PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Philosophy for optical filter optimization

Richard Juday

Richard D. Juday, "Philosophy for optical filter optimization," Proc. SPIE 10296, 1999 Euro-American Workshop Optoelectronic Information Processing: A Critical Review, 102960A (2 June 1999); doi: 10.1117/12.365907

SPIE.

Event: Euro-American Workshop on Optoelectronic Information Processing, 1999, Colmar, France

A philosophy for optical filter optimization

Richard D. Juday

Mail Code ER5, NASA Johnson Space Center, Houston TX 77058

ABSTRACT

I espouse an approach to optimizing filters on arbitrary SLMs that is different from some approaches to be found in the current literature. The method begins with selecting a metric by which to judge the operation of the correlator. The metric is to be based on observable features, not on inaccessible internal states of the correlator. That metric is optimized by choice of the implemented filter — significantly, the optimization is done under restriction to realizable filter values, not under other restrictions such as to unit filter energy or matching phase inappropriately. A necessary condition of optimality is that the partial derivative of the metric with respect to allowed changes in the realizable filter value is zero. The ramifications of these precepts are examined and examples are shown.

Keywords: Optical correlator, optimal filter.

1. INTRODUCTION

A great deal of work for the optimization of optical correlation filters is done without, in my opinion, two important considerations. One is that a filter should be optimized in terms of only such things we are able to observe physically about the performance of the correlator with the filter mounted in it. The second is that the resulting filter will perforce be implemented on the coding domain of the filter spatial light modulator. In these substandard approaches, a filter — ordinarily complex and having no particular limitation on its values — is often computed that, if it were realizable, would in some fashion be optimal for the problem it is directed towards. Then one mapping method or another is used to stick the computed fully complex filter onto the coding domain. A sterling example is that the complex matched filter (CMF) optimizes signal to noise ratio in the presence of additive noise whose power spectral density is known, and then one may cause the phase of the resulting signal to be maintained in realizing the filter. There might be some justification for any step in that process, but the combined effect is not the best that can be obtained with a more global view. Ordinarily we are not working in an environment in which the filter SLM permits access to an area of the unit disk, but instead with a discrete-set or curvilinear subset of the unit disk. So even by scaling and phasing the CMF, we are not able to realize it fully by jamming it, shape unaltered, into the coding domain. (We refer to that quality that is unaltered by multiplication by a complex scalar as the filter's "shape".)

In the second step of the traditional approach the computed complex filter is converted to coding domain values by some method or other. For a good while the watchword was to maintain the phase of the computed ideal complex filter¹, perhaps adding a constant phase offset or a spectral region in which the filter was set to zero². The result of maintaining phase during this step, called a phase-only filter (POF), has enjoyed considerable success, and there are circumstances of metric objective and available filter values in which it is demonstrably optimal to maintain the phase while setting the magnitude to its passive maximum of unity. I must add, however, that I am frankly

puzzled as to exactly what a “phase-only filter” actually is. Any filter that is built on a modulator whose only variation is phase is, one would suppose, a POF, and yet the literature is not refulgent with POFs computed except by matching the filter’s phase to that of a single reference object. Consider for a moment how to build a filter on an SLM that has coupled phase and amplitude behavior. I acknowledge here the argument that a matched-phase filter guarantees that the maximum of the correlation intensity is exactly centered in the correlation plane. For complicated reference objects, however, the decentration produced by maximizing intensity rather than matching phase is very small³.

In a move away from simply matching phase, Juday^{4,5}, then Farn and Goodman⁶, introduced maximum projection onto a realizable region as a way of maximizing correlation intensity for coupled modulators, as the shape of an ideal filter is converted to the realizable domain. Juday⁷ generalized that projection for single reference images and advanced metrics to include the minimum Euclidean distance (MED) projection for ideal filter values that were not “at infinity”, as the maximum-projection algorithm can be described. Laude and Réfrégier⁸ used the MED process as part of a realizable filter optimization that used a Lagrangian step in optimizing among multiple competing objectives for a metric. That is, they inserted a step into the process in which they compute an optimal filter [shape] that is constrained to unit energy while optimizing the tradeoff between differing metric goals. The unit-energy constrained filter is then mapped by MED onto the realizable domain. I contend that such a constraint to unit energy is at best unnecessary in achieving the optimal tradeoff. At worst it could deliver a sub-optimal filter by not delivering, as does my philosophical method, a confirmation that each frequency has its locally optimal value chosen in view of the combined effect of all other frequencies. Nowadays one finds a good deal of work in the literature in which filters are computed that are optimal in some sense, but the necessity of optical realizability is not taken into account. To cite some examples, a recent issue of Optical Engineering contains several papers^{9,10,11,12,13,14,15,16} on correlation pattern recognition in which there is a great deal of mathematics expended on filter computation without a word on optical realization of the filters. The one paper¹⁷ in the same issue that does take realizability into account assumes a phase-only constraint in projection onto a convex set – and its phase-only SLM with unit magnitude and 2π phase range is at least somewhat idealistic. More recently one finds the idea¹⁸ that a correlation filter (CF) is known, having originated in a tradeoff of noise robustness against sharpness of the correlation peak, without consideration given at that stage to correlation intensity. Then its values are mapped to the SLM’s available values in order to achieve the maximum intensity of the output correlation peak. I aver without detailed proof that this sequential approach will not generally achieve the optimal tradeoff between noise robustness and peak sharpness, since the shape of the filter that *did* best achieve that tradeoff is destroyed in the projection onto realizable values, and there is no analytical justification presented that the projection in any way maintains the optimality of that tradeoff. Nor, of course, will that approach achieve largest correlation intensity for the object, since we have inserted the tradeoff into the process sequence. Further, and again without detailed proof, I aver that the process, although it may well produce acceptable results, will not generally produce as satisfying a tradeoff among all three metric elements – sharpness, robustness, and intensity – as if they were simultaneously considered while choosing among realizable filter values. This simultaneity is addressed with the philosophical approach described here and justified in greater detail elsewhere^{7,19,20,21}.

Let me expand on the issue of metric observability in the philosophy I espouse, and contrast it against some other precepts. This is a philosophical and ruminative paper, so I

do not present laboratory results to back up each claim and statement I make. If you think the arguments might make sense, feel free to try them out; if not, then don't.

It makes no difference what happens inside a correlator if I am not observing it – and ordinarily I am not. Instead I observe only what is detected at the end of the process. An optical correlator is a quantum mechanical implement, and it is only after the energy has propagated through it and been detected that I have the information processed and ready to do further work with. Hence, I claim, it makes little sense to optimize such an internal state as the so-called optical efficiency (the fraction having as its denominator the light energy originating in the reference signal that has made it through the filter, and as its numerator the amount of that light that is directed into the central correlation spot). If I want a bright correlation spot in response to a given reference image, I should optimize exactly the correlation response. I have the tools to do that, under the philosophy described here. I admit the charm of the idea that optimizing the optical efficiency should produce spike-like correlations, but in laboratory practice²² I observe that a filter intended to optimize this version of optical efficiency (an inverse filter, if implemented in fully complex fashion) is very susceptible to noise. I observe further that correlations that have been optimized for their observable character such as intensity or signal to noise ratio are, in any practical sense, as narrow as those produced by a filter intended to optimize the optical efficiency. So, consonant with the philosophy and bolstered by my laboratory experience, I eschew filters designed to optimize the optical efficiency as defined above.

To be sure, often the general algorithm I discuss here will come down to computing a fully complex filter and then converting its values algorithmically into the coding domain. The philosophy I describe here, however, analytically develops very specific guidelines for choosing the shape of the ideal complex filter that is then represented in the coding domain by a projection process. In current literature those guidelines are not always followed in either the computation of the ideal complex filter or its conversion to realizable values. The shape is determined by the necessary condition of optimality to be described in the next Section, as it sets conditions on the values of search parameters and specifies the connection between the optimal realizable filter and the search parameters. The general algorithm is efficient in that it reduces the set of search parameters to a minimal set, and it is complete in that it provides confirmation that the result of the search is consistent with the necessary statement of optimality. Following the philosophy described here, our laboratory has been able to optimize some meaningful and complicated correlation pattern recognition metrics on arbitrary SLMs, and I shall show those results.

As a final introductory remark, I sometimes hear statements like “It is too complicated to do all these computations to optimize a filter, compared with just taking a signal, transforming it, extracting the phase, and matching it on my SLM.” I agree it is certainly more work to think about what observable metrics you would wish to optimize, to characterize your SLM carefully, and to do the fully optimizing computations. So we come down to a comparison of philosophies, and only you can judge what is correct for you: the ease and convenience of a rough demonstration, or having your system perform as well as it can.

2. A NECESSARY CONDITION OF OPTIMALITY

Regard the correlation process as producing an observable output that we shall adjudge according to some metric J . We carefully mark the distinction between a *filter*, a *filter*

drive, and a *filter SLM*. A filter drive is the set of controls (e.g. an array of bytes or array of voltages) that is applied to the filter SLM. The filter is the array of complex action (absorption and retardation expressed as a complex number) that the filter SLM wreaks on the transformed input image. Our control of the filter SLM is typically through a voltage (usually a physical continuum, even if addressed with one to eight bits) for each pixel. An SLM's operating curve is the range of complex values traced out as the control voltage runs through its own domain. (We often speak of a filter encoding domain, but at this stage the control is the domain, and the range is the set of filter values resulting from the drive values.) In terms of an input/output relationship, the metric is a function of the drive, and what happens within the correlator is largely unobservable. With this viewpoint, we can regard our problem as this: Optimize the metric by adjustment of the drive values. We can observe various things about the correlation intensity; we measure its two-dimensional values, and by feeding the correlator several instantiations of noise in the input – and several instantiations of the input object, too – we can get an idea of the variance of correlation intensity from those sources. Ideally we would optimize the ratio of correlation value at the location corresponding to the reference object to the largest correlation value appearing anywhere else in the plane. Unfortunately we can not (yet) handle that magnificent metric; we do not have an analytic expression for the value of such a secondary peak. Focusing attention onto the central correlation value has worked well, however, and all the work described here is done on that basis.

We suppose that there exists an optimal filter. Then the following must be true (if the derivative of filter value as a function of drive exists)... infinitesimal readjustments of any of the drive values produce zero change in the metric, to first order. (For suppose to the contrary that we could change a drive value and produce a change in the metric. Then one direction or the other of change in the real scalar drive value would increase the metric, and the filter we posited as optimal did not actually produce the highest value of the metric.) (This general argument applies except where the derivative of complex filter value as a function of drive does not exist, including end points of the operating curve. These conditions are taken care of in the discussion of metric gradient in Sect. 13 of Ref. 20.) That is, the partial derivative of the metric with respect to the drive is zero at all frequencies. From this simple statement that is a necessary condition of optimality we can infer the whole nature of a filter and optimize a widely ranging set of metrics on arbitrary SLMs.

We adopt the following nomenclature. The m -th frequency's value of the transformed signal is $S_m = A_m \exp(j\phi_m)$ and of the filter, $H_m = M_m \exp(j\theta_m)$. The filter is a function of applied drive, v . The central correlation field is $D = B \exp(j\beta) = \sum_k H_k S_k$ where the indicated sum is over all frequencies (we use the one-dimensional notation), and the correlation intensity is $I = B^2$. The power spectral density of the input noise is P_{nm} at the m -th frequency, and the input noise's contribution to variance is $\sigma_n^2 = \sum_k |H_k|^2 P_{nk} = \sum_k M_k^2 P_{nk}$. The metric, J , is some function of input object S , filter H , noise σ_n^2 , and possibly other quantities.

In addition to basing metrics on only those things that one can observe physically, my philosophy for filter optimization is summarized in the necessary condition for optimality of the filter. For each frequency (indexed by m),

$$\text{The philosophy: } \frac{\partial J}{\partial v_m} = 0 \quad (1)$$

From this simple-looking equation we derive a rich set of specifications for optimizing the filter. We can show how a frequency in the filter operates in conjunction with the filter values at all the other frequencies, and while doing that we can devise a reduced-dimensionality set of parameters to search for the optimal filter. Additionally, we have confirmation when finished with the search that the metric-optimizing parameter set causes Eq. (1) to hold. Although strictly speaking the process can be caught by local maxima in the function of metric vs. parameter set, we note that the process works well in practice, and we do not seem to get snarled in local maxima.

3. OPTIMIZATION OF SOME EXAMPLE METRICS

I shall now show three examples of metric optimization according to this philosophy. The first is signal to noise ratio when the noise is modeled as additive at the input plane and the noise is known only to its power spectral density. It optimizes the detection of a single reference object in a noisy background. The second is the Rayleigh quotient that optimizes the ratio of correlation values for two (possibly quite similar) objects at the center of the correlation plane. The third is the Fisher ratio for a group of training objects to be divided by correlation intensity into the "accept" and "reject" classes.

3.1 Signal to noise ratio

The signal to noise ratio is formulated as follows.

$$\text{SNR} = \frac{\left| \sum_k H_k S_k \right|^2}{\sigma_d^2 + \sum_k M_k^2 P_{nk}} \quad (2)$$

in which we see a term in the denominator, σ_d^2 , that represents noise in the correlation process that is irreducible by action of the filter. It is an experimentally derived term, and it has a profound effect on what the optimal filter is. In essence, its presence in the optimization process prevents one from building a very clean correlation signal that might come through the correlator with only microvolt correlation values, when one is dealing, say, with some millivolts of noise. If there is millivolt noise, then optimizing SNR will tend to produce correlation signals that dominate the noise of millivolts but at the expense of altering the shape of the realized filter.²³ Applying the necessary condition, Eq. (1), to the formula for SNR, one finds⁷ that the optimization comprises specifically and analytically the following steps.

1. Compute the ideal filter value at the m -th frequency,

$$H_m^I = \left\{ \frac{\sigma_d^2 + \sum_k M_k^2 P_{nk}}{\left| \sum_k H_k S_k \right|} \exp \left(j \text{Arg} \sum_k H_k S_k \right) \right\} \frac{S_m^*}{P_{nm}} \quad (3)$$

2. Use the realizable value that is closest (by Euclidean geometry in the complex plane) to H_m^I .

3. When we have computed all the filter values, verify that each meets the condition specified in Steps 1 and 2.

Well, we see that we have a problem. It seems that in order to compute the filter we have to know the filter first. We don't know the elements of the term in braces in Eq. (3) before computing the ideal filter value – and they are functions of the realized filter!

The saving circumstance is that the term in braces is not a function of m . The same complex value applies to all m . If there is an optimal filter that meets Eq. (1), then there exists such a complex factor as that in the braces, and all we have to do is find it. A major fact to notice is that the factor includes the effects of all frequencies as they influence what is optimal at the m -th frequency. Although the filter typically comprises tens of thousands of frequencies, their effects are coordinated through a single complex factor. The dimensionality of an optimizing search has been reduced by examining the consequence of the philosophy. This approach obviates such methods as annealing²⁴ where all frequencies are treated as independent. Instead we treat all frequencies alike that have the same spectral signal to noise ratio, S_m/P_{nm} . The algorithm, as a practical implication of Eq. (1), then becomes:

1) Choose a complex factor k .

2) For all m , compute $H_m^I = k \frac{S_m^*}{P_{nm}}$, then choose the closest realizable value for the m -th element of H .

3) Compute $J(k)=J(H(k))$; vary k to obtain $H^{\%}$ that maximizes J .

4) Perhaps optionally: Verify closure – that $k(H^{\%})$ is the value of k that maximized J , when the complex number in the braces in Eq. 3 is computed.

An interesting situation arises when P_{nm} is not zero. We consider how much each frequency can contribute to the metric. If realizable, H_m^I would contribute the maximum possible, and other values of H contribute less. Noting that $H_m=0$ contributes nothing to the metric, the interesting thing is that realizing $2H_m^I$ would also contribute nothing to the metric! Here's why... See Figure 1. There is a circularly symmetric falloff in the contribution that any frequency's filter value can make to the overall metric. If zero filter value contributes zero to the metric, then so also does twice the filter value, since it is at the same distance from the ideal value as is the origin. In Figure 2 we have laboratory results²² that show an improvement in useable correlation quality as one begins with the matched-phase filter having largest intensity; proceeds to the filter producing largest intensity; and finishes with the filter having the largest signal to noise ratio. See the reference for the $1/f$ noise added to the signal, the coupling between amplitude and phase of the filter SLM, and other information.

3.2 Rayleigh quotient

Now for the Rayleigh quotient. Suppose that we know we have one of two possibly very similar objects, that we know its centration, and that we wish to build a filter that will give a large correlation response to one of the objects but a comparatively small response to the other. Compared with the SNR case, this time we know more than the power spectral density of the object we wish to reject – we know its entire structure, including phase – and we can take advantage of that when computing the filter. Let S_1 and S_2 be the two signals. We slightly redefine the conventional Rayleigh quotient to include the physically observed, instrumentally introduced, but ineluctable and irreducible, noise.

$$RQ = \frac{\left| \sum_k H_k S_{1k} \right|^2}{\sigma_d^2 + \left| \sum_k H_k S_{2k} \right|^2} \quad (4)$$

Analytically following the consequences of applying Eq. (1) to Eq. (4), we find²⁰ that we must take a pair of complex constants k_1 and k_2 , then form the vector V in the complex plane of values for H_m ,

$$V = k_1 S_1 - k_2 S_2 \quad (5)$$

and choose the realizable filter value that has the largest component in the direction of V . Again we search over the complex parameter values (this time k_1 and k_2) to optimize RQ. The closure values for k_1 and k_2 are multiples, by a common real scalar, of

$$\left. \begin{aligned} k_1 &= B_1 \exp(j\beta_1) \\ k_2 &= B_2 \exp(j\beta_2) - \frac{B_1^2}{B_2^2 + \sigma_d^2} \end{aligned} \right\} \quad (6)$$

according to the terms developed for Eq. (11) of Ref.[20]. (We have subscripted the terms defined just above Eq. (1) to accommodate there being more than one reference object.) In Figure 3 we have laboratory results²⁵ showing that the Rayleigh quotient improves the discrimination between a pair of very similar images, compared with the poor discrimination resulting from building filters that just produce large correlation intensities. Optimizing the Rayleigh quotient causes a filter to reject the parts of the “accept” image that look like the “reject” image, whereas that same mutual character produces a substantial response when the filter is asked to produce as large a correlation intensity as possible, without regard to discrimination.

3.3 Fisher ratio

Optimizing the Fisher ratio (FR) is one step further along toward a truly “discriminant” filter. Such a filter would have a guaranteed minimal classification error. The FR is a long-standing metric in linear decision rules for pattern recognition. The FR measures the degree to which classes separate under the filter, normalized to the average variance within the classes. (See any text on statistical pattern recognition for more details.) I am unaware of any previous optimization of the Fisher ratio, based on observable optical quantities, restriction to realizable filter values, and incorporating scene noise and detection noise into the correlation intensity variance. I have optimized the FR as a function of the correlation filter, when the values of the filter are restricted to an arbitrary subset of the complex plane. I build the FR from optically observable quantities, in consonance with the philosophy’s dicta.

$$J = \frac{(\langle I \rangle_\Lambda - \langle I \rangle_\Psi)^2}{\sigma_\Lambda^2 + \sigma_\Psi^2} \quad (7)$$

in which the “accept” class is designated by Λ and the “reject” class by Ψ , the central correlation intensity is I , and the variances for the intensities of the classes are in the denominator. Optimizing the FR proceeds from the philosophy. The journal paper detailing the optimization of the FR is currently in peer review, so the details are not given here. However, I include initial simulation results from filter code that is in an incomplete stage of development. In Figure 4 we see eight images out of the sixteen that

the Fisher ratio optimizing filter has been trained to accept. (These are simulation data; all other data shown in this paper are taken in the optical laboratory.) The background is a difficult one for a correlator to reject, since it is highly structured and quite bright.

4. SOME GENERAL REMARKS

Coupling is your friend. This is true particularly in noisy situations. When noise is present, a great number of frequencies can have the spectral energy of the signal swamped by the noise power spectral energy. Ideal filter values at those frequencies would ideally have very small amplitude. The farther a realizable spectral filter value is from its ideal value, the more poorly the whole filter performs. Thus we often find that there is a cluster of ideal filter values near zero, and it behooves the operating curve of the filter to pass near zero to accommodate them. For a considerable while in the filter optimization industry it seemed as though coupled amplitude and phase behavior was the mark of a deficient modulator. Indeed there are a number of circumstances in which optimizing a filter calls for large amplitude and large phase range (the noiseless Rayleigh quotient comes to mind). But when we are attempting to defeat spectrally varying noise that can be modeled as additive in the input plane, we need amplitude control of the filter. We have the tools to use coupling to our advantage now, and we should use it.

Acknowledgements

The examples shown here are a group effort from the Hybrid Vision Laboratory at the Johnson Space Center in Houston. Stan Monroe, Tim Fisher, Shane Barton, Colin Soutar, Michael Morelli, Michael Rollins, and Ivan Spain are major contributors. T.-H. Chao at JPL kindly provided the images used with the Fisher ratio filter. Prof. B.V.K. Vijaya Kumar has been a highly valued colleague in these matters over the years.

References

- ¹ Joseph L. Horner and Peter D. Gianino, "Phase-only matched filtering", *Applied Optics* **23**, 812-816 (1984).
- ² B.V.K. Vijaya Kumar and Zouhir Bahri, "Phase-only filters with improved signal to noise ratio", *Applied Optics* **28**, 250-257 (1989).
- ³ B.V.K. Vijaya Kumar, Richard D. Juday, and Daniel W. Carlson, "Bias in correlation peak location", *Proc. SPIE* **1701** (1992).
- ⁴ Richard D. Juday, "Optical correlation with a cross-coupled spatial light modulator", *Spatial Light Modulators and Applications*, 1988 Technical Digest Series, vol. 8, Optical Society of America (1988).
- ⁵ Richard D. Juday, "Correlation with a spatial light modulator having phase and amplitude cross coupling", *Applied Optics* **28**, 4865-4869 (1989).
- ⁶ Michael W. Farn and Joseph W. Goodman, "Optimal maximum correlation filter for arbitrarily constrained devices", *Applied Optics* **28**, 3362-3366 (1989).
- ⁷ Richard D. Juday, "Optimal realizable filters and the minimum Euclidean distance principle", *Applied Optics* **32**, 5100-5111 (1993).
- ⁸ V. Laude and Ph. Réfrégier, "Multicriteria characterization of coding domains with optimal Fourier SLM filters", *Applied Optics* **33**, 4465-4471 (1994).

- ⁹ Fazlollahi, Amir H., and Bahram Javidi, "Optimum receivers for pattern recognition with nonoverlapping target and receiver noise", Optical Engineering **36** (10), 2633-2641 (1997).
- ¹⁰ Mahalanobis, Abhijit, and B. V. K. Vijaya Kumar, "Optimality of the maximum average correlation height filter for detection of targets in noise", Optical Engineering **36** (10), 2642-2648 (1997).
- ¹¹ Laude, Vincent, and Stephane Formont, "Bayesian target location in images", Optical Engineering **36** (10), 2649-2659 (1997).
- ¹² Guérault, Frédéric, Laurent Signac, François Goudai, and Philippe Réfrégier, "Location of target with random gray levels in correlated background with optimal processors and preprocessings", Optical Engineering **36** (10), 2660-2670 (1997).
- ¹³ Javidi, Bahram, Wenlu Wang, and Guanshen Zhang, "Composite Fourier-plane nonlinear filter for distortion-invariant pattern recognition", Optical Engineering **36** (10), 2690-2696 (1997).
- ¹⁴ Fisher, John W., III, and Jose C. Principe, "Recent advances to nonlinear minimum average correlation energy filters", Optical Engineering **36** (10), 2694-2709 (1997).
- ¹⁵ Hassebrook, Laurence G., Michael E. Lhamon, Mao Wang, and Jyoti P. Chatterjee, "Postprocessing of correlation for orientation estimation", Optical Engineering **36** (10), 2710-2718 (1997).
- ¹⁶ Casasent, David, and Satoshi Ashizawa, "Synthetic aperture radar detection, recognition, and clutter rejection with new minimum noise and correlation energy filters", Optical Engineering **36** (10), 2729-2736 (1997).
- ¹⁷ Shamir, Joseph, "Adaptive pattern recognition correlators", Optical Engineering **36** (10), 2675-2689 (1997).
- ¹⁸ Demoli, Nazif, Alexander Hirsch, Sven Kruger, Gunther Wernicke, Hartmut Gruber, and Mathias Senoner, "Optimization in mapping of correlation filters in a liquid crystal display based frequency plane correlator", Optical Engineering **38**, 1058-1064 (1999).
- ¹⁹ B.V.K. Vijaya Kumar, Daniel W. Carlson, and Abhijit Mahalanobis, "Optimal trade-off synthetic discriminant function filters for arbitrary devices", Optics Letters **19** 1556-1558 (1994).
- ²⁰ Richard D. Juday, "Generalized Rayleigh quotient approach to filter optimization", JOSA-A **15**(4), 777-790 (April 1998).
- ²¹ Richard D. Juday, "Optimal Fisher ratio for an arbitrary spatial light modulator", Optics Letters (submitted).
- ²² Richard D. Juday, R. Shane Barton, and Stanley E. Monroe, Jr., "Experimental optical results with MEDOF, coupled modulators, and quadratic metrics", Optical Engineering **38**, 302-312 (1999).
- ²³ B. V. K. Vijaya Kumar, Richard D. Juday, and P. Karivaratha Rajan, "Saturated filters", JOSA A **9**, 405-412 (1992).
- ²⁴ Richard D. Juday and Brian J. Daiuto, "Relaxation method of compensation in an optical correlator", Optical Engineering **26**, 1094-1101 (1987).

-
- ²⁵ Richard D. Juday, J. Michael Rollins, Stanley E. Monroe, Jr., and Michael V. Morelli, "Optical results with Rayleigh quotient discrimination filters", Proc. SPIE **3715**, 60-70 (April 1999).

Figure 1. The ideal filter value contributes the most possible to the metric. All filter values on the circle passing through the origin, and centered on the ideal value, contribute zero to the metric. That includes, of course, twice the ideal value. Outside the “zero contribution” circle, all contributions are negative.

Figure 2. Sections and images of optical lab results for filters that match phase (with greatest intensity under that restriction), maximize intensity, and maximize SNR on a highly coupled SLM (see references for fuller description including noise environment). The intensity filter creates by far the brightest image, the SNR filter the cleanest. The filter optimizing PCE (not shown) was essentially as noisy as the matched-phase filter. These are laboratory data.

Figure 3. A tank; the tank with a portion of its structure inverted; the difference between the two tanks.