FROM BIOMEDICAL IMAGING TO URBAN DATA MINING: THEORY OF SIGNAL REPRESENTATIONS

Jelena Kovačević

Departments of ECE and BME, Carnegie Mellon University, Pittsburgh, PA, USA

ABSTRACT

This paper presents the author's personal path through the signal representations of the past three decades, from the early days and excitement that surrounded the advent of wavelets and associated multiresolution representations, to the present day foray into graph signal processing and data mining. It is a tribute to Dr. John Cozzens of the NSF and his vision and support for the development of the field.

Index Terms— Signal representations, multiresolution, graph signal processing

1. PATH THROUGH SIGNAL REPRESENTATIONS

Signal representations are at the heart of signal processing. They offer a key to understanding signals and their properties, ways to model them, look at them in different domain, process them and mine them. A priori knowledge or application-domain insight offer further opportunities to help design blocks that build signal representations.

This paper presents an entirely personal path through signal representations over the last three decades describing why and how of each with no attempt at completeness. It is meant to illustrate choices the author made and personal thoughts about the field.

The author's entry into the field coincided with the enormous excitement that the advent of wavelets generated [1, 2]. The idea of contrasting the global behavior description by the Fourier techniques to the localization trade off wavelets offered generated a whole slew of new ideas and gave us, in signal processing, a new way to think about signal representations. The notions of redundancy, localization, adapting to the signal at hand, and processing hierarchy, among others, all came into play. In the work on local bases, multiple descriptions, and frames, the above concepts played a role of design constraints in building signal representations with structure.

In another paper in this session [3], Rebecca Willett talks about a new class of signal representations that facilitate novel inference methods. These include compressed sensing and sparse coding among others and draw upon signal processing, machine learning and statistics to allow for more flexible and adapted representations of complex data. This is exactly what drew us to into our latest foray into signal representations — graph signal processing [4, 5], as a means of representing data with complex structure.

In breaking with tradition, the author would like to acknowledge Dr. Cozzens here as the bulk of the work was supported through his dedication and vision at NSF in one form or another. Rare is the one who dares to aim higher without looking for low-hanging fruit; John has been the one who was willing to promote a community of ideas and encourage research in this more theoretical area. While not always motivated by applications, signal representations described here and in [3], are at the heart of a number of today's practical systems, from compression standards to biomedical imaging algorithms. The author's thanks go to John for the support of the group's research as well as pedagogical efforts [6, 7, 8].

That work started in the late 80s with work on wavelets and multiresolution representations. Thinking about concepts of redundancy, localization, hierarchy of representation building blocks, among others, lead the author to expand into work on multiple descriptions, local bases, and frames. In the early 2000s, the author moved to Carnegie Mellon University and embarked on a decade of work on biomedical image representation and mining; one of the keystones of that work was a new multiresolution supervised classification framework, which employed both bases and frames. That framework was refined through a number of different applications on a number of different imaging modalities, from fluorescent microscopy to digital histopathology. In a conversation with civil engineering colleagues, a project was born to use the same framework for monitoring structural health of bridges by collecting vibration data from passing vehicles and classifying information about the bridge health. To solve the issue of not enough labeled data as labeling bridges as healthy or not is an expensive process requiring either human visual inspection or sensor installation, the author's group used semi-supervised learning framework and graphs [9] and was lead to the nascent topic of graph signal processing. The current focus of the group is on sampling, recovery, localization and representations of graph signals with applications to urban data mining; thus the title of the paper.

2. WAVELETS AND MULTIRESOLUTION REPRESENTATIONS

The author's path through signal representations started with her PhD work [10]. This was the time when the initial papers by Mallat [1] and Daubechies [11] came out and generated excitement due to their link to filter banks and subband coding systems. The idea that what was then known as octave-band filter banks when iterated to infinity would lead to continuous-time wavelet bases was illuminating. Similarly to the spectral understanding of the Fourier representations for different signal domains — functions or sequences, infinitely supported or finitely supported with circular extension (periodic), we started understanding various flavors of wavelet representations — functions or sequences, infinitely or finitely supported.

Initially nonredundant (bases), these ideas extended to redundant ones (frames). The notion of localization also played a role; while localized Fourier representations such as the short-time Fourier transform were known, they suffered from fixed time/frequency localization at all frequencies. Wavelet representations were offering a trade-off of fine time localization at high frequencies and global view at low frequencies. The next generalization was to allow for signal-adapted multiresolution representations — wavelet packets [12].

This more general class of multiresolution representations, wavelet packets, spurred work on arbitrary tilings of the time-frequency plane (leading to signal-adapted representations) and local orthogonal bases due to their use in audio and image coding. These are intimately related since local orthogonal bases can be used to efficiently construct flexible bases with arbitrary tilings of the time-frequency plane.

The author's work in that period followed a few threads: (1) Understanding how to build multidimensional nonseparable wavelet bases and associated filter banks [13, 14] due to great interest in image and video processing, HDTV representation and coding and others [14, 15]. (2) The second was motivated by the goal of analyzing the signal into unequal subbands (such as in acoustics); in that case, rational sampling factors have to be allowed in a filter bank [16]. (3) This thread related local orthogonal bases and their use in image coding [17, 18] as well as construction of arbitrary tilings of the time-frequency plane [19, 20].

3. MULTIPLE DESCRIPTION TRANSFORMS

This work was motivated by the fact that a large fraction of the information that flows across networks is useful even in a degraded condition; examples include speech, audio, still images and video. When this information is subject to packet losses or retransmission is impossible due to real-time constraints, superior performance with respect to total transmitted rate, distortion, and delay may sometimes be achieved by adding redundancy to the bit stream rather than repeating lost packets.

In multiple description coding [21, 22], the data is broken into several streams with some redundancy among the streams. When all the streams are received, one can guarantee low distortion at the expense of having a slightly higher bit rate than a system designed purely for compression. On the other hand, when only some of the streams are received, the quality of the reconstruction degrades gracefully, which is very unlikely to happen with a system designed purely for compression. This illustrates the concepts of redundancy and no hierarchy imposed on the different building blocks (unlike, for example, in the case of wavelets where the building blocks capturing global behavior of the signal are of higher importance).

The author's work on this topic followed a few threads: (1) Generalized multiple description coding with correlated transforms [23] provided a general framework for multiple description transform coding. (2) Quantized frame expansions with erasures provided redundancy in multiple description coding through redundant representations (frames) [24]. (3) Multiple description lattice vector quantization adds another technique for multiple description transform coding [25, 26]. (4) Multiple descriptions for audio and image coding [27, 28].

4. FRAMES

Frames are redundant representations that have become popular over the recent years. It is the idea of removing doubt translated from our daily lives into signal representations. Given a signal, we represent it in another system, typically a basis, where its characteristics are more readily apparent in the transform coefficients (for example, wavelet-based compression). However, these representations are typically nonredundant, and thus corruption or loss of transform coefficients can be fatal. In comes redundancy; we build a safety net into our representation so that we can avoid those fatal disasters. The redundant counterpart of a basis is called a frame.

It is generally acknowledged that frames were born in 1952 in the paper by Duffin and Schaeffer [29]. Despite being over half a century old, frames gained popularity only in the last two decades, due mostly to the work of the three wavelet pioneers — Daubechies, Grossman and Meyer [30]. Frame-like ideas, that is, building redundancy into a signal expansion, can be seen in pyramid coding, quantization, denoising, robust transmission, CDMA systems, multiantenna code design, segmentation, classification, prediction of epileptic seizures, restoration and enhancement, motion estimation, signal reconstruction, coding theory, operator theory and quantum theory and computing, among others.

While frames are often associated with wavelet frames, it is important to remember that frames are more general than that. Wavelet frames possess structure; frames are redundant representations that only need to represent signals in a given

space with a certain amount of redundancy.

The author's work on this topic followed a few threads: (1) Finite-dimensional frame families [31, 32, 33] and (2) frames for applications such as robust transmission, wireless biometrics and biomaging [34, 35, 36].

5. GRAPH SIGNAL PROCESSING

As mentioned before, the author's encounter with graph signal processing occurred because of the need for semi-supervised multiresolution classification applied to bridge health monitoring; the original supervised multiresolution classification framework was developed in the context of biomedical imaging. Among semi-supervised approaches, graph-based ones are often used, because they are able to represent a given dataset with complex graph structure and allow unlabeled signals to provide distribution information. We proposed an adaptive graph filter for semi-supervised classification that allows for classifying unlabeled as well as unseen signals and for correcting mislabeled signals. This adaptive graph filter extends the applications of the then emerging area of discrete signal processing on graphs to classification.

Signal processing on graphs is a theoretical framework inspired by algebraic signal processing that generalizes classical discrete signal processing from regular domains, such as lines and rectangular lattices, to arbitrary, irregular domains commonly represented by graphs [37, 5, 4]. This topic has garnered a fair amount of enthusiasm [38, 39, 40]; numerous special sessions at conferences and even specialized workshop such as the one at University of Pennsylvania in May of 2016.

Recent additions to the toolbox, including some by the author, consist of sampling of graph signals [41, 42, 43], recovery of graph signals [44, 9, 45], representations for graph signals [41, 46, 47], uncertainty principles on graphs [48, 49], graph-based transforms [50, 51, 52] and community detection and clustering on graphs [53, 40, 54], among others.

6. REPRODUCIBLE RESEARCH

Finally, in parallel to the work described above, an important initiative, supported by NSF, evolved in the signal processing community to ensure that our research is reproducible. Starting with an early paper by Barni and Perez-Gonzales [55], a topic was introduced that was percolating in other areas [56, 57]. Predecessor ideas seem to originate with Knuth [58]; Claerbout was one of the pioneers of the reproducible research movement [58], many others followed [59, 60, 61] including the author [62, 63].

Reproducible research refers to the idea that in computational sciences, the ultimate product is not a published paper but rather the entire environment used to produce the results in the paper (data, software, etc.). While it might sound natural and obvious, only in the last decade has the idea gained

prominence, culminating in the statements on encouraging reproducible research now included in the information for authors in SPS publications.

7. CONCLUSIONS

Dr. John Cozzens has been key in supporting the work on signal representations for the last three decades. This paper is meant as a personal account and a grateful acknowledgment by the author of the impact he has made on the signal processing community.

8. REFERENCES

[1] S. Mallat, "Multiresolution approximation and wavelets," Tech. Rep., Philadelphia, PA, Sept. 1987.

[2] I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, PA, 1992.

[3] R. Willett, "Signal representations in modern signal processing," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, New Orleans, LA, Mar. 2017.

[4] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, pp. 83–98, May 2013.

[5] A. Sandryhaila and J. M. F. Moura, "Big data processing with signal processing on graphs," *IEEE Signal Process. Mag.*, vol. 31, no. 5, pp. 80–90, Sept. 2014.

[6] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*, Prentice Hall, Englewood Cliffs, NJ, 1995.

[7] M. Vetterli, J. Kovačević, and V. K. Goyal, *Foundations of Signal Processing*, Cambridge University Press, Cambridge, 2014.

[8] J. Kovačević, V. K. Goyal, and M. Vetterli, *Fourier and Wavelet Signal Processing*, Cambridge University Press, Cambridge, forthcoming.

[9] S. Chen, F. Cerda, P. Rizzo, J. Bielak, J. H. Garrett, and J. Kovačević, "Semi-supervised multiresolution classification using adaptive graph filtering with application to indirect bridge structural health monitoring," *IEEE Trans. Signal Process.*, vol. 62, no. 11, pp. 2879–2893, June 2014.

[10] J. Kovačević, *Filter Banks and Wavelets: Extensions and Applications*, Ph.D. thesis, Columbia University, 1991.

[11] I. Daubechies, "Orthonormal bases of compactly supported wavelets," *Commun. Pure Appl. Math.*, vol. 41, pp. 909–996, Nov. 1988.

[12] R. R. Coifman, Y. Meyer, S. Quake, and M. V. Wickerhauser, "Signal processing and compression with wavelet packets," Tech. Rep., Yale University, 1991.

[13] J. Kovačević and M. Vetterli, "Nonseparable multidimensional perfect reconstruction filter banks and wavelet bases for $\mathcal{R}^n$," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 533–555, Mar. 1992.

[14] J. Kovačević and M. Vetterli, "The commutativity of up/downsampling in two dimensions," *IEEE Trans. Inf. Theory*, vol. 37, no. 3, pp. 695–698, May 1991.

[15] J. Kovačević and M. Vetterli, "FCO sampling of digital video using perfect reconstruction filter banks," *IEEE Trans. Image Process.*, vol. 2, no. 1, pp. 118–122, Jan. 1993.

[16] J. Kovačević and M. Vetterli, "Perfect reconstruction filter banks with rational sampling rate changes," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Toronto, May 1991, vol. III, pp. 1785–1788.

[17] J. Kovačević, D. J. LeGall, and M. Vetterli, "Image coding with windowed modulated filter banks," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Glasgow, May 1989, vol. III, pp. 1949–1952.

[18] R. Bernardini and J. Kovačević, "Local orthogonal bases i: Construction," *Multidim. Syst. Signal Process., sp. iss. Wavelets Multiresolution Signal Process.*, vol. 7, no. 4, pp. 331–370, July 1996.

[19] C. Herley, J. Kovačević, K. Ramchandran, and M. Vetterli, "Tilings of the time-frequency plane: Construction of arbitrary orthogonal bases and fast tiling algorithms," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3341–3359, Dec. 1993.

[20] R. Bernardini and J. Kovačević, "Arbitrary tilings of the time-frequency plane using local bases," *IEEE Trans. Signal Process.*, vol. 47, no. 8, pp. 2293–2304, Aug. 1999.

[21] M. T. Orchard, Y. Wang, V. Vaishampayan, and A. R. Reibman, "Redundancy rate-distortion analysis of multiple description coding using pairwise correlating transforms," in *Proc. IEEE Int. Conf. Image Process.*, Santa Barbara, CA, Oct. 1997.

[22] S. D. Servetto, V. A. Vaishampayan, and N. J. A. Sloane, "Multiple description lattice vector quantization," in *Proc. Data Compr. Conf.*, Snowbird, UT, Mar. 1999.

[23] V. K. Goyal and J. Kovačević, "Generalized multiple description coding with correlated transforms," *IEEE Trans. Inf. Theory*, vol. 47, no. 6, pp. 2199–2224, Sept. 2001.

[24] V. K. Goyal, J. Kovačević, and J. A. Kelner, "Quantized frame expansions with erasures," *Appl. Comput. Harmon. Anal.*, vol. 10, no. 3, pp. 203–233, May 2001.

[25] J. A. Kelner, V. K. Goyal, and J. Kovačević, "Multiple description lattice vector quantization: Variations and extensions," in *Proc. Data Compr. Conf.*, Snowbird, UT, Mar. 2000, pp. 480–489.

[26] V. K. Goyal, J. A. Kelner, and J. Kovačević, "Multiple description vector quantization with a coarse lattice," *IEEE Trans. Inf. Theory*, vol. 48, no. 3, pp. 781–788, Mar. 2002.

[27] V. K. Goyal, J. Kovačević, R. Arean, and M. Vetterli, "Multiple description transform coding of images," in *Proc. IEEE Int. Conf. Image Process.*, Chicago, IL, Oct. 1998, pp. 2349–2352.

[28] G. Schuller, J. Kovačević, F. Masson, and V. K. Goyal, "Robust low-delay audio coding using multiple descriptions," *IEEE Trans. Speech Audio Process.*, vol. 13, no. 5, pp. 1014–1024, Sept. 2005.

[29] R. J. Duffin and A. C. Schaeffer, "A class of nonharmonic Fourier series," *Trans. Am. Math. Soc.*, vol. 72, pp. 341–366, 1952.

[30] I. Daubechies, A. Grossman, and Y. Meyer, "Painless nonorthogonal expansions," *J. Math. Phys.*, vol. 27, pp. 1271–1283, Nov. 1986.

[31] J. Kovačević and A. Chebira, "An introduction to frames," *Found. Trends Signal Process.*, vol. 2, no. 1, pp. 1–94, 2008.

[32] J. Kovačević and A. Chebira, "Life beyond bases: The advent of frames (Part I)," *IEEE Signal Process. Mag.*, vol. 24, no. 4, pp. 86–104, July 2007.

[33] J. Kovačević and A. Chebira, "Life beyond bases: The advent of frames (Part II)," *IEEE Signal Process. Mag.*, vol. 24, no. 5, pp. 115–125, Sept. 2007.

[34] P. L. Dragotti, J. Kovačević, and V. K. Goyal, "Quantized oversampled filter banks with erasures," in *Proc. Data Compr. Conf.*, Snowbird, UT, Mar. 2001, pp. 173–182.

[35] P. H. Hennings-Yeomans, J. Thornton, J. Kovačević, and B. V. K. V. Kumar, "Wavelet packet correlation methods in biometrics," *Appl. Opt., sp. iss. Biometric Recogn. Syst.*, vol. 44, no. 5, pp. 637–646, Feb. 2005.

[36] A. Chebira, Y. Barbotin, C. Jackson, T. E. Merryman, G. Srinivasa, R. F. Murphy, and J. Kovačević, "A multiresolution approach to automated classification of protein subcellular location images," *BMC Bioinformatics*, vol. 8, no. 210, 2007.

[37] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.

[38] S. Segarra, A. G. Marques, G. Leus, and A. Ribeiro, "Reconstruction of graph signals through percolation from seeding nodes," *IEEE Trans. Signal Process.*, 2015. Submitted.

[39] A. G. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Sampling of graph signals with successive local aggregations," *IEEE Trans. Signal Process.*, vol. 64, no. 7, pp. 1832–1843, Apr. 2016.

[40] X. Dong, P. Frossard, P. Vandergheynst, and N. Nefedov, "Clustering on multi-layer graphs via subspace analysis on Grassmann manifolds," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 905–918, Feb. 2014.

[41] X. Zhu and M. Rabbat, "Approximating signals supported on graphs," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Kyoto, Mar. 2012, pp. 3921–3924.

[42] A. Anis, A. Gadde, and A. Ortega, "Efficient sampling set selection for bandlimited graph signals using graph spectral proxies," *IEEE Trans. Signal Process.*, vol. 64, no. 14, pp. 3775–3789, July 2015.

[43] S. Chen, R. Varma, A. Sandryhaila, and J. Kovačević, "Discrete signal processing on graphs: Sampling theory," *IEEE Trans. Signal Process.*, vol. 63, no. 24, pp. 6510–6523, Dec. 2015.

[44] S. Chen, A. Sandryhaila, J. M. F. Moura, and J. Kovačević, "Signal recovery on graphs: Variation minimization," *IEEE Trans. Signal Process.*, vol. 63, no. 17, pp. 4609–4624, Sept. 2015.

[45] S. K. Narang, Akshay Gadde, and Antonio Ortega, "Signal processing techniques for interpolation in graph structured data," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Vancouver, May 2013, pp. 5445–5449.

[46] D. Thanou, D. I. Shuman, and P. Frossard, "Learning parametric dictionaries for signals on graphs," *IEEE Trans. Signal Process.*, vol. 62, pp. 3849–3862, June 2014.

[47] S. Chen, T. Li, R. Varma, A. Singh, and J. Kovačević, "Signal representations on graphs: Tools and applications," *IEEE Trans. Signal Process.*, Mar. 2016. Submitted.

[48] A. Agaskar and Y. M. Lu, "A spectral graph uncertainty principle," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4338–4356, July 2013.

[49] M. Tsitsvero, S. Barbarossa, and P. D. Lorenzo, "Signals on graphs: Uncertainty principle and sampling," *arXiv:1507.08822*, 2015.

[50] S.l K. Narang, G. Shen, and A. Ortega, "Unidirectional graph-based wavelet transforms for efficient data gathering in sensor networks," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Dallas, TX, Mar. 2010, pp. 2902–2905.

[51] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Appl. Comput. Harmon. Anal.*, vol. 30, pp. 129–150, Mar. 2011.

[52] D. I. Shuman, B. Ricaud, and P. Vandergheynst, "Vertex-frequency analysis on graphs," *Appl. Comput. Harmon. Anal.*, February 2015.

[53] N. Tremblay and P. Borgnat, "Graph wavelets for multi-scale community mining," *IEEE Trans. Signal Process.*, vol. 62, pp. 5227–5239, Oct. 2014.

[54] P.-Y. Chen and A.O. Hero, "Local Fiedler vector centrality for detection of deep and overlapping communities in networks," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Florence, 2014, pp. 1120–1124.

[55] M. Barni and F. Perez-Gonzales, "Pushing science into signal processing," *IEEE Signal Process. Mag.*, July 2005.

[56] S. Fomel and G. Hennenfent, "Reproducible computational experiments using scons," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Honolulu, HI, Apr. 2007, vol. 4, pp. 1257–1260.

[57] R. J. LeVeque, "Wave propagation software, computational science, and reproducible research," in *Proc. Int. Congr. of Mathematicians*, 2006.

[58] D. E. Knuth, "Literate programming," *Comput. J.*, vol. 27, pp. 97–111, 1984.

[59] C. Pouzat, "Reproducible Neurophysiological Data Analysis Web Site," http://www.biomedicale.univ-paris5.fr/SpikeOMatic/Compendium.html.

[60] J. Buckheit and D. L. Donoho, *Wavelets and Statistics*, vol. 103, chapter Wavelab and reproducible research, pp. 55–81, Springer-Verlag, 1995.

[61] C. Greyer, "Sweave Demo Web Site," http://www.stat.umn.edu/~charlie/Sweave/.

[62] J. Kovačević, "How to encourage and publish reproducible research," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Honolulu, HI, Apr. 2007, vol. IV, pp. 1273–1276.

[63] P. Vandewalle, J. Kovačević, and M. Vetterli, "Reproducible research in signal processing: What, why and how," *IEEE Signal Process. Mag.*, vol. 26, no. 3, pp. 37–47, May 2009.