

Designing Adaptive Sensing Policies for Meteorological Phenomena via Spectral Analysis of Radar Images

Bruno Castro da Silva, Andrew G. Barto, Jim Kurose

Technical Report UM-CS-2012-006

Computer Science Department
University of Massachusetts Amherst
{bsilva, barto, kurose}@cs.umass.edu

March 6, 2012

1 Abstract

In this work we consider the problem of determining efficient scanning strategies for adaptive networks of radars used to monitor meteorological events. While most meteorological radars traditionally scan 360 degrees, repeatedly, more accurate readings can be obtained if adaptive sector scanning is performed. We propose a scanning strategy based on analyzing eigenimages computed given a set of observed radar images. These eigenimages allow us to identify spatial regions with high radar measurement variance, and thus provide us with the necessary information to reconfigure radar resources towards regions where the storms' states seem to be changing the fastest. These are exactly the regions where newer and more accurate readings are usually necessary. We show that performing adaptive reading in this manner provides gains over the traditional sit-and-spin approach, and also over a baseline random scanning strategy. We also show that there is a difficult trade-off between the minimum size of storm that one might want to track, and the capability of saving scanning resources by not focusing on irrelevant sectors.

2 Introduction

Typical meteorological radars are usually configured to always scan 360 degrees. This is the case, for example, of the radars used by the National Weather Service NEXRAD system [9]. Although this type of radar provides a comprehensive

reading of the atmosphere, it might waste resources by scanning regions where no interesting phenomena are occurring, and therefore might not be able to obtain high-resolution measurements of the sectors that are in fact relevant.

Recently, a new type of radar technology has been proposed in order to construct small, low-power but agile radars that can perform sector scanning. This type of technology is being developed, for example, by the Collaborative Adaptive Sensing of the Atmosphere (CASA) Engineering Research Center [14]. Radars that can scan just a subsector of the atmosphere might outperform sit-and-spin radars. The reason for this is that given a fixed amount of time, scanning smaller sectors of the sky implies that more time can be spent obtaining readings from that sector, and that thus more accurate measurements are possible. By contrast, a radar that always scans 360 degrees will have to rotate faster in order to still cover the complete space in the same pre-determined fixed amount of time, and therefore will devote less sensing energy for detecting existing phenomena in a given region.

In this work, we study how to determine good scanning strategies for a network of radars; that is, how to configure a set of radars so that the quality of readings over all meteorological phenomena is maximized. There are several factors that should be considered when choosing a scanning strategy. The speed of rotation of the radar matters, as mentioned before, but we should also take into account the expected number of decision epochs before a sector is re-scanned. This is important because we want to make sure that even sectors that now seem uninteresting will be periodically re-scanned in order to make sure that no meteorological phenomena have appeared in the meantime. Another factor to be considered is whether to optimize myopically, that is, taking into account just the observations up to the current moment, or to optimize over expected future observations. This work will develop an approach for myopic decisions, although later on we discuss a simple extension that could allow for optimizations with a lookahead.

The scanning strategy proposed in this work is based on spectrally analyzing a sequence of observed radar images. Specifically, we construct a set of basis images (sometimes called eigenimages) that correspond to the eigenvectors of the covariance matrix of a fixed number of past radar images. The first such eigenimages describe a subspace of the original space of images that accounts for most of the observed variance over radar measurements. Usually these eigenvectors are used in order to find a lower-dimensional representation of the data, by projecting the original high-dimensional points (such as radar images) onto the bases of eigenvectors. We, on the other hand, analyze the eigenvectors' components directly, and use that information to construct relevant features of the storms. This type of analysis allows us to identify regions of the radar's sectors in which significant changes in the meteorological phenomena's states have been observed. These regions are exactly the ones for which newer and more accurate readings are usually necessary, and towards which the next measurement efforts should be directed.

For each possible radar joint action (i.e., for each possible way in which we can configure the set of radars) we specify a *quality metric*, which reflects how well that specific joint action allows the radars to scan for meteorologi-

cal phenomena. By using Boltzmann exploration with a specifically calculated temperature parameter, it is also possible to take into account some properties that meteorologists find desirable. For example, it is usually desirable that each radar sector be re-scanned on average at least once every 2 minutes.

We present comparisons of our method with (i) the scan strategy traditionally used in real-life, in which radars simply rotate 360 degrees all the time, and (ii) a baseline random scanning strategy, and show that our approach performs better in a variety of situations. We also show that there is a difficult trade-off between the minimum size of storm that one might want to track, and the capability of saving scanning resources by not focusing on irrelevant sectors.

The rest of this paper is organized as follows. Section 3 presents related work; Section 4 defines the Radar Scheduling Problem; Section 5 reviews the method of Principal Component Analysis, which will be used in our approach. Sections 6 and 7 present our approach and some results, and in Section 8 we summarize this work, discuss its limitations and outline possible future extensions.

3 Related Work

Several previous researchers have studied strategies for obtaining meaningful adaptive measurements of storms, and in particular for efficiently tracking and predicting their spatial motion and evolution. One traditional approach has been to identify and represent storms by means of their centroids, and then to track those centroids over time [3, 4]. Once storms are correctly detected, the change in the position of their centroids can be used as a basis for estimating their future states. This type of technique works well for predicting well-organized and isolated storms, since for these it is easy to assign a meaningful centroid. However, whenever storms split and merge, evolve over time in a complicated manner, or simply start growing or decaying, these methods do not perform satisfactorily. An approach based on tracking centroids is currently being used in the CASA's closed-loop operation [5]. The SCIT algorithm [4] is another example of such approach. Finally, an algorithm that detects centroids by running K-means on radar images is also presented in [6, 7].

Another type of approach used in order to perform tracking of meteorological phenomena consists in using correlation-based models [1, 12]. In this case, a set of observations is treated as a sequence of images, and then the local-area cross-correlation between every two successive images is computed. This technique enables small-scale storm analysis, but relies on a good choice for the size of the window to be analyzed: if the window is too small, the motion field might be incorrectly estimated due to storm growth and decay; if it is too large, local features might not be well represented.

In this work we want to avoid having to identify each storm individually, and also to avoid having to keep track of centroids. Centroids, although computationally efficient, are not particularly good representations for complex storms whose shape changes over time or that might split or merge with other storms. Correlation-based techniques are also not ideal since, as mentioned before, it

is not always easy to identify the right scale for analysis. Our work tries to bridge this gap by providing an approach that does not rely on a representation based on single centroids, and that does not assume any knowledge about the temporal evolution of storms, their spatial shape, or the dynamics that guide their evolution.

As will be mentioned in Section 6, our approach is based on analyzing the Principal Components of a set of radar images. Previous works such as [13] also used this type of spectral decomposition, but only in order to find a lower-dimensional representation of the measurements, and in order to obtain a short-term linear prediction of the next states of the phenomena.

4 The Radar Scheduling Problem

In this section we formulate the multi-radar scheduling problem; this formulation is based on the one presented in [8]. We assume that there is a constant number of radars, that each one is located at a fixed position in space, and that each has a given range. Notice that the radars might have overlapping footprints. At each decision epoch, each radar can choose from a set of *scan actions*. A scan action determines the size of the sector to be scanned, its start angle, its end angle, and the angle of elevation. As [8], we do not consider elevation angles. We assume that a radar can scan either a 90° , 180° , 270° , or 360° sector, and that the starting angle is limited to be either 0° , 90° , 180° , or 270° . These assumptions imply 13 actions per radar¹. Our goal is to use spectral analysis over a set of past radar images in order to obtain relevant features that help us to determine which scan actions to use, so as to obtain good quality measurements of the existing storms.

Before we proceed to present the Radar Scheduling Problem, it is important to highlight some properties of meteorological radars. Specifically, one important property of such radars is that the smaller the sector scanned, the better the resulting quality of collected data. The reason for this is that given a fixed amount of time, scanning smaller sectors of the sky implies that more time can be spent obtaining readings from that sector, and that thus more accurate measurements are possible. Another property of such radars is that any scan strategy that focuses on smaller segments of the sky incurs greater risk of missing new phenomena which might have arrived in the meantime in sectors that are not being scanned. Considering these two properties, we define an *effective strategy* to be one that chooses scan actions that provide a good balance between such conflicting requirements; that is, one that tries to maximize the accuracy of the readings while minimizing the probability of missing new phenomena.

We will now describe a quality function U which, for each storm in the environment, associates a real number describing how well the current scan action measures the properties of the corresponding storm. Such a quality

¹There are 16 possible actions, but four of those – specifically the ones that scan the whole 360° but which start at different angles, are equivalent, and thus we can disregard three of them.

function will be used later on to determine how noisy the measurements obtained by the scan action are, and consequently how effective the scanning strategy is. The quality function for a given storm and joint action will also be used to define a *cost function*, which is what we will try to optimize in this work.

In order to compute the quality with which an action measures a storm, we consider several properties of the radars, and of the specific storm being scanned. Specifically, we consider the overall storm's spatial shape, the location of the radars, the speed of rotation of the radars, and their range. Let s_r be a scan action of radar r . Let S be a set specifying a joint action, i.e., a scan action s_r for each individual radar. The quality $U_p(S)$ with which S scans a storm p is given by

$$U_p(S) = \max_{s_r \in S} U_p(s_r)$$

$$U_p(s_r) = F_c(c(p, s_r)) \times \left[\beta F_d(d(r, p)) + (1 - \beta) F_w\left(\frac{w(s_r)}{360}\right) \right]$$

where

- $c(p, s_r)$ = the area of p covered by the scanning action s_r ;
- $d(r, p)$ = the normalized distance from the center of mass of p to the location of radar r ;
- $w(s_r)$ = the area of space covered by scanning action s_r ; this depends on the range of r and on the start and end angles of s_r ;
- β = a tunable parameter that balances the relative importance of p being close to r , and the negative effects of high rotation speeds of r .

$U_p(s_r)$ is the quality obtained for scanning phenomenon p using radar r and scan action s_r . The functions $F_c(\cdot)$, $F_w(\cdot)$ and $F_d(\cdot)$ are adapted from [10] and were derived from discussions with radar meteorologists. F_c is a function that captures the effect on the quality function due to partially scanning a storm; F_w represents the negative effects of high radar rotation speeds; and F_d represents the negative effect on quality due to scanning storms that are far away from the radar. Examples of these functions, as used in [8], are shown in Figure 1.

Also, in order to represent the negative effect on the observation quality of storms that were not re-scanned in the last decision epoch, we decay their observation quality by a fixed amount κ :

$$U_p = \max(U_p - \kappa, 0) \quad \text{for all storms } p \text{ that were not observed in the last decision epoch.}$$

Given any strategy for determining a joint action S at a given decision epoch, we can compute its *cost*. The cost of a strategy at a given time step depends




Figure 1: Step functions for the computation of U_p , as used in [8].

on how well each storm is measured by the collective action of the set of radars being used. We first compute the quality U_p for each existing storm p , and then, based on that, we compute the current cost $C(t)$ of the strategy:

$$C(t) = \left(\sum_{i=1}^{N_{obs}} \sum_{j=1}^{\|d\|} \|d_{ij}^{obs} - d_{ij}\| \right) + \left((N - N_{obs})P \right)$$

where N is the true number of existing storms at time t , N_{obs} is the number of storms effectively observed by the last joint action, $\|d\|$ is the number of storm attributes being measured by the radars, d_{ij} is the true value of the j -th attribute of storm i , d_{ij}^{obs} is the measured value of the j -th attribute of storm i , and P is a penalty factor for missing a storm.

The value of the observed attributes d_{pj}^{obs} are calculated by adding zero-mean Gaussian noise to the true attributes d_{pj} of a storm p . The higher the quality of observation of p , the lower the noise. Specifically, the variance of the Gaussian noise affecting each attribute j of a storm p depends on the current scanning quality of p :

$$d_{pj}^{obs} = d_{pj} + N \left(0, \left[(1 - U_p) \frac{V_j^{max}}{\rho} \right]^2 \right)$$

where V_j^{max} corresponds to either the average value of attribute j (in case of velocities or radii), or to the largest positive value that that attribute can take. The parameter ρ is a scaling factor; as $\frac{1}{\rho}$ increases, so too does the amount of noise added to the true value of the attribute. The variance of the noise is not only used to perturb the true values of the attributes of the storm, but also in order to define the penalty P for missing a storm. We assume that all unobserved storms are actually observed with quality zero, and that therefore each of its attributes j will be on average $\frac{V_j^{max}}{\rho}$ away from its true value. Summing this quantity over all attributes of a storm gives the following penalty, for a given unobserved storm:

$$P = \sum_{j=1}^{\|d\|} \frac{V_j^{max}}{\rho}.$$

In the rest of this work, all performance measurements will be made by observing how the cost C of a scanning strategy changes over time, given a scanning strategy. The lower the cost, the better the strategy.

5 Principal Component Analysis

Principal Component Analysis is a statistical technique used mainly for dimensionality reduction and for feature extraction. The central idea is to find a linear projection of the original data onto a subspace that preserves most of the variance observed in the original data. This orthogonal projection defines what is known as a *principal subspace*.

Let $\{x_n\}$ be a set of N data points, each with dimensionality D . We want to find a projection of these points onto a space with dimensionality $M < D$ while maximizing the variance of the projected data. Let us start by assuming that we are trying to find the projection onto a *one-dimensional* subspace. The extension of the results shown below to $M > 1$ can be proved by induction.

Since $M = 1$, we can define the projection by using only one D -dimensional vector, which we call u_1 . We can assume that u_1 is a unit vector, since we are only interested in its direction, not its magnitude. Because u_1 is a unit vector, the projection of a data point x_n onto the one-dimensional space is simply given by $u_1^T x_n$. Let \bar{x} be the mean value of $\{x_n\}$. The variance of the projected data is then given by

$$\begin{aligned} V_{proj} &= \frac{1}{N} \sum_{n=1}^N \left(u_1^T x_n - u_1^T \bar{x} \right)^2 \\ &= \frac{1}{N} \sum_{n=1}^N \left(u_1^T (x_n - \bar{x}) \right)^2 \\ &= \frac{1}{N} \sum_{n=1}^N \left(u_1^T (x_n - \bar{x}) \right) \left(u_1^T (x_n - \bar{x}) \right)^T \\ &= \frac{1}{N} \sum_{n=1}^N \left(u_1^T (x_n - \bar{x}) \right) \left((x_n - \bar{x})^T u_1 \right) \\ &= u_1^T \frac{1}{N} \left(\sum_{n=1}^N (x_n - \bar{x})(x_n - \bar{x})^T \right) u_1 \\ &= u_1^T S u_1 \end{aligned}$$

where S is the covariance matrix of the data:

$$S = \frac{1}{N} \sum_{n=1}^N (x_n - \bar{x})(x_n - \bar{x})^T.$$

Remember that our goal is to find a basis vector u_1 that maximizes the projected variance of the projected data, $u_1^T S u_1$. We first have to transform this problem into a constrained maximization problem. Specifically, if we do not somehow encode the fact that u_1 is constrained to be a unit vector, maximizing the variance would be achieved by simply using a vector u_1 that points in the right direction, but that has infinite magnitude. We perform this transformation by introducing a Lagrange multiplier, which we denote λ_1 . Since u_1 is unitary, and therefore $u_1^T u_1 = 1$, we can find the best vector u_1 by making an unconstrained maximization of

$$u_1^T S u_1 + \lambda_1 (1 - u_1^T u_1).$$

By setting the derivative of this equation to zero, we obtain

$$S u_1 = \lambda_1 u_1$$

which we can then left-multiply by u_1^T to obtain

$$u_1^T S u_1 = \lambda_1$$

These two equations imply both that u_1 must be an eigenvector of the covariance matrix S , and that the vector u_1 that maximizes the projected variance is the one with largest eigenvalue λ_1 . In case $M > 1$, we can define additional principal components u_2, u_3, \dots, u_n by applying the same reasoning as above, but starting with the data after projecting it onto u_1 . In this manner, each additional principal component is responsible for “explaining” part of the variance that was not captured by the simple subspace defined by earlier basis vectors.

6 Contribution

Our approach for defining a good scanning strategy is based on analyzing the first eigenvector, E_1 , of the covariance matrix of a set of radar images. This eigenvector, as will be shown, is the one that encodes most of the information about what are the regions of highest variability in the radar images. Specifically, for a given number k of past radar images, each being $n \times m$ pixels reporting reflectivity, we can construct a matrix M where the columns correspond to images²; this gives us a matrix with dimensions $(nm) \times k$. By running

²Note, however, that nothing in our approach *requires* that the pixels of the images encode reflectivity; dealing with other types of radar readings is left as future work.

PCA on the data stored in this matrix we obtain a set of eigenvectors which describe a new space in which to represent the radar observations, and in which it is easier to characterize and analyze the variability in reflectivity of each storm. Specifically, the set of axis defining this new space can be ordered in a way as to indicate what are the attributes of the observed measurements (in this case, pixels reporting reflectivity) that vary the most.

Intuitively, each nm -dimensional eigenvector provided by PCA can be reshaped as a $n \times m$ image. This type of visualization creates what is usually called an *eigenimage* (more details in Section 6.2). These eigenimages form a basis that spans the whole set of observed radar images. The eigenimage with highest eigenvalue is specially interesting to us, since the value of its components, or pixels, indicates regions of highest variability over the past radar images. These regions are exactly the ones for which newer and more accurate readings are usually necessary, and towards which the next scanning actions should be directed.

Before we present representative examples of these eigenimages, and describe the way in which they are used to define an efficient scanning strategy, let us first describe the domain in which we ran our simulations, and also the manner in which we control the process that generates simulated radar readings.

6.1 Description of the domain and simulation process

The domain used in this work corresponds to a rectangular region of 90km by 60km. We use two radars with overlapping footprints, each with a 30km range. At any given time we allow at most a given fixed number of storm cells in the environment. New storms are created by a process described below. Whenever the center of mass of a storm is no longer within the range of any radar, that cell is removed from the simulation.

Storms appear in the environment at a given time probabilistically, based on a stochastic model of rainfall in which cells arrive according to a spatio-temporal Poisson process [2, 11]. Each decision epoch in our domain lasts 30 seconds. In order to compute how many new storm cells will potentially be added during each 30-second decision epoch, we sample a Poisson random variable with rate

$$\lambda\eta\delta_a\delta_t$$

where the parameters used above are the same as the ones modeled in [11]. Specifically, $\lambda = 0.075$ storms/ km^2 , and $\eta = 0.006$ cells/minute. From our radar setup we have $\delta_a = 90 \times 60 km^2$, and because each decision epoch takes 30 seconds, $\delta_t = 0.5$ minutes.

Each storm in our simulation is described by an arbitrarily rotated ellipse, and has the following 7 attributes:

- x coordinate;

- y coordinate;
- \dot{x} , the velocity along the x axis;
- \dot{y} , the velocity along the y axis;
- r_M , the ellipse's major radius;
- r_m , the ellipse's minor radius;
- α , the ellipse's rotation angle with respect to the x axis.

Note that meteorological radars are typically not able to measure these attributes directly; therefore, such attributes will only be for defining and simulating the storms' dynamics over time, but will *not* be given directly to our analysis algorithm. The input to our algorithm will correspond to a sequence of *simulated radar images*, all of which are constructed based on the simulation of the storms' dynamics, and on the previously mentioned values. Notice also that some attributes, such as velocity, are not directly encoded in the simulated radar images, although they might, to some extent, be inferred if one observes a sequence of images and tries to estimate the position and shape of each storm.

For each new storm being added to the simulation, we sample its attributes from the following distributions. The x and y coordinates are uniformly sampled, as well as the storm's angle of rotation. The storm's radii are drawn from a Gaussian distribution $N(\mu_R, \sigma_R^2)$. The horizontal and vertical velocities are sampled from a Gaussian distribution, whose parameters are based on real-life measurements, as follows. As in [8], 39 existing storm cell tracks from the National Severe Storms Laboratory were used in order to estimate the distribution of storm velocities. Each track consists of a series of measurements of latitude and longitude of storms' centers; by computing the difference in latitude and longitude between successive pairs of points, it is possible to fit the resulting latitude/longitude velocities using Gaussian distributions. Moreover, it is possible to convert each degree of latitude/longitude to its equivalent in kilometers, and thus to find the parameters that describe the distribution of horizontal and vertical velocities. [8] computed these values and found that the latitude (or x) velocity has mean 9.1 km/h and standard deviation of 35.6 km/h, and that the longitude (or y) velocity has mean of 16.7 km/h and standard deviation of 28.8 km/h. In order to obtain a new storm's (x, y) velocity, we sample from such a Gaussian distribution.

Based on the two building blocks described above (the Poisson process that creates new storms, and the uniform and Gaussian distributions used to instantiate their parameters), we are able to simulate the dynamics of a set of storms over time. Given such a simulated storm track, we can then generate readings that represent a simulated radar image of the track. *These simulated radar images will be used as input to all scanning algorithms.* An example of such a simulated radar image is shown in Figure 2. Circles delimit the range of each radar, and ellipses correspond to currently active storms.




Figure 2: Example of a simulate radar image generated based on a simulated track. In this simulation, a maximum of 10 storms was allowed to be active in the enviroment at any given time.

6.2 Analysis of eigenimages

In order to determine possible interesting regions to be scanned, we analyse the *components* of the first eigenvector obtained by running PCA over a set of past observed radar images. Before we define exactly what we mean by a component of an eigenvector, let us revise some of the concepts presented in the previous subsection, and also introduce some new ideas. First of all, let us note that if we take k past radar images, each being composed by $n \times m$ pixels reporting reflectivity, and organize them in the columns of a matrix M , we end up with a matrix whose dimensions are $(nm) \times k$. As described in Section 5, the method of Principal Component Analysis (PCA) works by computing the eigenvectors of S , the covariance matrix of the data stored in M . This covariance matrix is a $(nm \times nm)$ -dimensional matrix whose rank is at most k ; therefore, from basic linear algebra we can conclude that each eigenvector of S is a nm -dimensional vector, and that there are at most k orthogonal eigenvectors. We note also that because each original radar image is composed by $n \times m$ pixels, it is possible to interpret them as points in a nm -dimensional space, where each axis of that space represents a pixel. To give a concrete example, let us suppose that we were dealing with images formed by just 2 pixels; in this case, we could visualize each such image as a point in a 2-dimensional plane, in which the values along each of the two axes would represent the value of each of the two pixels. Since we are dealing with much larger images, however, which might in general be composed by $n \times m$ pixels, we have to interpret them not as points in a 2-D plane, but as points in a larger nm -dimensional vector space. Considering this fact, and the previously mentioned observation that eigenvectors of S are also nm -dimensional vectors, we can conclude that by reshaping such eigenvectors as $n \times m$ matrices, we might visualize them as $n \times m$ images. These images are typically called *eigenimages*. Note also that because each eigenvector of S is a nm -dimensional vector, it is in practice represented as a tuple of nm numbers; we call each of those numbers a *component* of the eigenvector. When viewing an eigenvector as an image, each of these components directly corresponds to

a pixel in the corresponding eigenimage. As will be discussed in the following Sections, great part of this work is based on the analysis of these components.

Finally, we note that because *i)* S , the covariance matrix of the radar images, has at most k eigenvectors, and *ii)* these eigenvectors define a new complete space that spans the whole set of k radar images, we can in principle reconstruct (or re-express) any of original radar images as a linear combination of such eigenvectors. In other words, the set of eigenimages computed by PCA spans the whole set of original radar images. Because of the properties of the eigenvectors computed by PCA, it is the case that the first such eigenvectors (i.e., the ones with highest eigenvalues, up to a threshold) can be used to define a subspace that encodes most of the observed variance across images. In fact, we observe that in practice much fewer than k eigenimages are typically necessary to reconstruct any of the original radar images, and thus we conclude that much fewer than k eigenvectors are sufficient to encode all the regions of the images where significant variability has been observed.

When analyzing eigenimages, we observe that each one of them tends to encode different aspects of the variability across images. To give a concrete example, suppose we were analyzing human faces, instead of radar images; this application of PCA is usually known as eigenfaces. The first eigenimage in an eigenface application tends to encode the region of the eyes, mouth and hair, since these are the parts of human faces that vary the most across individuals. These regions are usually characterized by pixels with very bright or very dark values. The subsequent eigenimages in an eigenface application (i.e., the ones with progressively smaller eigenvalues) tend to encode regions of the face they vary less, such as eyebrows and the hairline. These less significant eigenimages encode finer and finer features of the images, and in practice any eigenimage with a very small eigenvalue contributes very little to describing the overall, global regions of high variability across faces. In the case of eigenimages formed by analyzing radar images, on the other hand, we expect the first eigenimages to encode regions of the space where most of the change in reflectivity is occurring. In fact, as will be described shortly, the first eigenimage in our analysis was usually capable of encoding the movement, over time, of the fringes (or boundaries) of the storms, and thus allowed us to identify from where and to where storms seemed to be moving in the short term. This type of information is very important to our application, since it gives us hints regarding where the radars should focus on in order to obtain newer and more accurate readings of the evolving storms.

When constructing the matrix M , one needs to decide how many radar images, k , will be used. The larger the number of past observations is used, the more of an overall big picture we get regarding which spatial regions present high measurement variability. However, by using a large number of past observations we also get less precise indicators about which regions have undergone *recent* high variance measurements. In our experiments we have used a window of time of 10 past radar images, which means that our approach decides for a new scan action every 30 seconds, based on observations from the past 5 minutes. Using fewer images than that did not, empirically, provide enough information to describe the short-term changes in the state of the storms; using more images,

on the other hand, had the downside of “remembering” too much information from the past, and thus wrongly characterizing regions where storms used to be a long time ago as being regions where significant variability has recently occurred. Note that because we use 10 past radar images in each step of the decision-taking process, during the first 5 minutes of simulation we don’t have yet enough data to construct the required eigenimages. Therefore, during this initial period of the process we employ a simple sit-and-spin strategy.

We now present, in Figures 3, 4 and 5, examples of eigenimages generated based on our simulated radar images. In Figure 3 we show the first three eigenimages computed over a set of 100 past radar images; these images correspond to the simulation of 50 minutes of real time. The simulation used to generate the tracks allowed up to 5 storms to be active in the environment at any given time. A similar set of eigenimages, but now for a simulation allowing up to 2 active storms, is shown in Figure 4. Figure 4 shows us in a clearer way that in general very bright and very dark regions of the eigenimages tend to represent the regions of the scanned space in which most of the variability has occurred. Typically, dark regions of the first eigenimage tend to encode parts of the space where radar measurements went from reading a storm, to reading just background noise (i.e., regions where there was once a storm, but which then moved somewhere else). Brighter regions of the first eigenimage, on the other hand, tend to encode regions of the space where storms are arriving at. Long-term eigenimages such as the ones presented in Figures 3 and 4 provide us with a big picture regarding the overall storms’ structure, that is, about the way in which they tend to move in the long run. However, this type of long-term eigenimage has the disadvantage of also recording areas of high variability due to storms that might not be active anymore. Therefore, such images are not particularly useful for deciding a good *current* scan action.



Figure 3: First three eigenimages generated by analyzing the last 100 radar observations. A maximum of 5 storms was allowed to be active in the environment at any given time. Very bright and very dark regions encode places where high variance has been observed, and are related to the direction in which storms are moving.

In Figure 5, on the other hand, we present the first three eigenimages for a simulation involving up to 5 storms but considering just the *past 10 minutes of radar observations*. We notice that by using a smaller window of time, short-term storm motion features can be observed. In this case, bright and dark regions in the first eigenimage tend to represent the fringes, or boundaries, of the storms. The rest of the eigenimages represent regions of the images where less variability has been observed, and in general do not seem to encode features




Figure 4: First three eigenimages generated by analyzing the last 100 radar observations. A maximum of 2 storms was allowed to be active in the environment at any given time. Very bright and very dark regions encode places where high variance has been observed, and are related to the direction in which storms are moving.

that are directly useful for characterizing a good scanning strategy. In Section 8 we discuss a possible way in which other eigenimages, and the entropy over eigenvalues, could be used in order to provide us with further useful information.



Figure 5: First three eigenimages generated by analyzing the last 20 radar observations, or 10 minutes of real time observations. A maximum of 5 storms was allowed to be active in the environment at any given time. This type of eigenimage provides a good way of extracting short-term motion features based on the past observed radar images.

6.3 Action space and action quality

As previously mentioned, in this work we focus primarily on analyzing the first eigenimage. We choose to use just the first eigenimage because it is the one that encodes most of the information about what are the regions where the highest variabilities in reflectivity have been observed. The first eigenimage also seems to encode features that represent the fringes, or boundaries, of the storms. These are particularly useful for our purposes since they allow us to identify from where and to where storms seem to be moving in the short term, and thus give important indications regarding where the radars should focus on in order to obtain newer and more accurate readings of the evolving storms.

We will now formally define the set of actions that are available to each radar, and describe how to evaluate their qualities. We start by identifying the regions of an image that can be scanned by each radar r . We also identify which individual regions of an image can be covered by each possible scan action of

each radar r . For instance, when the leftmost radar in Figure 2 performs a scan whose start angle is 0 degrees, and whose end angle is 180 degrees, the scanned region corresponds to the top half of the leftmost circle shown in that Figure. Another example of this type of scanned region, implied by a specific radar action, is shown in Figure 6.

In what follows, we denote the spatial region of an image I that is scanned by an action s_r by $I(s_r)$. Let E_1 be the first eigenvector obtained when running PCA over a fixed window of past radar images. Let $A = \{E_1(s_r)\}$ be the set of all possible image regions scanned by each individual action s_r of radar r . For each action s_r , we compute its *quality* based on the intensity of the components of the region described by $E_1(s_r)$. That is, the quality of an action s_r is based on the intensity of the components of the region of the eigenimage that would be scanned by the action s_r .

If we assume that the set of active storms usually do not occupy the whole space all the time, then it is reasonable to assume that the most frequent component intensity present in the eigenimage indicates regions where no moving phenomena have recently occurred. In other words, by considering the intensities of components of an eigenimage E_1 , and by taking the mode of such intensities (or by analyzing its histogram), we can try to identify what we call the “background variance” over a set of images. The background variance empirically indicates regions of the scanning space in which the radars have very likely *not* scanned any recent moving phenomena, such as evolving storms. In Figure 6 we present an eigenimage, the same eigenimage but with background variance removed, and also the region $E_1(s_r)$ scanned by an action s_r of the rightmost radar³, in case s_r corresponds to the action whose scanned sector goes from 0° to 90° . In part (c) of Figure 6 we can see that all regions with intensity different than zero (i.e., all those not represented in black) correspond to the *fringes of a moving storm*. In other words, these regions encode and highlight only the regions of the space in which radar measurements indicate important changes in the state of the storms.




Figure 6: (a) Original eigenimage; (b) eigenimage with background variance removed; (c) example of region covered by a scanning action of the rightmost radar, in case the corresponding scanned sector goes from 0° to 90° .

Let $\|E_1(s_r)\|$ denote the size, in number of pixels, of the region scanned by

³not shown.

the action s_r . Let $H(E_1(s_r))$ be the number of non-zero elements in the region scanned by that action. Based on these values, we define the quality $Q(s_r)$ of a scanning action s_r to be a value proportional to the amount of the scanned region that contains features indicating high variability in the storm’s intensity, and penalized by the size of the scanned region:

$$Q(s_r) = \max\left(H(E_1(s_r)) - \tau\|E_1(s_r)\|, 0\right)$$

where τ is a tunable parameter that represents the fact that an action’s quality depends not only on the amount of the scanned region that contains features indicating high variability in the storm’s intensity, but also, in an inversely proportional manner, on the size of the region scanned. This is so because in order to scan large regions of the space, a radar has to rotate faster, and thus obtains less precise measurements. The penalty weight τ is probably the most important parameter in our approach, since it regulates how advantageous it is to expand the scanned area and to collect more information, but at the expense of losing radar resolution. The specific value of τ to be used is usually related to the minimum size of storm that one still wants to be able to detect. The importance of this parameter will be further discussed in Section 7.

Now let us describe how our approach decides which scan actions to execute. Let $s_r^1, s_r^2, \dots, s_r^p$ be the set of all p actions available to radar r . We choose the scan action of radar r probabilistically, according to a Boltzmann distribution constructed based on the quality of the actions of that radar. Specifically, the probability $\pi(s_r^i)$ of choosing action s_r^i is given by

$$\pi(s_r^i) = \frac{e^{Q(s_r^i)/T}}{\sum_{j=1}^p e^{Q(s_r^j)/T}}$$

where T is a temperature parameter that regulates how likely it is that we choose an exploratory action rather than the greedy action.

One important aspect of this type of probabilistic approach for picking actions is that it allows us to very easily cope with some requirements that meteorologists usually have. It can be desirable, for instance, to guarantee that each radar sector is scanned on average at least once every 2 minutes. In our simulations, each radar’s scan space was divided into 4 equally sized disjoint sectors. Each of the 13 actions available to a radar can then be seen as scanning each possible sequence of adjacent sectors. In order to guarantee that on average every sector is scanned at least once every K decision epochs, let us analyze the simplified case when we consider only scanning individual sectors. Let us imagine that we want to guarantee that one given sector, i , is scanned at least once every K decision epochs. Let us also assume that because all actions being considered scan regions of same size, τ can be set to zero. In the worst case, $H(E_1(s_r^i)) = 0$, that is, the sector of interest, i , has no indication of varying phenomena, and all other 3 sectors have maximum value for H : $H(E_1(s_r^j)) = 1$, for all $j \neq i$. In this case, the probability of scanning sector i is

$$\begin{aligned}
\pi(s_r^i) &= \frac{e^{Q(s_r^i)/T}}{\sum_{j=1}^p e^{Q(s_r^j)/T}} \\
&= \frac{e^{\frac{0}{T}}}{3e^{\frac{1}{T}} + e^{\frac{0}{T}}} \\
&= \frac{1}{3e^{\frac{1}{T}} + 1}.
\end{aligned}$$

If we want to lower bound the probability of scanning the sector i by a constant probability ϵ , we can solve the above equation for T :

$$\frac{1}{3e^{\frac{1}{T}} + 1} > \epsilon$$

and therefore

$$T > \frac{1}{\ln\left(\frac{1-\epsilon}{3\epsilon}\right)}.$$

In the worst case, the sector of interest, i , will continue to have zero quality even during future decision epochs, and in that case we can find the expected number of decision epochs until i is first scanned (or re-scanned) by using the properties of geometric distributions. Specifically, we might use a geometric distribution in order to analyze the expected number of Bernoulli trials, each with probability of success equal to ϕ , before the first success occurs. The expected number of trials until the first success occurs is given by $\frac{1}{\phi}$. Therefore, applying this result to our case, in which we want the expected number of decision epochs before i is re-scanned to be on average K , we find that $\frac{1}{\epsilon} = K$. Substituting this in the equation for the lower bound on T gives us

$$T > \frac{1}{\ln\left(\frac{1-\frac{1}{K}}{3\frac{1}{K}}\right)}.$$

The same type of argument could we developed for the case when we do not scan only individual sectors, and for when τ is not zero.

Notice that after having used the above mentioned strategies for choosing the scan action s_r for each radar r , the joint action S is then simply constructed as the set containing all individual actions, that is, $\{s_r\}$, for all r .

7 Results

We now present some results in which we compare our technique to the traditionally used sit-and-spin strategy, and to a baseline random strategy. As mentioned in Section 4, all scanning strategies will be evaluated based on the scanning costs in which they incur.

Let us first describe some specific details about how we organized the domain used for these comparisons. In our simulations we considered a window of 10 past radar images, in order to compute the eigenimages. As described in Section 4, there are 13 actions available for each radar, and therefore a total of 13^2 possible joint actions. The parameter ρ , which affects how noisy the observations can be, was set to $\rho = 20$. That way, vertical attributes of a storm could be measured up to 3 km off their correct location, for example. Storms were created according to the previously described Poisson process, and velocities were drawn from the presented Gaussian distributions. The major and minor radii of each storm were drawn either from $N(10, 4^2)$, in case we wanted to generate a simulation involving large storms, or from $N(5, 2^2)$, in case we were interested in generating a simulation involving only small storms.

Figure 7 presents the typical results that were obtained when simulating an intermediate number of maximum active storms, and when considering large storms. This Figure presents results corresponding to a simulation of 50 minutes of real time observations, in which we allowed up to 5 active storms at a time. Storms' radii were drawn from $N(10, 4^2)$. The performance gain that was obtained by our technique is a direct consequence of that fact that it successfully manages to focus its scan efforts only on the regions where some storm variability has been observed.

In case we allow many more active storms at a time (10 storms, as depicted in Figure 2), then the space quickly becomes crowded with meteorological phenomena. In this case, the best policy is usually very close to simply scanning the whole 360 degrees all the time, since there are phenomena to be observed everywhere and because skipping any one sector could possibly imply not observing several storms. Missing several storms would, of course, have a big negative impact on the scanning cost of the strategy. However, whenever it *is* possible for our approach to avoid scanning one or more sectors, it takes advantage of that fact. Nonetheless, we notice that, as a general trend, adding more storms makes the performance gain provided by our approach smaller, in comparison to sit-and-spin. This observation can be verified in Figure 8.

When, on the other hand, we allow the simulation to have only very few active storms at a time, we observe a similar outcome to the one described last. When just a few storms are active, it is very likely that one of them will be partially visible by at least one given scan action. In other words, it is likely that at some point in time a storm will *not* be completely contained in just one sector. Let s_r the the action of radar r that observes a storm just partially. Since we do not have access to the information of whether the measured variability in the corresponding region was due to a complete (and small) storm moving around, or due to a small portion of a larger, and partially observed storm,




Figure 7: Scanning cost for each of the scanning techniques. The results were averaged over 10 runs, using the same simulated track. The major and minor radii of the storms were sampled from $N(10, 4^2)$, allowing for relatively large storms. A maximum of 5 storms was allowed at a time, and the penalty τ was set to 0.25. We can see that our approach outperforms the sit-and-spin technique because it is able to save scanning resources by not focusing on irrelevant sectors.




Figure 8: Scanning cost for each of the scanning techniques. The results were averaged over 10 runs, using the same simulated track. The major and minor radii of the storms were sampled from $N(10, 4^2)$, allowing for relatively *large* storms. A maximum of 10 storms was allowed at a time, and the penalty τ was set to 0.1. We can see that our approach still performs better than sit-and-spin, but that the advantage is now not so evident, since the best policy in this case is very close to sit-and-spin.

how should we evaluate the action s_r ? The answer to this question will define whether or not we scan the sector; if we believe the variability to have been caused by a small portion of a storm that is mostly contained in another sector, then we should not waste resources by scanning the current sector. Otherwise, we should.

The bias on how we treat these situations is primarily affected by the penalty weight, τ . For small values of τ , there is almost no extra cost for scanning additional regions, and therefore the radar will tend to be optimistic and scan whatever regions contain variability that might indicate the presence of (even small) storms. In that case, the chance of being penalized by missing a storm is decreased, but the quality of measurements becomes lower due to the increase in the radar's rotation speed. In case τ is not small, the radar will tend to be agnostic as to whether small regions of variability should be scanned or not. The specific value of τ to be chosen is directly related to the size of the smallest storm that we want to be able to detect. Larger values of τ might decrease the scanning cost, since the radar would hopefully not scan irrelevant sectors, but could also eventually cause large penalties for missing small storms. This effect is further discussed in Section 7.1.

In Figure 9 we present the costs of scanning a system with very few active storms at a time – just two. In order to minimize the chance of missing storms, we had to set τ to a smaller value. The side-effect of this decision, however, was that because the radar now tended to perform optimistic scans, our strategy got closer to simple sit-and-spin. In fact, Figure 9 shows that even though the performance of our approach was still better than sit-and-spin, the margin of advantage was not so large as in the case where more storms were populating the space.

7.1 The problem with small storms

The problem of finding the value of τ that best balances possible improvements in the scanning cost, while at the same time minimizing the chance of missing storms, depends primarily on the size of the smallest storm that we want to be able to detect.

In case τ is not set appropriately, the scan strategy could end up ignoring possibly small storms, and even though that decision could provide temporary gains in performance, it could also cause big penalties in case actual small storms were mistaken for noise. This effect is clearly depicted in Figure 10, in which the performance of our approach was worse than simple sit-and-spin because our strategy repeatedly mistook small storms for noise.

8 Discussion

We have proposed a new method for determining efficient scanning strategies for adaptive networks of radars, based on spectrally analyzing a sequence of

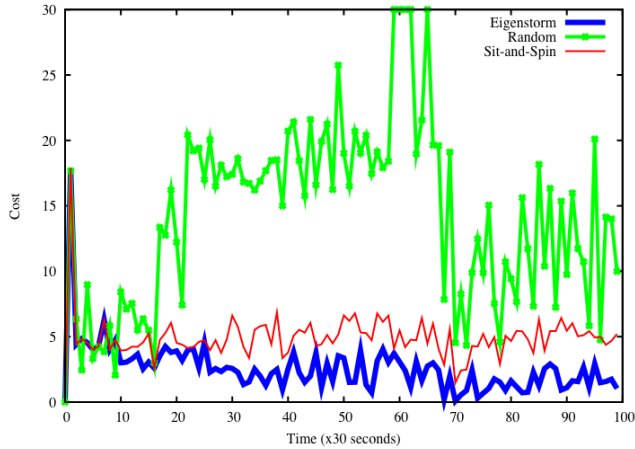


Figure 9: Scanning cost for each of the scanning techniques. The results were averaged over 10 runs, using the same simulated track. The major and minor radii of the storms were sampled from $N(10, 4^2)$, allowing for relatively *large* storms. A maximum of 2 storms was allowed at a time, and the penalty τ was set to 0.1. We can see that our approach still performs better than sit-and-spin, but that the advantage is now not so evident, since we had to set τ to a relatively small value in order for the radar not to confuse partially observed storms with noise. Small values of τ usually make our approach tend towards a sit-and-spin behavior.

radar images. Our approach has some advantages over previous works: it does not assume that the current or maximum number of storms is known; it also does not assume that the storms have been identified or labeled in any way, and neither that information about their centroids is available. We also do not assume anything about the spatial shape of the storms, although in our experiments we have restricted their shapes to only arbitrarily rotated ellipses.

Our approach seems to work well both when compared to the traditionally used sit-and-spin strategy, and also to a baseline random strategy. When there are many storms in the environment, in a way that the whole space is almost completely full with meteorological phenomena, our method successfully avoids scanning irrelevant sectors whenever possible. However, the performance gain in this case, when compared to sit-and-spin, is not very large, because sit-and-spin is indeed the best policy. When the number of storms is intermediate, however, our methods performs clearly better.

Whenever scanning and tracking small storms is important, our method suffers with the difficult problem of correctly setting a good value for the penalty weight, τ . Specifically, it is not trivial to balance the need of observing small storms while at the same time successfully ignoring noise and irrelevant partially observed storms. One possible way of dealing with this limitation could be to increase the size of the action space, since then each radar would be able to decide in a more refined way which regions to scan. However, in case this




Figure 10: Scanning cost for each of the scanning techniques. The results were averaged over 10 runs, using the same simulated track. The major and minor radii of the storms were sampled from $N(5, 2^2)$, allowing for relatively *small* storms. A maximum of 5 storms was allowed at a time, and the penalty τ was set to 0.1. We can see that in the presence of relatively small storms, it is easy to mistake an actual storm for noise or for a partially observed storm mostly located in another sector. Whenever our approach made this mistake, the scanning cost increased due to the penalties for missing storms.

technique were applied to noisy images, it is not clear whether a simple sweep through the whole sector would not be more efficient than constantly having to alternate between scanning very small regions, and skipping others.

We see three possible direct extensions to this work. The first one is related to complementing the method by analyzing additional eigenvectors. In the current work, we have considered just the eigenimage with higher eigenvalue, since it accounts for most of the variance. However, other eigenvectors might encode important features that describe the evolution of the storm at different scales. Another interesting extension could be to analyze the entropy of the eigenvalues. By computing the entropy of the eigenvalues up to the current decision epoch, and by studying how the entropy changes over time, it should be possible to get some evidence about how well-structured the storm cells are. In case the storm cells are not well-structured, either because their spatial shape evolve in complicated ways or because they move through possibly non-linear paths, the entropy will be large. On the other hand, in case the storms move in a predictable way and are more or less static in terms of their shapes, the entropy will be small.

Finally, and maybe more importantly, this work could be extended to consider non-myopic decisions. Currently, all actions are based only on past observations. We could certainly make use of the bases given by PCA in order to find a lower-dimensional representation of the radar images, and then use this new representation as the global state of the system. Based on this new

representation for the state, and assuming that the temporal evolution of the environment might be well approximated by a linear system on the new features, we could predict the next radar images by using Least Squares Regression. This would allow our technique to be extended for non-myopic decisions, which could improve on the current quality of the scanning strategy.

References

- [1] E. S. Chornoboy, A.M. Marlin, and J.P. Morgan. Automated storm tracking for terminal air traffic control. *Lincoln Laboratory Journal*, 7:427–448, 1994.
- [2] D. R. Cox and Valerie Isham. A simple spatial-temporal model of rainfall. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 415(1849):317–328, 1988.
- [3] M. Dixon and G. Wiener. TITAN: Thunderstorm identification, tracking, analysis, and nowcasting: a radar-based methodology. *Journal of Atmospheric and Oceanic Technology*, 10(6):785–797, 1993.
- [4] J. T. Johnson, Pamela L. MacKeen, Arthur Witt, E. De Wayne Mitchell, Gregory J. Stumpf, Michael D. Eilts, and Kevin W. Thomas. The storm cell identification and tracking algorithm: An enhanced WSR-88D algorithm. *Weather and Forecasting*, 13:263–276, 1998.
- [5] J. Kurose, E. Lyons, D. McLaughlin, D. Pepyne, B. Philips, D. Westbrook, and M. Zink. An end-user-responsive sensor network architecture for hazardous weather detection, prediction and response. 2006.
- [6] V. Lakshmanan. *A hierarchical, multiscale texture segmentation algorithm for real-world scenes*. PhD thesis, University of Oklahoma, 2001.
- [7] V. Lakshmanan, R. Rabin, and V. Debrunner. Multiscale storm identification and forecast. *Journal of Atmospheric Research*, pages 367–380, 2003.
- [8] V. Manfredi and J. Kurose. Comparison of myopic and lookahead scan strategies for meteorological radars. Technical Report 2006-62, University of Massachusetts at Amherst, December 2006.
- [9] V. Manfredi and J. Kurose. Scan strategies for adaptive meteorological radars. In *Advances in Neural Information Processing Systems 21*, 2007.
- [10] D. Pepyne. Defining and optimizing utility in NetRad, a collaborative adaptive sensor network for hazardous weather detection. Technical report, CASA, 2006.
- [11] I. Rodrigues-Iturbe and P. Eagleson. Mathematical models of rainstorm events in space and time. *Water Resources Research*, 23(1):181–190, 1987.
- [12] M. M. Wolfson, B. E. Forman, R. G. Hallowell, and M.P. Moore. The growth and decay storm tracker. In *8th Conference on Aviation*, pages 58–62, 1999.

- [13] G. Xu and V. Chandrasekar. Statistical modeling for spatiotemporal radar observations and its applications to nowcasting. In *IEEE Proceedings of the International Conference on Geoscience and Remote Sensing Symposium*, pages 2635–2638, 2006.
- [14] M. Zink, D. Westbrook, S. Abdallah, B. Horling, V. Lakamraju, E. Lyons, V. Manfredi, J. Kurose, and K. Hondl. Meteorological command and control: An end-to-end architecture for a hazardous weather detection sensor network. In *Mobisys Workshop on End-to-End, Sense-and-Respond Systems, Applications, and Services*, 2005.