Radar Signal Processing with Antenna Arrays via Maximum Likelihood*

A. Lee Swindlehurst
Dept. of Elec. & Comp. Engineering
Brigham Young University
Provo, UT 84602
USA
swindle@ee.byu.edu

Petre Stoica
Systems & Control Group
Uppsala University
SE-751 03 Uppsala
SWEDEN
ps@syscon.uu.se

Abstract

This paper considers two different models for coherent radar data in which a single target is observed by an array of antennas in the presence of noise, clutter, and jamming. One model is parameterized in terms of the target signal's amplitude and direction of arrival (DOA), while the other uses an unstructured "spatial signature" instead. The maximum likelihood (ML) estimator and generalized likelihood ratio detection test for the unstructured model are derived, and compared with the standard solutions for the structured DOA-based model. Using the extended invariance principle, it is further shown how asymptotically optimal estimates of the target signal's amplitude, DOA, and doppler can be obtained from the unstructured ML parameter estimates. Several advantages of the two-step unstructured approach are noted, including reduced computational cost and greater robustness to calibration errors.

1. Introduction

A standard problem in radar signal processing involves (1) detecting a target in the presence of noise, ground clutter, and electronic countermeasures (e.g., jamming), and (2) determining the target's position, speed, and type. These latter quantities are usually inferred from the target signal's direction of arrival (DOA), time of arrival, amplitude, and doppler frequency. Even when only a single target is present, this detection and estimation problem is quite daunting due to the large volume of the parameter space that must be searched. For this reason, computational ease and efficiency are of the utmost importance.

Standard solutions to these problems involve the use of classical space-time filters followed by some type of threshold detection test (e.g., see [2, 1, 7] and references therein). Maximum likelihood (ML) approaches have not been extensively considered, mainly because they are perceived to

be too computationally complex, and because there is some reluctance in accepting the required modeling assumptions. Recent work has shown that for a single target in Gaussian interference with unknown spatial covariance, the ML solution can be obtained via a 2-D search over target DOA and doppler, followed by a generalized likelihood ratio test (GLRT) [3, 6].

The primary goal of this paper is to demonstrate that, by suitably reparameterizing the problem, the ML solution can be obtained in a computationally efficient manner. In particular, if the array is uniform and linear, only a single 1-D search for the target doppler frequency is required. A simpler GLRT for detection is also proposed, and the combined detection and estimation procedure that results is shown to be significantly more robust in situations where the array is not precisely calibrated. The key to the proposed techniques is the use of an unstructured "spatial signature" to describe the response of the array to the target signal, rather than the signal amplitude and DOA parameters directly. A similar modeling approach has recently been taken in research related to blind source separation. Here, however, we use the so-called *extended invariance principle* (EXIP) [4] in a two-step approach, where the ML estimates of the unstructured parameters are used to find the target DOA and amplitude. Only the main ideas behind the proposed approach are presented in this paper; for more details, the reader should consult [5].

2. Radar Data Model

Suppose an m -element antenna array transmits a series of N pulses using a certain set of transmit weights. The transmit weights are usually chosen to steer a beam in some nominal direction, say θ_t . The output of each receiving antenna is sampled following each pulse, the time difference between each sampling instant and the pulse transmission corresponding to the distance, or range, between the array and a potential scattering source. The data collected for a given set of transmit weights is referred to as a *coherent processing interval*, or CPI. During a given CPI, a typical radar system collects N samples from each range "bin,"

*This work was supported by the Office of Naval Research under contract N00014-96-1-0934, and by the Swedish National Board for Technical Development (NUTEK).

and interrogates them for the presence/absence of a target. If present, the DOA, doppler frequency, and radar cross section of the detected target are estimated. This procedure is continually repeated for every range bin from numerous CPIs in an attempt to provide constant surveillance over a wide area.

A target present in a particular range bin during some CPI may be modeled as producing the following baseband vector signal (after pulse compression and demodulation) [7, 3, 6]:

$$\mathbf{x}(t) = b_0 \mathbf{a}(\theta_0) e^{j\omega_0 t} + \mathbf{e}(t) \in \mathcal{C}^m, \quad t = 1, \dots, N, \quad (1)$$

where b_0 is the complex amplitude of the signal, ω_0 is the doppler shift due to relative motion of the target, $\mathbf{a}(\theta_0)$ is the response of the array to a plane wave arriving from azimuthal direction θ_0 , and $\mathbf{e}(t)$ represents interference due to clutter, receiver noise, and jamming. We choose to model $\mathbf{e}(t)$ as a stationary, Gaussian random process over the N samples of the CPI, with zero-mean and arbitrary spatial and temporal color. The Gaussian hypothesis is made for the convenience of the ML approach presented later, but it is not critical in that the ML method derived under this assumption provides accurate estimates for many other data/noise distributions as well.

The generic problem considered in this paper may be succinctly stated as follows:

Given a collection of samples $\mathbf{x}(1), \dots, \mathbf{x}(N)$ as in (1), determine whether or not a target is present, and if present, estimate the parameters $\boldsymbol{\eta}_0 = \{\text{Re}\{b_0\}, \text{Im}\{b_0\}, \omega_0, \theta_0, \mathbf{Q}\}$.

While the above problem statement implies that detection precedes estimation, most solutions to the problem operate in reverse: the parameters are estimated as though a target were present, and then a statistical test is performed to compare the likelihood of the noise-only and signal-plus-noise models.

In the discussion above, we have written the array response explicitly in terms of θ_0 . Of course, to estimate θ_0 , the function $\mathbf{a}(\theta)$ must be known for all θ of interest. This requires an experimental or analytical calibration of the array, but such calibration procedures are never entirely accurate, due to various array perturbations. In some cases, these perturbations can cause the actual and calibrated array responses to be quite different, which can lead to a significant loss in detection and estimation performance. Another difficulty with the parameterization of (1) is the non-linear dependence of the data on the parameters θ_0 and ω_0 . As a result, most algorithms for estimating θ_0 and ω_0 require some type of non-linear 2-D search procedure that may be too time consuming to implement.

The above problems can be alleviated through the use of a less structured but still identifiable model. The price paid is of course a loss of performance in situations where the model of (1) holds, but that loss can be compensated for by the approach to be described in Section 3. In particular, consider the following model for the data, where the signal amplitude and array response have been combined into a single unstructured vector term:

$$\mathbf{x}(t) = \boldsymbol{\alpha}_0 e^{j\omega_0 t} + \mathbf{e}(t) \quad t = 1, \dots, N. \quad (2)$$

The vector $\boldsymbol{\alpha}_0$, sometimes referred to as the "spatial signature" of the target, takes the place of $b_0 \mathbf{a}(\theta_0)$, and in this model is simply treated as a set of m unknown complex constants. In the following sections, we discuss the use of (2) together with (1) in developing estimation and detection algorithms that are computationally and statistically efficient, and robust to array imperfections.

3. Maximum Likelihood Estimation

To begin with, we assume that $\mathbf{e}(t)$ is temporally white, with unknown spatial covariance $\mathbf{Q}$. The extension to the case of temporally colored clutter and interference will be considered below. The maximum likelihood solution for the unstructured model is derived under this assumption in [5]. Minimization of the log-likelihood with respect to $\mathbf{Q}$ gives

$$\hat{\mathbf{Q}}_u = \frac{1}{N} \sum_{t=1}^N (\mathbf{x}(t) - \boldsymbol{\alpha} e^{j\omega t}) (\mathbf{x}(t) - \boldsymbol{\alpha} e^{j\omega t})^*, \quad (3)$$

where the subscript u is used to explicitly denote an estimate obtained using the unstructured model. The concentrated criterion resulting from substitution of (3) into the log-likelihood expression may be simplified as follows:

$$\begin{aligned} V(\boldsymbol{\alpha}, \omega) &= \log \left| \frac{1}{N} \sum_{t=1}^N (\mathbf{x}(t) - \boldsymbol{\alpha} e^{j\omega t}) (\mathbf{x}(t) - \boldsymbol{\alpha} e^{j\omega t})^* \right| \\ &= \log \left| (\boldsymbol{\alpha} - \mathbf{y}(\omega)) (\boldsymbol{\alpha} - \mathbf{y}(\omega))^* + \hat{\mathbf{R}} - \mathbf{y}(\omega) \mathbf{y}^*(\omega) \right| \end{aligned} \quad (4)$$

where $\mathbf{y}(\omega)$ is the Fourier transform of $\mathbf{x}(t)$, and $\hat{\mathbf{R}}$ is the sample covariance of $\mathbf{x}(t)$. From (4), it is clear that the solution for $\boldsymbol{\alpha}$ is simply

$$\hat{\boldsymbol{\alpha}}_u = \mathbf{y}(\omega). \quad (5)$$

Some further simple calculations [5] show that the unstructured ML estimate of ω is given by

$$\hat{\omega}_u = \arg \max_{\omega} \mathbf{y}^*(\omega) \hat{\mathbf{R}}^{-1} \mathbf{y}(\omega). \quad (6)$$

Equation (6) shows that the maximum likelihood frequency estimate for this problem is calculated by first spatially pre-whitening the data, and then finding the largest peak in the magnitude of the Fourier transform averaged over the m channels. The derivation above also implies that in principle, array calibration data is not needed in detecting the presence or absence of a target in the radar data set. Given the estimate $\hat{\omega}_u$ from (6), one could develop a statistical test for detection based on the size of $\|\hat{\mathbf{R}}^{-\frac{1}{2}} \hat{\boldsymbol{\alpha}}_u\| = \|\hat{\mathbf{R}}^{-\frac{1}{2}} \mathbf{y}(\hat{\omega}_u)\|$. If the array is imprecisely calibrated, such a test may outperform the GLRT based on estimates of the target DOA. This issue is further studied in Section 4 and the simulations.

In cases where the array is perfectly calibrated, one might assume that the above ML solutions for $\mathbf{Q}$ and ω will not be as accurate as those obtained with a model that uses the calibration data. As outlined in the following theorem, this is not the case, at least for "large enough" N :

Theorem 1 Given $\hat{\mathbf{Q}}_u$, $\hat{\alpha}_u$, and $\hat{\omega}_u$ obtained from (3), (5), and (6), asymptotically (in N) efficient estimates of the radar target parameters in (1) can be found as follows:

$$\hat{\mathbf{Q}}_e = \hat{\mathbf{Q}}_u \quad (7)$$

$$\hat{\omega}_e = \hat{\omega}_u \quad (8)$$

$$\hat{b}_e = \frac{\mathbf{a}^*(\hat{\theta}_e) \hat{\mathbf{Q}}_u^{-1} \hat{\alpha}_u}{\mathbf{a}^*(\hat{\theta}_e) \hat{\mathbf{Q}}_u^{-1} \mathbf{a}(\hat{\theta}_e)} \quad (9)$$

$$\hat{\theta}_e = \arg \max_{\theta} \frac{|\mathbf{a}^*(\theta) \hat{\mathbf{Q}}_u^{-1} \hat{\alpha}_u|^2}{\mathbf{a}^*(\theta) \hat{\mathbf{Q}}_u^{-1} \mathbf{a}(\theta)}. \quad (10)$$

Proof: The proof uses the EXIP [4], and is given in [5].

The above theorem also demonstrates how the estimated spatial signature may be used to determine the target amplitude and DOA, and that these parameter estimates are asymptotically as accurate as those that would have been obtained using the fully structured model with array calibration. As detailed in [5], there are two main advantages to using the two-step approach implied by Theorem 1. The first is that two 1-D searches are required instead of a 2-D search. If the (nominal) array is uniform and linear, then the 1-D search for θ can be further eliminated by using a rooting technique. The second advantage is that, as demonstrated in Section 5, the criterion in (10) can be easily regularized to reduce sensitivity to calibration errors. It is not clear how such a regularization could be applied in the structured model case.

Since $\mathbf{y}(\omega)$ is just the DFT of $\mathbf{x}(t)$, the criterion of (6) may be efficiently evaluated at the frequencies $\omega_k = 2\pi k/N$ for $k = 0, \dots, N-1$ using a bank of m FFTs. The value of ω_k on the FFT grid that maximizes (6) can serve as an initial estimate for a local search using the computationally efficient zoom or chirp FFT algorithms. In Section 4 we will see that, in any case, the development of the detection criterion relies on using (6) evaluated on a grid with spacing equal to $2\pi/N$.

Extension to Non-White Clutter

When the interference term is dominated by clutter, a common approach [7] is to assume that an estimate of both the temporal and spatial statistics of $\mathbf{e}(t)$ is available from data in adjacent target-free range bins and CPIs. However, forming and inverting an estimate of the full space-time covariance of $\mathbf{e}(t)$ typically involves hundreds of correlation measurements, a large amount of secondary data, and significant computation. In this paper, we propose a computationally efficient alternative. In particular, we suggest a two-stage operation: (1) estimating a vector autoregressive (VAR) model for the temporal variation of the clutter using secondary data from adjacent range and CPI cells, and (2) using the estimated VAR model to pre-whiten the data prior to processing. The computational advantage of this approach stems from the fact that the model for the temporal statistics of the clutter are embedded in a relatively small number of AR parameters, rather than the comparatively large number of temporal covariance lags that may be necessary for

the methods of [7]. Additionally, the proposed algorithms estimate the spatial covariance $\mathbf{Q}$ along with the target parameters, and thus no prior estimate of $\mathbf{Q}$ is required.

Suppose we have at our disposal a VAR filter of the form

$$\mathcal{H}(z^{-1}) = \mathbf{I} + \sum_{l=1}^L \mathbf{H}_l z^{-l} \quad (11)$$

such that $\mathcal{H}(z^{-1})\mathbf{e}(t) = \epsilon(t)$ is temporally white. The estimation of $\mathcal{H}(z^{-1})$ from the secondary data is discussed in [5]. If we pre-process the array data with $\mathcal{H}(z^{-1})$ prior to detection and estimation of the target parameters, the filtered array output will be of the form

$$\mathbf{x}_f(t) = \mathcal{H}(z^{-1})\mathbf{x}(t) = b_0 \mathcal{H}(e^{-j\omega_0})\mathbf{a}(\theta_0)e^{j\omega_0 t} + \epsilon(t). \quad (12)$$

Defining the filtered spatial signature to be $\alpha_f = b_0 \mathcal{H}(e^{-j\omega_0})\mathbf{a}(\theta_0)$, we have

$$\mathbf{x}_f(t) = \alpha_f e^{j\omega_0 t} + \epsilon(t), \quad (13)$$

and we see that the unstructured ML estimator can be directly applied to $\mathbf{x}_f(t)$ with no change. Theorem 1 can then be applied to obtain the following estimates of b and θ :

$$\hat{b}_f = \frac{\mathbf{a}^*(\hat{\theta}_f) \mathcal{H}^*(e^{-j\hat{\omega}_f}) \hat{\mathbf{Q}}_f^{-1} \hat{\alpha}_f}{\mathbf{a}^*(\hat{\theta}_f) \mathcal{H}^*(e^{-j\hat{\omega}_f}) \hat{\mathbf{Q}}_f^{-1} \mathcal{H}(e^{-j\hat{\omega}_f}) \mathbf{a}(\hat{\theta}_f)} \quad (14)$$

$$\hat{\theta}_f = \arg \max_{\theta} \frac{|\mathbf{a}^*(\theta) \mathcal{H}^*(e^{-j\hat{\omega}_f}) \hat{\mathbf{Q}}_f^{-1} \hat{\alpha}_f|^2}{\mathbf{a}^*(\theta) \mathcal{H}^*(e^{-j\hat{\omega}_f}) \hat{\mathbf{Q}}_f^{-1} \mathcal{H}(e^{-j\hat{\omega}_f}) \mathbf{a}(\theta)} \quad (15)$$

where $\hat{\mathbf{Q}}_f$ is the estimate of the spatial covariance of $\epsilon(t)$ obtained from (3) using the filtered data and the estimates $\hat{\alpha}_f$ and $\hat{\omega}_f$.

4. Detection

In this section, we discuss the binary hypothesis testing problem that arises under the unstructured model when determining whether or not a target is present in a given data set. In particular, it is necessary to decide which of the following two hypotheses is valid for $\mathbf{x}(t)$, $t = 1, \dots, N$:

$$\begin{aligned} H_0 : \quad & \mathbf{x}(t) = \mathbf{e}(t) \\ H_1 : \quad & \mathbf{x}(t) = \alpha_0 e^{j\omega_0 t} + \mathbf{e}(t). \end{aligned}$$

The most common approach to solving this type of problem involves first estimating the signal parameters as if a target were present, and then comparing the likelihood of H_1 (with the true parameters replaced by estimates) to that of H_0 . This is referred to as a generalized likelihood ratio test (GLRT).

In [5], it is shown that to achieve a desired probability of false alarm P_f , the following detection procedure should be used:

1. Compute a threshold λ_T by solving the equation

$$F(\lambda_T) = (1 - P_f)^{\frac{1}{N}},$$

where $F(\cdot)$ is the cumulative distribution function for a chi-squared random variable with $2m$ degrees of freedom:

$$F(\lambda) = P\{\rho \leq \lambda \mid \rho \sim \chi^2(2m)\}. \quad (16)$$

2. Compute the sample covariance $\hat{\mathbf{R}}$, and calculate the DFT $\mathbf{y}(\omega_k)$ at the frequencies $\omega_k = 2\pi k/N$ for $k = 0, \dots, N-1$.
3. Find the DFT frequency ω_k that maximizes the vector periodogram $\mathbf{y}^*(\omega)\hat{\mathbf{R}}^{-1}\mathbf{y}(\omega)$; call this frequency $\hat{\omega}_u$.
4. Decide that H_1 is true if

$$\lambda_u = -2N \log \left(1 - \mathbf{y}^*(\hat{\omega}_u)\hat{\mathbf{R}}^{-1}\mathbf{y}(\hat{\omega}_u) \right) > \lambda_T,$$

otherwise assume that H_0 holds.

The GLRT for the unstructured model does not require any calibration information, and thus is unaffected by the presence of array perturbations. As shown in the next section, this can provide a considerable performance advantage even when the size of the calibration errors is not large. Even in the absence of such perturbations, the unstructured GLRT may be useful in providing a more rapid "initial scan" of the environment prior to application of the structured model. For example, one could apply the unstructured GLRT with a relatively high false alarm rate, and then use the complete model to process data sets that showed an initial detection. In the examples presented below, it will be seen that, even with perfect calibration data, the use of the unstructured GLRT suffers a loss of only about 2 dB.

5. Simulation Results

We present below the results of several simulations comparing the detection and estimation performance of the structured and unstructured ML estimators. The following parameters were common to all of the simulations:

- The array was composed of $m = 12$ nominally identical, uniformly spaced, unit-gain antennas separated by one-half wavelength.
- If present, the calibration errors were generated as Gaussian random variables independent and identically distributed from antenna to antenna. The term *calibration error* is used below to refer to the standard deviation of the perturbation at each sensor.
- The target was located at $\theta_0 = 0^\circ$ (broadside), with amplitude $b_0 = 1$ and doppler shift $\omega_0 = 2\pi/10$.

Figure 1. DOA Estimation Performance and CRB versus Number of Snapshots, 0 dB SNR

- The spatial covariance of the interference was set at $\mathbf{Q} = \sigma^2 \tilde{\mathbf{Q}}$ for various σ , where $\tilde{\mathbf{Q}}$ is a randomly chosen positive definite matrix scaled and regularized so that its determinant is unity. The same $\tilde{\mathbf{Q}}$ was used in all of the simulations; its maximum and minimum eigenvalues were 2.36 and 0.095, corresponding to a condition number of about 25. The SNR was defined to be $-20 \log \sigma$ dB, as the target amplitude and nominal antenna gains were both unity. For simplicity, $\mathbf{e}(t)$ was chosen to be temporally white.
- The probability of false alarm was set to $P_f = 10^{-6}$ for the detection tests.

In what follows, the term "unstructured GLRT" (or simply UGLRT) will be used to refer to the GLRT obtained for the unstructured array model. UML will refer to the process of estimating the target parameters via the unstructured ML algorithm and the EXIP. A corresponding meaning is to be attached to the acronyms SGLRT and SML.

The purpose of the first example we consider is to "validate" the asymptotic arguments presented earlier. In particular, we provide insight into the size of N required to achieve the asymptotic regime. The results presented here are indicative of numerous simulation runs. The empirical behavior of the SML and UML DOA estimates is compared with the CRB in Figure 1, which plots RMS error versus N for a case with 0 dB SNR and an ideal array (no calibration errors). Each data point on the plots is calculated from the results of 1000 independent experiments. It is clear from the figures that there is essentially no difference between the quality of the SML and UML estimates for all values of N . Thus, for these cases at least, the asymptotic equivalence of the two methods is apparent with very little data. Somewhat more data is required for their performance to reach the CRB, but there is little excess error for $N \geq 30$.

Figure 2. Probability of Detection versus SNR for various Pointing Errors.

Figure 2 compares the detection probabilities of the SGLRT and UGLRT at various SNR (500 trials at each SNR) for a case involving an array with a calibration error of 0.1 (note that the unperturbed array is assumed to have unit gain sensors). A different perturbation was applied to the array at each trial, so the curves in Figure 2 represent the detection probability averaged over uniform linear arrays with this level of calibration error. The three curves for the SGLRT show its performance for "pointing" errors (DOA estimation errors) of 0, 1, and 2 degrees. A pointing error is a result of an error in the DOA estimate, due (1) to interference and array perturbations, or (2) evaluation of the SGLRT at $\theta = \theta_t$, which is common in practice. Figure 2 demonstrates that the SGLRT is very sensitive to such errors. No such difficulty is observed with the UGLRT of course since it does not use DOA information. Even with no pointing error, we see that the performance of the SGLRT deteriorates at high SNR, since the presence of the calibration errors reduces the relative height of the peak value of the test statistic λ evaluated over ω . At higher SNR, there is less of a chance that the noise will artificially push the peak over the threshold. This deterioration is more evident with pointing errors, since the peak has been reduced even further.

We study next the estimation performance of SML and UML for the same scenario described above. The only difference in this example is that the calibration error for the array was increased to 0.2, and a strong jammer was added to the data. The jammer was located at a DOA of 5° , and the jammer-to-signal power ratio (JSR) was held fixed at 20 dB. The jammer was assumed to broadcast temporally white noise, so its net statistical effect is to make a large rank one contribution to $\mathbf{Q}$. The target parameters were estimated by maximizing the SML and UML criteria on a grid with spacing $2\pi/N_F$ in ω . The frequency resolution was increased by choosing $N_F > N$ and zero-padding the data to the required length. In each case, N_F was chosen large

Figure 3. Doppler Estimation Performance versus SNR, Strong Jammer Present.

enough so that the error due to the grid quantization would be several times smaller than the error predicted by the CRB. Figure 3 shows the RMS doppler estimation error for this example. The UML method is several times more accurate than SML at high SNR. The presence of the calibration error reduces the peak of the 2-D SML criterion at the true θ and ω , and produces a local maximum that is far removed from them. The UML technique does not suffer from this problem since at high SNR it is able to accurately estimate ω without being affected by the array perturbation. With an accurate estimate of ω , the error in the UML DOA estimate is correspondingly reduced. The difference between the methods is smaller at low SNR, but not insignificant.

References

- [1] S. Haykin, J. Litva, and T. Shepherd, editors. *Radar Array Processing*. Springer-Verlag, Berlin, 1993.
- [2] S. Haykin and A. Steinhardt, editors. *Adaptive Radar Detection and Estimation*. John Wiley & Sons, New York, NY, 1992.
- [3] A. Steinhardt. "Adaptive Multisensor Detection and Estimation". In S. Haykin and A. Steinhardt, editors, *Adaptive Radar Detection and Estimation*, pages 91--160, New York, NY, 1992. John Wiley & Sons, Inc.
- [4] P. Stoica and T. Soderstrom. "On Reparametrization of Loss Functions Used in Estimation and the Invariance Principle". *Signal Processing*, 17:383--387, August 1989.
- [5] A. Swindlehurst and P. Stoica. "Maximum Likelihood Methods in Radar Array Signal Processing". *Proc. of the IEEE (to appear)*, 1998.
- [6] M. Viberg, B. Ottersten, and O. Erikmat. "A Comparison of Model-Based Detection and Adaptive Sidelobe Cancelling for Radar Array Processing". Technical Report CTH-TE-13, Chalmers University of Technology, April 1994.
- [7] J. Ward. "Space-Time Adaptive Processing for Airborne Radar". Technical Report TR-1015, MIT Lincoln Labs, Dec. 1994.