

Compressed Sensing Radar Amid Noise and Clutter

Peter B. Tuuk; S. Lawrence Marple, Jr.
Georgia Tech Research Institute
Sensors and Electromagnetic Applications Laboratory
Atlanta, GA 30332
peter.tuuk@gtri.gatech.edu

Abstract—Adaptive radar processing has been successful in downward looking radars that must detect moving targets in the midst of clutter returns. The performance of compressed sensing techniques in the presence of clutter is explored herein and compared to existing adaptive radar processing methods, including Space-Time Adaptive Processing (STAP), via Monte Carlo exploration of detection performance. Finally, we propose extensions to standard ℓ_1 optimization techniques to account for known interference covariance matrix statistics. These extensions out-perform current compressed sensing techniques, out-perform the fully-sampled, non-adaptive matched filter estimate, and approach the performance level of the fully-sampled STAP estimate.

I. INTRODUCTION

MAIN-BEAM clutter may be neglected in some radar applications. But in many airborne and surveillance applications strong returns from ground and other fixed objects obscure target energy and cannot be neglected. A more energetic waveform can bring targets out of noise interference but it illuminates target and clutter alike, doing nothing for signal-to-clutter ratio. These ground returns can be much stronger than those of the target making detection difficult. But the clutter exhibits structure that can be described by an interference covariance matrix representing the relationship between interference in the various measurement cells. The interference covariance matrix can be used to build a filter that minimizes clutter while preserving target energy. Little prior work exists on techniques for incorporating the covariance information into a compressed sensing problem solution. This work describes such a technique and shows that it improves the probability of detection of targets in the midst of strong clutter returns.

A. Space-Time Adaptive Processing

Space-Time Adaptive Processing (STAP) is an adaptive radar signal processing technique that has proven useful in the filtering of structured interference like clutter and jammers [1], [2]. It typically operates on a data cube that is sampled in fast-time (sample-to-sample) to calculate range, slow-time (pulse-to-pulse) to calculate Doppler frequency and thus radial velocity, and space (element-to-element) to calculate angle of arrival [3].

To analyze and operate on this data cube it is reshaped into a vector $\mathbf{y} \in \mathbb{C}^n$. We model these measurements as consisting of signal and clutter $\{\mathbf{x}, \mathbf{c}\} \in \mathbb{C}^n$ components that are mapped from the environment onto the measurements by a square matrix

$\mathbf{S} \in \mathbb{C}^{n \times n}$. The sampling rates in each dimension define bins in each of the sampled dimensions. The intersection of these bins forms a grid of voxels over the three-dimensional observation space. The vector \mathbf{x} is the indicator vector that contains the true reflectivity of targets at each grid location – meaning it contains mostly zeros at all those gridpoints that contain no target. And \mathbf{c} is a clutter vector containing the geometry- and terrain-dependent clutter reflectivity at the grid locations. We model these as a constant gamma-distributed random variable as in [4].

In the real world targets are not confined to a discrete set of locations nor is clutter fully resolved by the sampling rate in either range or cross-range. We use this simpler model for analysis and computation cognizant that it is limited in this way. When we stimulate and test this model we do so with data generated by higher resolution and higher fidelity models that include these off-grid targets, unresolved returns, and other realistic non-linearities.

Each column of \mathbf{S} is a steering vector to a bin in the range-angle-Doppler grid. Much of the work in STAP and other adaptive radar literature uses the term “steering vector” to mean a strictly angle-Doppler dictionary element. But radars are becoming more fully digitized and analog-to-digital converters continue to move closer to the individual antenna elements; pulse compression is increasingly performed digitally [5]. Thus we mean the steering vector to represent the all-dimension dictionary element. In short, the i -th column of \mathbf{S} is the return one expects to receive from a target in the i -th voxel.

White Gaussian noise $\mathbf{n} \in \mathbb{C}^n$ also enters the measurements. The measurement can be decomposed into signal and interference components:

$$\mathbf{y} = \mathbf{y}_s + \mathbf{y}_i = (\mathbf{S}\mathbf{x}) + (\mathbf{S}\mathbf{c} + \mathbf{n}) = \mathbf{S}(\mathbf{x} + \mathbf{c}) + \mathbf{n}. \quad (1)$$

STAP techniques estimate the interference covariance matrix $\mathbf{R} = \mathbb{E}[\mathbf{y}_i \mathbf{y}_i^H]$ from training data and use its inverse to generate a whitening filter: $\mathbf{W} = \mathbf{R}^{-1}\mathbf{S}$. Computing \mathbf{R}^{-1} directly through sample matrix inversion can be expensive. Other methods exploit the low-rank structure of the clutter to estimate \mathbf{R}^{-1} from an eigen-decomposition of the sampled data. In either case the constructed filter maximizes signal-to-interference-plus-noise ratio (SINR) when applied to the measured data:

$$\hat{\mathbf{x}}_{stap} = \mathbf{W}^H \mathbf{y} = \mathbf{S}^H \mathbf{R}^{-1} \mathbf{y}. \quad (2)$$

This is as compared to the adjoint, or matched filter, estimate which does not make use of the interference covariance




Fig. 1: The adjoint estimate $\hat{\mathbf{x}}_{adj}$ is shown reshaped into the range-angle-Doppler cube. These three views are summations of the estimated intensity along each of the three dimensions. The marker indicates the true target location in each view.




Fig. 2: The STAP estimate $\hat{\mathbf{x}}_{stap}$ is shown reshaped into the range-angle-Doppler cube. These three views are summations of the estimated intensity along each of the three dimensions. The marker indicates the true target location in each view.

information:

$$\hat{\mathbf{x}}_{adj} = \mathbf{S}^H \mathbf{y}. \quad (3)$$

Figures 1 and 2 illustrate the ability of the STAP filter to reduce the contribution of clutter in the estimate while maintaining target detectability in uncluttered regions. The line of clutter slices through the third projection in the figures due to the interaction of the sensor platform motion with the stationary ground clutter. In the adjoint estimate that clutter line appears as a strong ridge in the angle-Doppler plane making the target invisible. In the STAP estimate the clutter has been nulled by the adaptive filter and the target can be easily located. These techniques find common use in downward looking airborne radar when attempting to detect targets that exhibit radial motion.

B. Compressed Sensing

Compressed sensing is a recently developed theory by which some under-determined systems of linear equations can be shown to be solvable. If the problem to be solved obeys certain conditions an exact or approximate solution can be found with a rigorously bounded degree of certainty [6], [7].

The formulation of the radar sensing problem presented in (1) makes application of compressed sensing techniques a natural step. Compressed sensing operates under two assumptions [8]:

- 1) Sparsity: The signal to be reconstructed has few non-zero elements or can be represented in some basis by few non-zero elements
- 2) Incoherence: The measurement model is incoherent with the sparsifying basis

Sparsity is evident in the formulation of (1). The target indicator vector \mathbf{x} is a sparse vector as the number of true

targets in the observation space is small relative to the number of gridpoints that have been laid in that space. In many realistic applications this will hold true. For instance, in an airspace monitoring radar most points in the observable sky do not contain an airplane but those that do compose the support set of the non-zero elements of a sparse vector. Other applications like synthetic aperture radar (SAR) in which images are formed of an area of land, do not exhibit native sparsity and may require representation in some other sparsifying basis. Because of the large extent of the observed clutter, downward looking MTI radar of the sort we consider herein does not yield a native sparsity. However, by using appropriate filters the desirable scene can be separated from the undesirable clutter and the sparsity can be found in the scene subspace.

Incoherence must then be addressed. If \mathbf{y} is the set of Nyquist-sampled measurements then it represents all the (bandlimited) electromagnetic information passing over the measurement aperture during the period of observation (a single coherent processing interval). The compressed measurements we model as another vector

$$\mathbf{z} = \mathbf{C}\mathbf{y}$$

where $\mathbf{z} \in \mathbb{C}^m$ and $\mathbf{C} \in \mathbb{C}^{m \times n}$. As indicated above, these compressive measurements must be incoherent with the sparsifying basis \mathbf{S} . Theory shows that random measurement matrices are highly likely to be incoherent with any specifying basis.

The exact content of the measurement matrix \mathbf{C} will depend on the measurement process it describes. In pulse-Doppler radar this process may introduce incoherence in fast-time by mixing incoming signals with pseudo-random modulation sequences before sampling slowly [9], it may introduce incoherence in slow-time by staggering the pulse repetition interval, it may introduce incoherence in the spatial domain using an adaptive measurement array. These ideas will not be broached in this work. We simply construct \mathbf{C} as a random matrix filled with independent and identically distributed (i.i.d.) random variables that take on the values ± 1 with equal likelihood.

A simple estimate of the target vector from the compressed measurements \mathbf{z} can be computed by performing a compressed adjoint which applies the matched filter to that measured data.

$$\hat{\mathbf{x}}_{cadj} = \mathbf{S}^H \mathbf{C}^H \mathbf{z}.$$

This method has low computational cost but does not necessarily yield a sparse solution.

A better estimate of the target vector can be computed by solving a convex linear program such as an ℓ_1 -regularized least-squares

$$\hat{\mathbf{x}}_{cs} = \arg \min_{\mathbf{x}} \|\mathbf{z} - \mathbf{C}\mathbf{S}\mathbf{x}\|_2^2 + \tau \|\mathbf{x}\|_1 \quad (4)$$

where the first term of the minimization objective is the Euclidean norm of the residual that enforces fidelity to measured data and the second term is the one-norm of the estimate that promotes sparsity in the solution. The parameter τ weights these sometimes competing priorities.

The field of adaptive radar and specifically STAP has been well-researched over the last 20 years. Compressed sensing has, itself, seen much work in the past five years. There has been little treatment of the intersection of these two fields.

The concepts of reduced-rank and reduced-dimension STAP have something to offer to this topic [1]. By exploiting the low-rank nature of the interference these techniques reduce the computational and training burden that full-rank STAP techniques impose. Additional work from a sparse-estimation perspective has taken place recently. One paper that presents an ℓ_1 -regularized MTI is given by Yang, et. al., [10]. This technique assumes sparsity in the rank of the interference subspace relative to the number of system degrees of freedom. In [11] the authors operate in the angle-Doppler domain and solve a compressed sensing optimization problem only over the portion of the plane that is judged to be outside the clutter ridge. Our work differs from those by optimizing in range-angle-Doppler space and by using the covariance matrix to adaptively reduce clutter and other structured interference.

D. Proposed Extension

None of the mentioned compressed sensing solution methods, as described, are equipped to handle structured interference. If \mathbf{R} is the covariance matrix of the fully sampled interference then the covariance matrix of the interference in the compressed domain can be expressed as $\mathbf{R}_c = \mathbf{C}\mathbf{R}\mathbf{C}^H$. The simplest way to obtain an estimate of the scene from the compressed measurements that also includes the covariance information is

$$\hat{\mathbf{x}}_{cstap} = \mathbf{S}^H \mathbf{C}^H \mathbf{R}_c^{-1} \mathbf{z}.$$

This is analogous to the $\hat{\mathbf{x}}_{stap}$ solution defined in (2) in that the matched filter is post-multiplied by the covariance matrix inverse to whiten the interference. To compute this estimate one incurs the cost of inverting the covariance matrix (or parametrically estimating the inverse) but gains a good deal of clutter suppression. However, this technique does not favor sparse solutions.

In addition to this technique, we propose a covariance-aware CS (CA CS) that accounts for structured interference in a compressed sensing framework. As in STAP methods this is done via the interference covariance matrix inverse. The original ℓ_1 regularized least squares problem (4) can be modified as follows:

$$\hat{\mathbf{x}}_{cacs} = \arg \min_{\mathbf{x}} (\mathbf{z} - \mathbf{C}\mathbf{S}\mathbf{x})^H \mathbf{R}_c^{-1} (\mathbf{z} - \mathbf{C}\mathbf{S}\mathbf{x}) + \gamma \|\mathbf{x}\|_1. \quad (5)$$

A fuller treatment of this work can be found in [12].

The first term maintains data-fidelity in the clear entries while the second term promotes sparsity. The first term in the objective function penalizes deviations from the measurements in interference-free regions but less so in interfered regions. This term is akin to a Mahalanobis distance. The second term penalizes large entries in the solution. Thus the entries in the interfered region that are unconstrained by the first term are driven to zero by the second term. Here the γ term serves the same purpose as τ in (4): balancing the weight of the sparsity promoting ℓ_1 norm against the fidelity-preserving ℓ_2 norm.

A. ASPEN

To test against realistic data we simulate measurements in the Adaptive Sensor Prototyping ENvironment (ASPENTM) tool developed at Georgia Tech Research Institute's Sensors and Electromagnetic Applications Laboratory. ASPEN is a high fidelity clutter modeling simulation that supports flexible sensor definitions. We model a 32-element uniform linear array, 32-pulse coherent processing interval, and 128-sample range window, yielding a data cube of 131,000 samples. The outputs consist of data (\mathbf{y}) measured at the Nyquist sampling rate, true target locations (from which we derive \mathbf{x}), interference covariance (\mathbf{R}) as well as simulation parameters like the transmitted waveform and relevant geometry.

B. Linear Model

The ASPEN measurement model is non-linear, breaking the assumption in (1): targets can exist off grid locations and nearby targets interfere in non-linear ways. But to compute a CS estimate a linear model is required. Thus we developed a forward model, i.e. \mathbf{S} , that approximates the ASPEN simulation.

This model describes simple propagation phenomena and can be represented explicitly as a matrix. Each column of the matrix \mathbf{S} is the return from a target at the corresponding range-angle-Doppler position in space. The size of the matrix grows rapidly as the dimension of the sample space and search space increase. For a problem of size $32 \times 32 \times 128$ as mentioned above the matrix has 1.7×10^{10} elements. Therefore, the most feasible, efficient way to implement this model is not by explicitly storing the matrix but by performing discrete Fourier transforms along the angle and velocity dimensions and a convolution in the range dimension. For computational gains, these can be implemented using the Fast Fourier Transform (FFT). By this means we reduce required storage and processing time considerably.

As in any real estimation problem, the true measurement model differs some from the model used for processing. Though this linear model does not exactly match the more realistic simulation it provides theoretical and computational tractability while maintaining fidelity to the reference and serving the purpose of estimating target parameters.

C. Solution Methods

Various estimates of the target scene can then be computed from the measured data. We compute the matched filter (3), STAP estimate (2), and CS and CA CS estimates using the Templates for First-Order Conic Solvers (TFOCS) algorithm [13] and implementation [14]. This toolset implements a number of state-of-the-art solution techniques for large-scale convex problems and provides a flexible framework for the solution of alternative problem formulations.

A. Scoring Metrics

We are interested in the ability of a radar system to identify the locations of targets in a field of noise and structured

interference. Therefore we will present results in terms of probability of false alarm (P_{fa}) and probability of detection (P_d) rather than the more standard (in CS literature at least) Euclidean error norm. Also, we use another performance measure, detection quantile Q_d , to look at the consistency of performance.

Consider t targets placed arbitrarily in the observation area. The scene description vector $\mathbf{x} \in \mathbb{C}^n$ contains mostly zeros, but at those entries that correspond to the grid-points closest to the true target locations entries are non-zero values. Call these locations $\mathcal{P} = \{p_1 \dots p_t\}$. By some technique $\hat{\mathbf{x}}$ is produced which is an estimate of \mathbf{x} . For a given probability of false alarm P_{fa} a detection threshold D_{th} can be computed: let $\hat{\mathbf{x}}_s$ contain all the elements of $|\hat{\mathbf{x}}|$ sorted in increasing order, then $D_{th} = \hat{\mathbf{x}}_s(\lceil P_{fa}n \rceil)$. Finally P_d is the fraction of elements of the estimate at the true target locations that have magnitude greater than or equal to the threshold: $P_d = \text{frac}(|\hat{\mathbf{x}}(\mathcal{P})| \geq D_{th})$

Detection quantile is, roughly, the P_{fa} required to achieve a perfect P_d . Define $Q_d = \text{frac}(|\hat{\mathbf{x}}| > \min(|\hat{\mathbf{x}}(\mathcal{P})|))$. If $Q_d = 0$ it implies that the true target position was the highest-amplitude bin in the estimate.

B. Signal Quality Definitions

Detection results are only interesting when presented in the context of a known signal quality. Definitions of signal-to-noise ratio (SNR), signal-to-clutter ratio (SCR), and signal to interference ratio (SIR) are slippery and vary with usage. Furthermore, these values can be measured at either the input or the output of the signal processing chain. In this work we define these functions as described in this section. Again let \mathcal{P} refer to the t true target locations in \mathbf{x} and \mathcal{A} refer to all locations in \mathbf{x} . And consider measurements collected such that $\mathbf{y} = \mathbf{A}(\mathbf{x} + \mathbf{c}) + \mathbf{n}$ and $\hat{\mathbf{x}}$ produced by some technique to estimate \mathbf{x} .

Define input SNR as:

$$\frac{\max(|\mathbf{Ax}|)^2}{\text{var}(\mathbf{n})}$$

Define input SCR as:

$$\frac{\max(|\mathbf{Ax}|)^2}{\text{var}(\mathbf{Ac})}$$

Define input SIR as:

$$\frac{\max(|\mathbf{Ax}|)^2}{\text{var}(\mathbf{Ac} + \mathbf{n})}$$

On the input side, where we use the $\max(\cdot)$ function we do so to represent the modulus of the complex waveform that is received. If the waveform returns from more than one target overlap (before pulse compression) the modulus of the overlapping segment may constructively or destructively interfere based on the relative phases of the targets.

On the output side of the receiver, pulse compression or some other estimation technique has been performed and the power for each individual target can be measured independently.

In the event of multiple true targets, we take the arithmetic mean of the corresponding SIR values. Define output SIR as:

$$\frac{\text{mean}(\hat{\mathbf{x}}(\mathcal{P})^2)}{\text{var}(\hat{\mathbf{x}}(\mathcal{A} \setminus \mathcal{P}))}$$

C. CS and CA CS in Clutter

To test the effectiveness of the various techniques in the presence of clutter we develop curves demonstrating performance over a range of signal quality. We are primarily interested in the ability to detect targets in the midst of the clutter.

In particular, we vary the SCR while holding all other parameters constant. At each SCR point we calculate an average probability of detection P_d . The detection statistics are dependent on the particular realization of the random noise, clutter, and target locations so we average the results over many trials to develop performance curves for each technique over the SNR domain. This result is shown in Figure 3.

The sparsity favoring aspect of the CS solution is as likely to identify strong clutter returns as it is the true target returns. This is shown in Figure 3 by the green lines that underperform the matched filter (adjoint) solution.

By contrast the proposed covariance-aware compressed sensing (CA CS) described in (5) shows success beyond that of the compressed sensing estimates. The results of this technique are shown in Figure 3 by the blue lines. Of particular interest is the $20\times$ undersampled solution that achieves nearly the same detection performance as the STAP solution. By using the detection performance covariance information this technique surpasses the performance of the fully-sampled (but covariance-ignorant) matched filter. The fully sampled adaptive STAP filter provides the gold standard.

Figure 4 shows the performance as measured by the detection quantile. An interesting comparison can be drawn between Figures 3 and 4. When judging by P_d as in Figure 3 the $40\times$ undersampled CA CS solution outperforms the fully-sampled adjoint. However when judging by Q_d , as defined in Section III-A and shown in Figure 4 the fully-sampled adjoint shows superior performance to the $40\times$ undersampled CA CS estimate. This shows that the adjoint solution often assigns a high amplitude to the correct bin but doesn't as often put it among the *highest* bins to be counted as a detection. The $20\times$ undersampled outperforms both of these estimates as measured by both P_d and Q_d .

IV. CONCLUSION

A. Review

Structured interference, including clutter, presents a different challenge to estimation techniques than does white noise. This structure can be exploited to improve performance. While estimation of sparse vectors in the presence of white noise has been well-studied in compressed sensing literature, estimation in structured noise has not. We show that compressed sensing techniques can be robust to structured interference and propose a method for exploiting the interference structure to improve detection statistics.




Fig. 3: The CA CS method sub-samples the data just as the standard CS method does, however it takes into account the covariance matrix that describes the interference structure. By doing so it improves the probability of detection over the CS case as well as beyond the fully sampled, matched filter case that does not use the covariance information. These results are shown with a probability of false alarm of .01.

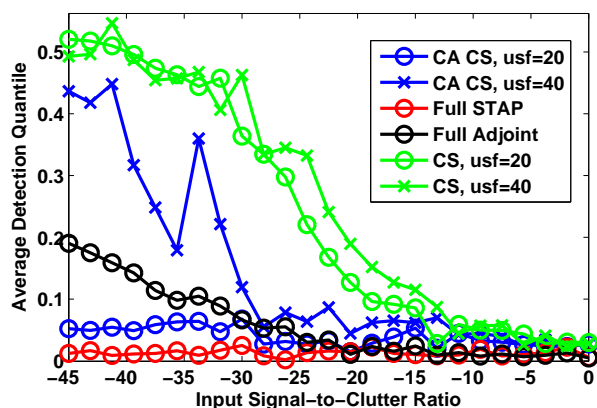


Fig. 4: A related evaluation criterion, the detection quantile Q_d measures the average ranking of the true-target bin among all the bins in the estimate. A Q_d of zero is perfect. These results show that again, the fully-sampled STAP estimate performs better than all other techniques. Also, the $20\times$ undersampled CA CS estimate is shown to consistently achieve performance near that of the fully-sampled STAP and better than the fully-sampled adjoint. In these results the $20\times$ sampled CA CS estimate is shown to be inferior to the fully-sampled adjoint, in contrast to the performance as measured by the P_d shown in Figure 3.

B. Future Work

A number of points remain to be investigated. Future work should proceed along many dimensions:

- The results presented here test the performance of the CA CS algorithm in the presence of clutter; by simulating jammer returns these techniques could be tested in the presence of this other common structured interference type.
- The CA CS technique must be refined with respect to selection of the balancing parameters and estimation of covariance matrix from compressed measurements.
- Theoretical bounds that describe the number of measurements required for a given sparsity level and signal quality would guide theorists and designers alike.
- To be convincing, these techniques must be paired with a more detailed description of the hardware required to collect the types of measurements required in a CS receiver as well as the size, weight, and power gains that could be made by designing to a CS specification.

There remains much progress to be made but the work presented here shows that CS radar systems can perform at a high level even in the presence of clutter.

REFERENCES

- [1] W. Melvin, "A STAP overview," *Aerospace and Electronic Systems Magazine, IEEE*, vol. 19, pp. 19–35, January 2004.
- [2] J. Guerci, *Space-Time Adaptive Processing for Radar*. Norwood, MA: Artech House, 2003.
- [3] M. A. Richards, *Fundamentals of Radar Signal Processing*. New York: McGraw-Hill, 2005.
- [4] D. Shnidman, "Radar detection in clutter," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 41, pp. 1056–1067, July 2005.
- [5] Z. Yingxi, H. Zhiming, and Z. Zhulin, "Design of the high-powered digital pulse compression real-time processing system based on ADSP-TS203," in *Radar, 2006. CIE '06. International Conference on*, pp. 1–4, October 2006.
- [6] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [7] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, pp. 1289–1306, April 2006.
- [8] E. Candès and J. Romberg, "Sparsity and incoherence in compressive sampling," *Inverse Problems*, vol. 23, pp. 969–985, April 2007.
- [9] J. Tropp, J. Laska, M. Duarte, J. Romberg, and R. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," *Information Theory, IEEE Transactions on*, vol. 56, pp. 520–544, January 2010.
- [10] Z. Yang, R. de Lamare, and X. Li, "L1-regularized STAP algorithms with a generalized sidelobe canceler architecture for airborne radar," *Signal Processing, IEEE Transactions on*, vol. 60, pp. 674–686, February 2012.
- [11] I. Selesnick, S. Pillai, K. Y. Li, and B. Himed, "Angle-doppler processing using sparse regularization," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pp. 2750–2753, March 2010.
- [12] P. B. Tuuk and S. L. Marple Jr., "Compressed sensing radar amid noise and clutter using interference covariance information," *submitted for publication*.
- [13] S. Becker, E. Candès, and M. Grant, "Templates for convex cone problems with applications to sparse signal recovery," *Mathematical Programming Computation*, vol. 3, pp. 165–218, 2011.
- [14] S. Becker, E. Candès, and M. Grant, "TFOCS templates for first-order conic solvers," April 2012. <http://tfocs.stanford.edu/>.