

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 24-01-2017	2. REPORT TYPE Final Report	3. DATES COVERED (From - To) 28-Nov-2012 - 27-Nov-2016
---	--------------------------------	---

4. TITLE AND SUBTITLE Final Report: Structured and Collaborative Geometric Signal Models for Big Data Analysis: Theory and Applications in Image, Video and Audio RA 3.1.1	5a. CONTRACT NUMBER W911NF-13-1-0011
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER 611102

6. AUTHORS Guillermo Sapiro	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Duke University C/O Office of Research Support 2200 W. Main St., Ste. 710 Durham, NC 27705 -4677	8. PERFORMING ORGANIZATION REPORT NUMBER
---	--

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211	10. SPONSOR/MONITOR'S ACRONYM(S) ARO
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) 63134-MA.53

12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited
--

13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT Efficient data modeling is critical for ill-posed problems in signal reconstruction, classification, and identification. This project develops a fundamental framework for structured and collaborative design of such signal models. The underlying foundations are derived via a combination of theories from sparse modeling, Gaussian Mixture Models, and Principal Component Analysis. Efficient computational approaches are an intrinsic part
--

15. SUBJECT TERMS Signal models, information, fundamental theory, image and video processing, optimization, learning, technology transfer
--

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Guillermo Sapiro
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 919-660-5252

Report Title

Final Report: Structured and Collaborative Geometric Signal Models for Big Data Analysis: Theory and Applications in Image, Video and Audio RA 3.1.1

ABSTRACT

Efficient data modeling is critical for ill-posed problems in signal reconstruction, classification, and identification. This project develops a fundamental framework for structured and collaborative design of such signal models. The underlying foundations are derived via a combination of theories from sparse modeling, Gaussian Mixture Models, and Principal Component Analysis. Efficient computational approaches are an intrinsic part of the project as well. In addition to theoretical and computational questions, particular applications here addressed include anomaly detection in videos with dynamic background, collaborative object detection and classification, source identification and separation, activity clustering in video, and dimensionality reduction onto physical meaningful spaces. The combination of multimodalities and incorporation of prior knowledge and side information will be considered as well. Scenarios with significantly undersampled or missing data are addressed as well as part of the proposed frameworks. The theoretical foundations include the development of a novel approach of collaborative compressed sensing and the incorporation of information theory tools into the sparse modeling world. Finally, in order to address critical big data challenges, on-line and real time optimization techniques will be developed. The project is carried out in close collaboration with Department of Defense researches, including frequent visits and technology transfers.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
01/24/2017	49 J. Sokolic, R. Giryes, M. Rodriguez, and G. Sapiro. Robust Large Margin Deep Neural Networks, IEEE TSP, (): . doi:
01/24/2017	36 E. Elhamifar, G. Sapiro, and S. Sastry. Dissimilarity-based sparse subset selection, IEEE Trans. Pattern Analysis Machine Intelligence, (01 2016): 0. doi:
07/14/2016	45 Cecilia Aguerrebere, Mauricio Delbracio, Alberto Bartesaghi and Guillermo Sapiro. Fundamental Limits in Multi-image Alignment, IEEE Trans. Signal Processing, (): . doi:
TOTAL:	3

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

We presented at numerous conferences and universities, and always the work presented reflects publications (published, submitted, in progress).

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
01/24/2017	46 J. Lezama, Q. Qiu, G. Sapiro. Not Afraid of the Dark: NIR-VIS Face Recognition via Cross-spectral Hallucination and Low-rank Embedding, CVPR 2017. 24-JUN-17, Hawai. : ,
01/24/2017	48 M. Tepper and G. Sapiro. Nonnegative Matrix Underapproximation for Robust Multiple Model Fitting, CVPR 2017. 24-JUN-17, Hawai. : ,
01/24/2017	47 Qixiang Ye, Tianliang Zhang, Qiang Qiu, Baochang Zhang, Jie Chen, and Guillermo Sapiro. Self-learning Scene-specific Pedestrian Detectors using a Progressive Latent Model, CVPR 2017. 24-JUN-17, Hawai. : ,
04/28/2016	38 M. Tepper, G. Sapiro. A SHORT-GRAPH FOURIER TRANSFORM VIA PERSONALIZED PAGERANK VECTORS, IEEE ICASSP. 25-MAR-06, . : ,
07/11/2014	18 A. Bronstein, G. Sapiro, P. Sprechmann. SUPERVISED NON-EUCLIDEAN SPARSE NMF VIA BILEVEL OPTIMIZATION WITHAPPLICATIONS TO SPEECH ENHANCEMENT, 2014 4th Joint Workshop on Hands-free Speech Communication and Microphone Arrays (HSCMA). 15-MAY-15, . : ,
09/05/2013	11 Mariano Tepper, Guillermo Sapiro. FAST L1 SMOOTHING SPLINES WITH AN APPLICATION TO KINECT DEPTH DATA, IEEE Image Processing (ICIP). 15-SEP-13, . : ,
TOTAL:	6

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
04/28/2016 39.00	Jiayi Huang , Qiang Qiu , Guillermo Sapiro, Robert Calderbank. Discriminative Robust Transformation Learning, NIPS 2015. 09-DEC-15, . . . ,
04/28/2016 44.00	Mariano Tepper, Alasdair Newson, Pablo Sprechmann, Guillermo Sapiro. Multi-temporal foreground detection in videos, 2015 IEEE International Conference on Image Processing (ICIP). 27-SEP-15, Quebec City, QC, Canada. : ,
04/28/2016 43.00	Qiang Qiu, Guillermo Sapiro, Jordan Hashemi. Cross-modality pose-invariant facial expression, 2015 IEEE International Conference on Image Processing (ICIP). 27-SEP-15, Quebec City, QC, Canada. : ,
04/28/2016 42.00	Mark Draelos, Qiang Qiu, Alex Bronstein, Guillermo Sapiro. Intel realsense = Real low cost gaze, 2015 IEEE International Conference on Image Processing (ICIP). 27-SEP-15, Quebec City, QC, Canada. : ,
04/28/2016 41.00	A. Newson, M. Tepper, G. Sapiro. Low-Rank Spatio-Temporal Video Segmentation, BMVC 2015. 01-SEP-15, . . . ,
04/28/2016 40.00	Jiayi Huang , Qiang Qiu , Robert Calderbank , Guillermo Sapiro. Geometry-aware Deep Transform, ICCV 2015. 15-DEC-15, . . . ,
07/11/2014 14.00	J. Masci, M. Bronstein, A. Bronstein, P. Sprechmann, and G. Sapiro. Sparse similarity-preserving hashing, International Conference Learning Representations. 16-APR-14, . . . ,
07/11/2014 22.00	Kimberly Carpenter,, Pablo Sprechmann, Marcelo Fiori,, Robert Calderbank,, Helen Egger,, Guillermo Sapiro. QUESTIONNAIRE SIMPLIFICATION FOR FAST RISK ANALYSIS OF CHILDREN MENTALHEALTH, ICASSP 2014. 05-MAY-14, . . . ,
07/11/2014 21.00	E. Elhamifar, G. Sapiro, A. Yang, S. Sastry. A Convex Optimization Framework for Active Learning, CVPR 2014. 20-JUN-13, . . . ,
07/11/2014 20.00	M. Tepper, G. Sapiro. ALL FOR ONE, ONE FOR ALL: CONSENSUS COMMUNITY DETECTION IN NETWORKS, ICASSP 2014. 04-MAY-14, . . . ,
07/11/2014 19.00	M. Tepper, G. Sapiro. Ants Crawling to Discover the Community Structure in Networks, 18th Iberoamerican Congress on Pattern Recognition. 15-NOV-13, . . . ,
07/11/2014 17.00	Pablo Sprechmann, Roei Litman, Tal Ben Yakar, Alex Bronstein, Guillermo Sapiro. Efficient Supervised Sparse Analysis and Synthesis Operators, NIPS 2013. 05-DEC-13, . . . ,
07/11/2014 16.00	M. Fiori,, P. Sprechmann, , J. Vogelstein, , P. Muse, , and G. Sapiro. Robust Multimodal Graph Matching: Sparse Coding Meets Graph Matching, NIPS 2013. 05-DEC-13, . . . ,

- 07/11/2014 15.00 Xin Yuan, Patrick Llull, Xuejun Liao, Jianbo Yang, Guillermo Sapiro, David J. Brady and Lawrence Carin. Low-Cost Compressive Sensing for Color Video and Depth, IEEE Computer Vision Pattern Recognition (CVPR). 20-JUN-14, . . . ,
- 07/12/2013 6.00 Pablo Sprechmann, Alex Bronstein, Michael Bronstein, Guillermo Sapiro. LEARNABLE LOW RANK SPARSE MODELS FOR SPEECH DENOISING, IEEE ICASSP. . . . ,
- 07/12/2013 10.00 Amir Asiaee Taheri, Mariano Tepper, Arindam Banerjee, Guillermo Sapiro. If You are Happy and You Know It... Tweet, ACM Conference on Information and Knowledge Management. . . . ,
- 07/20/2015 24.00 Qiang Qiu, Andrew Thompson, Robert Calderbank, Guillermo Sapiro. REPRESENTATION USING THE WEYL TRANSFORM, International Conference Learning Representations. 07-MAY-15, . . . ,
- 07/20/2015 26.00 Qianf Qiu, Guillermo Sapiro. LEARNING COMPRESSED IMAGE CLASSIFICATION FEATURES, IEEE ICIP. 27-OCT-14, . . . ,
- 07/20/2015 25.00 Jiaji Huang^L, Qiang Qiu^L, Robert Calderbank^L, Miguel Rodriguesy, Guillermo Sapiro. ALIGNMENT WITH INTRA-CLASS STRUCTURE CAN IMPROVE CLASSIFICATION, IEEE ICASSP. 19-APR-15, . . . ,
- 09/05/2013 12.00 Zhongwei Tang, Mariano Tepper, Guillermo Sapiro. Reflective Symmetry Detection by RectifyingRandomized Correspondences, BMVC 2013 - British Machine Vision Conference. 09-SEP-13, . . . ,
- 09/09/2013 13.00 Tal Ben Yakar, Roei Litman, Pablo Sprechmann, Alex Bronstein, Guillermo Sapiro. BILEVEL SPARSE MODELS FOR POLYPHONIC MUSICTRANSSCRIPTION, International Society for Music Information Retrieval Conferenc2. 15-NOV-13, . . . ,

TOTAL: 21

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

- | <u>Received</u> | <u>Paper</u> | |
|-----------------|--------------|---|
| 07/12/2013 | 1.00 | Patrick Llull, Xuejun Liao, Xin Yuan, Jianbo Yang, David Kittle., Lawrence Carin, Guillermo Sapiro, and David J. Brady. Coded aperture compressive temporal imaging, Optics Express (01 2013) |

TOTAL: 1

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

~~Patent on Computer Vision for Behavioral Analysis in Children~~

Patent for Burst Deblurring (licensed to Adobe)

Patent for Fast Face Recognition (licensed to SmartFaceID)

Patents Awarded

N. Harel and G. Sapiro,

~~``Methods and systems for high-resolution brain image pipeline and database program," U.S. Patent 9,412,076, July 2016.~~

X. Bai, J. Wang, and G. Sapiro,

``Methods and apparatus for dynamic color flow modeling,"

U.S. Patent 8,682,063, March 2014.

Awards

[2012:] Best Poster Award, P. Sprechmann, A. Bronstein, and G. Sapiro, ``Real-time online singing voice separation from monaural recordings using robust low-rank modeling," International Society for Music Information Retrieval Conference, Porto, October 2012.

[2013] SIAM Fellow.

[2013] Best Paper Award, P. Lull, X. Liao, X. Yuan, J. Yang, D. Kittle, L. Carin, G. Sapiro, and D. J. Brady, ``Compressive sensing for video using a passive coding element," Imaging and Applied Optics Congress, Arlington, VA, June 2013.

[2013] International Society in Magnetic Resonance in Medicine Recognition, Top 5 cited articles in Magnetic Resonance in Medicine '10, I. Aganj, C. Lenglet, G. Sapiro, E. Yacoub, K. Ugurbil, and N. Harel, ``Reconstruction of the orientation distribution function in single and multiple shell q-ball imaging within constant solid angle" Magnetic Resonance in Medicine 64:2, pp. 554-566, 2010.

[2013] IEEE Fellow.

[2013] Science Advisory Board, Institute for Computational and Experimental Research in Mathematics (ICERM), Brown University.

[2014] Plenary Speaker, European Signal Processing Conference (EUSIPCO), September 2014.

[2014-2017] Member, National Academies' Board on Mathematical Sciences and their Applications (BMSA).

[2015] Plenary Speaker, CIARP, December 2015.

[2015-2016] Distinguished Israel Pollack Lecturer, Technion, Haifa.

STUDENTS AWARDS

[2016] Junior Scientist Award, SIAM Image Science, Mauricio Delbracio

[2016:] Winner, Latin American Doctoral Thesis Contest - CLEI 2016, Marcelo Fiori.

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Jordan Hashemi	0.50	
Zhuoqing Chang	0.50	
Pablo Sprechmann	0.50	
Kim Jinyoung	0.50	
FTE Equivalent:	2.00	
Total Number:	4	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Mauricio Delbracio	0.50
Cecilia Aguerreberer	0.50
Qiang Qiu	1.00
Jose Lezama	0.00
Mariano Tepper	0.50
FTE Equivalent:	2.50
Total Number:	5

Names of Faculty Supported

NAME	PERCENT SUPPORTED	National Academy Member
Guillermo Sapiro	0.10	
Qiang Qiu	0.50	
FTE Equivalent:	0.60	
Total Number:	2	

Names of Under Graduate students supported

NAME	PERCENT SUPPORTED	Discipline
Liv Kelley	0.00	ECE
Owen Smith	0.00	ECE
Yuansong Feng	0.00	ECE
Yitaek Hwang	0.00	ECE
Sophia Park	0.00	ECE
FTE Equivalent:	0.00	
Total Number:	5	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 4.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 4.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 2.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 4.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 1.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

NAME	
Zhuoqing Chang	
Jordan Hashemi	
Total Number:	2

Names of personnel receiving PHDs

NAME	
F. Lecumberry	
p. Sprechmann	
L. Yatziv	
A. Castrodad	
M. Fiori	
J. Kim	
Total Number:	6

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Steven Espinosa	0.10
FTE Equivalent:	0.10
Total Number:	1

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

During the full years of this award we have achieved all our initial goals and beyond. This is fully described in the papers we have submitted, and a few examples include:

- * Theoretical foundations of learning;
- * Connections between geometry, metric learning, and DNN;
- * State-of-the-art computational algorithms for numerous optimization problems, from non-negative matrix factorization to dictionary learning and sparse modeling;
- * State-of-the-art face recognition (technology transfer);
- * State-of-the-art video deblurring (technology transfer);
- * Technology transfers to DoD, industry, NIH;
- * Development of apps for mental and developmental health;
- * Development of Coursera on-line class with over 100,000 students;
- * Leading algorithms for network inference and matching of multimodal data.

These are just samples of the contributions, each one resulting in numerous papers, technology transfers, and students training.

Technology Transfer

Video and image deblurring (Adobe)

Fast Face Recognition (SmartFaceID)

Subspace learning (ARL)

Face recognition (IR) (ARL)

Video background analysis (NGA)

Fast optimization (NGA, internally awarded best technology transfer of the year)

Subspace methods (NGA)

Sample of Last Year Technical Accomplishments

Guillermo Sapiro

Duke University

January 2017

We describe now a few examples of works from the last year of this award; additional details are available in the body of the report and in the published literature.

1 Self-Learning via Progressive Latent Models

In this work, a self-learning approach is proposed towards solving scene-specific pedestrian detection problem without any human annotation involved. The self-learning approach is deployed as progressive steps of object discovery, object enforcement, and label propagation. In the learning procedure, object locations in each frame are treated as latent variables that are solved with a progressive latent model (PLM). Compared with conventional latent models, the proposed PLM incorporates a spatial regularization term to reduce ambiguities in object proposals and to enforce object localization, and also a graph-based label propagation to discover harder instances in adjacent frames. With the difference of convex (DC) objective functions, PLM can be efficiently optimized with a concaveconvex programming and thus guaranteeing the stability of self-learning. Extensive experiments demonstrate that even without annotation the proposed self-learning approach outperforms weakly supervised learning approaches, while achieving comparable performance with transfer learning and fully supervised approaches.

2 IR Face Recognition

Surveillance cameras today often capture NIR (near infrared) images in low-light environments. However, most face datasets accessible for training and verification are only collected in the VIS (visible light) spectrum. It remains a challenging problem to match NIR to VIS face images due to the different light spectrum. Recently, breakthroughs have been made for VIS face recognition by applying deep learning on a huge amount of labeled VIS face samples. The same deep learning approach cannot be simply applied to NIR face recognition for two main reasons: First, much limited NIR face images

are available for training compared to the VIS spectrum. Second, face galleries to be matched are mostly available only in the VIS spectrum. In this work, we propose an approach to extend the deep learning breakthrough for VIS face recognition to the NIR spectrum, without retraining the underlying deep models that see only VIS faces. Our approach consists of two core components, cross-spectral hallucination and low-rank embedding, to optimize respectively input and output of a VIS deep model for cross-spectral face recognition. Cross-spectral hallucination produces VIS faces from NIR images through a deep learning approach. Low-rank embedding restores a lowrank structure for faces deep features across both NIR and VIS spectrum. We observe that it is often equally effective to perform hallucination to input NIR images or low-rank embedding to output deep features for a VIS deep model for cross-spectral recognition. When hallucination and lowrank embedding are deployed together, we observe significant further improvement; we obtain state-of-the-art accuracy on the CASIA NIR-VIS v2.0 benchmark, without the need at all to re-train the recognition system.

The code has been shared with ARL.

3 Generalization Error of Neural Networks

The generalization error of deep neural networks via their classification margin is studied in this work. Our approach is based on the Jacobian matrix of a deep neural network and can be applied to networks with arbitrary nonlinearities and pooling layers, and to networks with different architectures such as feed forward networks and residual networks. Our analysis leads to the conclusion that a bounded spectral norm of the networks Jacobian matrix in the neighbourhood of the training samples is crucial for a deep neural network of arbitrary depth and width to generalize well. This is a significant improvement over the current bounds in the literature, which imply that the generalization error grows with either the width or the depth of the network. Moreover, it shows that the the recently proposed batch normalization and weight normalization re-parametrizations enjoy good generalization properties, and leads to a novel network regularizer based on the networks Jacobian matrix. The analysis is supported with experimental results on the MNIST and CIFAR-10 datasets.

The work is now being extended to invariant learning.

4 Foundations of Multi-Image Alignment

The performance of multi-image alignment, bringing different images into one coordinate system, is critical in many applications with varied signal-to-noise ratio (SNR) conditions. A great amount of effort is being invested into developing methods to solve this problem. Several important questions thus arise, including: Which are the fundamental limits in multi-image alignment performance? Does having access to more images improve the alignment? Theoretical bounds provide a fundamental benchmark to compare methods and can help establish whether improvements can be made. In this work, we tackle the problem of finding the performance limits in image registration when multiple shifted and noisy observations are available. We derive and analyze the Cramer-Rao and Ziv-Zakai lower bounds under different statistical models for the underlying image. We show the existence of different behavior zones depending on the difficulty level of the problem, given by the SNR conditions of the input images. The analysis we present here brings further insight into the fundamental limitations of the multi-image alignment problem.

5 A Universal Perspective on Deep Learning

Three important properties of a classification machinery are: (i) the system preserves the core information of the input data; (ii) the training examples convey information about unseen data; and (iii) the system is able to treat differently points from different classes. In this work we show that these fundamental properties are satisfied by the architecture of deep neural networks. We formally prove that these networks with random Gaussian weights perform a distance-preserving embedding of the data, with a special treatment for in-class and out-of-class data. Similar points at the input of the network are likely to have a similar output. The theoretical analysis of deep networks here presented exploits tools used in the compressed sensing and dictionary learning literature, thereby making a formal connection between these important topics. The derived results allow drawing conclusions on the metric learning properties of the network and their relation to its structure, as well as providing bounds on the required size of the training set such that the training examples would represent faithfully the unseen data. The results are validated with state-of-the-art trained networks.

6 Large Scale Nonnegative Matrix Factorization

Nonnegative matrix factorization (NMF) has an established reputation as a useful data analysis technique in numerous applications. However, its usage in practical situations is undergoing challenges in recent years. The fundamental factor to this is the increasingly growing size of the datasets available and needed in the information sciences. To address this, in this work we propose to use structured random compression, that is, random projections that exploit the data structure, for two NMF variants: classical and separable. In separable NMF (SNMF) the left factors are a subset of the columns of the input matrix. We present suitable formulations for each problem, dealing with different representative algorithms within each one. We show that the resulting compressed techniques are faster than their uncompressed variants, vastly reduce memory demands, and do not encompass any significant deterioration in performance. The proposed structured random projections for SNMF allow to deal with arbitrarily shaped large matrices, beyond the standard limit of tall-and-skinny matrices, granting access to very efficient computations in this general setting. We accompany the algorithmic presentation with theoretical foundations and numerous and diverse examples, showing the suitability of the proposed approaches. The software was released to the community.

7 Sparse Subset Selection

Finding an informative subset of a large collection of data points or models is at the center of many problems in computer vision, recommender systems, bio/health informatics as well as image and natural language processing. Given pairwise dissimilarities between the elements of a source set and a target set, we consider the problem of finding a subset of the source set, called representatives or exemplars, that can efficiently describe the target set. We formulate the problem as a row-sparsity regularized trace minimization problem. Since the proposed formulation is, in general, NP-hard, we consider a convex relaxation. The solution of our optimization finds representatives and the assignment of each element of the target set to each representative, hence, obtaining a clustering. We analyze the solution of our proposed optimization as a function of the regularization parameter. We show that when the two sets jointly partition into multiple groups, our algorithm finds representatives from all groups and reveals clustering of the

sets. In addition, we show that the proposed framework can effectively deal with outliers. Our algorithm works with arbitrary dissimilarities, which can be asymmetric or violate the triangle inequality. To efficiently implement our algorithm, we consider an Alternating Direction Method of Multipliers (ADMM) framework, which results in quadratic complexity in the problem size. We show that the ADMM implementation allows to parallelize the algorithm, hence further reducing the computational time. Finally, by experiments on real-world datasets, we show that our proposed algorithm improves the state of the art on the two problems of scene categorization using representative images and time-series modeling and segmentation using representative models.

8 Data Representation and the Weyl Transform

The Weyl transform is introduced as a rich framework for data representation. Transform coefficients are connected to the WalshHadamard transform of multiscale autocorrelations, and different forms of dyadic periodicity in a signal are shown to appear as different features in its Weyl coefficients. The Weyl transform has a high degree of symmetry with respect to a large group of multiscale transformations, which allows compact yet discriminative representations to be obtained by pooling coefficients. The effectiveness of the Weyl transform is demonstrated through the example of textured image classification.

9 On Graph Matching

Graph matchingaligning a pair of graphs to minimize their edge disagreements has received wide-spread attention from both theoretical and applied communities over the past several decades, including combinatorics, computer vision, and connectomics. Its attention can be partially attributed to its computational difficulty. Although many heuristics have previously been proposed in the literature to approximately solve graph matching, very few have any theoretical support for their performance. A common technique is to relax the discrete problem to a continuous problem, therefore enabling practitioners to bring gradient-descent-type algorithms to bear. We prove that an indefinite relaxation (when solved exactly) almost always discovers the optimal permutation, while a common convex relaxation almost always fails to discover the optimal permutation. These theoretical results suggest

that initializing the indefinite algorithm with the convex optimum might yield improved practical performance. Indeed, experimental results illuminate and corroborate these theoretical findings, demonstrating that excellent results are achieved in both benchmark and real data problems by amalgamating the two approaches.

10 Estimating Dynamic Graphs

The interest in problems related to graph inference has been increasing significantly during the last decade. However, the vast majority of the problems addressed are either static, or systems where changes in one node are immediately reflected in other nodes. In this work we address the problem of mobility graph estimation, when the available dataset has an asynchronous and time-variant nature. We present a formulation for this problem consisting on an optimization of a cost function having a fitting term to explain the observations with the dynamics of the system, and a sparsity promoting penalty term, in order to select the paths actually used. The formulation is tested on two publicly available real datasets on US aviation and NY taxi traffic, showing the importance of the problem and the applicability of the proposed framework.

11 Fourier on Graphs

The short-time Fourier transform (STFT) is widely used to analyze the spectra of temporal signals that vary through time. Signals defined over graphs, due to their intrinsic complexity, exhibit large variations in their patterns. In this work we propose a new formulation for an STFT for signals defined over graphs. This formulation draws on recent ideas from spectral graph theory, using personalized PageRank vectors as its fundamental building block. Furthermore, this work establishes and explores the connection between local spectral graph theory and localized spectral analysis of graph signals. We accompany the presentation with synthetic and real-world examples, showing the suitability of the proposed approach.

12 Synthesis Based Gaze Analysis

Gaze analysis has gained much popularity over the years due to its relevance in a wide array of applications, including human-computer interaction, fa-

tigue detection, and clinical mental health diagnosis. However, accurate gaze estimation from low resolution images outside of the lab (in the wild) still proves to be a challenging task. The new Intel low-cost RealSense 3D camera, capable of acquiring sub-millimeter resolution depth information, is currently available in laptops, and such technology is expected to become ubiquitous in other portable devices. In this work, we focus on low-cost, scalable and real time analysis of human gaze using this RealSense camera. We exploit the direct measurement of eye surface geometry captured by the RGB-D camera, and perform gaze estimation through novel synthesis-based training and testing. Furthermore, we synthesize different eye movement appearances using a linear approach. From each 3D eye training sample captured by the RealSense camera, we synthesize multiple novel 2D views by varying the view angle to simulate head motions expected at testing. We then learn from the synthesized 2D eye images a gaze regression model using regression forests. At testing, for each captured RGB-D eye image, we first repeat the same synthesis process. For each synthesized image, we estimate the gaze from our gaze regression model, and factor-out the associated camera/head motion. In this way, we obtain multiple gaze estimations for each RGB-D eye image, and the consensus is adopted. We show that this synthesis-based training and testing significantly improves the precision in gaze estimation, opening the door to true low-cost solutions.

13 Additional Dissemination

1. Coursera Image Processing Class
2. Autism&Beyond App
3. *Future Directions in Compressed Sensing and the Integration of Sensing and Processing- What Can and Should We Know by 2030?* Report on the Workshop funded by the Basic Research Office, Office of the Assistant Secretary of Defense for Research & Engineering
4. Invited talks