

ADAPTIVE RATE COMPRESSIVE SENSING FOR BACKGROUND SUBTRACTION

Garrett Warnell¹, Dikpal Reddy², and Rama Chellappa¹

¹ Department of Electrical and Computer Engineering; Center for Automation Research, UMIACS; University of Maryland, College Park, MD

² Department of Electrical Engineering and Computer Science; University of California, Berkeley, CA

{warnellg,rama}@umiacs.umd.edu dikpal@berkeley.edu

ABSTRACT

We study the problem of adaptive compressive sensing (CS) of a time-varying signal with slowly changing sparsity and rapidly varying support. We are specifically interested in visual surveillance applications such as background subtraction and tracking. Classical CS theory assumes prior knowledge of signal sparsity in order to determine the number of sensor measurements needed to ensure adequate signal reconstruction. However, when dealing with time-varying signals such as video, prior information regarding the exact sparsity may be difficult to obtain. Assuming a sensor that is able to take an adaptive number of compressive measurements, we present an algorithm based on cross validation that quantitatively evaluates the current measurement rate and adjusts it as needed.

Index Terms— Compressive Sensing, Opportunistic Sensing, Background Subtraction

1. INTRODUCTION

Visual surveillance is burdened with a large amount of data that is used very selectively. For example, a surveillance system placed in a remote area with the goal of detecting intruders will often observe an inactive scene, yet record or transmit the same amount of data during these periods as if the scene were active. This is wasteful, yet it is not immediately clear how to approach the problem since the periods of activity and inactivity are unknown in advance. If such information were available *a priori*, a good approach would be to collect data only during times of activity. Such a strategy falls into the category of *opportunistic sensing* (OS) [1]: a methodology that aims to dynamically adjust sensing system parameters to the state of the environment.

In addressing the problem above, it is apparent that the system must make a dynamic decision regarding scene activity. However, such a decision can be made only if real-time data to that same effect is available. This means that either the video sensor itself must remain active at all times, or some secondary modality (e.g., a motion detector) must collect the actionable information. Using an additional modality comes at the cost of adding more hardware to the system, and the data may be unreliable due to factors such as limited range compared to a visual sensor. Therefore, we consider the case where a single video sensor continuously collects data in order to perform both standard visual surveillance tasks and the determination of the state of the environment. Compressive sensing lends itself to such a problem since each of the collected measurements are influenced by the entire signal. Therefore, even a small number of measurements can provide the system with information regarding

the overall state of the signal. Thus, we will assume our video sensor to be a CS camera, e.g. the single-pixel camera [2].

We specifically address the problem of background subtraction, and accomplish this task using compressive measurements as has been previously studied [3]. However, the existing approach is unable to adapt the system data rate to scene activity. Figure 1 shows the drawbacks of such a scheme: too few measurements result in a poor signal estimate, yet too many is wasteful since the estimate is not better for them. In this paper, we propose a method that adaptively chooses the number of compressive measurements collected based on cross validation (CV) theory in CS [4, 5]. Our contribution (a) extends the use of CV techniques to the estimation of time-varying signals, and (b) provides a practical algorithm for adaptively changing the data rate of the system in response to scene activity.

1.1. Related Work

Recent literature has investigated allowing the number of compressive measurements to change during CS measurement and decoding. Malioutov et. al [6] propose a method for estimating the CS reconstruction error directly from the measurements. Using this estimate, one can decide if additional measurements are warranted for an improved estimate of the signal of interest. For time-varying signals with slowly changing support, Vaswani [7] addresses the problem of adapting measurement rate to the signal structure. Using an estimate of the signal support provided by standard CS techniques, a reduced number of measurements can be used to estimate the signal using a Kalman filter on that subset of coefficients. Neither of these methods accomplishes what our proposed algorithm aims to do: adaptively select the measurement rate for a time-varying signal with rapidly varying support, e.g., a foreground signal with target motion.

2. BACKGROUND

2.1. Compressive Sensing

For $\mathbf{x} \in \mathbb{R}^N$, CS collects $M \ll N$ linear measurements via a measurement matrix $\Phi \in \mathbb{C}^{M \times N}$ ($\mathbf{y} = \Phi\mathbf{x}$). If it is known that \mathbf{x} belongs to the class of sparse signals and Φ is constructed appropriately, it is possible to deduce the value of \mathbf{x} from \mathbf{y} [8]. One condition which must be satisfied in order for this to occur is that the value of M must be large enough with respect to the number of significant coefficients in \mathbf{x} . In this paper, we adaptively choose M for a time-varying \mathbf{x} in order to reduce the overall amount of data collected.

This work was partially supported by Army Research Office MURI Grant W911NF-09-1-0383.




Fig. 1. Single frame in the Convoy2 dataset. (a) and (b) are the background and foreground plus background, respectively. (c) is the true foreground, with sparsity $s = 4095$. (d) and (f) are the foreground estimates as a result of using too few and too many compressive measurements, respectively ($k = 200$ and $k = 9810$). (e) is the foreground estimate using an appropriate number of compressive measurements ($k = 4095$). Note that overestimating s leads to wasted measurements, as they provide no discernible improvement in reconstruction quality.

2.1.1. Compressive Sensing with Cross Validation

Let \mathbf{x} be reconstructed via $\hat{\mathbf{x}} = \Delta(\mathbf{y}, \Phi)$ for some CS decoder Δ . In this work, Δ will represent ℓ_1 -minimization, i.e., if $\mathbf{y} = \Phi\mathbf{x}$, then Δ provides a point estimate of \mathbf{x} as

$$\hat{\mathbf{x}} = \Delta(\mathbf{y}, \Phi) = \arg \min_{\mathbf{z}} \|\mathbf{z}\|_1 \text{ subject to } \Phi\mathbf{z} = \mathbf{y} \quad .$$

Ward [4] estimates the reconstruction error $\|\mathbf{x} - \hat{\mathbf{x}}\|$ using a technique that is operationally similar to CV. A second CS measurement matrix $\Psi \in \mathbb{C}^{r \times N}$, referred to as a cross validation matrix, is used in parallel with Φ to sense \mathbf{x} . These CV measurements are given by $\mathbf{y}_\Psi = \Psi\mathbf{x}$. For a given accuracy parameter ϵ , selecting Ψ with $r = O(\epsilon^{-2})$ suffices to ensure with high probability that

$$(1 - \epsilon)\|\mathbf{y}_\Psi - \Psi\hat{\mathbf{x}}\| \leq \|\mathbf{x} - \hat{\mathbf{x}}\| \leq (1 + \epsilon)\|\mathbf{y}_\Psi - \Psi\hat{\mathbf{x}}\| \quad ,$$

i.e., $\|\mathbf{y}_\Psi - \Psi\hat{\mathbf{x}}\|$ is used to bound the CS reconstruction error.

Boufounos et. al [5] present an earlier version of this idea as a method for increasing the speed of iterative CS decoding procedures. At each iteration, the reconstruction error is approximated as $\|\mathbf{y}_\Psi - \Psi\hat{\mathbf{x}}\|$, which is used as a stopping criterion for the decoder.

Our method also utilizes a CV procedure. At each time instant, we collect a fixed number of measurements and assume no further access to the signal. After standard CS decoding, we estimate the reconstruction error and use it to determine if more or fewer measurements are necessary for sensing subsequent frames.

2.2. Compressive Background Subtraction

The problem of using compressive measurements of images to perform background subtraction has been considered by Cevher et. al. [3]. A vectorized image \mathbf{x}_t is assumed to consist of both foreground, \mathbf{f}_t , and background, \mathbf{b}_t , i.e.,

$$\mathbf{x}_t = \mathbf{f}_t + \mathbf{b}_t \quad . \quad (1)$$

A background-adaptive method for estimating \mathbf{f} from compressive measurements of \mathbf{x} is proposed: $\hat{\mathbf{f}}_t = \Delta(\Phi(\mathbf{x}_t - \mathbf{b}_t), \Phi)$, where \mathbf{b}_t is known via an estimation and update procedure.

Our method differs in that we assume a static background (i.e., $\mathbf{b}_t = \mathbf{b}$ for all t), but allow for a varying number of compressive measurements to be collected at each t . The background constancy assumption is clearly restrictive, but we use it to demonstrate the viability of a variable measurement rate method while leaving the development of a background-adaptive method for future work. To

realize a varying measurement rate, we require a different measurement matrix at each time instant, i.e. $\Phi_t \in \mathbb{C}^{M_t \times N}$. To form Φ_t , we first construct $\Phi \in \mathbb{C}^{N \times N}$ via standard CS measurement matrix construction techniques, such as selection from the set of partial Fourier matrices or by drawing entries from a Gaussian or Bernoulli distribution. We then define Φ_t by taking the first M_t rows of Φ and column-normalizing the result.

Our background model assumes that $\mathbf{y}^b = \Phi\mathbf{b}$ is a multivariate Gaussian random variable, i.e.

$$\mathbf{y}^b \sim \mathcal{N}(\mu^b, \Sigma) \quad . \quad (2)$$

Assuming we obtain L independent observations of \mathbf{y}^b , $\{\mathbf{y}_j^b\}_{j=1}^L$, we find the maximum likelihood (ML) estimate of μ^b as $\hat{\mu}^b = \frac{1}{L} \sum_{j=1}^L \mathbf{y}_j^b$.

During subsequent observation, we wish to estimate the compressive measurements of the foreground: $\mathbf{y}_t^f = \Phi_t \mathbf{f}_t$. Using (1) and (2) we see $\mathbf{y}_t^f \sim \mathcal{N}(\mathbf{y}_t - \mu_t^b, \Sigma)$, where $\mu_t^b \in \mathbb{C}^{M_t}$ is formed by retaining only the first M_t components of μ^b and rescaling to account for the different column-normalization factors in Φ and Φ_t . Thus, we can obtain the ML estimate of \mathbf{y}_t^f as

$$\hat{\mathbf{y}}_t^f = \mathbf{y}_t - \mu_t^b \quad . \quad (3)$$

From this point forward, we shall use \mathbf{y}_t^f in discussion with the understanding that $\hat{\mathbf{y}}_t^f$ is used in computation.

3. ADAPTIVE RATE COMPRESSIVE SENSING FOR BACKGROUND SUBTRACTION

We consider the following scenario: a single CS camera observes an area of interest, and we are ultimately concerned only with objects constituting the foreground of the scene. We further assume that this camera is such that the number of compressive measurements it collects at any given time may be dynamically controlled (i.e., “on the fly”). The problem considered in this paper is that of how to adaptively choose this number of measurements so as to limit the data rate of the sensor while simultaneously maintaining enough information such that we are able to faithfully reconstruct the time-varying foreground images.

Let $\{\mathbf{x}_t\}_{t=0}^\infty$ denote the vectorized intensity values of gray scale images in the video sequence of interest, where the index t is a discrete time index. We assume each \mathbf{x}_t adheres to the signal model (1) with a static background. Using the procedure outlined in Section 2.2, we obtain an estimate of \mathbf{y}_t^f from the compressive measurements $\mathbf{y}_t = \Phi_t \mathbf{x}_t$. In a similar fashion, we also estimate cross validation

measurements of the foreground, $\Psi \mathbf{f}_t$, via a $r \times N$ CV matrix Ψ and a corresponding background estimate $\hat{\mu}_\Psi^b$. The value of r is constant with respect to time and is negligible compared to N (in our experiments, less than two percent).

Typically the foreground signal \mathbf{f}_t is sparse in the spatial domain. Denote the number of significant coefficients it contains (i.e., its *sparsity*) by s_t . We assume that the video sampling rate is much faster than any changes in target appearance, and thus that the quantity s_t is slowly-varying with respect to t . However, due to target motion, the support of these coefficients may vary rapidly.

The purpose of our adaptive rate algorithm is to adjust the number of measurements the system collects with respect to time. Thus, we wish to select M_t such that $\mathbf{y}_t^f \in \mathbb{C}^{M_t}$ is of sufficient dimension to ensure that $\hat{\mathbf{f}}_t = \Delta(\mathbf{y}_t^f, \Phi_t)$ is a reliable estimate of the foreground. The method we present for doing so involves a loose estimate of s_t . We denote this estimate by k_t , and it will be known to the system before sensing at time t begins. Using k_t , our algorithm selects an appropriate M_t based on standard CS results. For example, if Φ is selected such that entries are drawn from a Gaussian distribution, it has been shown that choosing $M_t = O(k_t \log(N/k_t))$ suffices for adequate reconstruction of signals with k_t significant coefficients.

3.1. Measurement Rate Selection

The quality of $\hat{\mathbf{f}}_t$ depends on the relationship between s_t and M_t . However, since we do not know s_t prior to sensing, it is impossible to use it when selecting M_t . Cevher et. al. [3] mitigate this problem by assuming an upper bound s such that $s \geq s_t$ for all t , which is used to set $M_t = M$ for all t . Under these conditions, Δ will always yield reliable estimates, and a constant compression ratio of M/N is achieved. However, this approach is wasteful since s may be much greater than s_t for many values of t (e.g., consider the difference between frames with many targets and frames with few or no targets). As shown in Figure 1, collecting extra measurements due to an over-estimate of s_t does not significantly reduce the reconstruction error. Therefore, for those values of t , M_t can be selected to be smaller than M , and we can achieve a better compression ratio without impacting the quality of the foreground estimate.

Our adaptive rate approach accomplishes this by assuming $s_t = k_t$ before sensing, and using the observations collected to adjust k_t for the next time instant. We are only concerned with whether k_t is greater than or less than s_t . This can be posed as a hypothesis test:

$$\begin{aligned} \mathbf{H}_0 : k_t &\geq s_t \\ \mathbf{H}_1 : k_t &< s_t \end{aligned}$$

Since we want k_t to be close to s_t , if \mathbf{H}_0 (\mathbf{H}_1) is determined to be true at time t , then we select k_{t+1} such that $k_{t+1} < (>) k_t$.

We propose to make this determination based on the quality of the best possible estimate Δ could hope to provide given M_t measurements. We quantify this using the ℓ_2 residual of the best k_t -term approximation to \mathbf{f}_t . We will denote this quantity by $\sigma_{k_t}(\mathbf{f}_t)$, i.e.

$$\sigma_{k_t}(\mathbf{f}_t) = \min_{|\mathbf{z}| \leq k_t} \|\mathbf{f}_t - \mathbf{z}\|_2 \quad (4)$$

which is minimized by a \mathbf{z} that matches \mathbf{f}_t in the k_t components with the largest magnitude and is zero elsewhere.

We will assume that the insignificant coefficients of \mathbf{f}_t are distributed as independent, univariate Gaussian random variables with zero mean and identical variances σ_b^2 . Under \mathbf{H}_0 , the $N - k_t$ terms

Algorithm 1 ARCS for Background Subtraction

Require: $\Phi, \Psi, k_t, \hat{\mu}^b, \hat{\mu}_\Psi^b, l_0 \leq 1, l_1 \geq 1$
 Form Φ_t and $\hat{\mu}_t^b$ according to k_t
 Obtain sensor measurements: $\mathbf{y}_t = \Phi_t \mathbf{x}_t, \Psi \mathbf{x}_t$
 Compute ML estimates: $\hat{\mathbf{y}}_t^f = \mathbf{y}_t - \hat{\mu}_t^b, \Psi \mathbf{f}_t = \Psi \mathbf{x}_t - \hat{\mu}_\Psi^b$
 Estimate foreground: $\hat{\mathbf{f}}_t = \Delta(\hat{\mathbf{y}}_t^f, \Phi)$
 Evaluate decision rule (7)
if \mathbf{H}_0 **then**
 Set $k_{t+1} = l_0 k_t$
else
 Set $k_{t+1} = l_1 k_t$
end if

neglected by the \mathbf{z} that minimizes (4) are all insignificant coefficients and thus

$$\mathbb{E}[\sigma_{k_t}(\mathbf{f}_t)] = \sqrt{(N - k_t)\sigma_b^2} \quad (5)$$

where \mathbb{E} denotes expectation. Under \mathbf{H}_1 , the minimizing \mathbf{z} neglects significant coefficients of \mathbf{f}_t , and thus we expect $\mathbb{E}[\sigma_{k_t}(\mathbf{f}_t)]$ to be larger. Motivated by the above, we suggest the following heuristic decision rule:

$$\sigma_{k_t}(\mathbf{f}_t) \underset{\mathbf{H}_1}{\overset{\mathbf{H}_0}{\leq}} c \sqrt{(N - k_t)\sigma_b^2}, \quad (6)$$

for some constant c .

In order to evaluate (6), we need to know the value of $\sigma_{k_t}(\mathbf{f}_t)$. It is clear that we cannot directly compute it, since we do not know the true value of \mathbf{f}_t . However, we can obtain an upper bound for it by using Ward's CV technique [4]. Under assumptions satisfied by our system, the following inequality will hold with high probability: $\sigma_{k_t}(\mathbf{f}_t) \leq \|\mathbf{f}_t - \hat{\mathbf{f}}_t\|_2 \leq (1 + \epsilon) \|\Psi(\mathbf{f}_t - \hat{\mathbf{f}}_t)\|_2$, where Ψ is known and defined as in Section 2.1.1. $\Psi \hat{\mathbf{f}}_t$ is obtained by synthetically applying Ψ to the foreground estimate, and $\Psi \mathbf{f}_t$ is available per previous discussion. Using the previous bound in (6), we form the following decision rule:

$$\|\Psi(\mathbf{f}_t - \hat{\mathbf{f}}_t)\|_2 \underset{\mathbf{H}_1}{\overset{\mathbf{H}_0}{\leq}} c_0 \sqrt{(N - k_t)\sigma_b^2} \quad (7)$$

for some constant c_0 .

The fact that $\|\Psi(\mathbf{f}_t - \hat{\mathbf{f}}_t)\|_2$ is merely an upper bound for $\sigma_{k_t}(\mathbf{f}_t)$ might lead to deciding \mathbf{H}_1 for some t for which \mathbf{H}_0 is actually true, i.e., a false alarm. While false alarms may lead to more measurements, they will not negatively impact the quality of the estimate (e.g., see Figure 1). Further, in practice, it has been observed that $\|\Psi(\mathbf{f}_t - \hat{\mathbf{f}}_t)\|_2$ remains sufficiently close to $\sigma_{k_t}(\mathbf{f}_t)$ such that the number of extra measurements required by these false alarms is minimal.

3.2. Adaptive Rate Compressive Sensing

When assembled, the concepts above yield a strategy which we term *adaptive rate compressive sensing* (ARCS). ARCS provides a means by which to adaptively adjust the measurement rate of a CS system measuring a time-varying signal with dynamic sparsity. We summarize the procedure for background subtraction in Algorithm 1.

4. EXPERIMENTAL RESULTS

To demonstrate the effectiveness of ARCS, we tested the algorithm on the Convoy2 dataset. This dataset, collected on Spesutie Island,

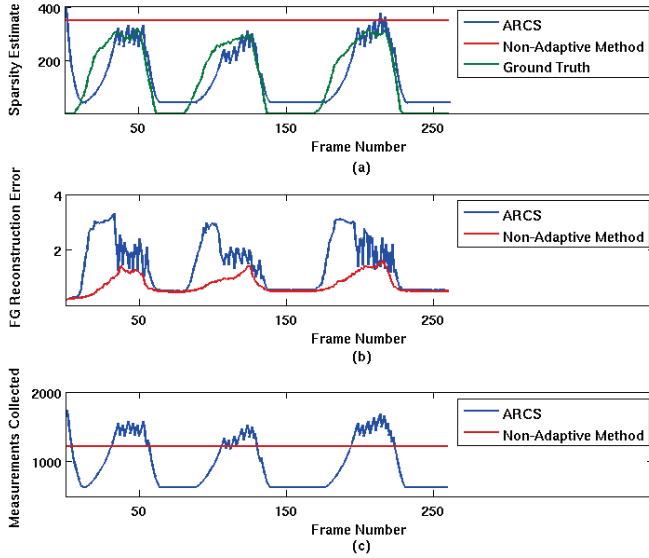


Fig. 2. Comparison between ARCS and a non-adaptive method. (a) Foreground sparsity estimates for each frame, including ground truth values. Note that the non-adaptive estimate upper bounds the ground truth for all frames. (b) ℓ_2 foreground reconstruction error. (c) Number of measurements required. Note the measurements savings provided by ARCS for most frames, and its ability to track the dynamic foreground sparsity.

consists of a video recorded by a single stationary camera. Vehicles enter and exit the field of view over time, which gives rise to a dynamic foreground sparsity (an excerpt can be seen in Figure 1). Since the video was recorded with a conventional imaging system, we simulated CS camera measurements via pre-multiplication with Φ and Ψ in software. For our experiments, Φ was a randomly row-permuted discrete Fourier transform matrix, and entries of Ψ were drawn from a Bernoulli distribution over the discrete set $\{0, 1\}$.

In Figure 2, we compare ARCS to a non-adaptive technique in which the signal sparsity is upper bounded for all t . After background estimation, both ARCS and a non-adaptive technique were used to observe a video sequence consisting of 260 frames. We initialized $k_1 = 400$, and selected $l_0 = 0.8$, $l_1 = 1.1$. Values of $c_0 = 1.2$, and $\sigma_b^2 = 1.5445 \times 10^{-4}$ yielded sufficient results.

Figure 3 depicts a qualitative comparison of ARCS and a non-adaptive method, this time using the PETS 2000 dataset [9]. We show the foreground reconstructions obtained from each technique. Each method was initialized to use an insufficient number of measurements for an accurate foreground reconstruction. While a static method is unable to compensate, the ARCS method is able to successfully recover from this poor initialization by increasing its measurement rate.

From the experiments, we see that ARCS is able to successfully track the value of s_t and adjust the system measurement rate accordingly. For upward swings in s_t , such as when a new target enters the frame, the reconstruction error becomes significantly higher than a non-adaptive scheme until the ARCS estimate is able to compensate. It should also be noted that when s_t approaches its maximum value, ARCS actually collects more measurements than the best non-adaptive scheme. This is due to the overhead required for the CV measurements obtained using Ψ . However, in a real scenario, it may be impossible to determine an appropriate upper bound for s_t in advance. Therefore, this bound may be much looser than the one selected for the experiments presented here.




Fig. 3. Comparison of ARCS and non-adaptive method for two frames in the PETS 2000 video sequence. (a), (b), and (c) show the ground truth, non-adaptive, and ARCS frames (respectively) at t_1 . (d), (e), and (f) show the frames from each at t_2 , where ARCS was able to adapt its measurement rate in order to obtain a better estimate.

5. CONCLUSION

We have presented ARCS, an on-line algorithm that adaptively controls the number of compressive measurements collected while sensing a time-varying signal with dynamic sparsity. We have experimentally validated this algorithm for the application of background subtraction in video surveillance, where ARCS provides significant data savings over non-adaptive approaches. ARCS is also able to adapt to poor initial information regarding the signal sparsity, whereas a non-adaptive scheme is unable to do so.

The practical effectiveness of ARCS in the scenarios described above justifies several topics of future investigation. Among these are the replacement of (7) with a rule more optimal in a decision-theoretic sense and the exploration of other rate-adjustment techniques once this decision is made. Further work could also investigate background-adaptive modifications. We are also interested employing other sources of side-information other than what is provided by Ψ , such as that provided by object tracking.

6. REFERENCES

- [1] "MURI - Opportunistic sensing," <http://dsp.rice.edu/muri31>.
- [2] M.F. Duarte, M.a. Davenport, D. Takhar, J.N. Laska, K.F. Kelly, and R.G. Baraniuk, "Single-Pixel Imaging via Compressive Sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 83–91, Mar. 2008.
- [3] V. Cevher, A. Sankaranarayanan, M. Duarte, D. Reddy, R. Baraniuk, and R. Chellappa, "Compressive sensing for background subtraction," *European Conf. on Computer Vision*, pp. 155–168, 2008.
- [4] Rachel Ward, "Compressed Sensing With Cross Validation," *IEEE Transactions on Information Theory*, vol. 55, no. 12, pp. 5773–5782, 2009.
- [5] P. Boufounos, M.F. Duarte, and R.G. Baraniuk, "Sparse signal reconstruction from noisy compressive measurements using cross validation," in *IEEE 14th Workshop on Statistical Signal Processing*, 2007, pp. 299–303.
- [6] D. M. Malioutov, S. R. Sanghavi, and A. S. Willsky, "Sequential Compressed Sensing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 435–444, Apr. 2010.
- [7] N. Vaswani, "Kalman filtered compressed sensing," in *15th IEEE International Conf. on Image Processing*, 2008, pp. 893–896.
- [8] Emmanuel Candès, "Compressive sampling," in *Proceedings of the International Congress of Mathematicians*, Jan. 2006, vol. 17.
- [9] "PETS 2000 Benchmark Data," <ftp://ftp.pets.rdg.ac.uk/pub/PETS2000>.