Spring 2009, Aaron Lanterman

ECE 6279: Spatial Array Processing Homework 9

Due: Wed 4/22/09 at the *start* of class - homeworks turned in later in the hour may be penalized at my discretion.

You are welcome to discuss approaches to the problems and solutions to difficulties you encounter with one another and with others outside the class. You can and should learn from each other as much as, and even more than, you learn from me. However, your solutions should be your own work and should be written up by yourself; feel free to discuss things, but don't be looking at someone else's paper when you are writing your solution. It's too easy to freeload that way and not learn anything. See the class website for more guidelines.

Looking at solutions to homeworks and quizzes of previous offerings of ECE6279 is expressly forbidden.

If you cannot make a class, please make arrangements to get your homework to me ahead of time.

1 Required Problems

1. In Lecture 30, we will derive an expression for the Cramér-Rao bound on the parameter γ under the "deterministic" Gaussian model, where γ is the "electrical angle" specified by $\gamma = (2\pi d/\lambda)\sin\phi$ (where $\theta = 0$) for the case of an equally-spaced linear array of M elements with spacing d lying along the x-axis. Here is the exact result here:

$$CRB(\gamma) = \frac{6}{L(SNR)} \frac{1}{M(M^2-1)}$$

   SNR denotes the signal power to noise power ratio. The expression acts as you might expect; achieving a better SNR or adding more elements gives you a lower CRB. Interestingly, that the CRB on the electrical angle is not actually a function of the true electrical angle.

   In this problem, we'll consider an alternate formulation where we take $\phi = \pi/2$, and consider the azimuth angle θ. Hence, in this problem, we will let $\gamma = (2\pi d/\lambda)\cos\theta$.

   (a) We'd like a CR bound in terms of θ instead of γ. At the end of the Lecture 31, I showed you two formulas for computing the CR bound on a parameter θ (I happened to use ϕ in the example in lecture, but the same idea applies) that's defined by a functional mapping $\theta = f(\gamma)$, given the CR bound on γ. Using these formulas, compute the CR bound on θ using the CR bound on γ given above for an M-element linear array. Try both formulas, and make sure you get the same answer with each!

   (b) What happens to the CRB on θ as $\theta \to 0°$?

(c) Suppose I told you that there are estimators for θ *that beat the CRB computed in (b).* (If you think about your answer for part (b) a bit, and realize that the greatest error you can ever have in estimating an angle is 180°, you'll see that this must be true.) What does this tell you about those estimators? (I'm looking for a short answer here - something like two or three words.)

2. The extremely simple formula for the CRB used above required two simplifying assumptions: a linear array assumption and a single-target assumption. In this problem, we will keep using a linear array, but we will explore the CRB for a two-target problem. To do this more complicated computation, you will need to pull out MATLAB.

   Consider a linear array of 32 equally spaced elements (with half-wavelength spacing) along the x-axis, centered around the origin (although, as usual, it doesn't really matter where it is centered for this problem.) We will characterize things in terms of the electrical angle γ. Recall the CRB on γ is constant in the single-target case; we will see if this is still true in the two-target case. Use MATLAB to do the hard work; just code up the Fisher information matrix formula on slide 9, without trying to do any serious symbolic simplification on your own, or else you will go insane.

   Assume we only have one snapshot, and that $\sigma^2 = 1$. Let the true γ_2 for source 2 correspond to a real angle of $\phi_2 = \pi/4$. Plot the square root of the CRB on γ_1 for source 1 vs. the true γ_1 for γ_1 corresponding to a sweep of ϕ_1 between $-\pi/2 \leq \phi \leq \pi/2$. This will be the upper-left entry of the inverse of the Fisher information matrix. Make plots for two cases: $s_1 = 1$ and $s_2 = 0.2$ (s_2 much weaker), and $s_1 = 1$ and $s_2 = 1$ (both the same strength).

   Comment on interesting effects you observe. How does the CRB change with the strength of s_2? How does it change as γ_1 approaches γ_2? (Note: You may need to tweak the plot if something numerically strange goes on at certain points. Or, you might not.)

3. Let's pull out our 9-element cross array one last time, the one you used in many previous homeworks. So as usual, you can probably take code you already have and just tweak it to do this problem.

   Let's use 150 snapshots to form our covariance estimate. (There will still be some variability, so you may want to run your experiments a few times to watch overall trends and make sure you didn't get just one "lucky" case. You need only turn in one set of plots per subpart, though.)

   Let's have two sources, one at $\phi = 45°$ and $\theta = 30°$, and another $\phi = 45°$ and $\theta = 60°$. Let's let each have power 1/4 (so you'll want to multiply your steering vectors by 1/2 when making the "seen by array" variables), and let's make the noise have power 2. Notice we're really pushing our algorithms hard now; the signals are pretty weak compared to the noise.

   Let's experiment with the model order estimation procedures we discussed in Lecture 32. As usual, please provide printouts of your code. (You need not provide separate printouts for minor tweaks - i.e. when you change the noise level, don't bother printing out a whole other page just to show the code with that one change).

(a) On the same graph, plot (1) the loglikelihood given in class for the model order estimation problem (the expression that has the geometric mean in the numerator and the artihmetic mean in the denominator), (2) the penalized likelihood using the AIC penalty, (3) the penalized likelihood using the MDL penalty, all as a function of N_s, the assumed number of sources. For what N_s do the different curves take their maximum? Did AIC estimate the correct model order? What about MDL?

(b) If your model order estimators got the correct model order in part (a), slowly jack up the noise power from 2 until you start regularly seeing errors in the model order estimates. At what noise power do the criteria start to have problems? (There isn't an exact number here that's the "right" answer. Just experiment until you get something in the "ballpark.") Give a plot like in part (a) for one of these cases where both AIC and MDL get the model order wrong. Do they overestimate or underestimate the number of sources?