From *SIAM News*, Volume 43, Number 7, September 2010

Imaging Science

Learning Dictionaries and Sparse Image Representation

By Guillermo Sapiro

One of my first impressions of the 2010 SIAM Conference on Imaging Science, held in Chicago in April, was its high quality. Adel Faridani of Oregon State University and Stacey Levine of Duquesne University co-chaired the organizing committee that produced the outstanding program. The second thing that impressed me about the conference was its diversity, which is easily observed in the list of minisymposia as well as in the topics covered by the invited plenary speakers: learning theory (Amnon Shashua, and in part myself), geometry and dynamics (Alain Trouvé), optimization and numerical analysis (Gabriele Steidl), inverse problems, in particular as applied to seismic imaging (William Symes), and compressed sensing applied to astrophysics (Jean-Luc Starck). Both great mathematics and creative applications abounded in Chicago. Funding agencies and national labs were represented at the meeting, as speakers, minisymposium organizers, and attendees. As typical for SIAM meetings, diversity was in evidence at all levels.

I had the honor of giving one of the six plenary talks—on dictionary learning for sparse image representation. In the basic underlying model, natural images, and signals in general, have a sparse decomposition in some learned and adapted dictionary. This means that some combination of a small number of "atoms" from the (learned) dictionary leads to an efficient representation of the signal. Use of dictionaries adapted not only to the data but also to the task, rather than off-the-shelf dictionaries, significantly improves many image and video processing tasks, including image enhancement and classification.

Following a brief survey of examples from some of my work with Julien Mairal (whose public-domain code for sparse image representation is available at http://www.di.ens.fr/willow/SPAMS/), Miki Elad (who introduced me to the topic), Francis Bach, Jean Ponce, and Andrew Zisserman, my talk was centered on structured sparse models, in which the atoms used for signal representation cannot be freely selected from the dictionary, and the selection follows an (often pre-defined) structure. It turns out that adding structure to sparse models is critical if we are to further improve the state of the art, to speed up the computations, and to address novel problems.

As one example of structured sparse modeling, I described work with Guoshen Yu (University of Minnesota) and Stéphane Mallat (Ecole Polytechnique). In this work we learn a dictionary composed of blocks, each block being a PCA, or principal components analysis, of the overlapping image patches using the particular PCA, followed by linear model selection and coding procedures. In this way, we obtain computationally efficient block sparsity (one block, or PCA, per image patch is used at a time), modeling the image as a mixture of Gaussians and connecting structured sparse modeling with techniques from expectation maximization, collaborative filtering, manifold learning, and mixture models. Figure 1 shows an example of image interpolation (inpainting) done with this technique.* The dictionary, the collection of PCAs or the Gaussian mixture, is learned solely from the image to be processed. This is also the case in work of Larry Carin and collaborators at Duke University, who are developing exciting approach-

Figure 1. *Inpainting: Techniques from structured sparse modeling were used to recover 80% of the pixels missing from the original image (left).*

*Details can be found at http://www.cmap.polytechnique.fr/~yu/research/SSMS/demo.html, along with information on other image enhancement applications and connections to the literature (see also http://arxiv.org/abs/1006.3056 for the full report).

es to sparse modeling via non-parametric Bayesian techniques. Stan Osher and colleagues at UCLA are developing an interesting alternative approach for this application, based on concepts of non-local total variation (inspired in part by the pioneering work of Buades, Coll, and Morel on non-local means).

My second example of structured sparse modeling was drawn from work on collaborative hierarchical sparse coding with Pablo Sprechmann and Ignacio Ramirez of the University of Minnesota and Yonina Eldar of the Technion, Israel (currently visiting Stanford). In this work multiple signals are simultaneously coded to share blocks of atoms (the collaborative sparse block coding part), while allowed to have different sparsity patterns inside each block; the result is a hierarchical code natural for handling examples like those shown in Figure 2 (details can be found at http://arxiv.org/abs/1006.1346).

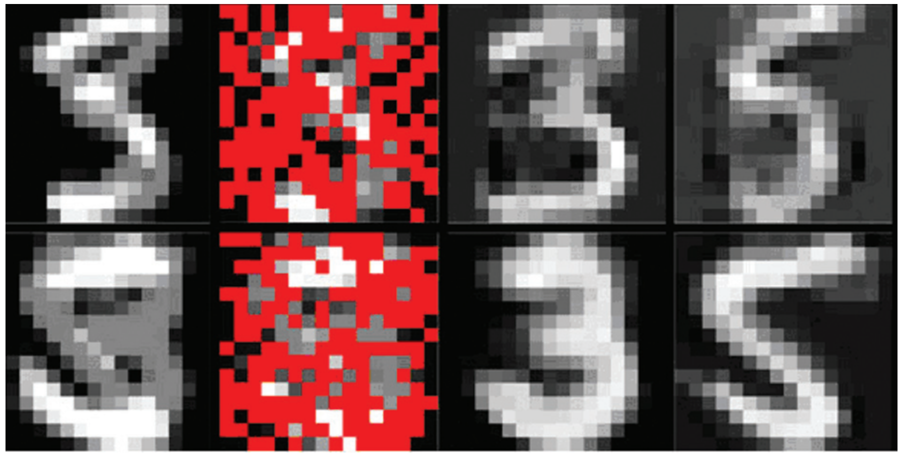

Figure 2. *Simultaneous restoration of the digits 3 and 5 from different combinations of the digits. Shown for the two examples, from left in each row, are the original mixture, the mixture with missing data (in red), and the recovered digits.*

In these examples, multiple combinations of two digits are simultaneously restored from data of which 60% is missing. Because all the images are combinations of the same two digits, their codes share the same blocks (those with dictionary elements appropriate to those digits). At the same time, because the particular images of the digits combined in a sample are not the same, each digit (image) is permitted to use different atoms from the blocks. The proposed algorithm simultaneously finds the correct blocks (class/object identification) and the atoms inside the blocks for the proper reconstruction. Related work (without the collaborative component) has been done by Jerome Friedman, Trevor Hastie, and Robert Tibshirani of Stanford University. Interested readers will also find theoretical and practical contributions to structured sparse modeling in the work of Rodolphe Jenatton, Jean-Yves Audibert, and Francis Bach.

These are just two examples of numerous efforts and new results on sparse modeling and dictionary learning presented at the SIAM conference—and a very small sample overall of the topics covered at the meeting. We live in a very exciting moment in image analysis, and SIAM is playing a leading role. We hope to see many readers in two years at the next meeting!

Acknowledgments

I thank my collaborators in the sparse modeling endeavors, in particular the graduate students and postdocs for teaching me so much and for making the work so much fun.

Guillermo Sapiro is a professor in the Department of Electrical and Computer Engineering at the University of Minnesota.