Introduction to Graph Signal Processing

An intuitive and accessible text explaining the fundamentals and applications of graph signal processing. Requiring only an elementary understanding of linear algebra, it covers both basic and advanced topics, including node domain processing, graph signal frequency, sampling, and graph signal representations, as well as how to choose a graph. Understand the basic insights behind key concepts and learn how graphs can be associated with a range of specific applications across physical, biological, and social networks, distributed sensor networks, image and video processing, and machine learning. With numerous exercises and Matlab examples to help the reader put knowledge into practice, and a solutions manual available online for instructors, this unique text is essential reading for graduate and senior undergraduate students taking courses on graph signal processing, signal processing, information processing, and data analysis, as well as researchers and industry professionals.

Antonio Ortega is a professor of electrical and computer engineering at the University of Southern California, and a fellow of the IEEE.

Introduction to Graph Signal Processing

ANTONIO ORTEGA
University of Southern California

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre,
New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of
education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781108428132
DOI: 10.1017/9781108552349

First published 2022

A catalogue record for this publication is available from the British Library.

Library of Congress Cataloging-in-Publication Data
Names: Ortega, Antonio, 1965– author.
Title: Introduction to graph signal processing / Antonio Ortega, University of Southern California.
Description: New York, NY : Cambridge University Press, 2021. |
Includes bibliographical references and index.
Identifiers: LCCN 2021038900 | ISBN 9781108428132 (hardback)
Subjects: LCSH: Signal processing. | BISAC: TECHNOLOGY & ENGINEERING /
Signals & Signal Processing | TECHNOLOGY & ENGINEERING / Signals & Signal Processing
Classification: LCC TK5102.9 .O77 2021 | DDC 621.382/2–dc23
LC record available at https://lccn.loc.gov/2021038900

ISBN 978-1-108-42813-2 Hardback

To Mineyo and Naoto

Contents

Preface

This book grew out of my own research with students and collaborators, and evolved into a set of notes developed for a special topics course on graph signal processing at the University of Southern California. I have taught several versions of this material as a standalone class as well as in combination with advanced topics in signal processing (wavelets and extensions, dictionary representations and compressed sensing). In the meantime, much has changed since the first time I offered this class in Fall 2013. Graph signal processing (GSP) has grown as a field, understanding of specific problems has improved, and more applications have been considered. Both my experience in teaching this material and the evolution of the field informed the choices I made in writing this book.

First, while this is a book about graph signal processing, my goal has been to make the book accessible to readers who do not have any signal processing background. The main assumption I make is that readers will have taken an elementary linear algebra course. Appendix A provides a review of elementary concepts, presented from a signal representation perspective, and can be used to review this material.

Second, from my experience teaching this material to masters and PhD students, and from feedback I received from undergraduates who read early drafts of the book, I think the main challenge for those studying this topic for the first time is to understand the basic insights that make it possible to *use* the concepts. Without those insights, even if the mathematical ideas are understood, it is difficult to apply them. For this reason, I have spent time trying to develop intuition about the key concepts, lingering on ideas that are likely to be obvious to active researchers but hoping that they will prove to be useful to those new to the area. In short, whether or not I have succeeded in these goals, the word "Introduction" certainly deserves its place on the title.

Finally, by design, this book does not aim to provide a comprehensive and detailed survey of all recent research. In a rapidly evolving field this is difficult to do: much has been published but competing methods have not been compared, connections have not been made, the dust hasn't quite settled yet. I have chosen to summarize, classify and make connections where possible, but many details and approaches are left out. For advanced topics covered in this book (sampling, graph learning, signal representations) there are recent overview papers that provide a more detailed literature survey. Similarly, GSP is being used in a growing number of applications, and only a small subset of those are introduced in this book.

How to use this book There are of course many different ways to use the material in this book for teaching. I describe some possible scenarios.

- For a semester-long advanced undergraduate class, Chapter 1 can be followed by the first three sections of Appendix A, and then the rest of the chapters. The remaining sections in Appendix A can be introduced before tackling Chapter 5. Additional programming assignments can be included to help students with limited `Matlab` experience.
- I have used this material as part of a semester-long graduate level class covering other topics (e.g., filterbanks, various types of wavelets and compressed sensing) and colleagues at other universities have used it in a similar way. For this type of course, the material in Appendix A can be used as an introduction to both GSP and to other advanced signal processing topics.
- Finally, a semester-long GSP class for advanced graduate students can complement the material in this book with reading from recent literature and advanced research oriented projects.

Exercises Since Chapters 1–3 and the two appendices cover basic concepts they contain exercises, while for Chapters 4–7, which deal with more advanced topics, students could be asked to read some of the published literature and work on a class project.

`Matlab` examples There are sections in Appendix B corresponding to Chapters 1–6, so this appendix can be used to complement each of those chapters, allowing students to get a more hands-on experience through `Matlab` code examples. Alternatively, Appendix B can be used as a standalone introduction to `GraSP`.

The `GraSP` toolbox is freely available at `https://www.grasp-toolbox.org/`. The source code for all examples and supplementary materials is available on the book's web page (`http://www.graph-signal-processing-book.org/`).

Acknowledgments

Interest in graph signal processing (GSP) has grown out of research on multiple applications where observed signals can be associated with an underlying graph. The journey that led me to completing this book started, more than 10 years ago, with the study of the representation and compression of signals captured by sensor networks. I often get asked about good applications for GSP, as if the results in this field had emerged out of purely theoretical research. In fact, much of what is described here, and in particular most of my own work on this topic, has been motivated by practical applications of GSP. While applications are important as a motivation, the reader should not be looking here for detailed solutions to problems that arise in specific applications. Instead, the main goal of this book is to develop the mathematical tools and insights that can allow us to think about some of these problems in terms of the processing of signals on graphs.

This book would not have been possible without the help and support of many people. First and foremost I would like to thank all of my current and former PhD students and postdoctoral fellows at the University of Southern California (USC). A brief note in these acknowledgments can hardly do justice to the importance of their contributions. Among my former PhD students I would like to thank in particular Alexandre Ciancio and Godwin Shen, who focused on transforms for sensor networks and developed methods for transforms over trees; Sunil Narang, who extended these ideas to graphs and introduced critically sampled graph filterbanks; Wooshik Kim, Yung-Hsuan (Jessie) Chao, Hilmi Egilmez and Eduardo Pavez, who developed new graph constructions and studied a number of image and video applications; Jiun-Yu (Joanne) Kao and Amin Rezapour, who took GSP methods into new and interesting application domains; and Akshay Gadde and Aamir Anis, who developed new methods for graph signal sampling and its application to machine learning. Ongoing and recent work with some of my current students, including Pratyusha Das, Keng-Shih Lu, Ajinkya Jayawant and Sarath Shekkizhar, as well as other work with Shay Deutsch, Alexander Serrano, Yoon Hak Kim, Lingyan Sheng, Sungwon Lee and Yongzhe Wang, also contributed to my research in this area. Finally, I would like to acknowledge the contributions of undergraduate and graduate students from various universities who visited USC and collaborated with my group, and in particular Eduardo Martínez-Enríquez from Universidad Carlos III, Madrid, and a series of students from Universitat Politècnica de Catalunya, Barcelona, and in particular Xavier Perez-Trufero, Apostol Gjika, Eduard Sanou, Javier Maroto, Victor González and Jùlia Barrufet, as well as David Bonet, who also provided comments on the manuscript.

Special thanks go to Benjamin Girault. Benjamin started the development of the GraSP Matlab toolbox during his PhD studies and has continued to grow, develop, and maintain it since joining USC as a postdoc. Without Benjamin's contribution, including Appendix B, a comprehensive introduction to GraSP, and numerous code examples, this book would be a lot less useful to students and practitioners. Benjamin's contributions go far beyond this one chapter. He has been a close collaborator, his ideas have fundamentally shaped several parts of this book and he has also read multiple chapter drafts, some of them more than once, and sometimes early versions that were far from being ready. His comments on various versions of the manuscript were thoughtful and detailed, and always prompted me to go deeper into the material.

Throughout my time at USC I have been fortunate to work with many outstanding colleagues in the school of engineering and beyond. Among those who have collaborated more directly in research related to this book, I would like to mention Salman Avistimehr, Bhaskar Krishnamachari, Urbashi (Ubli) Mitra, Shrikanth (Shri) Narayanan and Cyrus Shahabi, as well as some of their students and postdocs, among others Sundeep Pattem, Mahesh Sathiamoorthy, Eyal En Gad, Aly El Gamal, Basak Guler, Nicolò Michelusi and Marco Levorato.

One of the great pleasures of academic research is having the opportunity to collaborate with researchers around the world, both remotely and through some very enjoyable visits. Some of these collaborations have had a major role in shaping my research in this area. During several visits to École polytechnique fédérale de Lausanne (EPFL), discussions with Pascal Frossard, Pierre Vandergheynst and David Shuman led to an overview paper and to additional collaboration with Thomas Maugey, Xiaowen Dong and Dorina Thanou, among others. It was during an extended sabbatical visit to National Institute of Informatics (NII) in Tokyo, Japan, that I started drafting this book, and I would like to thank my long-time collaborator Gene Cheung (then at NII, now at York University), and his students Wei Xu, Jiahao Pang and Jin Zeng for many productive discussions and a very fruitful collaboration. A short visit to the University of New South Wales (UNSW) in Sydney, Australia, and many discussions with David Taubman are also acknowledged. Since 2015 I have also visited regularly the Tokyo University of Agriculture and Technology (TUAT), and various parts of the manuscript were developed there. I am very grateful to TUAT and my hosts Yuichi Tanaka and Toshihisa Tanaka and their students, in particular Akie Sakiyama and Koki Yamada, for hosting me and for their collaboration. While most of this collaboration has been remote, I would also like to thank Vincent Gripon, and his students Carlos Lassance and Myriam Bontonou, at Institut Mines-Télécom (IMT) Atlantique and David Tay at Deakin University for their collaboration.

My work in this area has benefited greatly from collaboration and funding from several companies, and in particular I would like to thank Anthony Vetro, Dong Tian and Hassan Mansour at Mitsubishi Electric Research Labs (MERL), Phil Chou, Debargha Mukherjee, Yue Chen at Google and Amir Said at Qualcomm.

My research in this area has been supported by multiple sources over the years and this funding is most gratefully acknowledged. In particular, I would like to acknowledge multiple grants from the National Science Foundation and funding from the National

Aeronautics and Space Administration (NASA), from the Defense Advanced Research Project Agency (DARPA) and several companies, including Samsung, LG Electronics, Google, Mitsubishi Electric, Qualcomm, Tencent and KDDI. My summer and sabbatical visits were funded by the Japan Society for the Promotion of Science, NII, UNSW, TUAT and EPFL.

A first version of this manuscript was based on class notes for a special topics course on GSP at USC, and portions of successive drafts were used in teaching a class that included GSP along with wavelets. I am thankful to students in those classes for their questions and comments. Several students who took my undergraduate linear algebra class gamely volunteered to test my hypothesis that this material could be made accessible to undergraduate students with only basic linear algebra. Time will tell whether this is possible, or a good idea, but I am grateful to them, in particular Alex Vilesov, Reshma Kopparapu, Pengfei Chang, Lorand Cheng and Keshav Sriram, for their comments and questions and Catherine (Cami) Amein for her comments and for her wonderful cover illustration.

Several people helped me with detailed comments on various parts of the manuscript. In addition to Benjamin Girault, I would like to thank Yuichi Tanaka, David Tay, Hilmi Egilmez and Eduardo Pavez. Baltasar Beferull-Lozano and Gene Cheung both used the material in their teaching and provided comments. I am also thankful to the team at Cambridge University Press, including Julie Lancashire, Julia Ford, and Sarah Strange, for their inexhaustible patience and their support throughout the years it took for this book to be completed, Susan Parkinson for her detailed reading and many helpful suggestions and Sam Fearnley for managing the final production.

As of the time of this writing, 2020 and 2021 have been often described as "interesting" and "unusual" years. While not much else has gone according to plan in the last year, I can at least look at this book as a small (if late) positive outcome. I thank my wife Mineyo and our son Naoto for supporting me as I completed this project over the last few years, for pretending to believe my repeated (and highly unreliable) claims that the book was almost done and for making this year of work from home interesting and unusual, but in a very good way.

Notation

Vectors are written in boldface lowercase, $\mathbf{x}$, while matrices are in capital boldface, $\mathbf{A}, \mathbf{B}$. For graph matrices we use calligraphic letters if the matrix is normalized, so for example we would write $\mathcal{L}$ instead of $\mathbf{L}$. A summary of other specific notation is given below.

$S_1 \cap S_2$	Intersection of sets S_1 and S_2
$S_1 \cup S_2$	Union of sets S_1 and S_2
$\lvert S \rvert$	Number of elements in set S
$\mathbf{I}$	Identity matrix
$\mathbf{J}$	Exchange matrix – all ones on the anti-diagonal
$\mathbf{1}$	Vector with all entries equal to 1
$\mathbf{A}$	Adjacency matrix
$\mathbf{B}$	Incidence matrix
$\mathbf{E}$	Matrix of self-loop weights
$\mathbf{L}$	Combinatorial graph Laplacian
$\mathcal{A}$	Symmetric normalized adjacency matrix
$\mathcal{L}$	Symmetric normalized Laplacian
$\mathcal{Q}$	Row normalized adjacency matrix
$\mathcal{P}$	Column normalized adjacency matrix
$\mathcal{T}$	Random walk Laplacian
$\mathbf{S}$	Sample covariance matrix
$\mathbf{Q}$	Precision matrix
$i \sim j$	Nodes i and j are connected
$\mathcal{I}(A)$	Indicator function (1 if A is true, 0 otherwise)
$\mathbf{Z}$	Generic one-hop fundamental graph operator
$\Delta_{\mathbf{S}}(\mathbf{x}) = \mathbf{x}^{\mathsf{T}}\mathbf{S}\mathbf{x}$	Variation operator
$\mathcal{N}(i)$	Set of nodes connected to node i
$\mathcal{N}_k(i)$	Set of nodes in the k-hop neighborhood of node i
$p_{\min}(\mathbf{Z})$	Minimal polynomial of $\mathbf{Z}$
$p_{\mathrm{c}}(\mathbf{Z})$	Characteristic polynomial of $\mathbf{Z}$
$p_x(\mathbf{Z})$	Minimal polynomial of $\mathbf{Z}$ for vector $\mathbf{x}$