An Ideal Observer for Appearance: Reconstruction from Samples

David H. Brainard, Department of Psychology

University of California, Santa Barbara, CA 93106

Draft of December, 1995

Note: This report provides additional technical detail related to the Bayesian demosaicing method outlined in

> Brainard, D. H. (1994) Bayesian method for reconstructing color images from trichromatic samples. *Proceedings of the IS&T 47th Annual Meeting, Rochester, NY,* 375-379.

and

> Brainard, D. H. and Sherman, D. (1995) Reconstructing images from trichromatic samples: from basic research to practical applications. *Proceedings of the IS&T/SID Color Imaging Conference: Color Science, Systems, and Applications, Scottsdale, AZ,* , 4-10.

If you would like to cite this work, please use the following:

> Brainard, D. H. (1995). "An Ideal Observer for Appearance: Reconstruction from Samples," UCSB Vision Labs Tech Report 95-1, Department of Psychology, UC Santa Barbara, Santa Barbara, CA. http://color.psych.ucsb.edu/brainard/papers/bayessampling.pdf.

Methods

Reconstruction from Samples

Image representation. I represent an image by the discrete function $i(x_i, y_j, \lambda_k)$. The coordinates x_i $(1 \le i \le N_{rows})$ and y_j $(1 \le j \le N_{cols})$ represent N_{rows} N_{cols} evenly spaced sample locations on a rectangular grid. This grid should be much denser than the sampling array being studied. The λ_k $(1 \le k \le N_{wls})$ represent the different color bands in the image. The λ_k represent evenly spaced sample wavelengths throughout the visible spectrum, so that the entire function $i(x_i, y_j, \lambda_k)$ represents the spectrum at each image location. The use of discrete representations for both spatial position and wavelength has been discussed extensively elsewhere.[1-3]

I use the N_{rows} N_{cols} N_{wls} dimensional column vector $\mathbf{i}$ to represent the function $i(x_i, y_j, \lambda_k)$. The n^{th} entry of $\mathbf{i}$ is $i(x_i, y_j, \lambda_k)$ with

$$\begin{aligned} i &= 1 + ((n-1) \text{ modulo } N_{rows}) \\ j &= 1 + (\lfloor (n-1)/N_{rows} \rfloor \text{ modulo } N_{cols}) \\ k &= 1 + \lfloor (n-1)/(N_{rows} N_{cols}) \rfloor \end{aligned} \qquad (1)$$

This choice of indexing enumerates all of the values of the discrete function $i(x_i, y_j, \lambda_k)$.

Sampling. Consider the relation between the response of a single receptor and the image data. Each receptor is completely characterized by a polychromatic receptive field, which specifies how strongly it responds to light from each color band at each location in the image. The receptive field of a single sensor may be described by a function $s(x_i, y_j, \lambda_k)$. The response r of the receptor is given by

$$r = \sum_{i,j,k} s(x_i, y_j, \lambda_k) i(x_i, y_j, \lambda_k) + e \qquad (2)$$

where e is a random variable representing sensor noise. Note that this receptive field specification allows us to handle any optical blurring that precedes spatial samping, arbitrary sampling geometry, and the fact that different photoreceptors have different spectral sensitivities.

We can rewrite Eq. 2 as

$$r = \mathbf{s}\,\mathbf{i} + e \tag{3}$$

where the $N_{rows}\,N_{cols}\,N_{wls}$ dimensional row vector **s** is the vector representation of $s(x_i, y_j, \lambda_k)$.

Suppose there are N_{rec} sensors in the overall camera array, enumerated by the index m. The responses of all the sensors can be represented by a single N_{rec} dimensional column vector **r** whose m^{th} entry is the response of the m^{th} sensor. Let **S** be the N_{rec} by $N_{rows}\,N_{cols}\,N_{wls}$ matrix whose m^{th} row is the row vector that represents the receptive field of the m^{th} receptor. Then

$$\mathbf{r} = \mathbf{S}\mathbf{i} + \mathbf{e} \tag{4}$$

where **e** is an N_{rec} dimensional column vector representing receptor noise. In this paper, I assume that **e** is multivariate Normal with mean $\mathbf{u}_e$ and covariance $\mathbf{K}_e$.

Eq. 4 describes the sampling process. The goal of a reconstruction algorithm is to provide an inverse mapping, $\hat{\mathbf{i}} = f(\mathbf{r})$, where $\hat{\mathbf{i}}$ is an estimate of the original image **i**. In this sense, reconstruction from samples is an inverse problem. It is a linear inverse problem because, in the absence of noise, **r** is a linear function of **i**. It is an underdetermined inverse problem because the number of samples is typically much smaller than the number of image parameters to be estimated. In this paper, I adopt a Bayesian approach to the finding an estimator $f(\mathbf{r})$. Other ways to approach underdetermined linear inverse problems have been extensively discussed (see e.g. Pratt[1] or Menke[4]).

Bayesian approach. The Bayesian approach to inverse problems is very simple (see e.g. Berger[5]). The problem is defined in terms of two probability distributions and a loss function.

The first probability distribution is the prior $p(\mathbf{i})$. The prior specifies the likelihood of encountering each possible image. It captures what is known about images before the receptor data are examined. Prior information is used in most reconstruction methods. For example, low-pass filtering (i.e. sinc interpolation) is based on the prior assumption that the image contains no frequency components above the Nyquist limit of the array. What differs in the Bayesian framework is that the prior information is expressed explicitly as a probability distribution. I discuss the construction of sensible prior distributions below.

The second probability distribution $p(\mathbf{r}|\mathbf{i})$ is the likelihood of the sensor data conditional on the image being viewed. The likelihood is really a specification of the imaging model. Here $p(\mathbf{r}|\mathbf{i})$ is determined by Eq. 3 and we write

$$p(\mathbf{r}|\mathbf{i}) = N(\mathbf{S}\mathbf{i} + \mathbf{u}_e, \mathbf{K}_e). \quad (5)$$

Given the prior and the likelihood, we may compute the posterior distribution $p(\mathbf{i}|\mathbf{r})$. The posterior specifies how likely any given image is, conditional on the observed receptor responses. The posterior is computed using Bayes' rule:

$$p(\mathbf{i}|\mathbf{r}) = C\ p(\mathbf{r}|\mathbf{i})p(\mathbf{i}). \quad (6)$$

In this expression, C is a normalizing constant. In this paper we will not be concerned with the exact value of C, as only the relative shape of the posterior influences our choice of $\hat{\mathbf{i}}$.

The loss function $L(\hat{\mathbf{i}}, \mathbf{i})$ specifies the cost of estimating $\hat{\mathbf{i}}$ when the underlying image is $\mathbf{i}$. Given a loss function and the posterior, we may compute the Bayes risk

$$R(\hat{\mathbf{i}}|\mathbf{r}) = E_{p(\mathbf{i}|\mathbf{r})}\{L(\hat{\mathbf{i}}, \mathbf{i})\} \quad (7)$$

where $E_{p(\mathbf{i}|\mathbf{r})}$ denotes the expected value with respect to the posterior. The Bayes risk specifies the expected loss for each choice of **i**, conditional on **r**. We choose the reconstructed image as the particular $\hat{\mathbf{i}}$ that minimizes the Bayes risk. In this paper, I only consider the squared error loss functions. The simplest example of a squared error loss function is $L(\hat{\mathbf{i}}, \mathbf{i}) = \| \mathbf{i} - \hat{\mathbf{i}} \|^2$.

Specifying the prior. To find an explicit Bayesian estimator, it is necessary to specify the prior distribution $p(\mathbf{i})$. I use linear models to construct prior distributions. I assume that each image can be expressed as a weighted sum of N_{basis} basis images. Let the basis images be represented by the columns of a $N_{rows} N_{cols} N_{wls}$ by N_{basis} basis matrix **B**. I assume that every image **i** may be expressed as a weighted sum of these columns, so that

$$\mathbf{i} = \mathbf{B}\,\mathbf{w} . \tag{8}$$

Images constrained by Eq. 8 are said to lie within the linear model defined by **B**. The N_{basis} dimensional vector **w** contains the linear model weights.

The linear model constraint by itself does not express a probability distribution. To turn Eq. 8 into a statement about probability, I assume that **w** is multivariate Normal with mean $\mathbf{u}_w$ and covariance $\mathbf{K}_w$, so that $p(\mathbf{w}) = N(\mathbf{u}_w, \mathbf{K}_w)$. This implies that

$$p(\mathbf{i}) = N(\mathbf{B}\mathbf{u}_w, \mathbf{B}\mathbf{K}_w\mathbf{B}^T) \tag{9}$$

The general form of Eq. 9 can capture two features thought to be true of natural images. First, the power spectra of natural images fall off fairly rapidly as a function of spatial frequency.[6,7] Second, there are correlations across wavelength.[7] Below I discuss specific choices for the basis images and the prior parameters $\mathbf{u}_w$ and $\mathbf{K}_w$.

The general form of Eq. 9 cannot, however, capture all of the structure in natural images. For example, Normal priors cannot express the fact that image intensities must be positive. None-the-less, these priors are more flexible than the assumptions of band-limited signals that have been used

in previous reconstruction algorithms and it seems worthwhile to investigate the performance possible when such priors are assumed. Future specification of better prior distributions may be incorporated within a general Bayesian framework.

Reconstruction. When the prior and the likelihood are both multivariate normal with known mean and covariance, so is the posterior. Moreover, the mean and covariance of the posterior may be computed in closed form. See Lee[8] for the derivation in the univariate case. The multivariate generalization is presented in the Appendix. The estimator that minimizes the Bayes risk for quadratic loss is exactly the posterior mean, and using this fact we may derive that

$$\hat{\mathbf{i}} = \mathbf{I}\,\mathbf{r} + \mathbf{i}_0 \tag{10}$$

with

$$\begin{aligned} \mathbf{I} &= \mathbf{B}\mathbf{K}_w(\mathbf{S}\mathbf{B})^T\left((\mathbf{S}\mathbf{B})\mathbf{K}_w(\mathbf{S}\mathbf{B})^T + \mathbf{K}_e\right)^{-1} \\ \mathbf{i}_0 &= \mathbf{B}\,\mathbf{u}_w - \mathbf{I}(\mathbf{S}\mathbf{B})\mathbf{u}_w - \mathbf{I}\,\mathbf{u}_e \end{aligned} \tag{11}$$

Eqs. 10 and 11 describe an affine estimator. The estimator is determined by parameters describing the sampling process ($\mathbf{S}, \mathbf{u}_e, \mathbf{K}_e$) and parameters describing the image population ($\mathbf{B}, \mathbf{u}_w, \mathbf{K}_w$). An advantage of the Bayesian approach is that it lets us design the reconstruction algorithm in terms of parameters that may (in principle) be measured or controlled. Note that for the special case of linear inverse problems, Normal priors, and squared error loss function, the Bayesian estimator of Eqs. 10 and 11 is identical to the discrete Wiener estimator described by Pratt for image restoration.[1] (Note that Pratt's published expression, Eq. 14.6-9, contains a typographical error for the sign of the contribution of $\mathbf{I}\,\mathbf{u}_e$ to $\mathbf{i}_0$.)

Relation to earlier methods.

Some insight the Bayesian estimator can be gained by considering the special case of a simple prior specified by $\mathbf{u}_w = \mathbf{0}$ and $\mathbf{K}_w = \mathbf{1}$ and no additive noise ($\mathbf{u}_e = \mathbf{0}, \mathbf{K}_e = \mathbf{0}$). (I use the symbol $\mathbf{0}$

to denote a zero vector or matrix of appropriate dimension and the symbol **1** to denote the identity matrix of appropriate dimension.) If the column dimension of **B** matches the row dimension of **S** and there are no degeneracies, then the matrix (**SB**) is non-singular and may be inverted. In this case, Eqs. 10 and 11 reduce to

$$\hat{\mathbf{i}} = \mathbf{B}(\mathbf{SB})^{-1}\mathbf{r}. \tag{12}$$

Consider the one-dimensional monochromatic case with no optical blurring and spatially regular sampling. If we choose the basis images to be sinusoids with spatial frequencies below the Nyquist limit of the sampling array, Eq. 12 reduces to a discrete version of low-pass or sinc interpolation. For spatially irregular two-dimensional sampling and with the basis images chosen to be regularly spaced two-dimensional sinc functions, Eq. 12 reduces to a discrete version of Yen interpolation.[9,10] If instead the basis images are band-limited sinc functions centered on the sampling points, Eq. 12 reduces to a discrete version of Chen-Allebach interpolation.[10,11]. Evidently the Bayesian formulation generalizes classical methods for reconstruction from samples. In its general form, it incorporates optical blurring, additive noise, and a richer description of image structure into the reconstruction process.

Specific Choice of Priors

Separability. To implement Eqs. 10 and 11, it is necessary to make specific choices of basis images and prior paramters. In this paper, I consider a restricted set of Normal priors that are separable in space and spectrum. A separable prior satisfies

$$p(\mathbf{i}) = p(x_i, y_j, \lambda_k) = p_{xy}(x_i, y_j)\, p_{\lambda}(\lambda_k) \tag{13}$$

Preliminary measurements of natural images suggest that separability holds approximately though not exactly.[7] Assuming separability simplifies the process of selecting a prior over images. It also

allows computational savings which make the numerical implementation of the reconstruction method feasible.

To construct separable priors, I proceed as follows. Let $\mathbf{B}_{xy}$ be an N_{pixels} by $N_{xybasis}$ matrix, each of whose columns represents a monochromatic image. This matrix describes $N_{xybasis}$ spatial basis functions. We can form a prior over monochromatic images by placing a Normal distribution over weights $\mathbf{w}_{xy}$ for this linear model, so that $p(\mathbf{w}_{xy}) = N(\mathbf{u}_{xy},\mathbf{K}_{xy})$. Similarly, let $\mathbf{B}_\lambda$ be an N_{wls} by N_{cbasis} matrix, each of whose columns represents a function of wavelength. This matrix describes N_{cbasis} spectral basis functions. We can form a prior over spectra by placing a Normal distribution over weights $\mathbf{w}_\lambda$ for this linear model, so that $p(\mathbf{w}_\lambda) = N(\mathbf{u}_\lambda,\mathbf{K}_\lambda)$.

Given $\mathbf{B}_{xy}$, $\mathbf{B}_\lambda$, $\mathbf{u}_{xy}$, $\mathbf{K}_{xy}$, $\mathbf{u}_\lambda$. $\mathbf{K}_\lambda$, we form a basis matrix $\mathbf{B}$ by taking the Kronecker (or tensor) product of $\mathbf{B}_\lambda$ and $\mathbf{B}_{xy}$:

$$\mathbf{B} = \mathbf{B}_\lambda \quad \mathbf{B}_{xy} . \tag{14}$$

The Kronecker is illustrated in Fig. 1. It provides a compact description of the type of structure matrices that arise in computations with separable distributions. To form a prior over imges that is separable in space and spectrum, we choose a Normal distriubtion with mean $\mathbf{u}_\mathbf{w} = \mathbf{u}_\lambda \quad \mathbf{u}_{xy}$ and covariance $\mathbf{K}_\mathbf{w} = \mathbf{K}_\lambda \quad \mathbf{K}_{xy}$ over the weights $\mathbf{w}$ for the linear model specified by $\mathbf{B}$. The structure imposed by the use of the Kronecker product in defining $\mathbf{B}$, $\mathbf{u}_\mathbf{w}$, and $\mathbf{K}_\mathbf{w}$ guarantees that the resulting $p(\mathbf{i})$ will be separable.

Figure 1 about here.

Choice of spatial prior parameters. To explore the performance of the reconstruction algorithm, it is convenient to choose $\mathbf{B}_{xy}$, $\mathbf{u}_{xy}$, and $\mathbf{K}_{xy}$ so that the properties of the prior may be manipulated via a small number of parameters. I assume that the population of spatial patterns may be described by distribution that is separable in x and y, so that

$$\mathbf{B}_{xy} = \mathbf{B}_x \quad \mathbf{B}_y \tag{15}$$

with

$$\begin{aligned} \mathbf{u}_{xy} &= \mathbf{u}_x \quad \mathbf{u}_y \\ \mathbf{K}_{xy} &= \mathbf{K}_x \quad \mathbf{K}_y . \end{aligned} \tag{16}$$

I assume that both the x and y distributions are identical, and that both are discrete first order Gauss-Markov processes. Thus we may specify both x and y distributions by their common pixel mean u_{xy} , pixel variance σ_{xy}^2 , and correlation ρ_{xy} between neighboring pixels. From these, we can construct the mean vectors $\mathbf{u}_x = \mathbf{u}_y$ and covariance matrices $\mathbf{K}_x = \mathbf{K}_y$ for the x and y distributions using the formulae provided by Pratt[1] (see section 5.4). Together with the choice $\mathbf{B}_x = \mathbf{B}_y = \mathbf{1}$ these provide us with a full distribution through Eqs. 15 and 16. The dimensionality $N_{xybasis}$ is equal to the number of pixels, however, and is too large for practical computations. Thus we choose reduced bases $\mathbf{B}_x = \mathbf{B}_y$ by finding the $N_{xbasis} = N_{ybasis}$ eigenvectors of $\mathbf{K}_x = \mathbf{K}_y$ that correspond to the largest eigenvalues. Given these, we then use Eqs. 15 and 16 to obtain a full distribution with dimension $N_{xybasis} = N_{xbasis}\, N_{ybasis}$.

Choice of spectral prior parameters. I use a similar strategy to specify the spectral priors with a small number of parameters. I assume that the spectral distribution is a discrete first order Gauss-Markov processes. Thus we may specify both the distribution by its mean at each wavelength u_λ , variance at each wavelength σ_λ^2 , and correlation ρ_λ between neighboring wavelengths. From these, we can construct the mean vector $\mathbf{u}_\lambda$ and covariance matrix $\mathbf{K}_\lambda$ using the formulae provided by Pratt[1] (see section 5.4). Together with the choice $\mathbf{B}_\lambda = \mathbf{1}$ these provide us with a full distribution. The dimensionality $N_{\lambda basis}$ is equal to the number of sample wavelengths, however, and is too large for practical computations. Thus we choose a reduced basis $\mathbf{B}_\lambda$ by finding the $N_{\lambda basis}$ eigenvectors of $\mathbf{K}_\lambda$ that correspond to its largest eigenvalues and transform $\mathbf{u}_\lambda$ and $\mathbf{K}_\lambda$ appropriately.

Meaning of the prior parameters. Both the spatial and spectral priors are determined by the values of parameters u, σ^2, ρ, and N_{basis}. The parameters u_{xy} and u_λ determine mean image. In this paper, I set both of these parameters to 1 and investigate reconstructions of modulations around this mean. For fixed u, the parameters σ^2 and ρ determine how much information is provided by the prior. Very large values of σ^2 and very small values of ρ lead to relatively non-informative priors where all images are essentially equally likely. Smaller values of σ^2 imply that pixel values far from the mean are unlikely. Values of ρ close to one indicate that rapid variation in space or wavelength is unlikely. I analyzed the pixel distribution in a number of uncalibrated images and obtained a value for σ^2_{xy} of 0.25 (given u_{xy} =1). Pratt[1] suggests that a reasonable value for ρ_{xy} for natural images is about 0.95. To determine values for σ^2_λ and ρ_λ I analyzed a set of population of spectra derived from an analysis of the spectra of Munsell chips and CIE daylights. I obtained values of $\sigma^2_\lambda = 0.3$ (given u_λ =1) and ρ_λ =0.99 (for wavelength spacing of 10 nm). All of these values should be taken as ballpark figures at best. Ideally, we would set $N_{xybasis}$ equal to the number of pixels and $N_{\lambda basis}$ equal to the number of wavelength samples. As mentioned above, this is computationally prohibitive. The computations reported in this paper were done on 32 by 32 pixel blocks containing roughly 64 receptors. I generally used $N_{xybasis} = 64$ and $N_{\lambda basis} = 4$. These truncations increases the informativeness of the priors only slightly for the large values of ρ that seem appropriate for natural images and spectra.

Specific Choice of Imaging Model

To perform calculations, we must supply an actual imaging model. This consists of specifying the matrix **S** and the noise mean $\mathbf{u}_e$ and covariance matrix $\mathbf{K}_e$. The matrix **S** must incorporate optical blurring, the mosaic topography, and the cone spectral sensitivites.

Optics. To model the eye's optics, I used the formula provided by Williams et. al.[12] for the foveal monochromatic modulation transfer function of the human eye. From this I generated a polychromatic modulation transfer function that incorporated chromatic aberration by using the calculations reported by Marimont and Wandell[13]. I converted the modulation transfer function to

a point spread function numerically under the assumption that the point spread function is circularly symmetric.

Mosaic topography. Curcio and Sloan[14] provide a two parameter model for the packing geometry of the overall human cone mosaic. The model depends on a compression parameter c and a jitter parameter j. I used compression parameter c = 0.84 and jitter parameter j = 0.10 to generate the overall packing of the foveal mosaics.

For image blocks of 32 by 32 pixels, I constructed mosaics consisting of 64 cones. Of these, 38 were L cones, 20 were M cones, 6 were S cones. With this choice, the S cones are about 9% of the total cone population and the L to M cone ratio is 1.9:1. Both figures are consistent with current estimates.[references]. The packing geometry of the S cones was quasi-regular. It was obtained by constructing an S cone mosaic with the Curio/Sloan algorithm and then finding the nearest neighbors to the constructed points in the overall cone mosaic. The L and M cones were colored randomly from the remainder[Mollon ref].

Cone spectral sensitivities. I used the Smith-Pokorny estimates of the human cone spectral sensitivities.[refs]

Noise. I assumed that the noise had zero mean and was independent for each cone. The amount of photon noise depends on the mean luminance. I used noise whose standard deviation for each cone was 10% of that cones mean response to the mean image.

Acknowledgments

I thank A. Ahumada, D. Sherman, B.Wandell and D. Williams for helpful discussions. Initial results were presented at the 1989 and 1993 ARVO annual meetings.[15,16] and are described briefly in proceedings papers. The work was supported in part by NEI EY 10016.

Appendix A: Bayesian Estimator -- Linear model and Normal Prior

We start with a prior $p(\mathbf{i}) = N(\mathbf{u}_i, \mathbf{K}_i)$ and a likelihood $p(\mathbf{r}|\mathbf{i}) = N(\mathbf{Si} + \mathbf{u}_e, \mathbf{K}_e)$. We would like to derive that the posterior is Normal and find the mean and covariance of the posterior. From Bayes rule, we have

$$
\begin{aligned}
p(\mathbf{i}\,|\,\mathbf{r}) &\sim N(\mathbf{u}_i, \mathbf{K}_i) N(\mathbf{S}\mathbf{i} + \mathbf{u}_e, \mathbf{K}_e) \\
&\sim \exp\left(-\frac{1}{2} \left\| \mathbf{K}_i^{-1/2}(\mathbf{i} - \mathbf{u}_i) \right\|^2 \right) \exp\left(-\frac{1}{2} \left\| \mathbf{K}_e^{-1/2}(\mathbf{r} - (\mathbf{S}\mathbf{i} + \mathbf{u}_e)) \right\|^2 \right). \\
&\sim \exp\left(-\frac{1}{2} \left\| \mathbf{K}_i^{-1/2}(\mathbf{i} - \mathbf{u}_i) \right\|^2 \right) \exp\left(-\frac{1}{2} \left\| \mathbf{K}_e^{-1/2}(\mathbf{S}\mathbf{i} - (\mathbf{r} - \mathbf{u}_e)) \right\|^2 \right)
\end{aligned}
\qquad \text{(A- 1)}
$$

By completing the square of the sum of the exponents (see below), we derive

$$
p(\mathbf{i}\,|\,\mathbf{r}) \sim N(\mathbf{u}, \mathbf{K}) \,. \qquad \text{(A-2)}
$$

with

$$
\begin{aligned}
\mathbf{u} &= \mathbf{K}\left(\mathbf{K}_i^{-1}\mathbf{u}_i + \mathbf{S}^T\mathbf{K}_e^{-1}(\mathbf{r} - \mathbf{u}_e)\right) \\
\mathbf{K} &= \left(\mathbf{K}_i^{-1} + \mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}
\end{aligned}
\,. \qquad \text{(A-3)}
$$

The best MAP or MMSE estimator for the **i** is the posterior mean, we write

$$
\begin{aligned}
\hat{\mathbf{i}} = &\left(\mathbf{K}_i^{-1} + \mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{r} \;\ldots \\
&- \left(\mathbf{K}_i^{-1} + \mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{u}_e \;\ldots \\
&+ \left(\mathbf{K}_i^{-1} + \mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\mathbf{K}_i^{-1}\mathbf{u}_i
\end{aligned}
\,. \qquad \text{(A-4)}
$$

Now we would like to simplify this expression. We can use the identity

$$
\begin{aligned}
\mathbf{I} = &\left(\mathbf{K}_i^{-1} + \mathbf{S}^T\mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\mathbf{S}^T\mathbf{K}_e^{-1} = \ldots \\
&\mathbf{K}_i\mathbf{S}^T\left(\mathbf{S}\,\mathbf{K}_i\mathbf{S}^T + \mathbf{K}_e\right)^{-1}
\end{aligned}
\,. \qquad \text{A-5)}
$$

This identity is quoted by Pratt[1] (page 412) in his derivation of the discrete Weiner estimator. It may be derived from his identity 5.1-11 on page 123 as follows. Let the matrix **F** be defined by

$$\mathbf{F} = \left[\begin{array}{c|c} -\mathbf{K}_e & \mathbf{S} \\ \hline \mathbf{S}^T & \mathbf{K}_i^{-1} \end{array}\right] \tag{A-6}$$

and note that $\mathbf{F}$ and thus $\mathbf{F}^{-1}$ are symmetric. In this case, the two off-diagonal blocks of the expression for $\mathbf{F}^{-1}$ in 5.1-11 are related by the transpose operation. By substituting for our particular $\mathbf{F}$ and simplifying, we get the desired result. Identity 5.1-11 itself may be verified by direct block matrix multiplication.

Using this result and substituting into A-4, we get

$$\hat{\mathbf{i}} = \mathbf{I}\mathbf{r} + \mathbf{i}_0 \tag{A-7}$$

where

$$\begin{aligned} \mathbf{i}_0 &= \left(\mathbf{K}_i^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\mathbf{K}_i^{-1}\mathbf{u}_i - \mathbf{I}\mathbf{u}_e \\ &= \left(\mathbf{K}_i^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\left(\mathbf{K}_i^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S} - \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\right)\mathbf{u}_i - \mathbf{I}\,\mathbf{u}_e \\ &= \left(\mathbf{K}_i^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\left(\mathbf{K}_i^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\right)\mathbf{u}_i \ldots \\ &\quad -\left(\mathbf{K}_i^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\right)^{-1}\mathbf{S}^T \mathbf{K}_e^{-1}\mathbf{S}\mathbf{u}_i - \mathbf{I}\mathbf{u}_e \\ &= \mathbf{u}_i - \mathbf{I}\mathbf{S}\,\mathbf{u}_i - \mathbf{I}\mathbf{u}_e \end{aligned} \quad . \tag{A-8}$$

Note that this agrees with Pratt's result (Eqs. 14.6-8 through 14.6-11, page 412) except for the sign associated with the term $\mathbf{I}\mathbf{u}_e$ in the expression for $\mathbf{i}_0$ (Eq. 14.6-9). Pratt's expression is incorrect. The sign was lost betrween the left and right and sides of Eq. 14.6-9. Pratt's derivation is bassed on different principles, however, as he minimizes the expected squared error directly.

Linear model priors

Often we are concerned with the case where the prior is induced by a prior over the weights of a linear model. In this case we have $\mathbf{u}_i = \mathbf{B}\,\mathbf{u}_w$ and $\mathbf{K}_i = \mathbf{B}\,\mathbf{K}_w\,\mathbf{B}^T$. Substituting, we arrive at

$$\hat{\mathbf{i}} = \mathbf{I}\,\mathbf{r} + \mathbf{i}_0 \tag{A-9}$$

with

$$\begin{aligned} \mathbf{I} &= \mathbf{B}\,\mathbf{K}_w\,(\mathbf{SB})^T\left((\mathbf{SB})\,\mathbf{K}_w(\mathbf{SB})^T + \mathbf{K}_e\right)^{-1} \\ \mathbf{i}_0 &= \mathbf{B}\,\mathbf{u}_w - \mathbf{I}\,\mathbf{SB}\,\mathbf{u}_w - \mathbf{I}\mathbf{u}_e \end{aligned} \tag{A-10}$$

where the notation is abused by continuing to use $\mathbf{I}$ and $\mathbf{i}_0$.

Completing the square

The following result is necessary for the derivation. Consider the expression

$$E = \left\| \mathbf{K}_1^{-1/2}(\mathbf{A}_1\mathbf{x} - \mathbf{u}_1) \right\|^2 + \left\| \mathbf{K}_2^{-1/2}(\mathbf{A}_2\,\mathbf{x} - \mathbf{u}_2) \right\|^2 \tag{A-11}$$

I show that if we let

$$\begin{aligned} \mathbf{K} &= \left(\mathbf{A}_1^T\mathbf{K}_1^{-1}\mathbf{A}_1 + \mathbf{A}_2^T\mathbf{K}_2^{-1}\mathbf{A}_2\right)^{-1} \\ \mathbf{u} &= \mathbf{K}\left(\mathbf{A}_1^T\,\mathbf{K}_1^{-1}\,\mathbf{u}_1 + \mathbf{A}_2^T\,\mathbf{K}_2^{-1}\,\mathbf{u}_2\right) \end{aligned} \tag{A-12}$$

then

$$E = \left\| \mathbf{K}^{-1/2}(\mathbf{x} - \mathbf{u}) \right\|^2 + C \tag{A-13}$$

where the constant C is independent of $\mathbf{x}$. The proof is brute force expansion of A-11.

$$
\begin{aligned}
\left\| \mathbf{K}^{-1/2}(\mathbf{x}-\mathbf{u}) \right\|^2 &= (\mathbf{x}-\mathbf{u})^{\mathrm{T}}\, \mathbf{K}^{-1}(\mathbf{x}-\mathbf{u}) \\
&= \mathbf{x}^{\mathrm{T}}\mathbf{K}^{-1}\, \mathbf{x} - 2\, \mathbf{x}^{\mathrm{T}}\mathbf{K}^{-1}\, \mathbf{u} + \mathbf{u}^{\mathrm{T}}\mathbf{K}^{-1}\, \mathbf{u} \\
&= \mathbf{x}^{\mathrm{T}}\mathbf{A}_1^{\mathrm{T}}\mathbf{K}_1^{-1}\mathbf{A}_1\, \mathbf{x} + \mathbf{x}^{\mathrm{T}}\mathbf{A}_2^{\mathrm{T}}\mathbf{K}_2^{-1}\mathbf{A}_2\, \mathbf{x} \ldots \\
&\quad - 2\, \mathbf{x}^{\mathrm{T}}\mathbf{K}^{-1}\, \mathbf{K}\left(\mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 + \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2\right) \ldots \\
&\quad + \left(\mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 + \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2\right)^{\mathrm{T}}\, \mathbf{K}\, \mathbf{K}^{-1}\, \mathbf{K}\left(\mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 + \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2\right) \\
&= \mathbf{x}^{\mathrm{T}}\mathbf{A}_1^{\mathrm{T}}\mathbf{K}_1^{-1}\mathbf{A}_1\, \mathbf{x} + \mathbf{x}^{\mathrm{T}}\mathbf{A}_2^{\mathrm{T}}\mathbf{K}_2^{-1}\mathbf{A}_2\, \mathbf{x} \ldots \\
&\quad - 2\, \mathbf{x}^{\mathrm{T}}\, \mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 - 2\, \mathbf{x}^{\mathrm{T}}\, \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2 \ \ldots \\
&\quad + C_1
\end{aligned}
\tag{A-14}
$$

with

$$
C_1 = \left(\mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 + \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2\right)^{\mathrm{T}}\, \mathbf{K}\, \mathbf{K}^{-1}\, \mathbf{K}\left(\mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 + \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2\right). \tag{A-15}
$$

Thus

$$
\begin{aligned}
\left\| \mathbf{K}^{-1/2}(\mathbf{x}-\mathbf{u}) \right\|^2 &= \mathbf{x}^{\mathrm{T}}\mathbf{A}_1^{\mathrm{T}}\mathbf{K}_1^{-1}\mathbf{A}_1\, \mathbf{x} + \mathbf{x}^{\mathrm{T}}\mathbf{A}_2^{\mathrm{T}}\mathbf{K}_2^{-1}\mathbf{A}_2\, \mathbf{x} \ldots \\
&\quad - 2\, \mathbf{x}^{\mathrm{T}}\, \mathbf{A}_1^{\mathrm{T}}\, \mathbf{K}_1^{-1}\mathbf{u}_1 - 2\, \mathbf{x}^{\mathrm{T}}\, \mathbf{A}_2^{\mathrm{T}}\, \mathbf{K}_2^{-1}\mathbf{u}_2 \ \ldots \\
&\quad + \mathbf{u}_1^{\mathrm{T}}\mathbf{K}_1^{-1}\mathbf{u}_1 + \mathbf{u}_2^{\mathrm{T}}\mathbf{K}_2^{-1}\mathbf{u}_2 \ldots \\
&\quad + C_2
\end{aligned}
\tag{A-16}
$$

with

$$
C_2 = C_1 - \mathbf{u}_1^{\mathrm{T}}\mathbf{K}_1^{-1}\mathbf{u}_1 - \mathbf{u}_2^{\mathrm{T}}\mathbf{K}_2^{-1}\mathbf{u}_2. \tag{A-17}
$$

Gathering terms appropriately, we have

$$
\begin{aligned}
\left\| \mathbf{K}^{-1/2}(\mathbf{x}-\mathbf{u}) \right\|^2 &= (\mathbf{A}_1\mathbf{x}-\mathbf{u}_1)^{\mathrm{T}}\mathbf{K}_1^{-1}\, (\mathbf{A}_1\mathbf{x}-\mathbf{u}_1) \ldots \\
&\quad + (\mathbf{A}_2\mathbf{x}-\mathbf{u}_2)^{\mathrm{T}}\mathbf{K}_2^{-1}\, (\mathbf{A}_2\mathbf{x}-\mathbf{u}_2) \ldots \\
&\quad + C_2
\end{aligned}
\tag{A-18}
$$

.

This proves the result if we let $C = -C_2$.

Estimation error

The expected squared estimation error, given an observed **r**, is given by

$$\varepsilon_{\mathbf{r}} = E_{\mathbf{r}} \left\{ (\mathbf{i}-\hat{\mathbf{i}})^T (\mathbf{i}-\hat{\mathbf{i}}) \right\} \tag{A-19}$$

where the expectation is taken over the posterior $p(\mathbf{i}|\mathbf{r})$. It is simple to show that this expectation is exactly the trace of the posterior covariance matrix, so that

$$\varepsilon_{\mathbf{r}} = \mathrm{tr}\{\mathbf{K}\} = \mathrm{tr}\left\{ \left(\mathbf{K}_{\mathbf{i}}^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1} \mathbf{S} \right)^{-1} \right\} \tag{A-20}$$

Conveniently, $\varepsilon_{\mathbf{r}}$ is independent of **r**, so that the expected reconstruction error ε given a random draw of **i** from the prior on **i** is given by the same expression

$$\varepsilon = \mathrm{tr}\left\{ \left(\mathbf{K}_{\mathbf{i}}^{-1} + \mathbf{S}^T \mathbf{K}_e^{-1} \mathbf{S} \right)^{-1} \right\}. \tag{A-21}$$

Estimating linear functions of i

The posterior for **i** contains everything we know about **i** after looking at the sensor responses **r**. Sometimes we may wish to estimate some function of **i**, rather than **i** itself. If the function is linear, so that the quantity we wish to estimate is **m** = **M i**, then the solution is easy. The posterior for **m** is obtained from the posterior for **i** by looking at the distribution induced by transforming by **M**. Since for our case the posterior is Normal over **i**, it is also Normal over **m**. We derive

$$\hat{\mathbf{m}} = \mathbf{M}\,\mathbf{I}\mathbf{r} + \mathbf{M}\,\mathbf{i}_0. \tag{A-22}$$

The corresponding expected reconstruction error is given by

$$\varepsilon_{\mathbf{M}} = \mathrm{tr}\left\{\mathbf{M}\left(\mathbf{K}_{\mathbf{i}}^{-1} + \mathbf{S}^{\mathsf{T}}\,\mathbf{K}_{\mathbf{e}}^{-1}\mathbf{S}\right)^{-1}\mathbf{M}^{\mathsf{T}}\right\} \qquad \text{(A-23)}$$

Note that this is also the expression for the expected value for a weighted error measure on **i**.

Sensor choice

Equations (A-22) and (A-23) are rather interesting. They allow us to compute the expected reconstruction error as a function of the sensor matrix **S**. If we are trying to design sensors for a known problem, then we can maximize our performance (in a squared error sense) if we choose the sensors to minimize $\varepsilon_{\mathbf{M}}$.

It is not obvious that there is an analytical expression for the gradient of $\varepsilon_{\mathbf{M}}$ with respect to the entries of S. If we sample the spectum in 10 nm steps, we could try searching over the sensor sensitivities to minimize the expected error. I think it would be hard to do in general, but possibly you would do it using a greedy algorithm for one sensor followed by a second and so forth.

References

1. Pratt, W. K. *Digital Image Processing*, (John Wiley & Sons: New York, 1978).

2. Nielsen, K. R. K. and Wandell, B. A., "Discrete analysis of spatial sensitivity models," *Journal of the Optical Society of America A* **5**, 743-755 (1988).

3. Brainard, D. H., "Colorimetry." In *Handbook of Optics: Volume 1. Fundamentals, Techniques, and Design*; M. Bass eds., (McGraw-Hill: New York, 1995).

4. Menke, W. *Geophysical data analysis: discrete inverse theory*, (Academic Press: San Diego, 1989).

5. Berger, T. O. *Statistical Decision Theory and Bayesian Analysis*, (Springer-Verlag: New York, 1985).

6. Field, D. J., "Relations between the statistics of natural images and the response properties of cortical cells," *Journal of the Optical Society of America A* **4**, 2379-2394 (1987).

7. Burton, G. J. and Moorehead, I. R., "Color and spatial structure in natural images," *Applied Optics* **26**, 157-170 (1987).

8. Lee, P. M. *Bayesian Statistics*, (Oxford University Press: London, 1989).

9. Yen, J. L., "On nonuniform sampling of bandwidth-limited signals," *IRE Transactions on circuit theory* **CT-3**, 251-257 (1956).

10. Yellott, J. I., Jr., "The photoreceptor mosaic as an image sampling device." In *Advances in Photoreception*; N. R. C. C. o. Vision eds., (National Academy Press: Washington, D.C., 1990).

11. Chen, D. S. and Allebach, J. P., "Analysis of error in reconstruction of two-dimensional signals from irregularly spaced samples," *IEEE Transactions on Acoustics, Speech, and Signal Processing* **ASSP-35**, 173-180 (1987).

12. Williams, D. R., Brainard, D. H., MacMahon, M. and Navarro, R., "Double pass and interferometric measures of the optical quality of the eye," *J. Opt. Soc. Am. A* (in press).

13. Marimont, D. H. and Wandell, B. A., "Chromatic aberration paper -- fill in reference.," *J. Opt. Soc. Am. A.* (in press).

14. Curcio, C. A. and Sloan, K. R., "Packing geometry of human cone photoreceptors: variation with eccentricity and evidence for local anisotropy," *Visual Neuroscience* **9**, 169-180 (1992).

15. Brainard, D. H., Wandell, B. A. and Poirson, A. B., "Discrete analysis of spatial and spectral aliasing," *Investigative Opthalmology and Visual Science, Supplement* **30**, 53 (1989).

16. Brainard, D. H. and Williams, D. R., "Bayes estimator for reconstruction from samples," *Investigative Ophthalmology and Visual Science, Supplement* **34**, 777 (1993).

Figures

$$\mathbf{B} = \begin{bmatrix} B_\lambda(1,1)\,\mathbf{B}_{xy} & B_\lambda(1,2)\,\mathbf{B}_{xy} & \cdots & B_\lambda(1,N_{\lambda basis})\,\mathbf{B}_{xy} \\ B_\lambda(2,1)\,\mathbf{B}_{xy} & B_\lambda(2,2)\,\mathbf{B}_{xy} & & \\ \cdots & & \cdots & \\ B_\lambda(N_{wls},1)\,\mathbf{B}_{xy} & & & B_\lambda(N_{wls},N_{\lambda basis})\,\mathbf{B}_{xy} \end{bmatrix}$$

Figure 1. The Kronecker product. The figure shows the basis matrix $\mathbf{B} = \mathbf{B}_\lambda \quad \mathbf{B}_{xy}$. This matrix has a block structure, where each block is a scaled version of the matrix $\mathbf{B}_{xy}$. Each individual block is scaled by the corresponding element of $\mathbf{B}_\lambda$. If the matrix $\mathbf{B}_\lambda$ has dimensions N_{wls} by $N_{\lambda basis}$ and the matrix $\mathbf{B}_{xy}$ has dimensions N_{pixels} by $N_{xybasis}$, then the matrix **B** has dimensions $N_{pixels}N_{wls}$ and column dimension $N_{xybasis}N_{\lambda basis}$. In general the Kronecker product of two arbitrary matrices **X** and **Y** is formed by replicating the matrix **Y**, with each replica scaled by the corresponding entry of **X**.