

SPRINGER BRIEFS IN COMPUTER SCIENCE

Harry Strange  
Reyer Zwiggelaar

# Open Problems in Spectral Dimensionality Reduction



Springer

# SpringerBriefs in Computer Science

## *Series editors*

Stan Zdonik  
Peng Ning  
Shashi Shekhar  
Jonathan Katz  
Xindong Wu  
Lakhmi C. Jain  
David Padua  
Xuemin Shen  
Borko Furht  
V. S. Subrahmanian  
Martial Hebert  
Katsushi Ikeuchi  
Bruno Siciliano

For further volumes:  
<http://www.springer.com/series/10028>

Harry Strange · Reyer Zwiggelaar

# Open Problems in Spectral Dimensionality Reduction

 Springer

Harry Strange  
Department of Computer Science  
Aberystwyth University  
Aberystwyth  
UK

Reyer Zwiggelaar  
Department of Computer Science  
Aberystwyth University  
Aberystwyth  
UK

ISSN 2191-5768

ISSN 2191-5776 (electronic)

ISBN 978-3-319-03942-8

ISBN 978-3-319-03943-5 (eBook)

DOI 10.1007/978-3-319-03943-5

Springer Cham Heidelberg New York Dordrecht London

Library of Congress Control Number: 2013956626

© The Author(s) 2014

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law. The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media ([www.springer.com](http://www.springer.com))

# Preface

The ability to simplify means to eliminate the unnecessary  
so that the necessary may speak.

—Hans Hofmann, 1880–1966

The last few years have seen a great increase in the amount of data available to scientists, engineers, and researchers from many disciplines. Datasets with millions of objects and hundreds, if not thousands, of measurements are now commonplace in areas such as image analysis, computational finance, bio-informatics, and astrophysics. This large volume of data does, however, come at a price, more often than not many computational techniques used to analyze these datasets cannot cope with such large data. Therefore, strategies need to be employed as a pre-processing step to reduce the number of objects, or measurements, whilst retaining important information inherent to the data. One of the key problems with such datasets is how to reduce the number of measurements, often referred to as dimensions, in such a way that the reduced set of measurements captures the main properties of the original data. Spectral dimensionality reduction is one such family of methods that has proven to be an indispensable tool in the data processing pipeline. In recent years, the area has gained much attention; thanks to the development of nonlinear spectral dimensionality reduction methods, often referred to as manifold learning algorithms.

Spectral dimensionality reduction methods can be broadly split into two categories; those that seek to maintain linear properties in the data, and those that seek to maintain nonlinear, manifold, properties. Both linear and nonlinear methods achieve the reduction in dimensionality through the careful construction of a feature matrix, the spectral decomposition of which gives rise to the reduced dimensionality dataset. Ever since the first nonlinear spectral dimensionality reduction methods were proposed over a decade ago, numerous algorithms and improvements have been proposed for the purpose of performing spectral dimensionality reduction. Although these algorithms may improve and extend existing techniques, there is still no gold standard technique. The reasons for this are many; however, one of the core problems with the area is that there are still many obstacles that need to be overcome before spectral dimensionality reduction that can be applied to a specific problem area. These obstacles, referred to herein

as *open problems*, have implications for those without a background in the area who wish to employ spectral dimensionality reduction to their problem domain.

Those wish to use spectral dimensionality reduction without prior knowledge of the field will immediately be confronted with questions that need answering; what parameter values to use? how many dimensions should the data be embedded into? how are new data points incorporated? what about large-scale data? For many, a search of the literature to find answers to these questions is impractical, as such, there is a need for a concise discussion into the problems themselves, how they affect spectral dimensionality reduction and how these problems can be overcome.

This book provides a survey and reference aimed at advanced undergraduate and postgraduate students as well as researchers, scientists, and engineers in a wide range of disciplines. Dimensionality reduction has proven useful in a wide range of problem domains, and so this book will be applicable to anyone with a solid grounding in statistics and computer science seeking to apply spectral dimensionality to their work.

## **Acknowledgments**

We both would like to thank the love and support of Jo, Diana, and our families. As well as to this, we would like to acknowledge both NISCHR and the Aberystwyth APRS funding schemes for partially supporting this work.

Aberystwyth, October 2013

Harry Strange  
Reyer Zwiggelaar

# Contents

<b>1</b>	<b>Introduction</b>	1
1.1	Goals of the Book	2
1.2	Relation to Existing Work	3
1.3	Outline	4
	References	4
<b>2</b>	<b>Spectral Dimensionality Reduction</b>	7
2.1	A General Setting for Spectral Dimensionality Reduction	7
2.2	Linear Spectral Dimensionality Reduction	9
2.2.1	Principal Components Analysis (PCA)	9
2.2.2	Classical Multidimensional Scaling (MDS)	10
2.3	Nonlinear Spectral Dimensionality Reduction	11
2.3.1	Isomap	12
2.3.2	Maximum Variance Unfolding (MVU)	13
2.3.3	Diffusion Maps	14
2.3.4	Locally Linear Embedding (LLE)	16
2.3.5	Laplacian Eigenmaps	17
2.3.6	Local Tangent Space Alignment (LTSA)	19
2.4	Kernel Formulation	20
2.5	Summary	21
	References	21
<b>3</b>	<b>Modelling the Manifold</b>	23
3.1	Overview of Neighbourhood Graph Construction	24
3.2	Building Neighbourhood Graphs	25
3.2.1	Optimised Neighbourhood Methods	26
3.2.2	Adaptive Estimation Methods	28
3.3	Topological and Multi-manifold Considerations	30
3.3.1	Manifolds with Loops	31
3.3.2	Multiple Manifolds	32
3.4	Noise and Outliers	34
3.4.1	Pre-processing Methods	35
3.4.2	Noise Handling Extensions	36

3.5	Summary . . . . .	38
	References . . . . .	38
<b>4</b>	<b>Intrinsic Dimensionality . . . . .</b>	<b>41</b>
4.1	Background . . . . .	41
4.2	Estimating Dimensionality: The Spectral Gap . . . . .	42
4.3	Estimating Dimensionality: Other Methods . . . . .	44
4.3.1	Global Estimators . . . . .	44
4.3.2	Local Estimators . . . . .	47
4.4	Choosing Techniques and Limitations . . . . .	49
4.5	Summary . . . . .	50
	References . . . . .	51
<b>5</b>	<b>Incorporating New Points . . . . .</b>	<b>53</b>
5.1	Natural Incorporation . . . . .	53
5.2	Out-of-Sample Extensions . . . . .	54
5.2.1	The Nyström Method . . . . .	55
5.2.2	Generalised . . . . .	57
5.2.3	Common Problems . . . . .	59
5.3	Pre-image . . . . .	60
5.4	Incremental Learning . . . . .	61
5.4.1	Incremental PCA . . . . .	62
5.4.2	Incremental Isomap . . . . .	62
5.4.3	Incremental LLE . . . . .	64
5.4.4	Incremental Laplacian Eigenmaps . . . . .	64
5.4.5	A Unified Framework . . . . .	65
5.5	Summary . . . . .	66
	References . . . . .	67
<b>6</b>	<b>Large Scale Data . . . . .</b>	<b>69</b>
6.1	Computational Complexity and Bottlenecks . . . . .	69
6.1.1	Complexity of Spectral Dimensionality Reduction Algorithms . . . . .	70
6.1.2	Common Bottlenecks . . . . .	73
6.2	The Nyström Method . . . . .	73
6.2.1	Sampling Methods . . . . .	74
6.3	Other Approximation Methods . . . . .	75
6.3.1	Isomap Variants . . . . .	75
6.3.2	Laplacian Eigenmap Variants . . . . .	76
6.3.3	Maximum Variance Unfolding Variants . . . . .	77
6.3.4	Diffusion Maps Variants . . . . .	77
6.4	Using Brute Force: GPU and Parallel Implementations . . . . .	78
6.5	Summary . . . . .	79
	References . . . . .	79



- 7 Postscript** ..... 83
  - 7.1 Digging Deeper: How do You Measure Success? ..... 83
  - 7.2 Beyond Spectral Dimensionality Reduction ..... 85
    - 7.2.1 Manifold Learning and ANNs ..... 85
    - 7.2.2 Beyond Dimensionality Reduction ..... 85
  - 7.3 Implementations ..... 86
  - 7.4 Conclusions ..... 87
  - 7.5 Future Paths ..... 88
  - 7.6 Summary ..... 88
- References ..... 89
  
- Index** ..... 91

# Chapter 1

## Introduction

**Abstract** A brief introduction to dimensionality reduction and manifold learning is provided and supported by a visual example. The goals of the book and its place in the literature is given, while the chapter is concluded by an outline of the remainder of the book.

**Keywords** Manifold learning · Spectral dimensionality reduction · Medical image analysis

Many problems in the machine learning, computer vision, and pattern recognition domains are inherently high-dimensional, that is, the number of measurements taken per observation is considered ‘high’. There are both practical and theoretical reasons for wanting to reduce the number of measurements to a more manageable amount. One such motivating factor is that visualising information that contains more than three dimensions is a near impossible task. Although methods have been presented to aid in multi-variate visualisation, humans are still limited to visualising data in three dimensions or less. As such, reducing the number of measurements, or dimensionality, of a dataset is an important process when seeking to visually analyse and understand the data. From a theoretical perspective, high-dimensional spaces have a number of properties that can pose real difficulties. The well known “curse of dimensionality” [1] and “empty space problem” [2] are two examples of such difficulties; as the number of dimensions increases, the properties often associated with 2 or 3-dimensional Euclidean spaces disintegrate and are replaced by strange and complex phenomena [3].

Such practical and theoretical motivations drive the need for automatic methods that can reduce the dimensionality of a dataset in an intelligent way. Spectral dimensionality reduction is one such family of methods. Spectral dimensionality reduction seeks to transform the high-dimensional data into a lower dimensional space that retains certain properties of the subspace or sub manifold upon which the data lies. This transformation is achieved via the spectral decomposition of a square symmetric