Research Summary: Namrata Vaswani (Iowa State University)

Namrata Vaswani is a Professor of Electrical and Computer Engineering, and (by courtesy) of Mathematics, at Iowa State University. She has been on the Iowa State faculty since Fall 2005. Her research interests lie at the intersection of statistical machine learning and data science, computer vision and medical imaging. She has graduated six Ph.D. and three M.S. students. She currently has a group of five Ph.D. students, one of whom will graduate in Spring 2018. Her work is or has been supported by seven NSF grants and two grants from industry. Her research has appeared in the premier journals in signal processing and information theory - IEEE Transactions on Signal Processing (TSP), Information Theory (T-IT), Image Processing (TIP), and on Pattern Analysis and Machine Intelligence (T-PAMI) – and, more recently, also in top machine learning conferences such as NIPS and AISTATS. Vaswani's work has been *cited over 3100 times*, she has an *h-index of 30*, and eight of her papers have at least a 100 citations. She has been recognized with the Harpole-Pentair Assistant Professorship and the Early Career Engineering Faculty Research Awards at Iowa State. In 2014, she received the *IEEE Signal Processing Society (SPS) Best Paper Award for a 2010 IEEE Transactions on Signal Processing paper (co-authored with my graduate student, Wei Lu) titled "Modified-CS: Modifying Compressive Sensing for Problems with Partially Known Support".* This award is given to only five to seven papers published in the last five years in any of the SPS Transactions or in JSTSP.

Vaswani is an *Area Editor for IEEE Signal Processing Magazine*; has served twice as an Associate Editor for TSP; and is the Lead Guest-Editor for a forthcoming Proceedings IEEE Special Issue on Rethinking PCA for Modern Datasets and a forthcoming Signal Processing Magazine feature cluster. She is also the current chair of the Women in Signal Processing Committee of SPS.

Vaswani's most interesting work consists of provably correct and practically useful online algorithms for the following two structured high-dimensional (big) data recovery problems – (a) dynamic compressive sensing (CS) and (b) dynamic robust principal component analysis (PCA). Online algorithms are needed for real-time applications, and even for offline applications, they are typically faster and need less storage compared to batch techniques. Most importantly, *her work shows that online algorithms provide a natural way to exploit temporal dependencies in a dataset without increasing algorithm complexity (speed or memory); and that exploiting such dynamics provably results in either reduced sample complexity or improved outlier tolerance or both.* The former implies proportionally reduced acquisition time for applications such as MRI where data is acquired one sample at a time. The latter implies increased robustness to difficult outliers such as persistent foreground occlusions in videos.

Dynamic Robust PCA (Robust Subspace Tracking) and PCA in Data-Dependent Noise

PCA is a widely used tool for dimension reduction. It is useful in a variety of applications ranging from exploratory data analysis to recommendation system design and video analytics. Given a matrix of data that is not too noisy, PCA is easily accomplished via singular value decomposition (SVD). While PCA is a relatively easy problem when the data is clean, it becomes much harder if data is corrupted by even a few outliers. This harder problem is called robust PCA. Outliers occur in practical applications for various reasons such as malicious or lazy users in recommendation system design, or foreground occlusions in video analytics. This important problem remained unsolved until the recent (2009) work of Candes, Wright, Li, and Ma posed it as a problem of decomposing a data matrix into the sum of a low rank matrix (clean data) and a sparse matrix (outliers). This new formulation led to a large volume of nice work on provably correct and practically useful robust PCA solutions.

However, all existing solutions make strong assumptions on either the number or the randomness of outliers. For practical applications such as video analytics, these translate to requiring that foreground occlusions (typically by moving objects) are either small sized and moving fast (outlier support changes a lot over time), or occur in randomly selected locations in each image. Both these are impractical requirements. Vaswani's work on this topic uses a novel insight to relax these requirements. She proposes to solve the dynamic robust PCA problem which can be understood as a time-varying extension of (static) robust PCA. It assumes that the true data lie in a slowly changing low-dimensional subspace, rather than a fixed one. In other contexts, this problem is also referred to as robust subspace tracking. *Vaswani's proposed solution, called* Recursive Projected Compressive Sensing (ReProCS), *is the first provable solution to this problem. It is also the first robust PCA solution that can provably tolerate a constant fraction of outlier per row without needing any assumptions on how the outlier support is distributed, as long as two simple extra assumptions hold – slow subspace change and a lower bound on most outlier magnitudes*. The former is a valid assumption for most static camera video backgrounds, while the latter essentially follows from the definition of an outlier as a large magnitude corruption. Under these two mild assumptions, ReProCS improves upon all previous work, all of which either needed outlier support be uniformly randomly distributed or needed a significantly tighter bound of order 1/r on this fraction. Here r is the dimension of the data subspace. For the practitioner, this means that foreground moving objects can be slow moving or occasionally static, and her solution would still work, while other methods will not. *Moreover, because it is an online approach, ReProCS is fast and has near-optimal memory complexity.* In today's big data age, memory complexity is the most important factor in determining the practical usability of an algorithm. Vaswani also demonstrates all these three advantages via application to real video analytics data, where she shows that, even in real video experiments involving videos with slow moving or large-sized foreground occlusions, ReProCS solution significantly outperforms existing algorithms, is at least two times faster, and needs significantly lesser memory.

The ReProCS solution framework is algorithmically very different from solutions for robust PCA that existed in earlier work. *This meant that proving the correctness of her approach was extremely challenging. New proof techniques were needed and there was no prior related work to guide this development.* At each time, ReProCS converts dynamic robust PCA into two simpler problems – projected CS and subspace update via PCA or incremental PCA. While PCA and incremental PCA are well studied problems, the specific problem encountered in case of ReProCS is that of PCA or incremental PCA in data-dependent noise. PCA has never been studied in this setting in any past work. Thus, in order, to obtain a provable guarantee for ReProCS, Vaswani first needed to analyze PCA in data-dependent noise. These results are themselves important because they provide the first finite sample (non-asymptotic) guarantees for PCA in non-isotropic and data-dependent noise. In many regimes of practical interest, she shows that this problem can be solved with near-optimal sample complexity.

Dynamic Compressive Sensing (CS) or Sparse Recovery

This work develops provably correct online algorithms for recovering a time sequence of approximately sparse signals from highly under-sampled measurements. In this problem, also referred to as dynamic CS, the signals are assumed to be sparse in some transform domain (sparsity basis or dictionary) and their sparsity patterns can change with time. Fast solutions to dynamic CS have important applications in dynamic imaging, e.g., in dynamic MRI for real-time medical applications such as interventional radiology or MRI-guided surgery.

To solve this problem, Vaswani *introduced the simple but very useful idea that, by exploiting slow support change, the problem can be converted into one of CS or sparse recovery with partial support knowledge*. This more general problem had a natural solution that she named Modified-CS. She has been able to both prove and experimentally demonstrate that Modified-CS achieves a significant sample complexity gain over static CS solutions because it exploits the dynamics inherently present in most practically occurring sparse signal sequences – slow support change, and sometimes also slow signal value change. *This result is significant because (i) it is the first exact recovery guarantee for the problem with partial support knowledge or, in fact, for CS with any kind of prior knowledge; (ii) it* *establishes that dynamic CS is a special case of this more general problem and hence also provides the first exact recovery guarantee for dynamic CS; and (iii) it opened up a new area and inspired much later work by other researchers*. It is rare that one needs to solve a sparse recovery problem where no prior knowledge about the signal support or values is available. This makes solutions for CS with partial support knowledge very valuable for practical purposes. For example, most frequency sparse signals will definitely contain a large number of low frequency components. Similarly, wavelet sparse images will have mostly nonzero entries in the low frequency subband, assuming most of the image is nonzero. Both these pieces of information can be converted into “partial support knowledge”.

Vaswani also studied Modified-CS for solving the dynamic CS problem in the presence of measurement noise*. This is the more realistic, but also the more challenging, setting* since an accurate recovery result for a single time instant does not guarantee that the error will not keep increasing over time. She has been able to prove stability under a simple and practical model on signal change over time, and weaker measurement model assumptions than existing solutions. A stability guarantee is critical for any online (recursive) algorithm since it ensures that the error does not blow up over time.

Vaswani’s original Modified-CS paper has been cited over 400 times and it received the *2014 IEEE Signal Processing Society Best Paper Award.* The entire body of work has been cited over 900 times.

Particle Filtering and Computer Vision

In even older work, Vaswani has made important contributions to efficient high dimensional particle filtering and computer vision (shape tracking and deformable contour tracking). Her algorithms relied on the simple but very useful idea that, in large class of high-dimensional tracking problems, most of the change occurs in only a few “key” dimensions. Her methods were developed in the context of deformable contour tracking from videos (needed for region-of-interest tracking in biological image sequences or for tracking moving objects in low-contrast videos), but are also applicable to various other complex tracking problems in computer vision, environment sensing and weather forecasting.

Future Plans: Dynamic Structured (Big) Data Recovery from Nonlinear Measurements – Low Rank Phase Retrieval

In both the problems that are described in detail above, the measurement model was linear. In recently started work, Vaswani is studying the low rank phase retrieval problem and its dynamic extensions. This refers to the problem of recovering a low rank matrix from magnitude-only measurements of linear projections of each column of the matrix. It finds applications in dynamic phaseless imaging problems such as dynamic Fourier ptychography or sub-diffraction imaging. She has developed a simple and fast iterative algorithm called LRPR to solve this problem. In preliminary dynamic Fourier ptychography experiments, this achieves a significant reduction in sample complexity (and hence data acquisition time or resources) compared with existing solutions which do not exploit the approximate low rank structure inherent in many slow changing image sequences. Her preliminary work on analyzing the method suggests that it should be also possible to mathematically prove this claim under simple assumptions. In ongoing work, she is also studying the dynamic version of this problem – how to exploit slow subspace change to either further reduce sample complexity or allow for robustness to outliers.

Low-rank phase retrieval is one instance of a dynamic structured high-dimensional data recovery problem when the measurements consist of element-wise nonlinearities. In future work, Vaswani plans to explore other instances of this problem, many of which occur in learning the weights for deep networks.