Color constancy. II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces

Michael D'Zmura and Geoffrey Iverson

Department of Cognitive Sciences and Institute for Mathematical Behavioral Sciences, University of California, Irvine, Irvine, California 92717

Received October 21, 1992; revised manuscript received April 8, 1993; accepted April 13, 1993

Our analysis of color constancy in a companion paper [J. Opt. Soc. Am A **10,** 2148 (1993)] provided an algorithm that lets one test how well linear color constancy schemes work. Here we present the results of applying the algorithm to a large parametric class of color constancy problems involving bilinear models that relate photoreceptoral spectral sensitivities, surface reflectance functions, and illuminant spectral power distributions. These results, supported by simulation and further analysis, provide a detailed classification of two-stage linear methods for recovering the spectral properties of reflectances and illuminants from reflected lights.

1. INTRODUCTION

We continue here a study of color constancy that we started in a companion paper.[1] We posed there a general problem concerning a p-chromatic visual system that uses a bilinear model with an m-dimensional linear model for illumination and an n-dimensional linear model for reflectances. When provided v views of s surfaces, can such a system recover uniquely the s surface reflectance functions and the v illuminant spectral power distributions?

For a recovery procedure to perform perfectly, for some choice of parameters $(p\,m\,n\,v\,s)$ it is necessary and sufficient that the bilinear model provide a one-to-one relationship between sets of lit surfaces and quantum catch data.[1] For two-stage linear recovery procedures, we expressed these conditions through a homogeneous system of polynomial equations, determined by a particular bilinear model, that has no nontrivial solution if and only if the model provides a recovery procedure that returns uniquely the spectral descriptions from quantum catch data. From these systems of equations we derived a model check algorithm.

The aim of the work presented here is to use the model check algorithm to help determine, for all choices of parameters $(p\,m\,n\,v\,s)$ that satisfy $p \geq m$ and $s \geq n$, where the number p of photoreceptoral types ranges from two through four, (1) whether there exist bilinear models that provide perfect recovery procedures, (2) whether the problem parameters are such that only imperfect recovery procedures are possible, or (3) whether the parameters are such that all recovery procedures are total failures. The inequalities $p \geq m$ and $s \geq n$ limit the problems to those for which one can devise a two-stage recovery procedure, such as those of Maloney and Wandell[2] and D'Zmura.[3]

A feasibility condition derived in the companion paper[1] excludes a number of problems from consideration. For a linear recovery procedure to have any chance of succeeding, it is necessary that the number Q of quantum catch data equal or exceed the number D of unknown descriptors of reflectances and illuminants to be recovered. If we take into account the scaling ambiguity of the recovered descriptors, the feasibility condition is $\mathsf{Q} \geq \mathsf{D} - 1$, or

$$sup \geq sn + vm - 1. \tag{1}$$

We examine here only problems with parameters $(p\,m\,n\,v\,s)$ that satisfy this feasibility condition [inequality (14) of Ref. 1].

Our methods for classifying feasible two-stage linear recovery procedures are threefold. The primary method is the numerical application of the model check algorithm to particular bilinear models. Two further methods include the numerical simulation of recovery with particular bilinear models, which can tell us whether a recovery procedure is a total failure, and theoretical analysis of particular problems. Entailments among problems and among bilinear models extend greatly the breadth of particular results found with these methods.

We apply these methods to dichromatic, trichromatic, and tetrachromatic systems. The generality of the methods lets us examine color constancy schemes that use only one view, such as those of Maloney and Wandell,[2] in addition to those that make use of the chromatic change in reflected lights provided by two or more views.[3] This work was described elsewhere in preliminary form.[4]

2. METHODS

The methods for classifying color constancy problems include numerical methods, which we use to test particular bilinear models, and analytic methods, which we use to derive results for a problem in a more general way.

A. Model Check Algorithm

Our chief method for showing that a particular problem can provide perfect color constancy schemes is to conduct successful checks of particular bilinear models with the problem's parameters by using the model check algorithm. This algorithm uses the known components of a particular bilinear model to construct a set of E homogeneous equations in U unknowns that are monomials of degree $m - v + 1$ in $n^2 - 1$ underlying variables. The quantities E

0740-3232/93/102166-15$06.00

and U are given by products of binomial coefficients [Eqs. (71) and (72) of Ref. 1]:

$$\underset{(p\geq m\geq v,\, n=s)}{\text{Number of Equations}} = \mathsf{E} = \binom{np-m}{m-v+1}\binom{m}{m-v+1}, \tag{2}$$

$$\underset{(p\geq m\geq v,\, n=s)}{\text{Number of Unknowns}} = \mathsf{U} = \binom{n^2+m-v-1}{m-v+1}. \tag{3}$$

If there is no solution (other than **0**) to the set of equations, considered as a set of linear homogeneous equations in the monomial unknowns, then the recovery procedure works perfectly to return unique descriptors when provided adequate data. For problems where $m \leq v$, it is necessary and sufficient for unique recovery that a bilinear model pass the model check. In problems where $m > v$, the model check algorithm provides a sufficient test of a particular bilinear model, and in such cases we infer nothing from the failure of a particular model to pass the test. For a check of a particular bilinear model to be conducted, it is necessary for E to equal or exceed U.[1]

The primary tool is the singular-value decomposition of the model check matrix **M** representing the generally overdetermined system of E equations in U unknowns, in double precision (64 bits), with use of the LAPACK routine dgesvd.[5] We take a plunge between adjacent values in the spectrum of ordered singular values of 5 log units or greater to indicate the presence of a nontrivial kernel and failure of the check. We take the absence of any difference between adjacent values greater than 1 log unit in magnitude to indicate that the model check matrix has full rank and that the tested model passes the check. These criteria classify the model check matrices readily for all but 3 color constancy problems out of a total of 64 problems examined algorithmically. We present exemplary spectra in the results section (Section 3).

Methods that examined the positive semidefinite symmetric matrices $\mathbf{S} = \mathbf{M}^T\mathbf{M}$ of matrix dimension U × U were used to help confirm the ranks of model check matrices deduced from singular values. The rank of a symmetrized check matrix **S** is necessarily identical to that of the original model check matrix **M**,[6] and we examined each matrix **S** in two ways. First, we calculated the LU decomposition of **S** and examined its spectrum of pivots and confirmed that the spectrum of pivots of **S** matched qualitatively the spectrum of singular values of **M** (ludcmp[7]). Second, we performed a Cholesky factorization (for positive definite symmetric matrices) and inverted **S**, when possible (dpotrf and dpotri[5]). Success here helps confirm that the corresponding matrix **M** has full rank.

The bulk of the work was conducted on a DECstation 5000/200 computer in double precision. In the largest cases, model checks were performed also on a Cray Y-MP computer in long double precision (128 bits).

B. Particular Bilinear Models

A successful check of a single bilinear model with a particular problem's parameters suffices to prove that the problem's parameters permit perfect color constancy schemes. However, for each problem we tested many models to lend reliability to the results.

Each bilinear model is formed of three components: (1) a set of p photoreceptoral spectral sensitivity functions, (2) a linear model for illumination comprising m basis functions, and (3) a linear model for reflectance with n basis functions. In Table 1 are listed the components that we tested for the ranges $2 \leq p \leq 4$, $2 \leq m \leq 4$, and $2 \leq n \leq 15$. The table states which models were tested for each choice of problem parameters ($p\,m\,n\,v\,s$). For example, the problem with parameters ($p\,m\,n\,v\,s$) = (3 3 3 2 3) involved tests of the $5 \times 4 \times 3 = 60$ models formed by taking each possible combination of the five trichromatic systems in column one, the four illumination models of dimension three in column two, and the three reflectance models of dimension three in column three (cf. Ref. 3).

Each function of wavelength was represented by a vector of 31 entries containing the function's values at 10-nm intervals over the range [400 nm, 700 nm]. Each such vector was normalized to have length one. These normalized vectors were then used in calculations of bilinear model matrix entries [cf. Eq. (8) of Ref. 1].

C. Models for Exemplary Spectra

Exemplary spectra of singular values of model check matrices **M** are presented in the Results section (Section 3). These are drawn from bilinear models that use Smith–Pokorny[8] photoreceptors, CIE daylight basis functions for illumination,[9,10] and Fourier basis functions for surface reflectance. The photoreceptoral systems that underlie the exemplary spectra for $p = 2, 3, 4$, are the Smith–Pokorny protanope, the Smith–Pokorny trichromat, and the Smith–Pokorny trichromat augmented by the V'_λ rod sensitivity,[10] respectively. Illumination models of dimension $m = 2, 3, 4$ are provided by the first two, the first three, and the first four CIE daylight basis functions, respectively. Finally, the reflectance models with dimension n ranging from 2 through 15 that underlie the exemplary spectra are the first n functions in the Fourier series decomposition of functions on the interval [400 nm, 700 nm] in the order 1, sin, cos... when 400 nm is shifted to zero.

D. Simulation

In several problems of interest, the model check algorithm provides fewer equations than unknowns or returns a failing result for a particular bilinear model and so is either inapplicable or inconclusive, respectively. By simulating recovery successfully, one can show that particular bilinear models with such a problem's parameters do not produce recovery procedures that are total failures.

We simulated recovery by using the two-stage procedure of D'Zmura,[3] which was generalized in the companion paper.[1] For a simulation with parameters ($p\,m\,n\,v\,s$), we generated randomly (1) 256 sets of v nonnegative and linearly independent illumination spectral power distributions described exactly by a particular linear model for illumination, for each illumination component in Table 1 of dimension m and (2) 256 sets of n linearly independent reflectance functions taking values within [0, 1], described exactly by a particular linear model for reflectance, for each reflectance component of dimension n in Table 1. Photoreceptoral responses were calculated for each choice

Table 1. Tested Bilinear Model Components

Photoreceptors	Illuminants	Reflectances
Two dimensions		
Smith–Pokorny[a] protanope	Judd *et al.*[b]	Cohen[c]
Smith–Pokorny deuteranope[a]	Dixon[d]	Parkkinen *et al.*[e]
Smith–Pokorny tritanope[a]	Fourier (1, sin)	Fourier (1, sin)
	Fourier (1, cos)	
Three dimensions		
Smith–Pokorny[a]	Judd *et al.*[b]	Cohen[c]
Hurvich and Jameson[f]	Dixon[d]	Parkkinen *et al.*[e]
CIE 10° observer[g]	Indoor (D_{65}, A, F_2)[g]	Fourier
Sony XC-007 CCD RGB	Fourier	
Sony XC-711 CCD RGB		
Four dimensions		
Smith–Pokorny + rod V'_λ[g]	Judd *et al.*[b]	Cohen[c]
Hurvich–Jameson + rod V'_λ[f]	Dixon[d]	Parkkinen *et al.*[e]
CIE 10° observer + rod V'_λ	Fourier	Fourier
Gaussians (2σ = 60 nm; m = 460, 520, 580, 640 nm)		
Five dimensions		Parkkinen *et al.*[e]
		Fourier
Six to fifteen dimensions		Fourier

[a]Ref. 8.
[b]Refs. 9 and 10.
[c]Ref. 11.
[d]Ref. 12.
[e]Ref. 13.
[f]Ref. 14.
[g]Ref. 10.

of p-chromatic system listed in Table 1 taken combinatorially with the illumination and the reflectance components. Each triple of photoreceptoral, illumination, and reflectance components was used to generate 256 sets of photoreceptoral responses, and the descriptors returned by the recovery algorithm were compared with the original descriptors. For instance, for the problem with parameters $(p\,m\,n\,v\,s) = (3\,3\,4\,2\,4)$, we tested recovery by each of 60 possible bilinear models (see Table 1) 256 times, for a total of 15,360 simulated recoveries.

Flawless recovery by a model over many simulation trials rules out the possibility that the model is a total failure and shows that the model is, at worst, a partial failure.

E. Analysis

We proved in the companion paper[1] that all bilinear models with parameters (2 2 2 2 2) are total failures, i.e., no algorithm can recover reflectance and illuminant descriptors properly. We showed this by examining the eigenstructure of matrices within the system of equations expressing the condition that sets of lit surfaces be related to quantum catch data in a one-to-one fashion. We analyze other color constancy problems in a similar fashion below, with the aim of determining whether, for all bilinear models with the problems' parameters, there are distinct sets of lit surfaces that provide identical quantum-catch data.

F. Entailments

The results of (1) applying the model check algorithm to demonstrate perfect recovery, (2) simulating recovery successfully to demonstrate that recovery does not fail totally, and (3) analyzing situations under which recovery breaks down are extended in their scope considerably by entailments that hold among color constancy problems.

Table 2 lists both positive (a)–(h) and negative (i)–(p) entailments among problems. Each positive entailment is a proposition of the following form: if there exists a bilinear model with parameters on the left that provides a perfect recovery procedure, then there exists a bilinear model with parameters on the right that also provides a perfect recovery procedure. The positive entailments

Table 2. Entailments

Positive		
	$(p\,m\,n\,v\,s) \Rightarrow (p+1\,m\,n\,v\,s)$	(a)
	$(p\,m\,n\,v\,s) \Rightarrow (p\,m\,n\,v+1\,s)$	(b)
	$(p\,m\,n\,v\,s) \Rightarrow (p\,m\,n\,v\,s+1)$	(c)
	$(p\,m\,n\,v\,s) \Rightarrow (p\,n\,m\,s\,v)$	(d)
	$(m > v)$ and $(p\,m\,n\,v\,s) \Rightarrow (p\,m-1\,n\,v\,s)$	(e)
	$(v > m)$ and $(p\,m\,n\,v\,s) \Rightarrow (p\,m\,n\,v-1\,s)$	(f)
	$(n > s)$ and $(p\,m\,n\,v\,s) \Rightarrow (p\,m\,n-1\,v\,s)$	(g)
	$(s > n)$ and $(p\,m\,n\,v\,s) \Rightarrow (p\,m\,n\,v\,s-1)$	(h)
Negative		
	$\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p-1\,m\,n\,v\,s)$	(i)
	$\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,m\,n\,v-1\,s)$	(j)
	$\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,m\,n\,v\,s-1)$	(k)
	$\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,n\,m\,s\,v)$	(l)
	$(m \geq v)$ and $\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,m+1\,n\,v\,s)$	(m)
	$(v \geq m)$ and $\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,m\,n\,v+1\,s)$	(n)
	$(n \geq s)$ and $\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,m\,n+1\,v\,s)$	(o)
	$(s \geq n)$ and $\sim(p\,m\,n\,v\,s) \Rightarrow \sim(p\,m\,n\,v\,s+1)$	(p)

(a)–(e) reproduce the entailments (74)–(78), respectively, stated in the companion paper.[1] To review: the entailment (a) of Table 2 states that if there exists a p-chromatic bilinear model with parameters ($p\,m\,n\,v\,s$) that provides perfect recovery, then there exists a ($p+1$)-chromatic bilinear model with parameters ($p+1\,m\,n\,v\,s$) that also works. Likewise, offering additional views (b), offering additional surfaces (c), transposing the model (d), and reducing the dimension of the illumination model (e), within limits, will not break a functioning recovery algorithm.

The entailment (e) is embellished by the further entailment (f): if there exists a bilinear model that functions perfectly with use of data from v views in excess of the m views required for revealing the m-dimensional subspace of quantum catch data grouped according to view, then there exists a bilinear model (e.g., the same one) that can make do with data from one fewer view (provided, of course, that the remaining illuminants continue to provide data that span the m-dimensional subspace). The entailments (g) and (h) reformulate the entailments (e) and (f), respectively, in terms of surfaces: entailments (g) and (h) are the transpositions of entailments (e) and (f), respectively, just as entailment (c) is the transposition of entailment (b). As for (g), if a bilinear model with parameters ($p\,m\,n\,v\,s$) provides perfect recovery, then certainly it can recover $n-1$ descriptors per surface from the same information (ignore one of the descriptors). Likewise, if there exists a bilinear model that functions perfectly with use of data from s surfaces in excess of the n surfaces required for revealing the n-dimensional subspace of quantum catch data grouped according to surface, then (h) there exists a bilinear model that can make do with data from one fewer surface (provided, of course, that the remaining surfaces continue to provide data that span the n-dimensional subspace).

Each positive entailment is a proposition that one can negate formally to provide a negative entailment. These negative entailments, (i)–(p), are listed in Table 2; their order corresponds to that of the respective positive entailments from which they are derived. The negative entailments have the following form: if there is no bilinear model with parameters on the left that provides a perfect recovery procedure, then there does not exist a bilinear model with parameters on the right that provides a perfect recovery procedure. It is evident that if there does not exist a bilinear model with parameters ($p\,m\,n\,v\,s$) that provides a perfect recovery procedure, then reducing the information available to the visual system by eliminating a photoreceptoral type (i), denying it a view (j), or removing a surface (k) cannot help but produce further models that fail. Likewise, model transposition (l) cannot help a failing recovery procedure.

The negative entailment (m) states that if no model is able to recover perfectly m descriptors per illuminant, then no model can recover perfectly $m+1$ descriptors per illuminant from the same information. In particular, any failing model with parameters ($p\,m\,n\,v\,s$), when augmented by a further basis function for illumination, will continue to fail when presented illuminants within the original m-dimensional subspace. The negative entailment (n) states that increasing the number v of views beyond the number m, which is the number required for defining uniquely the m-dimensional data subspace spanned by the quantum catch data vectors grouped according to view, will not help a failing model. The entailments (o) and (p) reformulate the entailments (m) and (n), respectively, in terms of surfaces. As for (o), if all bilinear models with parameters ($p\,m\,n\,v\,s$) fail to recover perfectly n descriptors per surface, then there is no model that can recover perfectly $n+1$ descriptors per surface from the same information. Finally, increasing the number of viewed surfaces beyond the number n, which is the number required for defining uniquely the n-dimensional data subspace spanned by the quantum catch data vectors from the surfaces, will not help a failing model (p).

3. RESULTS

We consider dichromatic (Subsection 3.A), trichromatic (Subsections 3.B and 3.C) and tetrachromatic (Subsections 3.D and 3.E) problems in turn.

A. Dichromacy

Figure 1 shows the results for two-stage linear recovery of spectral descriptions by dichromatic systems. The format of the diagram follows that of the corresponding figure in the preceding paper (Fig. 4 of Ref. 1) and is detailed in the caption. Shown immediately beneath the points, each of which represents the parameters of a particular problem

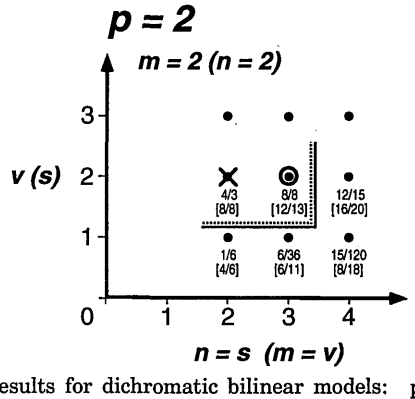

Fig. 1. Results for dichromatic bilinear models: problems involving square bilinear model matrices ($p = m = 2$). The horizontal axis marks the dimension n of the reflectance model, which is taken equal to the number s of surfaces, while the vertical axis marks the number v of views. The solid lines divide cases that satisfy the necessary condition $svp \geq sn + vm - 1$ [inequality (1)]. The number of quantum catch data $Q = svp$ and the number $D = sn + vm$ of spectral descriptors to be recovered are indicated for each problem by the bracketed pair [Q/D] beneath the appropriate point. Points that lie beneath and to the right of the solid lines, where $Q < D - 1$, fail the feasibility condition [inequality (1)] and so represent problems for which unique recovery is impossible. The dotted lines divide cases that satisfy the necessary condition for the test provided by the model check algorithm to be performed, namely, that $E \geq U$ [Eqs. (2) and (3)]. The pair E/U is shown directly beneath each point. In problems where $m \leq v$, such checks provide necessary and sufficient tests of whether particular bilinear models with the problem's parameters provide perfect recovery algorithms. The **X** marks the problem ($p\,m\,n\,v\,s$) = (2 2 2 2 2), for which recovery fails totally.[1] The circled point marks the problem (2 2 3 2 3), for which successful model checks show that there are perfect two-stage linear recovery algorithms with these problem parameters. By transposition [entailments (d) and (l) of Table 2], the result for each problem ($p\,m\,n\,v\,s$) also represents the result for the transposed problem ($p\,n\,m\,s\,v$), and the transposed parameters are indicated in parentheses at the top of the diagram and along its axes. See text for further discussion.

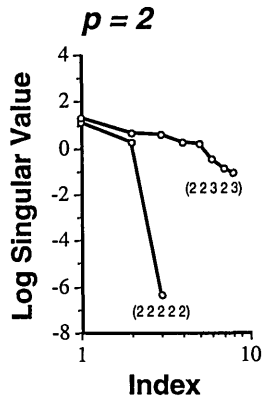

Fig. 2. Spectra of model check matrices for exemplary dichromatic bilinear models. Plotted on a log axis are the ordered singular values of the model check matrices for exemplary bilinear models that combine the Smith–Pokorny protanope,[8] the CIE daylight basis for illumination,[9,10] and the Fourier basis for reflectance (see Table 1). The model check matrix for the exemplary model with parameters $(p\,m\,n\,v\,s) = (2\,2\,2\,2\,2)$ has a kernel of dimension one and fails the check; the matrix for the exemplary model with parameters (2 2 3 2 3) has full rank and passes the check. See text for discussion.

(and its transpose), are (1) the number E of equations and the number U of unknowns provided by the model check algorithm in format E/U and, immediately beneath these, (2) the number Q of quantum catch data and the number D of descriptors in format [Q/D].

The case $(p\,m\,n\,v\,s) = (2\,2\,2\,2\,2)$ is marked with an **X** to indicate that all models with these parameters fail totally: for no such model is unique recovery possible.[1]

The only other interesting feasible dichromatic problem that permits two-stage recovery has parameters (2 2 3 2 3). Because the dimension of the illuminant model is equal to the number of views, one can apply a model check that expresses both necessary and sufficient conditions for perfect recovery. We checked models with all possible combinations of the following components: (1) the three Smith–Pokorny[8] dichromats listed in the left-hand (photoreceptoral) column of Table 1, (2) the four two-dimensional illuminant models in the middle (illumination) column, and (3) the three three-dimensional models for reflectance listed in the right-hand (reflectance) column. Checks of these $3 \times 4 \times 3 = 36$ models showed that their respective 8×8 model check matrices have full rank, so that each of these models functions perfectly. The recovery procedures use 12 quantum catch data to recover 13 spectral descriptors, up to a single scale. We place a circle around the corresponding point in Fig. 1 to indicate this perfect recovery result.

Spectra of ordered singular values of model check matrices for exemplary bilinear models with parameters (2 2 2 2 2) and (2 2 3 2 3) are shown in Fig. 2. For the parameters (2 2 2 2 2), the three singular values are those of the 4×3 model check matrix for the exemplary bilinear model formed by combining the Smith–Pokorny[8] protanope, the first two CIE daylight illumination basis functions,[9,10] and the first two Fourier basis functions, as per Subsection 2.C. The difference in magnitude between the second and the third singular values of over 6 log units shows that the matrix has a kernel of dimension one. Because the parameters of this problem satisfy $m = v$, the model check expresses necessary and sufficient conditions for perfect recovery. Finding that the model check matrix has a kernel of dimension one thus shows that this bilinear model provides a recovery procedure that fails, in agreement with analysis.[1] The other 35 bilinear models with parameters (2 2 2 2 2) that we checked had similar spectra.

Figure 2 also shows that the exemplary bilinear model with parameters (2 2 3 2 3) provides an 8×8 model check matrix of full rank. This single successful check shows that the problem with parameters (2 2 3 2 3) admits of models that function perfectly. The other 35 bilinear models with parameters (2 2 3 2 3) provided spectra similar to that shown.

We have not marked results for the points in the top row (three views) of Fig. 1: because a third view provides information that is (theoretically) in excess of that needed to reveal two-dimensional illuminants, the results for the points in the three-view row are identical to the results for the points immediately below them in the two-view row [entailments (b) and (n) of Table 2].

The parenthesized parameters that label Fig. 1 transpose the parameters of the problems in the diagram. Transposition [entailment (d) of Table 2] shows that the success of (2 2 3 2 3) implies the success of (2 3 2 3 2), which is the problem in which three-dimensional illuminants and two-dimensional reflectances are recovered from three views of two surfaces by a dichromatic system. The results for the problems (2 2 2 2 2), (2 2 3 2 3), and (2 3 2 3 2) exhaust the two-stage linear recovery possibilities for dichromacy.

It is well to note, at this point, that there are further linear recovery possibilities for dichromacy that are not pictured in Fig. 1. For instance, is it possible for a dichromatic system to recover three-dimensional spectral descriptions for both illuminants and reflectances? Although a comparison of the numbers of quantum catch data and unknown descriptors shows that it is feasible, the problem $(p\,m\,n\,v\,s) = (2\,3\,3\,3\,3)$ falls outside the scope of two-stage linear recovery procedures, because both m and n exceed p: the bilinear model matrices $\mathbf{B}_j, j = 1,2,3$, and their transposes $\mathbf{B}_i^t, i = 1,2,3$, have matrix dimensions 2×3 and so do not possess unique left inverses.

B. Trichromacy

Figures 3A and 3B show results for the trichromatic problems $(3\,3\,n\,v\,s)$ and $(3\,2\,n\,v\,s)$, respectively, in cases where $s \geq n$. The bracketed parameters label Figs. 3A and 3B for the transposed problems $(3\,m\,3\,v\,s)$ and $(3\,m\,2\,v\,s)$, respectively, in cases where $v \geq m$. The circled points mark problems for which perfect recovery is possible; these results are shown through application of the model check algorithm. The points with squares mark problems shown by simulation of recovery to provide, at worst, imperfect recovery procedures. The triangles, finally, mark problems shown by analysis to allow perfect recovery.

Figures 4A and 4B show the spectra of singular values for the model check matrices of exemplary bilinear models (Subsection 2.C) with parameters that correspond to the checkable problems of Figs. 3A and 3B, respectively. The spectra are shown, from bottom to top, in order of increasing view, and for a particular choice of view, in order of increasing dimension n for surface reflectance. We have

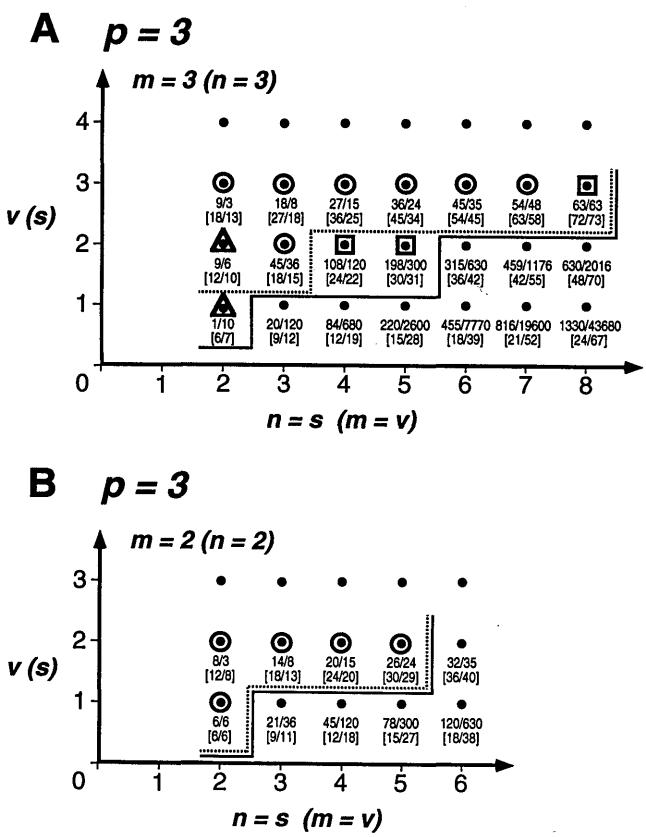

Fig. 3. Results for trichromatic bilinear models. A, The case of square bilinear model matrices ($p = m = 3$ or $p = n = 3$); B, the case of rectangular bilinear model matrices ($p = 3, m = 2$ or $p = 3, n = 2$). The triangles mark problems that are shown by analysis to provide perfect recovery algorithms. The squares mark problems that are shown by successful simulation of recovery to provide, at worst, imperfect recovery. As in Fig. 1, the circles mark problems that are shown by the model check algorithm to support perfect recovery procedures. See the caption for Fig. 1 and text for further details.

scaled the spectra to stagger their maximal singular values along the vertical axis at half log unit intervals.

We turn first to the problems of Fig. 3A (with exemplary spectra in Fig. 4A).

1. (3 3 n 2 s), $s \geq n$

The problem with parameters (3 3 2 2 2) produces bilinear models that fail the model check miserably. Each of the $5 \times 4 \times 3 = 60$ bilinear models with these parameters, drawn from Table 1, provides a 9×6 model check matrix with a kernel of dimension three, a feature displayed by the exemplary spectrum (Fig. 4A). We defer further consideration of this problem until Subsection 3.C, where we specify the conditions under which recovery procedures applied to models with these parameters and the related ones with parameters (3 3 2 1 2) work.

The model check algorithm provides a sufficient test of bilinear models with parameters (3 3 3 2 3), the case examined by D'Zmura.[3] Each of the 45×36 model check matrices of the $5 \times 4 \times 3 = 60$ bilinear models with these parameters (Table 1) is of full rank.

The model check algorithm provides no test for models with parameters (3 3 4 2 4) and (3 3 5 2 5). We simulated recovery numerically in these cases to establish that recovery by bilinear models with these parameters need not be a total failure (see Subsection 2.D). For the problem with parameters (3 3 4 2 4), we tested recovery by each of $5 \times 4 \times 3 = 60$ bilinear models drawn from Table 1 256 times, for a total of 15,360 successfully simulated recoveries. For the problem with parameters (3 3 5 2 5), we tested recovery by each of $5 \times 4 \times 2 = 40$ models drawn from Table 1 256 times, for a total of 10,240 successfully simulated recoveries. These findings rule out the possibility that all models with these parameters are total failures, and the problems are marked accordingly by squares in Fig. 3A.

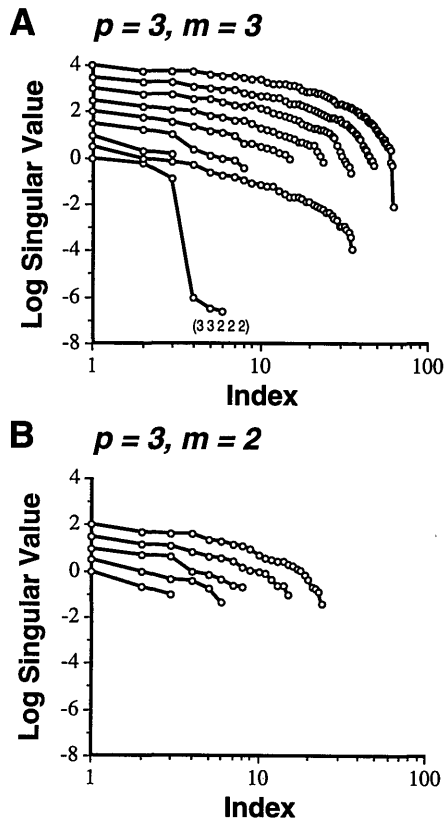

Fig. 4. Spectra of model check matrices for exemplary trichromatic bilinear models. A, The case of square bilinear model matrices ($p = m = 3$). The Smith–Pokorny trichromat[8] and the CIE daylight basis[9,10] were used in combination with Fourier reflectance models of dimension n ranging from two through eight. Plotted on a log axis are the ordered singular values of the model check matrices for these exemplary bilinear models. The spectra are shown, from bottom to top, in order of increasing view and, for a particular choice of view, in order of increasing dimension n for surface reflectance. We have scaled the spectra to stagger the maximal singular values along the vertical axis at half log unit intervals. The parameters of the exemplary models whose spectra are shown are, from bottom to top, (3 3 2 2 2), (3 3 3 2 3), (3 3 3 3 3), (3 3 4 3 4), (3 3 5 3 5), (3 3 6 3 6), (3 3 7 3 7), and (3 3 8 3 8). The model with parameters (3 3 2 2 2) provides a matrix with a kernel of dimension three and fails the model check; with the possible exception of the model with parameters (3 3 8 3 8), the remaining models provide matrices of full rank and so pass the model check. B, The case of rectangular bilinear model matrices ($p = 3, m = 2$). The Smith–Pokorny trichromat[8] and the first two CIE daylight basis functions[9,10] were used in combination with Fourier reflectance models of dimension n ranging from two through five. The spectra of the model check matrices for these exemplary bilinear models are ordered and staggered as in A; the parameters are, from bottom to top, (3 2 2 1 2), (3 2 2 2 2), (3 2 3 2 3), (3 2 4 2 4), and (3 2 5 2 5). All provide matrices of full rank and so pass the model check. See text for further discussion.

2. (3 3 n 3 s), $s \geq n$

Results for the problems with three views show that a trichromatic visual system can use three views to recover anywhere from two to eight reflectance descriptors per surface (Fig. 3A, $v = 3$; Fig. 4A, top seven spectra). We obtained these results by applying the necessary and sufficient tests of unique recovery provided by the model check algorithm to problems for which $m = v$. Starting at the left ($n = s = 2$) and working toward the right, these checks involved the following: for (3 3 2 3 2), determining that the $5 \times 4 \times 3 = 60$ bilinear models formed of the appropriate components of Table 1 provide 9×3 model check matrices of full rank; for (3 3 3 3 3), checking successfully the 18×8 matrices of $5 \times 4 \times 3 = 60$ models; for (3 3 4 3 4), checking successfully the 27×15 matrices of $5 \times 4 \times 3 = 60$ models; for (3 3 5 3 5), checking successfully the 36×24 matrices of $5 \times 4 \times 2 = 40$ models; for (3 3 6 3 6), checking successfully the 45×35 matrices of $5 \times 4 \times 1 = 20$ models; and for (3 3 7 3 7), checking successfully the 54×48 matrices of $5 \times 4 \times 1 = 20$ models.

For the problem with parameters (3 3 8 3 8), we checked the 63×63 matrices $\mathbf{M}$ for $5 \times 4 \times 1 = 20$ models and found equivocal spectra similar to the top spectrum shown in Fig. 4A. The symmetrized check matrices $\mathbf{S}$ were poorly conditioned (dsycon, Ref. 5). However, there were no sharp falloffs (of magnitude 5 log units or more, Subsection 2.A) in either the spectra of singular values or the spectra of pivots from the LU decomposition of the matrices $\mathbf{S}$. Furthermore, the Cholesky factorization and inversion routines for positive definite symmetric matrices returned putative inverses of the matrices $\mathbf{S}$, albeit inaccurate ones. Finally, simulation of recovery (Section 2.D) proceeded flawlessly. The problem with parameters (3 3 8 3 8) is the first of three problems for which the model check algorithm, although applicable, does not produce a clear answer concerning rank.

Again, results for the points in the top row (four views) of Fig. 3A are not indicated; by entailments (b) and (n) these are identical to the results for the points immediately below them.

3. (3 2 n v s), $s \geq n$

The results for two-dimensional models of illumination are shown in Fig. 3B; spectra for exemplary models are shown in Fig. 4B. There are five feasible problems, all of which can be checked with the model check algorithm. The results of the checks show that all five problems permit perfect recovery. They are taken in order of increasing n in what follows.

The problem (3 2 2 1 2) is represented by bilinear models for which the model check algorithm provides a sufficient check ($m > v$) through tests of whether 6×6 model check matrices are of full rank. These checks were performed on the $5 \times 4 \times 3 = 60$ models formed of the appropriate components listed in Table 1. All models pass the check, so that in all cases perfect recovery obtains. Note that the result for this problem follows by entailment (e) from the successful result for the problem (3 3 2 1 2), which we present in Subsection (3.C).

The perfect recovery result for (3 2 2 1 2) implies a perfect recovery result for the problem (3 2 2 2 2) by entailment (b) of Table 2; likewise, the result for (3 3 2 2 2), analyzed in Subsection 3.C, implies success for (3 2 2 2 2) through entailment (e). All the $5 \times 4 \times 3 = 60$ models with parameters (3 2 2 2 2) had 8×3 model check matrices of full rank.

Third, the problem (3 2 3 2 3) is the transpose of the problem (3 3 2 3 2), which was checked successfully (Fig. 3A). Indeed, success for (3 2 3 2 3) follows also through entailment (a) from the success of (2 2 3 2 3). Furthermore, the result for (3 3 3 2 3) implies success for (3 2 3 2 3) through entailment (e). Wholly unnecessary checks of the 45×36 model check matrices of the $5 \times 4 \times 3 = 60$ bilinear models with the parameters (3 2 3 2 3) nevertheless were performed, all with success.

Finally, necessary and sufficient model checks were performed successfully for the two problems (3 2 4 2 4) and (3 2 5 2 5). For (3 2 4 2 4), these involved checking that the 20×15 model check matrices for $5 \times 4 \times 3 = 60$ distinct bilinear models formed of the appropriate components from Table 1 have full rank; for (3 2 5 2 5) these involved checks of the ranks of 26×24 matrices for $5 \times 4 \times 2 = 40$ distinct bilinear models. The results for three views in Fig. 3B are, by entailments (b) and (n), identical to those for two views with all other parameters held constant.

The computational results show that a trichromatic visual system can recover perfectly as many as seven and possibly eight surface color descriptors from three views provided by distinct three-dimensional illuminants. It can recover three and possibly as many as five reflectance descriptors perfectly from two views, while from one view it can recover perfectly only two reflectance descriptors (see also Subsection 3.C).

There are further feasible trichromatic problems, not pictured in Fig. 3, that lie outside the scope of two-stage linear recovery. While the linear recovery results for the use of a single view are complete, there are several problems [e.g., (3 4 4 2 4)] in which the number of views is two that are not within the purview of two-stage linear recovery. Another class of problem is exemplified by (3 3 3 2 2),[15] in which the number of views and the number of surfaces are fewer than the dimensions for illumination and reflectance, respectively.

C. Recovery of Two Reflectance Descriptors

The two problems with parameters ($p\,m\,n\,v\,s$) = (3 3 2 1 2) and (3 3 2 2 2), marked by triangles in Fig. 3A, are feasible, yet the model check proves unable to determine whether there are models with these parameters that permit perfect recovery: (1) the model check cannot be performed on models with parameters (3 3 2 1 2), owing to an excess of unknowns U over equations E, and (2) the model check matrices for bilinear models with the parameters (3 3 2 2 2) have kernels of dimension three. While the latter finding is inconclusive, because model checks for cases where $m > v$ are sufficient but not necessary for showing unique recovery, it prompted us to analyze these two problems.

We show here that all such problems can provide perfect recovery procedures, despite the fact that there are always distinct sets of lit surfaces that provide identical quantum catch data when they are seen through bilinear models with parameters (3 3 2 1 2) or (3 3 2 2 2). This lack of a one-to-one relationship between sets of lit surfaces and quantum catch data accounts for the failure of the model checks with parameters (3 3 2 2 2). However, it need not be the case that the sets of lit surfaces that cause recovery

to fail are physically realizable. Through a judicious choice of bilinear model, one need never encounter in practice a set of lit surfaces that causes recovery to fail.

1. The Problem (3 3 2 1 2)

We consider first the problem with parameters $(p\,m\,n\,v\,s) = (3\,3\,2\,1\,2)$; the findings here generalize readily to that with parameters (3 3 2 2 2). Our analysis starts with equations that express the necessary and sufficient conditions for a bilinear model with parameters (3 3 2 1 2) to provide a one-to-one relationship between sets of lit surfaces and quantum catch data:

$$\mathbf{z} = (e_{11}\mathbf{I} + e_{12}\boldsymbol{\Gamma}_{12})\mathbf{a} = (e_{21}\boldsymbol{\Gamma}_{21} + e_{22}\mathbf{I})\mathbf{a}. \quad (4)$$

As described in the companion paper [cf. Eqs. (39)–(42) of Ref. 1], (1) the 3×1 column vectors $\mathbf{z}$ and $\mathbf{a}$ hold illuminant descriptors, (2) the identity matrix $\mathbf{I}$ and the gamma matrices $\boldsymbol{\Gamma}_{12}$ and $\boldsymbol{\Gamma}_{21}$ are 3×3 matrices, and (3) the gamma matrices are determined by bilinear model matrices according to $\boldsymbol{\Gamma}_{12} = \mathbf{B}_1^{-1}\mathbf{B}_2 = \boldsymbol{\Gamma}_{21}^{-1}$ [Eq. (40) of Ref. 1). Note that if one or both of the bilinear model matrices $\mathbf{B}_1$ and $\mathbf{B}_2$ are singular, then recovery fails; examining the conditions for perfect recovery that are expressed in Eq. (4) is sensible only if the model matrices, and hence the gamma matrices, are invertible. Finally, the variables e_{11}, e_{12}, e_{21}, and e_{22} relate two sets of surfaces seen under the illuminant $\mathbf{z}$ and the illuminant $\mathbf{a}$, respectively.

The system of Eq. (4) must force the single scaling family of solutions $e_{11} - e_{22} = e_{12} = e_{21} = 0$ if sets of lit surfaces are to bear a one-to-one relationship with quantum catch data.[1] Yet one can readily show that there are other solutions to this system. The details depend on the eigenstructure of the matrix $\boldsymbol{\Gamma}_{12}$. The case of two or more distinct eigenvalues is the most important from the practical point of view; each of the 60 distinct bilinear models for $p = 3$, $m = 3$, and $n = 2$ drawn from Table 1 has a gamma matrix $\boldsymbol{\Gamma}_{12}$ with this eigenstructure. Let us assume, then, that $\boldsymbol{\Gamma}_{12}$ has two distinct eigenvalues that are nonzero; call these λ_1 and λ_2 and suppose that they are real. To these correspond the two independent 3×1 eigenvectors $\boldsymbol{\varepsilon}_1$ and $\boldsymbol{\varepsilon}_2$, respectively. If we choose the illuminant described by $\mathbf{a}$ to lie within the span of the eigenvectors $\boldsymbol{\varepsilon}_1$ and $\boldsymbol{\varepsilon}_2$, then we can show the existence of nonscaling solutions to Eq. (4).

Take $\mathbf{a}$ to describe an illuminant of the form

$$\mathbf{a} = \alpha\boldsymbol{\varepsilon}_1 + \beta\boldsymbol{\varepsilon}_2, \quad (5)$$

where α and β are nonzero constants. Applying to this illuminant vector $\mathbf{a}$ the matrices $e_{11}\mathbf{I} + e_{12}\boldsymbol{\Gamma}_{12}$ and $e_{21}\boldsymbol{\Gamma}_{12}^{-1} + e_{22}\mathbf{I}$ of Eq. (4), we find that

$$(e_{11} - e_{22})(\alpha\boldsymbol{\varepsilon}_1 + \beta\boldsymbol{\varepsilon}_2) + e_{12}(\alpha\lambda_1\boldsymbol{\varepsilon}_1 + \beta\lambda_2\boldsymbol{\varepsilon}_2) - e_{21}(\alpha\boldsymbol{\varepsilon}_1/\lambda_1 + \beta\boldsymbol{\varepsilon}_2/\lambda_2) = 0. \quad (6)$$

Now the eigenvectors $\boldsymbol{\varepsilon}_1$ and $\boldsymbol{\varepsilon}_2$ are independent, so that from Eq. (6) one has

$$\alpha[(e_{11} - e_{22}) + \lambda_1 e_{12} - e_{21}/\lambda_1] = \beta[(e_{11} - e_{22}) + \lambda_2 e_{12} - e_{21}/\lambda_2] = 0, \quad (7)$$

and by virtue of $\alpha, \beta \neq 0$, the following two equations:

$$(e_{11} - e_{22}) + \lambda_1 e_{12} - e_{21}/\lambda_1 = 0$$
$$(e_{11} - e_{22}) + \lambda_2 e_{12} - e_{21}/\lambda_2 = 0. \quad (8)$$

The system of Eqs. (8) has the following family of solutions, which includes solutions other than the scaling solutions:

$$e_{12} = c,$$
$$e_{21} = -c\lambda_1\lambda_2,$$
$$e_{11} - e_{22} = -c(\lambda_1 + \lambda_2), \quad (9)$$

where c is an arbitrary real constant. These solutions are identical in form to those for the problem (2 2 2 2 2) [Eqs. (59) of Ref. 1].

Note that if the two distinct eigenvalues form a complex-conjugate pair, then the solutions for the variables e_{ij} are real [Eqs. (9)]; in such a case, one can choose the coefficients α and β appropriately to form an illuminant $\mathbf{a}$ with real descriptors.

The existence of the nonscaling solutions [$c \neq 0$ in Eqs. (9)] shows that there are distinct sets of surfaces that produce identical quantum catch data when viewed under illuminants lying within the span of two (distinct) eigenvectors of the gamma matrix $\boldsymbol{\Gamma}_{12} = \mathbf{B}_1^{-1}\mathbf{B}_2$. If the gamma matrix has, in fact, three real distinct eigenvalues λ_1, λ_2, and λ_3 that correspond to three linearly independent eigenvectors $\boldsymbol{\varepsilon}_1$, $\boldsymbol{\varepsilon}_2$, and $\boldsymbol{\varepsilon}_3$, then there are three sets of nonscaling solutions that are like those of Eqs. (9). Each of the three planes spanned by the three possible pairs of eigenvectors, namely, span[$\boldsymbol{\varepsilon}_1, \boldsymbol{\varepsilon}_2$], span[$\boldsymbol{\varepsilon}_1, \boldsymbol{\varepsilon}_3$], and span[$\boldsymbol{\varepsilon}_2, \boldsymbol{\varepsilon}_3$], represents a distinct set of illuminants for which recovery fails. Note that if the illuminant $\mathbf{a}$ is chosen to lie off one of these planes, then Eq. (4) has only the scaling solutions $e_{11} - e_{22} = e_{12} = e_{21} = 0$.

Figure 5 shows that bilinear models with parameters (3 3 2 1 2) can provide either imperfect recovery procedures (Fig. 5A) or perfect recovery procedures (Fig. 5B), depending on whether the planes of failure contain physically realizable (nonnegative) illuminants or do not, respectively. The cones in Fig. 5 mark the subset of illuminant descriptors that represent physically realizable light

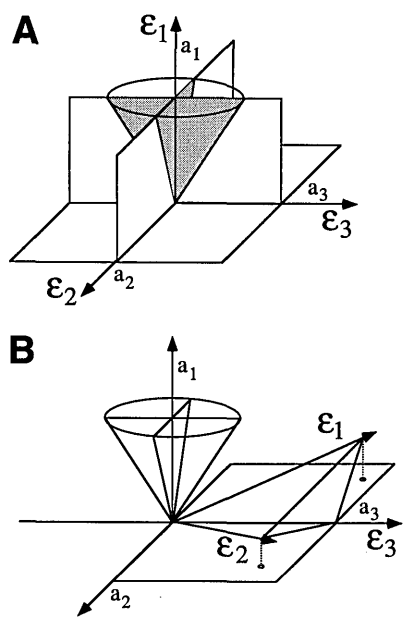

Fig. 5. Possible failure of bilinear models with parameters (3 3 2 1 2) or (3 3 2 2 2). See text for further discussion.

sources. To see that the cone is a reasonable representation, consider the three-dimensional Fourier model of illumination $\{1, \sin, \cos\}$, and the limits on modulation of descriptors a_2 and a_3 relative to the value of a_1 that are required for maintaining nonnegativity.[16]

If any one of the three planes of failure span$[\varepsilon_1, \varepsilon_2]$, span$[\varepsilon_1, \varepsilon_3]$, or span$[\varepsilon_2, \varepsilon_3]$ passes through the cone of physically realizable illuminants, then the recovery procedure is imperfect. The shaded areas in Fig. 5A represent physically realizable illuminants that cause the recovery procedure to fail.

If none of the three planes of failure contains (nonzero) physically realizable illuminants, then the recovery procedure is perfect. If the bilinear model has three eigenvectors that are like those shown in Fig. 5B, then in no case can an illuminant lying in a plane of failure spanned by two eigenvectors be realized physically. It is always possible for one to choose the eigenvectors of Γ_{12} so that the planes span$[\varepsilon_1, \varepsilon_2]$, span$[\varepsilon_1, \varepsilon_3]$, and span$[\varepsilon_2, \varepsilon_3]$ do not pass through the cone of physically realizable illuminants and to determine a bilinear model consistent with these eigenvectors. In conclusion, there exist bilinear models with parameters (3 3 2 1 2) that provide perfect recovery procedures.

The 60 models of Table 1 with parameters (3 3 2 1 2) provide examples of both situations that are represented in Fig. 5.

2. The Problems (3 3 2 2 2) and (3 3 2 3 2)

The eigenstructure of the matrix Γ_{12}, for some choice of bilinear model with parameters that satisfy $p = 3$, $m = 3$, and $n = 2$, does not depend on the number of views (or the number of surfaces), so that we have the following progression in results with increasing number of views: if one illuminant can be chosen to lie within a single plane spanned by a pair of eigenvectors, then recovery by a model with parameters (3 3 2 1 2) will fail, as shown in the preceding section. A similar derivation shows that if two linearly independent illuminants are chosen to lie within a single plane spanned by a pair of eigenvectors, then recovery by a model with parameters (3 3 2 2 2) will fail. However, it is impossible to choose three linearly independent illuminants that lie within a single plane spanned by a pair of eigenvectors, and the provision of perfect recovery procedures by the problem with parameters (3 3 2 3 2) is indicated by the success of the model check.

3. Problems of the Form (pp 2 1 2) and (pp 2 2 2)

The description of the sets of failure for the problem (3 3 2 1 2) [Eqs. (5)–(9)] can be generalized to problems with parameters (pp 2 1 2). There are various possible eigenstructures for the $p \times p$ matrix Γ_{12}. Yet if the matrix has at least two distinct eigenvalues, then there is at least one plane of failure in the p-dimensional space of illuminants, for which there are nonscaling solutions among surfaces that are like those described by Eqs. (9). If one supposes that the matrix has p distinct real eigenvalues, then there are $p(p-1)/2$ such planes of failure. With two views, but not three or more, it is possible to choose the illuminants to lie in a single such plane. Because the model check algorithm is insensitive to physical realizability, checks will always fail for models with parameters (pp 2 1 2) and (pp 2 2 2) yet succeed for (pp 2 3 2), etc. Determining whether the planes of failure for a particular bilinear model are positioned in such a way as to cause the recovery procedure to be imperfect is a programming problem not pursued here. Yet it seems likely that one can always construct a bilinear model such that the eigenvectors of Γ_{12} fall in positions that cause failing illuminants to be physically unrealizable. If this is the case, there exist models with these parameters for which recovery algorithms function perfectly.

D. Tetrachromacy

Figures 6A, 6B, and 6C show results for the tetrachromatic problems (4 4 n v s), (4 3 n v s), and (4 2 n v s), respectively, in cases where $s \geq n$. The bracketed parameters label the panels for the transposed problems (4 m 4 v s), (4 m 3 v s), and (4 m 2 v s), respectively, in cases where $v \geq m$. Figure 6D shows results for the tetrachromatic problems (4 4 n v s) and (4 3 n v s) in cases where $s \geq n > 8$ and completes the diagrams shown in Figs. 6A and 6B. An analytic result for (4 4 3 1 3) is derived in Subsection 3.E.

Figures 7A, 7B, 7C, and 7D show the spectra of singular values for the model check matrices of exemplary bilinear models with parameters that correspond to the checkable problems of Figs. 6A, 6B, 6C, and 6D, respectively. From bottom to top, these are shown in order of increasing view and, for a particular number of views, in order of increasing reflectance dimension. As in Fig. 4, we have scaled the spectra to stagger their maximal singular values along the vertical axis at half log unit intervals.

We discuss here more briefly the results found through computation and start with the problems of Fig. 6A.

1. (4 4 n 2 s), $s \geq n$

As suggested by the exemplary spectra of Fig. 7A (bottom three curves), the bilinear models with parameters (4 4 2 2 2), (4 4 3 2 3), and (4 4 4 2 4) failed their model checks. Because $m > v$ in these problems, such failure is not conclusive.

The problem (4 4 2 2 2) has the form (pp 2 2 2), and we know from earlier considerations (Subsection 3.C.3) that models with these parameters will never pass the check, because the check is insensitive to the criterion of physical realizability. While all models tested had model check matrices with kernels of dimension six, it may well be possible to position planes of failure so that they correspond to illuminants (viz., pairs of illuminants) that are not physically realizable. Similar comments hold for the problem (4 4 2 1 2).

The problem (4 4 3 2 3) gives rise to model check matrices with kernels of dimension one. This is a situation in which one can extend the model check algorithm. We examined, for each of $4 \times 3 \times 3 = 36$ models drawn from Table 1, whether the kernel's ray possessed values for the 120 monomial unknowns of degree $m - v + 1 = 3$ that were consistent with the initial equations in the underlying $n^2 - 1 = 8$ variables [cf. Eqs. (67)–(73) of Ref. 1]. The kernel ray of each tested model had four nonzero monomial unknowns that had values inconsistent with the original equations. This shows that there are bilinear models with parameters (4 4 3 2 3) that provide perfect recovery algorithms.

The model check algorithm, applied to the $4 \times 3 \times 3 = 36$ models of Table 1 with parameters (4 4 4 2 4), produces

880 × 680 matrices with kernels of dimension eight: rather disappointingly, the checks for this analytically difficult problem failed. We simulated recovery successfully 9216 times to show that bilinear models with these parameters provide recovery procedures that are, at worst, imperfect.

The surfeit of unknowns over equations in the model checks for the problems with parameters (4 4 5 2 5), (4 4 6 2 6), and (4 4 7 2 7) led us to simulate recovery in these cases. Successful recoveries in 6144, 3072, and 3072 trials, respectively, shows that there are bilinear models with these parameters that are, at worst, imperfect.

2. (4 4 n 3 s), s ≥ n

The model check algorithm was applied with success to all bilinear models drawn from Table 1 with parameters matching (4 4 2 3 2) through (4 4 8 3 8); this uniform result of perfect recovery is indicated by the circles of the third ($v = 3$) row of Fig. 6A. The seven middle curves of Fig. 7A show spectra for the exemplary bilinear models.

The results for three views continue in the middle row of Fig. 6D, labeled (4 4 *n* 3 *n*). The feasible problems (4 4 9 3 9), (4 4 10 3 10), and (4 4 11 3 11) possess more unknowns than equations in the model check algorithm, so we simulated recovery for bilinear models with these parameters. The successful simulations are indicated by the squares.

3. (4 4 n 4 s), s ≥ n

That the model checks for problems with parameters (4 4 2 4 2) through (4 4 8 4 8) (the fourth row of Fig. 6A) were successful is implied by entailment (b) and the successful results for three views (the circled points in the third row). The model checks were conducted, nevertheless, all with success. The top seven curves of Fig. 7A show the spectra for the corresponding exemplary bilinear models.

The results for four views are continued in Fig. 6D, in the top row labeled (4 4 *n* 4 *n*). The model checks for reflectance dimensions 9 through 14 were uniformly successful, but those for the problem (4 4 15 4 15) produced equivocal results (see the spectra for the corresponding exemplary bilinear models in Fig. 7D, top seven curves). Successful simulated recovery in 3072 trials shows that a tetrachromatic system can use four views to recover 15 reflectance descriptors per surface in a way that fails occasionally, at worst.

Again, the results for five views (top row of Fig. 6A) are identical to those for four views, by entailments (b) and (n) of Table 2, when all other parameters are held constant.

4. (4 3 n v s), s ≥ n

Figure 6B shows results for tetrachromatic visual systems with use of a three-dimensional model for illumination

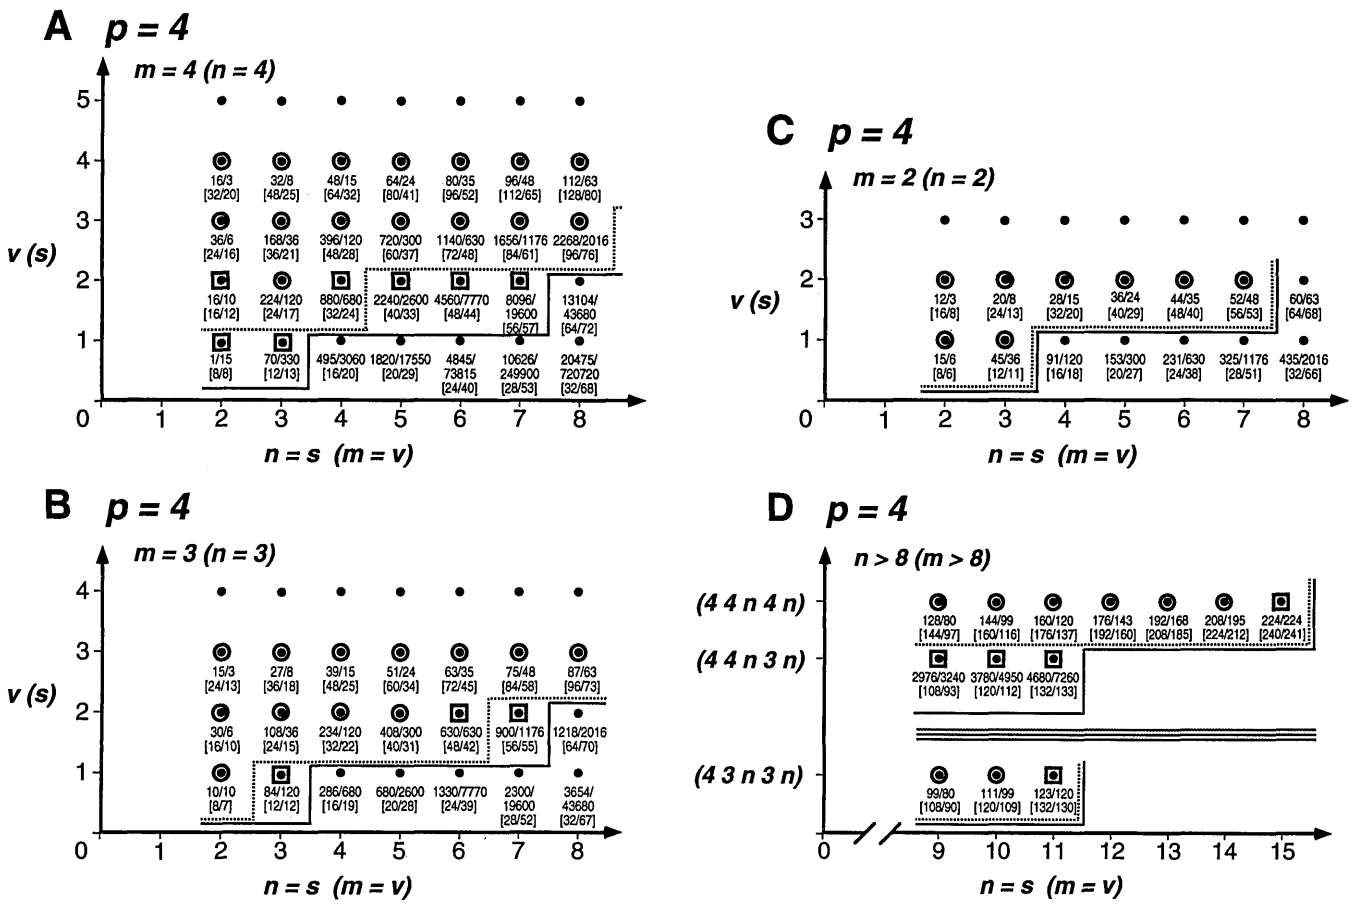

Fig. 6. Results for tetrachromatic bilinear models. A, Problems that involve square bilinear model matrices ($p = m = 4$ or $p = n = 4$); B,C, problems that involve rectangular bilinear model matrices (B: $p = 4, m = 3$ or $p = 4, n = 3$; C: $p = 4, m = 2$ or $p = 4, n = 2$); D, problems that involve the recovery of greater than eight reflectance descriptors per surface (or greater than eight descriptors per illuminant). See the captions for Figs. 1 and 3 and text for further discussion.

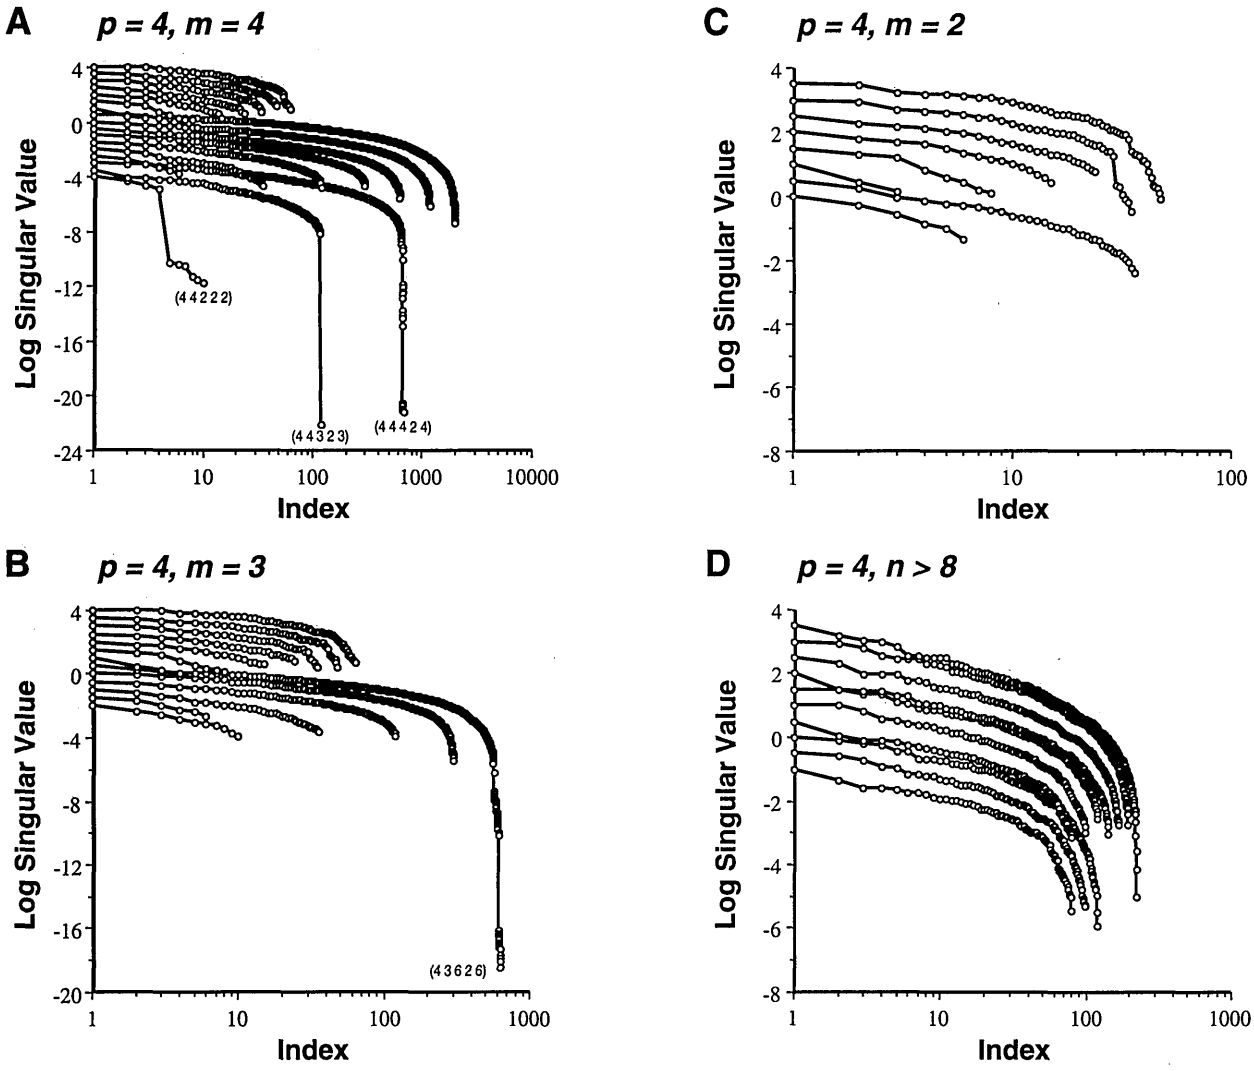

Fig. 7. Spectra of model check matrices for exemplary tetrachromatic bilinear models. The Smith–Pokorny[8] trichromat and the rod V'_λ photoreceptoral sensitivities, together with the CIE daylight basis,[9,10] were used in combination with Fourier reflectance models of dimension n ranging from two through fifteen for formation of exemplary bilinear models. Plotted on a log axis are the ordered singular values of the model check matrices for these exemplary bilinear models. The spectra are shown, from bottom to top, in order of increasing view, and for a particular choice of view, in order of increasing dimension n for surface reflectance. We have scaled the spectra to stagger the maximal singular values along the vertical axis at half log unit intervals. A, The case of square bilinear model matrices ($p = m = 4$, $n \leq 8$; see Fig. 6A). The parameters of the exemplary models whose spectra are shown are, from bottom to top, $v = 2$: (44222), (44323), and (44424); $v = 3$: (44232), (44333), (44434), (44535), (44636), (44737), and (44838); $v = 4$: (44242), (44343), (44444), (44545), (44646), (44747), and (44848). Note that all the models that use two views fail the model check. B, The case of rectangular bilinear model matrices ($p = 4$, $m = 3$, $n \leq 8$; see Fig. 6B). The spectra of the model check matrices for the exemplary bilinear models are ordered and staggered as in A; the parameters are, from bottom to top, $v = 1$: (43212); $v = 2$: (43222), (43323), (43424), (43525), and (43626); $v = 3$: (43232), (43333), (43434), (43535), (43636), (43737), and (43838). The exemplary models provide matrices of full rank and so pass the model check, with the exception of the model with parameters (43626). C, The case of rectangular bilinear model matrices ($p = 4$, $m = 2$, $n \leq 8$; see Fig. 6C). The spectra of the model check matrices for the exemplary bilinear models are ordered and staggered as in A; the parameters are, from bottom to top, $v = 1$: (42212), (42313); $v = 2$: (42222), (42323), (42424), (42525), (42626), and (42727). All provide matrices of full rank and so pass the model check. D, Problems that involve the recovery of greater than eight reflectance descriptors per surface (see Fig. 6D). From bottom to top, the parameters of the exemplary bilinear models whose spectra are exhibited are (43n3s): (43939), (43103 10), and (43113 11); (44n4s): (44949), (44104 10), (44114 11), (44124 12), (44134 13), (44144 14), and (44154 15). See text for further discussion.

(or, by transposition, a three-dimensional model for reflectance). A large number of these results follow from earlier results by entailment. For instance, the results for the problems (43232) through (43737) for three views follow, by entailment (a) of Table 2, from earlier results on the use of three views by trichromatic systems (Fig. 3A). Results for the problems (43232) through (43838) for three views follow by entailment (e) from earlier results, for four-dimensional models of illumination, on the use of three views by tetrachromatic visual systems (Fig. 6A). Model checks were performed, nevertheless, on all problems for which this was possible.

Starting with one view (Fig. 6B, bottom row), the model check algorithm shows that two reflectance descriptors can be recovered perfectly from one view when a three-dimensional model for illumination is used. This finding

implies, by entailment (b), the positive results for the problems (4 3 2 2 2), (4 3 2 3 2), etc., which are represented by the points lying above (4 3 2 1 2).

We are unable to prove or disprove perfect recovery by bilinear models with parameters (4 3 3 1 3); successful simulations of recovery show that the problem is not a total failure.

The model check algorithm shows that two views may be used to recover perfectly two through five reflectance descriptors (Fig. 6B, $v = 2$); this result is supported by the corresponding exemplary spectra of Fig. 7B (bottom four curves). These results imply, by entailment (b), positive results for the corresponding problems with three or more views.

Model checks for the problem (4 3 6 2 6), however, produce matrices with kernels of rank 18. We simulated recovery successfully 4096 times for models with these parameters. We did the same for models with the parameters (4 3 7 2 7), thus showing that a tetrachromatic visual system can use two views, in conjunction with a three-dimensional model for illumination, to recover, at worst imperfectly, seven reflectance descriptors per surface.

Checks for the bilinear models from Table 1 that use three views, with parameters ranging from (4 3 2 3 2) through (4 3 10 3 10), were uniformly successful. Exemplary spectra for $2 \leq n \leq 8$ are shown in Fig. 7B (top seven curves). The problems with parameters (4 3 9 3 9), (4 3 10 3 10), and (4 3 11 3 11) are represented in the bottom row, labeled (4 3 n 3 n), of Fig. 6D. Their exemplary spectra are shown in Fig. 7D (bottom three curves). As suggested by the exemplary spectrum, the model checks for the problem (4 3 11 3 11) produced equivocal results. Successful simulated recovery in 4096 trials shows that a tetrachromatic system with a three-dimensional model for illumination can use 3 views to recover 11 reflectance descriptors per surface in a way that, at worst, fails occasionally.

5. (4 2 n v s), $s \geq n$

The final diagram (Fig. 6C) shows results for problems in which a tetrachromatic visual system uses a two-dimensional model for illumination. The model checks were uniformly successful, as suggested by the exemplary spectra of Fig. 7C, showing that such a system can recover as many as three reflectance descriptors from one view and as many as seven reflectance descriptors from two views.

E. The Problem (4 4 3 1 3)

We conclude by showing that there are always distinct sets of lit surfaces that provide identical quantum catch data when seen through bilinear models with parameters (4 4 3 1 3). The uniqueness equations for this problem [cf. Eqs. (39)–(44) of Ref. 1] involve 4×4 gamma matrices. Suppose that Γ_{12}, in particular, has at least one real eigenvalue λ (and thus a corresponding real eigenvector $\boldsymbol{\varepsilon}$):

$$\Gamma_{12}\boldsymbol{\varepsilon} = \lambda\boldsymbol{\varepsilon}. \tag{10}$$

Take the illuminant that provides the single view for the problem (4 4 3 1 3), represented by 4×1 vector $\mathbf{a}$ of illuminant descriptors, to be this eigenvector $\boldsymbol{\varepsilon}$:

$$\mathbf{a} = \boldsymbol{\varepsilon}. \tag{11}$$

If $e_{13} = e_{23} = 0$, then one finds the following solutions for the variables $e_{ij}, i, j = 1, 2, 3$, which include nonscaling solutions:

$$\begin{aligned} e_{33} &= e_{11} + \lambda e_{12}, \\ e_{22} &= e_{11} + \lambda e_{12} - e_{21}/\lambda, \\ e_{31} &= -\lambda e_{32}. \end{aligned} \tag{12}$$

The above argument holds for any eigenvector of Γ_{12}, as well as for any eigenvector of the matrices Γ_{13} and Γ_{23}. If each of these three gamma matrices possesses four distinct real eigenvalues, then there are potentially 12 rays in the space of illuminant descriptors along each of which recovery fails.

These subspaces of illuminants on which recovery fails do not necessarily persist if the dimension of the illumination model is dropped to three, as in the problem (4 3 3 1 3).

In conclusion, we can describe illuminants that cause recovery procedures based on bilinear models with parameters (4 4 3 1 3) to fail. We do not know, however, whether it is possible to choose a bilinear model for which all of these illuminants are physically unrealizable (compare Subsection 3.C.1). The square for the problem (4 4 3 1 3) in Fig. 6A indicates successful simulation by us, in addition to that by Wandell and Maloney.[17] Whether there are models that afford perfect recovery for the problems (4 4 3 1 3) and (4 3 3 1 3) remains an open question.

4. DISCUSSION

Our goal has been to determine ways that a visual system can recover spectral descriptions of surface reflectance functions and illuminant spectral power distributions from reflected lights. Our analysis of color constancy has dwelt on the use, by two-stage linear recovery procedures, of one or more views of a set of surfaces, in which each view is provided by a distinct illuminant.

A prime impetus for this work was the problem of proving the perfect recovery result for the problem (3 3 3 2 3), discussed by D'Zmura,[3] in which a trichromatic system recovers, from two views, three descriptors for each surface and illuminant. While simulations of recovery had eliminated the possibility that this and several other problems were total failures, the question of whether there were situations under which recovery would break down went unanswered. The method of proof presented here and in the companion paper[1] generalized readily to further problems and led to the present attempt to determine complete results for dichromatic, trichromatic, and tetrachromatic visual systems.

Our primary result is that a visual system's ability to use bilinear models to recover spectral descriptions of lights and surfaces from reflected lights is not an isolated happenstance: successful two-stage linear recovery is the rule rather than the exception.

We have shown that the dimension of the recovered descriptions increases significantly with the number of views. Using a three-dimensional model for illumination, a trichromatic visual system can recover from one view two reflectance descriptors, from two views up to five reflectance descriptors, and from three views up to eight

reflectance descriptors. Furthermore, we have proved that, for trichromatic visual systems, recovery is perfect in cases in which two descriptors are recovered from a single view, in cases in which up to three descriptors are recovered from two views, and in cases in which as many as seven descriptors are recovered from three views.

The number of recovered descriptors also increases with increasing number of distinct photoreceptoral types in a way that is summarized in Fig. 6 of the companion paper.[1] Using two views, in particular, a dichromatic system can recover three descriptors, a trichromatic system can recover five descriptors, and a tetrachromatic system can recover seven descriptors per surface reflectance.

We have proved and refined an earlier claim, made by Maloney and Wandell[2] and Maloney[18] that a trichromatic visual system can recover two descriptors per surface reflectance from a single view. The proof (Subsection 3.C) depends on knowledge of the necessary and sufficient conditions for recovery to be unique: a bilinear model must provide a one-to-one relationship between sets of lit surfaces and quantum catch data.[1] Investigation of these conditions for the problem (3 3 2 1 2) shows that there are planes of failure in the space of illuminant descriptors. These planes may or may not pass through the subset of physically realizable illuminants, and perfect recovery depends acutely on the particular choice of bilinear model. For the problem (3 2 2 1 2), successful model checks for bilinear models with these parameters first prove the claim for the case of two-dimensional models of illumination.

We are unable to prove or disprove claims made by Maloney and Wandell,[2] Wandell and Maloney,[17] and Maloney[18] regarding the use of information from a single view by tetrachromatic visual systems. While we can produce illuminant descriptors, for any bilinear model, that will cause recovery to fail, we have neither (1) described completely the subspaces of failure within the space of illuminant descriptors nor (2) determined whether it is possible to construct a bilinear model with the property that these subspaces of failure do not pass through the subset of physically realizable illuminants. The problems (4 4 3 1 3) and (4 3 3 1 3) remain open.

Further open problems for trichromacy and for tetrachromacy include those marked by squares in Figs. 3 and 6, respectively. We have not determined whether such problems admit perfect recovery procedures or not; successful simulation of recovery indicates merely that a recovery procedure is not a total failure.

A. Alternative Methods of Proof

The model check algorithm provides a tractable test of nonlinear uniqueness problems, met when $m > v$, by linearizing a homogeneous system of polynomial equations. One then examines this linearized system numerically, for particular bilinear models, to provide proof of perfect recovery.

We have found no alternative method that is both feasible and general. As mentioned in the companion paper,[1] there is a generalized elimination procedure involving resultants that handles problems like ours in a general way, without performing linearization. Yet this method is utterly intractable.

A second method of proof is taken up by Iverson and D'Zmura.[19] This method examines the eigenstructure of gamma matrices, thus remaining accessible to intuition, and provides results for several trichromatic problems, including that with parameters (3 3 3 2 3). Generality remains elusive with this method, however.

B. Human Color Constancy?

The results lead one to consider their relevance to human visual processing. An immediate issue is raised by a large set of psychophysical experiments whose results suggest that the stability of surface color appearance under changing illumination is largely an illusion and that, studied properly, human color constancy is an inexact and overrated faculty. Into such a category fall the studies by McCann *et al.*,[20] by Arend and Reeves,[21] by Arend *et al.*,[22] and by Brainard and Wandell.[23] In the minority camp falls the view of the late Edwin Land, based largely on work with demonstrations.[24–26]

C. Surface Correspondence

The analysis suggests that the studies on which the majority view is based are incomplete. The changing color signals that allow surface color properties to be represented stably depend on a change in illumination. Now it is possible to devise situations in which such a change is performed in a way so that surface correspondence across two or more views cannot possibly be maintained by the visual system: this is a feature of the studies mentioned above. Yet it is far more natural for several surfaces to be seen under both the old and the new illuminants. The obvious retention of surface identity across visible changes in illumination is a feature of Land's demonstrations of color constancy with Mondrians. The analysis suggests that a critical requirement for color constancy is correspondence: that several surfaces remain visible and retain their identity when the illuminant is changed across space or time.

D. Low-Level Chromatic Adaptation and Eye Movements?

A leading explanation of the visual system's accommodation to unknown illumination conditions relies on the chromatic adaptation of low-level mechanisms.[27] In the simplest version of this hypothesis, the work is accomplished through von Kries adaptation, in which each photoreceptor's gain is determined (reciprocally) by its quantum catch input.[28] Eye movements let each photoreceptor adapt to the space-averaged light, which, if the gray-world assumption holds, matches the chromaticity of the illuminant.[29] Brainard and Wandell[30] have shown that two versions of Land's retinex scheme are asymptotically equivalent to this type of adaptation, which was considered early on[27] and is still considered[23,28,31–33] to be a primary mechanism of color constancy.

Cone-level adaptation, however, provides an inexact rendition of surface chromatic properties.[32] Departures from the gray-world assumption in particular scenes lead simply to further instability in surface color appearance. The time-course of adaptation of postreceptoral chromatic mechanisms appears too lengthy to account for the stability of surface color appearance under rapidly changing illumination conditions, although it is certain that such processes significantly influence color appearance under persistent changes in illumination.[23,33,34]

E. Detecting and Tracking Chromatic Change

It was recognized early on that low-level mechanisms likely are supplemented by high-level processes (e.g., Ref. 27). Determining the chromatic properties of surfaces and illuminants by using chromatic change is an intrinsically high-level process. Detecting and tracking chromatic change across space and time requires mechanisms that are inseparable in their chromatic and their spatiotemporal properties, in just the way that motion-sensitive mechanisms bear receptive fields that are inseparable in space and time.[35,36] Furthermore, the chromatic change of the light from a single surface, caused by changing its illumination, only loosely constrains possible physical interpretations. Mechanisms that are sensitive to local chromatic change may feed a further set of integrative mechanisms that track these motions in terms that parallel the physical properties of surfaces and their illumination.

Double-opponent mechanisms sensitive to changes in color across space are basic elements in accounting for simultaneous color contrast.[37] Electrophysiological studies of cortical mechanisms of color vision in macaque suggest that early mechanisms with separable receptive fields[38,39] feed an intermediate stage characterized by inseparable receptive fields[40] of the sort required for detection of local spatiochromatic motion. These are thought to feed further stages comprising mechanisms with spatially extended receptive fields that are not yet fully understood.[41–46] Mechanisms that track the time-varying chromatic change of patterns, suggested by the success of Land's demonstrations of color constancy, are less well studied.[24,25] The likelihood that slow, persistent changes in illumination drive passive mechanisms of adaptation better than mechanisms sensitive to transient chromatic change suggests that one work toward understanding color vision in a way that integrates lower and higher levels.

ACKNOWLEDGMENTS

We thank David Brainard, Carole Cicerone, Don Hoffman, Tarow Indow, Ram Kakarala, Brian Wandell, and Jack Yellott for helpful discussion. We also thank D. Frederick for his advice on numerical algorithms. This work was supported, in part, by National Eye Institute grant EY10014 to M. D'Zmura and by National Science Foundation grant DIR-9014278 to the Institute for Mathematical Behavioral Sciences, University of California, Irvine, R. D. Luce, director. We thank the University of California, Irvine, for an allocation of computer time and the San Diego Supercomputing Center for the use of their facilities.

REFERENCES

1. M. D'Zmura and G. Iverson, "Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces," J. Opt. Soc. Am. A **10,** 2148–2165, (1993).
2. L. T. Maloney and B. A. Wandell, "Color constancy: a method for recovering surface spectral reflectance," J. Opt. Soc. Am. A **3,** 29–33 (1986).
3. M. D'Zmura, "Color constancy: surface color from changing illumination," J. Opt. Soc. Am. A **9,** 490–493 (1992).
4. M. D'Zmura and G. Iverson, "Color structure from chromatic motion," in *Annual Meeting*, Vol. 23 of 1992 OSA Technical Digest Series (Optical Society of America, Washington, D.C., 1992), p. 51.
5. E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, A. DuCroz, S. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen, LAPACK *User's Guide* (Society for Industrial and Applied Mathematics, Philadelphia, Pa., 1992).
6. G. Strang, *Linear Algebra and Its Applications*, 3rd ed. (Harcourt Brace Jovanovitch, San Diego, Calif., 1988).
7. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C. The Art of Scientific Computing* (Cambridge U. Press, New York, 1988).
8. V. C. Smith and J. Pokorny, "Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm," Vision Res. **15,** 161–171 (1975).
9. D. B. Judd, D. L. MacAdam, and G. Wyszecki, "Spectral distribution of typical daylight as a function of correlated color temperature," J. Opt. Soc. Am. **54,** 1031–1040 (1964).
10. G. Wyszecki and W. S. Stiles, *Color Science. Concepts and Methods, Quantitative Data and Formulas*, 2nd ed. (Wiley, New York, 1982).
11. J. Cohen, "Dependency of the spectral reflectance curves of the Munsell color chips," Psychonom. Sci. **1,** 369–370 (1964).
12. E. R. Dixon, "Spectral distribution of Australian daylight," J. Opt. Soc. Am. **68,** 437–450 (1978).
13. J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, "Characteristic spectra of Munsell colors," J. Opt. Soc. Am. A **6,** 318–322 (1989).
14. L. M. Hurvich and D. Jameson, "Some quantitative aspects of an opponent-color theory. II. Brightness, saturation, and hue in normal and dichromatic vision," J. Opt. Soc. Am. **45,** 602–616 (1955).
15. M. Tsukada and Y. Ohta, "An approach to color constancy using multiple images," Proc. Third Int. Conf. Comput. Vis. **3,** 385–393 (1990).
16. M. H. Brill and T. Benzschawel, "Remarks on signal-processing explanations of the trichromacy of vision," J. Opt. Soc. Am. A **2,** 1794–1796 (1985).
17. B. A. Wandell and L. T. Maloney, "Color imaging process," U.S. patent 4,648,051 (March 3, 1987).
18. L. T. Maloney, "Computational approaches to color constancy," Stanford Applied Psychology Laboratory Tech. Rep. 1985-01 (Stanford University, Stanford, Calif., 1985).
19. G. Iverson and M. D'Zmura, "Criteria for color constancy in trichromatic bilinear models," UC Irvine Institute for Mathematical Behavioral Sciences Tech. Rep. 93-18 (University of California, Irvine, Irvine, Calif., 1993).
20. J. J. McCann, S. P. McKee, and T. H. Taylor, "Quantitative studies in retinex theory. A comparison between theoretical predictions and observer responses to the 'color Mondrian' experiments," Vision Res. **16,** 445–458 (1976).
21. L. Arend and A. Reeves, "Simultaneous color constancy," J. Opt. Soc. Am. A **3,** 1743–1751 (1986).
22. L. E. Arend, Jr., A. Reeves, J. Schirillo, and R. Goldstein, "Simultaneous color constancy: papers with diverse Munsell values," J. Opt. Soc. Am. A **8,** 661–672 (1991).
23. D. Brainard and B. A. Wandell, "A bilinear model of the illuminant's effect on color appearance," in *Computational Models of Visual Processing*, M. Landy and J. A. Movshon, eds. (MIT Press, Cambridge, Mass., 1991), pp. 171–186.
24. E. H. Land, "Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image," Proc. Natl. Acad. Sci. USA **80,** 5163–5169 (1983).
25. E. H. Land, "Recent advances in retinex theory," Vision Res. **26,** 7–21 (1986).
26. E. H. Land and N. W. Daw, "Colors seen in a flash of light," Proc. Natl. Acad. Sci. USA **48,** 1000–1008 (1962).
27. D. B. Judd, "Hue, saturation and lightness of surface colors with chromatic illumination," J. Opt. Soc. Am. **30,** 2–32 (1940).
28. J. A. Worthey and M. H. Brill, "Heuristic analysis of von Kries color constancy," J. Opt. Soc. Am. A **3,** 1708–1712 (1986).
29. G. Buchsbaum, "A spatial processor model for object colour perception," J. Franklin Inst. **310,** 1–26 (1980).

30. D. Brainard and B. A. Wandell, "An analysis of the retinex theory of color vision," J. Opt. Soc. Am. A **3,** 1651–1661 (1986).
31. D. I. A. MacLeod, "Receptoral constraints on colour appearance," in *Central and Peripheral Mechanisms of Colour Vision*, D. Ottoson and S. Zeki, eds. (Macmillan, London, 1985), pp. 103–116.
32. M. D'Zmura and P. Lennie, "Mechanisms of color constancy," J. Opt. Soc. Am. A **3,** 1662–1672 (1986).
33. M. Fairchild and P. Lennie, "Chromatic adaptation to natural and incandescent illuminants," Vision Res. **32,** 2077–2085 (1992).
34. S. Ahn and D. I. A. MacLeod, "Adaptation in the chromatic and luminance channels," Invest. Ophthalmol. Vis. Sci. Suppl. **31,** 109 (1990).
35. E. H. Adelson and J. Bergen, "Spatiotemporal energy models for the perception of motion," J. Opt. Soc. Am. A **2,** 284–299 (1985).
36. E. H. Adelson and J. Bergen, "The plenoptic function and the elements of early vision," in *Computational Models of Visual Processing*, M. Landy and J. A. Movshon, eds. (MIT Press, Cambridge, Mass., 1991), pp. 3–20.
37. C. R. Michael, "Color vision mechanisms in monkey striate cortex: dual-opponent cells with concentric receptive fields," J. Neurophysiol. **41,** 572–588 (1978).
38. P. Lennie, J. Krauskopf, and G. Sclar, "Chromatic mechanisms in striate cortex of macaque," J. Neurosci. **10,** 649–669 (1990).
39. D. Y. Ts'o and C. D. Gilbert, "The organization of chromatic and spatial interactions in the primate striate cortex," J. Neurosci. **8,** 1712–1727 (1988).
40. M. S. Livingstone and D. H. Hubel, "Anatomy and physiology of a color system in primate primary visual cortex," J. Neurosci. **4,** 309–356 (1984).
41. S. M. Zeki, "Colour coding in the cerebral cortex: the reaction of cells in monkey visual cortex to wavelengths and colours," Neuroscience **9,** 741–765 (1983).
42. S. M. Zeki, "Colour coding in the cerebral cortex: the response of wavelength-selective and colour-coded cells in monkey visual cortex to changes in wavelength composition," Neuroscience **9,** 767–781 (1983).
43. H. M. Wild, S. R. Butler, D. Carden, and J. J. Kulikowski, "Primate cortical area V4 important for colour constancy but not wavelength discrimination," Nature **313,** 133–135 (1985).
44. C. A. Heywood and A. Cowey, "On the role of cortical area V4 in the discrimination of hue and pattern in macaque monkeys," J. Neurosci. **7,** 2601–2617 (1987).
45. P. E. Haenny and P. H. Schiller, "State dependent activity in monkey visual cortex—I. Single cell activity in V1 and V4 on visual tasks," Exp. Brain Res. **69,** 225–244 (1988).
46. S. J. Schein and R. Desimone, "Spectral properties of V4 neurons in macaque," J. Neurosci. **10,** 3369–3389 (1990).