. 6, 1993.
- [140] M. Faria, J. Pinto, F. Py, J. Fortuna, H. Dias, R. Martins, F. Leira, T. A. Johansen, J. Sousa, and K. Rajan, "Coordinating UAVs and AUVs for Oceanographic Field Experiments: Challenges and Lessons Learned Experiments in UAV and AUV control for Coastal Oceanography," in *IEEE Int. Conf. Robotics and Automation, Hong Kong*, 2014.
- [141] D. A. Jessup, M. A. Miller, J. P. Ryan, H. M. Nevins, H. A. Kerkering, A. Mekebri, D. B. Crane, T. A. Johnson, and R. M. Kudela, "Mass Stranding of Marine Birds Caused by a Surfactant-Producing Red Tide," *PLOS ONE*, vol. 4, no. 2, 2009.
- [142] G. Johnsen, M. A. Moline, L. H. Pettersson, J. Pinckney, D. V. Pozdnayakov, E. S. Egeland, and O. M. Schofield, *Optical monitoring of phytoplankton bloom pigment signatures*, chapter 14, Cambridge University Press, 2011.
- [143] H. Dierssen, G. B. McManus, A. Chlus, D. Qiu, B.-C. Gao, and S. Lin, "Space station image captures a red tide ciliate bloom at high spectral and spatial resolution," *Proceedings of the National Academy of Sciences*, vol. 112, no. 48, 2015.
- [144] K. R. Turpie, S. Ackelson, T. Bell, H. Dierssen, J. Goodman, R. O. Green, L. Guild, E. Hochberg, V. V. Klemas, S. Lavender, C. M. Lee, P. Minnett, T. Moisan, F. Muller-Karger, J. Ortiz, S. Palacios, D. R. Thompson, and R. Zimmerman, "Global Observations of Coastal and Inland Aquatic Habitats," Tech. Rep., NASA, 2016.
- [145] A. G. C. Guerra, F. Francisco, J. Villate, F. Agelet, O. Bertolami, and K. Rajan, "On small satellites for oceanography: A survey," *Acta Astronautica*, vol. 127, 2016.
- [146] M. Guelman and F. Ortenberg, "Small satellite's role in future hyperspectral Earth observation missions," *Acta Astronautica*, vol. 64, no. 11, 2009.
- [147] T. D. Dickey and R. R. Bidigare, "Interdisciplinary oceanographic observations: the wave of the future," *Scientia Marina*, vol. 69, 2005.
- [148] M. Soukup, J. Gailas, D. Fantin, A. Jochemsen, C. Aas, P. J. Baeck, L. Benhadj, S. Livens, B. Delauré, M. Menenti, B. G. H. Gorte, S. E. Aria Hosseini, M. Esposito, and C. N. van Dijk, "HyperScout: Onboard Processing of Hyperspectral Imaging Data on a Nanosatellite," in *Small Satellites, System and Services Symposium (4S), Valletta, Malta*, 2016.